

---

# Building and developing GHC

The GHC Team

## Abstract

This Guide is primarily aimed at those who want to build and/or hack on GHC. It describes how to get started with building GHC on your machine, and how to tweak the settings to get the kind of build you want. It also describes the inner workings of the build system, so you can extend it, modify it, and use it to build your code.

The bulk of this guide applies to building on Unix systems; see Section 12, “Instructions for building under Windows” for Windows notes.

## Table of Contents

1. Getting the sources .....	2
2. Things to check before you start .....	3
3. Installing pre-supposed utilities .....	3
3.1. Tools for building parallel GHC (GPH) .....	5
4. Building from source .....	5
5. Quick start for GHC developers .....	5
6. Working with the build system .....	6
6.1. History .....	7
6.2. Build trees .....	7
6.3. Getting the build you want .....	8
6.4. The story so far .....	10
6.5. Making things .....	11
6.6. Bootstrapping GHC .....	12
6.7. Standard Targets .....	12
6.8. Using GHC from the build tree .....	14
6.9. Fast Making .....	14
7. The Makefile architecture .....	14
7.1. Debugging .....	14
7.2. A small example .....	15
7.3. Boilerplate architecture .....	16
7.4. The mk/boilerplate.mk file .....	17
7.5. Platform settings .....	18
7.6. Pattern rules and options .....	19
7.7. The main mk/target.mk file .....	20
7.8. Recursion .....	21
7.9. Way management .....	21
7.10. When the canned rule isn't right .....	22
8. Building the documentation .....	22
8.1. Tools for building the Documentation .....	22
8.2. Installing the DocBook tools .....	23
8.2.1. Installing the DocBook tools on Linux .....	23
8.2.2. Installing DocBook on FreeBSD .....	23
8.2.3. Installing from binaries on Windows .....	23
8.3. Configuring the DocBook tools .....	23
8.4. Building the documentation .....	23
8.5. Installing the documentation .....	24

9. Porting GHC .....	24
9.1. Booting/porting from C (.hc) files .....	24
9.2. Porting GHC to a new platform .....	25
9.2.1. Cross-compiling to produce an unregistered GHC .....	26
9.2.2. Porting the RTS .....	28
9.2.3. The mangler .....	29
9.2.4. The splitter .....	29
9.2.5. The native code generator .....	29
9.2.6. GHCi .....	29
10. Known pitfalls in building Glasgow Haskell .....	30
11. Platforms, scripts, and file names .....	31
11.1. Windows platforms: Cygwin, MSYS, and MinGW .....	31
11.1.1. MinGW .....	31
11.1.2. Cygwin and MSYS .....	31
11.1.3. Targeting MinGW .....	32
11.1.4. File names .....	32
11.1.5. Crippled <b>ld</b> .....	33
11.1.6. Host System vs Target System .....	33
11.2. Wrapper scripts .....	33
12. Instructions for building under Windows .....	34
12.1. Installing and configuring MSYS .....	34
12.2. Installing and configuring Cygwin .....	35
12.3. Configuring SSH .....	36
12.4. Other things you need to install .....	37
12.5. Building GHC .....	38
12.6. A Windows build log using Cygwin .....	39
Index .....	40

# 1. Getting the sources

You can get your hands on the GHC sources in two ways:

## Source distributions

You have a supported platform, but (a) you like the warm fuzzy feeling of compiling things yourself; (b) you want to build something “extra”—e.g., a set of libraries with strictness-analysis turned off; or (c) you want to hack on GHC yourself.

A source distribution contains complete sources for GHC. Not only that, but the more awkward machine-independent steps are done for you. For example, if you don't have **happy** you'll find it convenient that the source distribution contains the result of running **happy** on the parser specifications. If you don't want to alter the parser then this saves you having to find and install **happy**. You will still need a working version of GHC (version 5.x or later) on your machine in order to compile (most of) the sources, however.

## The darcs repository.

We make releases infrequently. If you want more up-to-the minute (but less tested) source code then you need to get access to our darcs repository.

Information on accessing the darcs repository is on the wiki: <http://hackage.haskell.org/trac/ghc/wiki/GhcDarcs>.

The repository holds source code only. It holds no mechanically generated files at all. So if you check out a source tree from darcs you will need to install every utility so that you can build all the derived files

from scratch.

## 2. Things to check before you start

Here's a list of things to check before you get started.

1. Disk space needed: from about 100Mb for a basic GHC build, up to probably 500Mb for a GHC build with everything included (libraries built several different ways, etc.).
2. Use an appropriate machine / operating system. GHC Platform Support [<http://hackage.haskell.org/trac/ghc/wiki/Platforms>] lists the currently supported platforms; if yours isn't amongst these then you can try porting GHC (see Section 9, "Porting GHC").
3. Be sure that the "pre-supposed" utilities are installed. Section 3, "Installing pre-supposed utilities" elaborates.
4. If you have any problem when building or installing the Glasgow tools, please check the "known pitfalls" (Section 10, "Known pitfalls in building Glasgow Haskell"). Also check the FAQ for the version you're building, which is part of the User's Guide and available on the GHC web site [<http://www.haskell.org/ghc/>].

If you feel there is still some shortcoming in our procedure or instructions, please report it.

For GHC, please see the bug-reporting section of the GHC Users' Guide [<http://www.haskell.org/ghc/docs/latest/set/bug-reporting.html>], to maximise the usefulness of your report.

If in doubt, please send a message to [<glasgow-haskell-bugs@haskell.org>](mailto:glasgow-haskell-bugs@haskell.org).

## 3. Installing pre-supposed utilities

Here are the gory details about some utility programs you may need; **perl**, **gcc** and **happy** are the only important ones. (PVM is important if you're going for Parallel Haskell.) The **configure** script will tell you if you are missing something.

GHC

GHC is required to build GHC, because GHC itself is written in Haskell, and uses GHC extensions. It is possible to build GHC using just a C compiler, and indeed some distributions of GHC do just that, but it isn't the best supported method, and you may encounter difficulties. Full instructions are in Section 9, "Porting GHC".

GHC can be built using either an earlier released version of GHC (currently 5.04 and later are supported), or bootstrapped using a GHC built from exactly the same sources. Note that this means you cannot in general build GHC using an arbitrary development snapshot, or a build from say last week. It might work, it might not - we don't guarantee anything. To be on the safe side, start your build using the most recently released stable version of GHC.

Perl

*You have to have Perl to proceed!* Perl version 5 at least is required. GHC has been known to tickle bugs in Perl, so if you find that Perl crashes when running GHC try updating (or downgrading) your Perl in-

stallation. Versions of Perl before 5.6 have been known to have various bugs tickled by GHC, so the configure script will look for version 5.6 or later.

For Win32 platforms, you should use the binary supplied in the Install-Shield (copy it to `/bin`). The Cygwin-supplied Perl seems not to work.

Perl should be put somewhere so that it can be invoked by the `#!` script-invoking mechanism. The full pathname may need to be less than 32 characters long on some systems.

GNU C (**gcc**)

Most GCC versions should work with the most recent GHC sources. Expect trouble if you use a recent GCC with an older GHC, though (trouble in the form of mis-compiled code, link errors, and errors from the `ghc-asm` script).

If your GCC dies with “internal error” on some GHC source file, please let us know, so we can report it and get things improved. (Exception: on x86 boxes—you may need to fiddle with GHC's `-monly-N-regs` option; see the User's Guide)

GNU Make

The GHC build system makes heavy use of features specific to GNU **make**, so you must have this installed in order to build GHC.

NB. it has been reported that version 3.79 no longer works to build GHC, and 3.80 is required.

Happy  
[<http://www.haskell.org/happy>]

Happy is a parser generator tool for Haskell, and is used to generate GHC's parsers.

If you start from a source tarball of GHC (i.e. not a darcs checkout), then you don't need Happy, because we supply the pre-processed versions of the Happy parsers. If you intend to modify the compiler and/or you're using a darcs checkout, then you need Happy.

Happy version 1.15 is currently required to build GHC. Grab a copy from Happy's Web Page [<http://www.haskell.org/happy/>].

Alex

Alex is a lexical-analyser generator for Haskell, which GHC uses to generate its lexer.

Like Happy, you don't need Alex if you're building GHC from a source tarball, but you do need it if you're modifying GHC and/or building a darcs checkout.

Alex is written in Haskell and is a project in the darcs repository. Alex distributions are available from Alex's Web Page [<http://www.haskell.org/alex/>].

autoconf

GNU autoconf is needed if you intend to build from the darcs sources, it is *not* needed if you just intend to build a standard source distribution.

Version 2.52 or later of the autoconf package is required. NB. version 2.13 will no longer work, as of GHC version 6.1.

**autoreconf** (from the autoconf package) recursively builds **configure** scripts from the corresponding `configure.ac` and `aclocal.m4` files. If you modify one of the latter files, you'll need **autoreconf** to re-build the corresponding `configure`.

**sed** You need a working **sed** if you are going to build from sources. The build-configuration stuff needs it. GNU sed version 2.0.4 is no good! It has a bug in it that is tickled by the build-configuration. 2.0.5 is OK. Others are probably OK too (assuming we don't create too elaborate configuration scripts.)

## 3.1. Tools for building parallel GHC (GPH)

**PVM version 3:** PVM is the Parallel Virtual Machine on which Parallel Haskell programs run. (You only need this if you plan to run Parallel Haskell. Concurrent Haskell, which runs concurrent threads on a uniprocessor doesn't need it.) Underneath PVM, you can have (for example) a network of workstations (slow) or a multiprocessor box (faster).

The current version of PVM is 3.3.11; we use 3.3.7. It is readily available on the net; I think I got it from `research.att.com`, in `netlib`.

A PVM installation is slightly quirky, but easy to do. Just follow the Readme instructions.

**bash:** Sadly, the **gr2ps** script, used to convert “parallelism profiles” to PostScript, is written in Bash (GNU's Bourne Again shell). This bug will be fixed (someday).

More tools are required if you want to format the documentation that comes with GHC. See Section 8, “Building the documentation”.

## 4. Building from source

“I just want to build it!”

No problem. This recipe should build and install a working GHC with all the default settings. (unless you're on Windows, in which case go to Section 12, “Instructions for building under Windows”).

```
$ autoreconf1
$ ./configure
$ make
$ make install
```

For GHC, this will do a 2-stage bootstrap build of the compiler, with profiling libraries, and install the results in the default location (under `/usr/local` on Unix, for example).

The `configure` script is a standard GNU `autoconf` script, and accepts the usual options for changing install locations and the like. Run `./configure --help` for a list of options.

If you want to do anything at all non-standard, or you want to do some development, read on...

## 5. Quick start for GHC developers

---

<sup>1</sup>not necessary if you started from a source tarball

This section is a copy of the file HACKING from the GHC source tree. It describes how to get started with setting up your build tree for developing GHC or its libraries, and how to start building.

#### Getting started with hacking on GHC

-----

So you've decided to hack on GHC, congratulations! We hope you have a rewarding experience. This file will point you in the direction of information to help you get started right away.

#### The GHC Developer's Wiki

-----

The home for GHC Developers, with information on accessing the latest sources, the bug tracker, and documentation on the code:

<http://hackage.haskell.org/trac/ghc>

In particular, the wiki contains the following pages of interest to new hackers:

##### Quick Start for developers

<http://hackage.haskell.org/trac/ghc/wiki/Building/Hacking>

This section on the wiki will get you up and running with a serviceable build tree in no time:

This is part of the "Building GHC" section of the wiki, which has more detailed information on GHC's build system should you need it.

##### The GHC Commentary

<http://hackage.haskell.org/trac/wiki/Commentary>

Notes on the internals and architecture of GHC.

#### Mailing lists

-----

Ask on [glasgow-haskell-users@haskell.org](mailto:glasgow-haskell-users@haskell.org) if you have difficulties. If you're working with the current CVS sources of GHC, then [cvs-ghc@haskell.org](mailto:cvs-ghc@haskell.org) might be a more appropriate (developers hang out here). See <http://www.haskell.org/mailman/listinfo> for subscription.

Happy Hacking! --The GHC Team

## 6. Working with the build system

This rest of this guide is intended for duffers like me, who aren't really interested in Makefiles and systems configurations, but who need a mental model of the interlocking pieces so that they can make them

work, extend them consistently when adding new software, and lay hands on them gently when they don't work.

## 6.1. History

First, a historical note. The GHC build system used to be called "fptools": a generic build system used to build multiple projects (GHC, Happy, GreenCard, H/Dir, etc.). It had a concept of the generic project-independent parts, and project-specific parts that resided in a project subdirectory.

Nowadays, most of these other projects are using Cabal [<http://www.haskell.org/cabal/>], or have faded away, and GHC is the only regular user of the fptools build system. We decided therefore to simplify the situation for developers, and specialise the build system for GHC. This resulted in a simpler organisation of the source tree and the build system, which hopefully makes the whole thing easier to understand.

You might find old comments that refer to "projects" or "fptools" in the documentation and/or source; please let us know if you do.

## 6.2. Build trees

If you just want to build the software once on a single platform, then your source tree can also be your build tree, and you can skip the rest of this section.

We often want to build multiple versions of our software for different architectures, or with different options (e.g. profiling). It's very desirable to share a single copy of the source code among all these builds.

So for every source tree we have zero or more *build trees*. Each build tree is initially an exact copy of the source tree, except that each file is a symbolic link to the source file, rather than being a copy of the source file. There are "standard" Unix utilities that make such copies, so standard that they go by different names: **ln**, **mkshadowdir** are two (If you don't have either, the source distribution includes sources for the X11 **ln**—check out `utils/ln`). See Section 6.4, "The story so far" for a typical invocation.

The build tree does not need to be anywhere near the source tree in the file system. Indeed, one advantage of separating the build tree from the source is that the build tree can be placed in a non-backed-up partition, saving your systems support people from backing up untold megabytes of easily-regenerated, and rapidly-changing, gubbins. The golden rule is that (with a single exception—Section 6.3, "Getting the build you want") *absolutely everything in the build tree is either a symbolic link to the source tree, or else is mechanically generated*. It should be perfectly OK for your build tree to vanish overnight; an hour or two compiling and you're on the road again.

You need to be a bit careful, though, that any new files you create (if you do any development work) are in the source tree, not a build tree!

Remember, that the source files in the build tree are *symbolic links* to the files in the source tree. (The build tree soon accumulates lots of built files like `Foo.o`, as well.) You can *delete* a source file from the build tree without affecting the source tree (though it's an odd thing to do). On the other hand, if you *edit* a source file from the build tree, you'll edit the source-tree file directly. (You can set up Emacs so that if you edit a source file from the build tree, Emacs will silently create an edited copy of the source file in the build tree, leaving the source file unchanged; but the danger is that you think you've edited the source file whereas actually all you've done is edit the build-tree copy. More commonly you do want to edit the source file.)

Like the source tree, the top level of your build tree must be (a linked copy of) the root directory of the GHC source tree. Inside Makefiles, the root of your build tree is called `$(GHC_TOP)`. In the rest of this document path names are relative to `$(GHC_TOP)` unless otherwise stated. For example, the file `mk/target.mk` is actually `$(GHC_TOP)/mk/target.mk`.

## 6.3. Getting the build you want

When you build GHC you will be compiling code on a particular *host platform*, to run on a particular *target platform* (usually the same as the host platform). The difficulty is that there are minor differences between different platforms; minor, but enough that the code needs to be a bit different for each. There are some big differences too: for a different architecture we need to build GHC with a different native-code generator.

There are also knobs you can turn to control how the software is built. For example, you might want to build GHC optimised (so that it runs fast) or unoptimised (so that you can compile it fast after you've modified it. Or, you might want to compile it with debugging on (so that extra consistency-checking code gets included) or off. And so on.

All of this stuff is called the *configuration* of your build. You set the configuration using a three-step process.

Step 1: get ready for configuration. NOTE: if you're starting from a source distribution, rather than darcs sources, you can skip this step.

Change directory to `$(GHC_TOP)` and issue the command

```
$ autoreconf
```

(with no arguments). This GNU program (recursively) converts `$(GHC_TOP)/configure.ac` and `$(GHC_TOP)/aclocal.m4` to a shell script called `$(GHC_TOP)/configure`. If **autoreconf** bleats that it can't write the file `configure`, then delete the latter and try again. Note that you must use **autoreconf**, and not the old **autoconf**! If you erroneously use the latter, you'll get a message like "No rule to make target 'mk/config.h.in'".

Some parts of the source tree, particularly libraries, have their own configure script. **autoreconf** takes care of that, too, so all you have to do is calling **autoreconf** in the top-level directory `$(GHC_TOP)`.

These steps are completely platform-independent; they just mean that the human-written files (`configure.ac` and `aclocal.m4`) can be short, although the resulting files (the **configure** shell scripts and the C header template `mk/config.h.in`) are long.

Step 2: system configuration. Runs the newly-created **configure** script, thus:

```
$ ./configure [args]
```

**configure**'s mission is to scurry round your computer working out what architecture it has, what operating system, whether it has the `vfork` system call, where **tar** is kept, whether **gcc** is available, where various obscure `#include` files are, whether it's a leap year, and what the systems manager had for lunch. It communicates these snippets of information in two ways:

- It translates `mk/config.mk.in` to `mk/config.mk`, substituting for things between “@” brackets. So, “@HaveGcc@”

will be replaced by “YES” or “NO” depending on what **configure** finds. `mk/config.mk` is included by every Makefile (directly or indirectly), so the configuration information is thereby communicated to all Makefiles.

- It translates `mk/config.h.in` to `mk/config.h`. The latter is `#included` by various C programs, which can thereby make use of configuration information.

**configure** takes some optional arguments. Use `./configure --help` to get a list of the available arguments. Here are some of the ones you might need:

`--with-ghc=path` Specifies the path to an installed GHC which you would like to use. This compiler will be used for compiling GHC-specific code (eg. GHC itself). This option *cannot* be specified using `build.mk` (see later), because **configure** needs to auto-detect the version of GHC you're using. The default is to look for a compiler named `ghc` in your path.

`--with-hc=path` Specifies the path to any installed Haskell compiler. This compiler will be used for compiling generic Haskell code. The default is to use `ghc`. (NOTE: I'm not sure it actually works to specify a compiler other than GHC here; unless you really know what you're doing I suggest not using this option at all.)

`--with-gcc=path` Specifies the path to the installed GCC. This compiler will be used to compile all C files, *except* any generated by the installed Haskell compiler, which will have its own idea of which C compiler (if any) to use. The default is to use `gcc`.

Step 3: build configuration.

Next, you say how this build of GHC is to differ from the standard defaults by creating a new file `mk/build.mk` *in the build tree*. This file is the one and only file you edit in the build tree, precisely because it says how this build differs from the source. (Just in case your build tree does die, you might want to keep a private directory of `build.mk` files, and use a symbolic link in each build tree to point to the appropriate one.) So `mk/build.mk` never exists in the source tree—you create one in each build tree from the template. We'll discuss what to put in it shortly.

And that's it for configuration. Simple, eh?

What do you put in your build-specific configuration file `mk/build.mk`? *For almost all purposes all you will do is put make variable definitions that override those in `mk/config.mk.in`.* The whole point of `mk/config.mk.in`—and its derived counterpart `mk/config.mk`—is to define the build configuration. It is heavily commented, as you will see if you look at it. So generally, what you do is

look at `mk/config.mk.in`, and add definitions in `mk/build.mk` that override any of the `config.mk` definitions that you want to change. (The override occurs because the main boilerplate file, `mk/boilerplate.mk`, includes `build.mk` after `config.mk`.)

For your convenience, there's a file called `build.mk.sample` that can serve as a starting point for your `build.mk`.

For example, `config.mk.in` contains the definition:

```
GhcHcOpts=-Rghc-timing
```

The accompanying comment explains that this is the list of flags passed to GHC when building GHC itself. For doing development, it is wise to add `-DDEBUG`, to enable debugging code. So you would add the following to `build.mk`:

```
GhcHcOpts += -DDEBUG
```

GNU **make** allows existing definitions to have new text appended using the “+=” operator, which is quite a convenient feature.

Haskell compilations by default have `-O` turned on, by virtue of this setting from `config.mk`:

```
SRC_HC_OPTS += -H16m -O
```

`SRC_HC_OPTS` means "options for HC from the source tree", where HC stands for Haskell Compiler. `SRC_HC_OPTS` are added to every Haskell compilation. To turn off optimisation, you could add this to `build.mk`:

```
SRC_HC_OPTS = -H16m -O0
```

Or you could just add `-O0` to `GhcHcOpts` to turn off optimisation for the compiler. See Section 5, “Quick start for GHC developers” for some more suggestions.

When reading `config.mk.in`, remember that anything between “@...@” signs is going to be substituted by **configure** later. You *can* override the resulting definition if you want, but you need to be a bit surer what you are doing. For example, there's a line that says:

```
TAR = @TarCmd@
```

This defines the Make variables `TAR` to the pathname for a **tar** that **configure** finds somewhere. If you have your own pet **tar** you want to use instead, that's fine. Just add this line to `mk/build.mk`:

```
TAR = mytar
```

You do not *have* to have a `mk/build.mk` file at all; if you don't, you'll get all the default settings from `mk/config.mk.in`.

You can also use `build.mk` to override anything that **configure** got wrong. One place where this happens often is with the definition of `GHC_TOP_ABS`: this variable is supposed to be the canonical path to the top of your source tree, but if your system uses an automounter then the correct directory is hard to find automatically. If you find that **configure** has got it wrong, just put the correct definition in `build.mk`.

## 6.4. The story so far

Let's summarise the steps you need to carry to get yourself a fully-configured build tree from scratch.

1. Get your source tree from somewhere (darcs repository or source distribution). Say you call the root directory `myghc` (it does not have to be called `ghc`).
2. (Optional) Use **lndir** or **mkshadowdir** to create a build tree.

```
$ cd myghc
$ mkshadowdir . /scratch/joe-bloggs/myghc-x86
```

(N.B. **mkshadowdir**'s first argument is taken relative to its second.) You probably want to give the build tree a name that suggests its main defining characteristic (in your mind at least), in case you later add others.

3. Change directory to the build tree. Everything is going to happen there now.

```
$ cd /scratch/joe-bloggs/myghc-x86
```

4. Prepare for system configuration:

```
$ autoreconf
```

(You can skip this step if you are starting from a source distribution, and you already have `configure` and `mk/config.h.in`.)

5. Do system configuration:

```
$ ./configure
```

Don't forget to check whether you need to add any arguments to `configure`; for example, a common requirement is to specify which GHC to use with `--with-ghc=ghc`.

6. Create the file `mk/build.mk`, adding definitions for your desired configuration options.

You can make subsequent changes to `mk/build.mk` as often as you like. You do not have to run any further configuration programs to make these changes take effect. In theory you should, however, say **make clean; make**, because configuration option changes could affect anything—but in practice you are likely to know what's affected.

## 6.5. Making things

At this point you have made yourself a fully-configured build tree, so you are ready to start building real things.

The first thing you need to know is that *you must use GNU **make***. On some systems (eg. FreeBSD) this is called **gmake**, whereas on others it is the standard **make** command. In this document we will always refer to it as **make**; please substitute with **gmake** if your system requires it. If you use a the wrong **make** you will get all sorts of error messages (but no damage) because the GHC **Makefiles** use GNU **make**'s facilities extensively.

To just build the whole thing, **cd** to the top of your build tree and type **make**. This will prepare the tree and build the various parts in the correct order, resulting in a complete build of GHC that can even be used directly from the tree, without being installed first.

## 6.6. Bootstrapping GHC

GHC requires a 2-stage bootstrap in order to provide full functionality, including GHCi. By a 2-stage bootstrap, we mean that the compiler is built once using the installed GHC, and then again using the compiler built in the first stage. You can also build a stage 3 compiler, but this normally isn't necessary except to verify that the stage 2 compiler is working properly.

Note that when doing a bootstrap, the stage 1 compiler must be built, followed by the runtime system and libraries, and then the stage 2 compiler. The correct ordering is implemented by the top-level `Makefile`, so if you want everything to work automatically it's best to start `make` from the top of the tree. The top-level `Makefile` is set up to do a 2-stage bootstrap by default (when you say `make`). Some other targets it supports are:

<code>stage1</code>	Build everything as normal, including the stage 1 compiler.
<code>stage2</code>	Build the stage 2 compiler only.
<code>stage3</code>	Build the stage 3 compiler only.
<code>bootstrap</code> , <code>bootstrap2</code>	Build stage 1 followed by stage 2.
<code>bootstrap3</code>	Build stages 1, 2 and 3.
<code>install</code>	Install everything, including the compiler built in stage 2. To override the stage, say <code>make install stage=n</code> where <code>n</code> is the stage to install.
<code>binary-dist</code>	make a binary distribution. This is the target we use to build the binary distributions of GHC.
<code>dist</code>	make a source distribution. Note that this target does “make distclean” as part of its work; don't use it if you want to keep what you've built.

The top-level `Makefile` also arranges to do the appropriate `make boot` steps (see below) before actually building anything.

The `stage1`, `stage2` and `stage3` targets also work in the `compiler` directory, but don't forget that each stage requires its own `make boot` step: for example, you must do

```
$ make boot stage=2
```

before `make stage2` in `compiler`.

## 6.7. Standard Targets

In any directory you should be able to make the following:

<code>boot</code>	does the one-off preparation required to get ready for the real work. Notably, it does <b>make depend</b> in all directories that contain programs. It also builds the necessary tools for compilation to proceed.
-------------------	--

Invoking the `boot` target explicitly is not normally necessary. From the top-level directory, invoking `make` causes `make boot` to be invoked in various subdirectories first, in the right order. Unless you really know what you are doing, it is best to always say `make` from the top level first.

If you're working in a subdirectory somewhere and need to update the dependencies, `make boot` is a good way to do it.

<code>all</code>	makes all the final target(s) for this Makefile. Depending on which directory you are in a “final target” may be an executable program, a library archive, a shell script, or a Postscript file. Typing <b>make</b> alone is generally the same as typing <b>make all</b> .
<code>install</code>	installs the things built by <code>all</code> (except for the documentation). Where does it install them? That is specified by <code>mk/config.mk.in</code> ; you can override it in <code>mk/build.mk</code> , or by running <b>configure</b> with command-line arguments like <code>--bindir=/home/simonpj/bin</code> ; see <code>./configure --help</code> for the full details.
<code>install-docs</code>	installs the documentation. Otherwise behaves just like <code>install</code> .
<code>uninstall</code>	reverses the effect of <code>install</code> (WARNING: probably doesn't work).
<code>clean</code>	Delete all files from the current directory that are normally created by building the program. Don't delete the files that record the configuration, or files generated by <b>make boot</b> . Also preserve files that could be made by building, but normally aren't because the distribution comes with them.
<code>distclean</code>	Delete all files from the current directory that are created by configuring or building the program. If you have unpacked the source and built the program without creating any other files, <code>make distclean</code> should leave only the files that were in the distribution.
<code>mostlyclean</code>	Like <code>clean</code> , but may refrain from deleting a few files that people normally don't want to recompile.
<code>maintainer-clean</code>	Delete everything from the current directory that can be reconstructed with this Makefile. This typically includes everything deleted by <code>distclean</code> , plus more: C source files produced by Bison, tags tables, Info files, and so on.  One exception, however: <code>make maintainer-clean</code> should not delete <code>configure</code> even if <code>configure</code> can be remade using a rule in the Makefile. More generally, <code>make maintainer-clean</code> should not delete anything that needs to exist in order to run <code>configure</code> and then begin to build the program.  After a <code>maintainer-clean</code> , a <code>configure</code> will be necessary before building again.

All of these standard targets automatically recurse into sub-directories. Certain other standard targets do not:

`depend` make a `.depend` file in each directory that needs it. This `.depend` file contains mechanically-generated dependency information; for example, suppose a directory contains a Haskell source module `Foo.lhs` which imports another module `Baz`. Then the generated `.depend` file will contain the dependency:

```
Foo.o : Baz.hi
```

which says that the object file `Foo.o` depends on the interface file `Baz.hi` generated by

compiling module Baz. The `.depend` file is automatically included by every Makefile.

Some Makefiles have targets other than these. You can discover them by looking in the Makefile itself.

## 6.8. Using GHC from the build tree

If you want to build GHC and just use it direct from the build tree without doing `make install` first, you can run the in-place driver script. To run the stage 1 compiler, use `compiler/stage1/ghc-inplace`, stage 2 is `compiler/stage2/ghc-inplace`, and so on.

Do *NOT* use `compiler/stage1/ghc`, or `compiler/stage1/ghc-6.xx`, as these are the scripts intended for installation, and contain hard-wired paths to the installed libraries, rather than the libraries in the build tree.

## 6.9. Fast Making

Sometimes the dependencies get in the way: if you've made a small change to one file, and you're absolutely sure that it won't affect anything else, but you know that **make** is going to rebuild everything anyway, the following hack may be useful:

```
$ make FAST=YES
```

This tells the make system to ignore dependencies and just build what you tell it to. In other words, it's equivalent to temporarily removing the `.depend` file in the current directory (where **mkdependHS** and friends store their dependency information).

A bit of history: GHC used to come with a **fastmake** script that did the above job, but GNU make provides the features we need to do it without resorting to a script. Also, we've found that fastmaking is less useful since the advent of GHC's recompilation checker (see the User's Guide section on "Separate Compilation").

# 7. The Makefile architecture

**make** is great if everything works—you type **make install** and lo! the right things get compiled and installed in the right places. Our goal is to make this happen often, but somehow it often doesn't; instead some weird error message eventually emerges from the bowels of a directory you didn't know existed.

The purpose of this section is to give you a road-map to help you figure out what is going right and what is going wrong.

## 7.1. Debugging

Debugging Makefiles is something of a black art, but here's a couple of tricks that we find particularly useful. The following command allows you to see the contents of any make variable in the context of the current Makefile:

```
$ make show VALUE=HS_SRCS
```

where you can replace `HS_SRCS` with the name of any variable you wish to see the value of.

GNU make has a `-d` option which generates a dump of the decision procedure used to arrive at a conclusion about which files should be recompiled. Sometimes useful for tracking down problems with su-

perfluous or missing recompilations.

## 7.2. A small example

To get started, let us look at the Makefile for an imaginary small program, `small`. Each program or library in the GHC source tree typically has its own directory, in this case we'll use `$(GHC_TOP)/small`. Inside the `small/` directory there will be a Makefile, looking something like this:

```
# Makefile for program "small"
TOP = ..
include $(TOP)/mk/boilerplate.mk

HS_PROG = small

include $(TOP)/target.mk
```

this Makefile has three sections:

1. The first section includes <sup>2</sup> a file of “boilerplate” code from the top level ). As its name suggests, `boilerplate.mk` consists of a large quantity of standard Makefile code. We discuss this boilerplate in more detail in Section 7.4, “The `mk/boilerplate.mk` file”.

Before the `include` statement, you must define the **make** variable `TOP` to be the top-level directory of the source tree, containing the `mk` directory in which the `boilerplate.mk` file is. It is *not* OK to simply say

```
include ../mk/boilerplate.mk # NO NO NO
```

Why? Because the `boilerplate.mk` file needs to know where it is, so that it can, in turn, include other files. (Unfortunately, when an included file does an `include`, the filename is treated relative to the directory in which **make** is being run, not the directory in which the included sits.) In general, *every file `f.o.mk` assumes that `$(TOP)/mk/f.o.mk` refers to itself*. It is up to the Makefile doing the `include` to ensure this is the case.

2. The second section defines the standard **make** variable `HS_PROG` (the executable binary to be built). We will discuss in more detail what the “standard variables” are, and how they affect what happens, in Section 7.7, “The main `mk/target.mk` file”.
3. The last section includes a second file of standard code, called `target.mk`. It contains the rules that tell **make** how to make the standard targets (Section 6.7, “Standard Targets”). Why, you ask, can't this standard code be part of `boilerplate.mk`? Good question. We discuss the reason later, in Section 7.3, “Boilerplate architecture”.

You do not *have* to include the `target.mk` file. Instead, you can write rules of your own for all the standard targets. Usually, though, you will find quite a big payoff from using the canned rules in `target.mk`; the price tag is that you have to understand what canned rules get enabled, and what they do (Section 7.7, “The main `mk/target.mk` file”).

In our example Makefile, most of the work is done by the two included files. When you say **make all**, the following things happen:

- **make** looks in the current directory to see what source files it can find (eg. `Foo.hs`, `Baz.c`), and from that it figures out what object files need to be built (eg. `Foo.o`, `Baz.o`). Because source files

<sup>2</sup> One of the most important things to note here is that `include` is not a `#include` directive. It is not found and used automatically, so including them from a program or library has to be done manually.

(see `EXCLUDED_SRCS` in Section 7.4, “The `mk/boilerplate.mk` file”).

- It uses a boilerplate pattern rule to compile `Foo.hs` to `Foo.o` using a Haskell compiler. (Which one? That is set in the build configuration.)
- It uses another standard pattern rule to compile `Baz.c` to `Baz.o`, using a C compiler. (Ditto.)
- It links the resulting `.o` files together to make `small`, using the Haskell compiler to do the link step. (Why not use `ld`? Because the Haskell compiler knows what standard libraries to link in. How did `make` know to use the Haskell compiler to do the link, rather than the C compiler? Because we set the variable `HS_PROG` rather than `C_PROG`.)

All `Makefile`s should follow the above three-section format.

## 7.3. Boilerplate architecture

Every `Makefile` includes a `boilerplate.mk` file at the top, and `target.mk` file at the bottom. In this section we discuss what is in these files, and why there have to be two of them. In general:

- `boilerplate.mk` consists of:
  - *Definitions of millions of **make** variables* that collectively specify the build configuration. Examples: `HC_OPTS`, the options to feed to the Haskell compiler; `NoFibSubDirs`, the sub-directories to enable within the `nofib` project; `GhcWithHc`, the name of the Haskell compiler to use when compiling GHC in the `ghc` project.
  - *Standard pattern rules* that tell **make** how to construct one file from another.

`boilerplate.mk` needs to be included at the *top* of each `Makefile`, so that the user can replace the boilerplate definitions or pattern rules by simply giving a new definition or pattern rule in the `Makefile`. **make** simply takes the last definition as the definitive one.

Instead of *replacing* boilerplate definitions, it is also quite common to *augment* them. For example, a `Makefile` might say:

```
SRC_HC_OPTS += -O
```

thereby adding “-O” to the end of `SRC_HC_OPTS`.

- `target.mk` contains **make** rules for the standard targets described in Section 6.7, “Standard Targets”. These rules are selectively included, depending on the setting of certain **make** variables. These variables are usually set in the middle section of the `Makefile` between the two `includes`.

`target.mk` must be included at the end (rather than being part of `boilerplate.mk`) for several tiresome reasons:

- **make** commits target and dependency lists earlier than it should. For example, `target.mk` has a rule that looks like this:

```
$(HS_PROG) : $(OBJS)
            $(HC) $(LD_OPTS) $< -o $@
```

If this rule was in `boilerplate.mk` then `$(HS_PROG)` and `$(OBJS)` would not have their final values at the moment **make** encountered the rule. Alas, **make** takes a snapshot of their current values, and wires that snapshot into the rule. (In contrast, the commands executed when the

rule “fires” are only substituted at the moment of firing.) So, the rule must follow the definitions given in the `Makefile` itself.

- Unlike pattern rules, ordinary rules cannot be overridden or replaced by subsequent rules for the same target (at least, not without an error message). Including ordinary rules in `boilerplate.mk` would prevent the user from writing rules for specific targets in specific cases.
- There are a couple of other reasons I've forgotten, but it doesn't matter too much.

## 7.4. The `mk/boilerplate.mk` file

If you look at `$(GHC_TOP)/mk/boilerplate.mk` you will find that it consists of the following sections, each held in a separate file:

`config.mk` is the build configuration file we discussed at length in Section 6.3, “Getting the build you want”.

`paths.mk` defines **make** variables for pathnames and file lists. This file contains code for automatically compiling lists of source files and deriving lists of object files from those. The results can be overridden in the `Makefile`, but in most cases the automatic setup should do the right thing.

The following variables may be set in the `Makefile` to affect how the automatic source file search is done:

`ALL_DIRS` Set to a list of directories to search in addition to the current directory for source files.

`EXCLUDED_SRCS` Set to a list of source files (relative to the current directory) to omit from the automatic search. The source searching machinery is clever enough to know that if you exclude a source file from which other sources are derived, then the derived sources should also be excluded. For example, if you set `EXCLUDED_SRCS` to include `Foo.y`, then `Foo.hs` will also be excluded.

`EXTRA_SRCS` Set to a list of extra source files (perhaps in directories not listed in `ALL_DIRS`) that should be considered.

The results of the automatic source file search are placed in the following make variables:

`SRCS` All source files found, sorted and without duplicates, including those which might not exist yet but will be derived from other existing sources. `SRCS` *can* be overridden if necessary, in which case the variables below will follow suit.

`HS_SRCS` all Haskell source files in the current directory, including those derived from other source files (eg. Happy sources also give rise to Haskell sources).

`HS_OBJS` Object files derived from `HS_SRCS`.

`HS_IFACES` Interface files (`.hi` files) derived from `HS_SRCS`.

`C_SRCS` All C source files found.

C\_OBJS            Object files derived from C\_SRCS.

SCRIPT\_SRCS     All script source files found (.lprl files).

SCRIPT\_OBJS     “object” files derived from SCRIPT\_SRCS (.prl files).

HSC\_SRCS        All hsc2hs source files (.hsc files).

HAPPY\_SRCS     All happy source files (.y or .hy files).

OBJS            the concatenation of \$(HS\_OBJS), \$(C\_OBJS), and \$(SCRIPT\_OBJS).

Any or all of these definitions can easily be overridden by giving new definitions in your Makefile.

What, exactly, does `paths.mk` consider a “source file” to be? It's based on the file's suffix (e.g. `.hs`, `.lhs`, `.c`, `.hy`, etc), but this is the kind of detail that changes, so rather than enumerate the source suffices here the best thing to do is to look in `paths.mk`.

`opts.mk`        defines **make** variables for option strings to pass to each program. For example, it defines `HC_OPTS`, the option strings to pass to the Haskell compiler. See Section 7.6, “Pattern rules and options”.

`suffix.mk`     defines standard pattern rules—see Section 7.6, “Pattern rules and options”.

Any of the variables and pattern rules defined by the boilerplate file can easily be overridden in any particular Makefile, because the boilerplate `include` comes first. Definitions after this `include` directive simply override the default ones in `boilerplate.mk`.

## 7.5. Platform settings

There are three platforms of interest when building GHC:

The *build*        The platform on which we are doing this build.

The *host*        The platform on which these binaries will run.

The *target*      The platform for which this compiler will generate code.

These platforms are set when running the `configure` script, using the `--build`, `--host`, and `--target` options. The `mk/config.mk` file defines several symbols related to the platform settings (see `mk/config.mk` for details).

We don't currently support `build` & `host` being different, because the build process creates binaries that are both run during the build, and also installed.

If `host` and `target` are different, then we are building a cross-compiler. For GHC, this means a compiler which will generate intermediate `.hc` files to port to the target architecture for bootstrapping. The libraries and stage 2 compiler will be built as `HC` files for the target system (see Section 9, “Porting GHC” for details).

More details on when to use `BUILD`, `HOST` or `TARGET` can be found in the comments in `config.mk`.

## 7.6. Pattern rules and options

The file `suffix.mk` defines standard *pattern rules* that say how to build one kind of file from another, for example, how to build a `.o` file from a `.c` file. (GNU **make**'s *pattern rules* are more powerful and easier to use than Unix **make**'s *suffix rules*.)

Almost all the rules look something like this:

```
% .o : % .c
    $(RM) $@
    $(CC) $(CC_OPTS) -c $< -o $@
```

Here's how to understand the rule. It says that *something.o* (say `Foo.o`) can be built from *something.c* (`Foo.c`), by invoking the C compiler (path name held in `$(CC)`), passing to it the options `$(CC_OPTS)` and the rule's dependent file of the rule `$<` (`Foo.c` in this case), and putting the result in the rule's target `$@` (`Foo.o` in this case).

Every program is held in a **make** variable defined in `mk/config.mk`—look in `mk/config.mk` for the complete list. One important one is the Haskell compiler, which is called `$(HC)`.

Every program's options are held in a **make** variables called `<prog>_OPTS`. the `<prog>_OPTS` variables are defined in `mk/opts.mk`. Almost all of them are defined like this:

```
CC_OPTS = \
    $(SRC_CC_OPTS) $(WAY$_way)_CC_OPTS) $(*_CC_OPTS) $(EXTRA_CC_OPTS)
```

The four variables from which `CC_OPTS` is built have the following meaning:

`SRC_CC_OPTS`: options passed to all C compilations.

: options passed to C compilations for way `<way>`. For example, `WAY_mp_CC_OPTS` gives options to pass to the C compiler when compiling way `mp`. The variable `WAY_CC_OPTS` holds options to pass to the C compiler when compiling the standard way. (Section 7.9, “Way management” dicusses multi-way compilation.)

`<module>_CC_OPTS`: options to pass to the C compiler that are specific to module `<module>`. For example, `SMap_CC_OPTS` gives the specific options to pass to the C compiler when compiling `SMap.c`.

`EXTRA_CC_OPTS`: extra options to pass to all C compilations. This is intended for command line use, thus:

```
$ make libHS.a EXTRA_HC_OPTS="-v"
```

## 7.7. The main `mk/target.mk` file

`target.mk` contains canned rules for all the standard targets described in Section 6.7, “Standard Targets”. It is complicated by the fact that you don’t want all of these rules to be active in every Makefile. Rather than have a plethora of tiny files which you can include selectively, there is a single file, `target.mk`, which selectively includes rules based on whether you have defined certain variables in your Makefile. This section explains what rules you get, what variables control them, and what the rules do. Hopefully, you will also get enough of an idea of what is supposed to happen that you can read and understand any weird special cases yourself.

. If `HS_PROG` is defined, you get rules with the following targets:

`HS_PROG` itself. This rule links `$(OBJS)` with the Haskell runtime system to get an executable called `$(HS_PROG)`.

`install` installs `$(HS_PROG)` in `$(bindir)`.

`C_PROG` is similar to `HS_PROG`, except that the link step links `$(C_OBJS)` with the C runtime system.

`LIBRARY` is similar to `HS_PROG`, except that it links `$(LIB_OBJS)` to make the library archive `$(LIBRARY)`, and `install` installs it in `$(libdir)`.

Some rules are “double-colon” rules, thus

```
install :: $(HS_PROG)
...how to install it...
```

GNU **make** treats double-colon rules as separate entities. If there are several double-colon rules for the same target it takes each in turn and fires it if its dependencies say to do so. This means that you can, for example, define both `HS_PROG` and `LIBRARY`, which will generate two rules for `install`. When you type **make install** both rules will be fired, and both the program and the library will be installed, just as you wanted.

## 7.8. Recursion

In leaf Makefiles the variable `SUBDIRS` is undefined. In non-leaf Makefiles, `SUBDIRS` is set to the list of sub-directories that contain subordinate Makefiles. *It is up to you to set `SUBDIRS` in the Makefile.* There is no automation here—`SUBDIRS` is too important to automate.

When `SUBDIRS` is defined, `target.mk` includes a rather neat rule for the standard targets (Section 6.7, “Standard Targets” that simply invokes **make** recursively in each of the sub-directories.

*These recursive invocations are guaranteed to occur in the order in which the list of directories is specified in `SUBDIRS`.* This guarantee can be important. For example, when you say **make boot** it can be important that the recursive invocation of **make boot** is done in one sub-directory (the include files, say) before another (the source files). Generally, put the most independent sub-directory first, and the most dependent last.

## 7.9. Way management

We sometimes want to build essentially the same system in several different “ways”. For example, we want to build GHC's `Prelude` libraries with and without profiling, so that there is an appropriately-built library archive to link with when the user compiles his program. It would be possible to have a completely separate build tree for each such “way”, but it would be horribly bureaucratic, especially since often only parts of the build tree need to be constructed in multiple ways.

Instead, the `target.mk` contains some clever magic to allow you to build several versions of a system; and to control locally how many versions are built and how they differ. This section explains the magic.

The files for a particular way are distinguished by munging the suffix. The “normal way” is always built, and its files have the standard suffices `.o`, `.hi`, and so on. In addition, you can build one or more extra ways, each distinguished by a *way tag*. The object files and interface files for one of these extra ways are distinguished by their suffix. For example, way `mp` has files `.mp.o` and `.mp.hi`. Library archives have their way tag the other side of the dot, for boring reasons; thus, `libHS_mp.a`.

A **make** variable called `way` holds the current way tag. *way is only ever set on the command line of **make*** (usually in a recursive invocation of **make** by the system). It is never set inside a `Makefile`. So it is a global constant for any one invocation of **make**. Two other **make** variables, `way_` and `_way` are immediately derived from  $\$(way)$  and never altered. If `way` is not set, then neither are `way_` and `_way`, and the invocation of **make** will build the “normal way”. If `way` is set, then the other two variables are set in sympathy. For example, if  $\$(way)$  is “mp”, then `way_` is set to “mp\_” and `_way` is set to “\_mp”. These three variables are then used when constructing file names.

So how does **make** ever get recursively invoked with `way` set? There are two ways in which this happens:

- For some (but not all) of the standard targets, when in a leaf sub-directory, **make** is recursively invoked for each way tag in  $\$(WAYS)$ . You set `WAYS` in the `Makefile` to the list of way tags you want these targets built for. The mechanism here is very much like the recursive invocation of **make** in sub-directories (Section 7.8, “Recursion”). It is up to you to set `WAYS` in your `Makefile`; this is how you control what ways will get built.
- For a useful collection of targets (such as `libHS_mp.a`, `Foo.mp_o`) there is a rule which recursively invokes **make** to make the specified target, setting the `way` variable. So if you say **make `Foo.mp_o`** you should see a recursive invocation **make `Foo.mp_o way=mp`**, and *in this recursive invocation the pattern rule for compiling a Haskell file into a .o file will match*. The key pattern rules (in `suffix.mk`) look like this:

```
%.${way_}o : %.lhs
    $(HC) $(HC_OPTS) $< -o $@
```

Neat, eh?

- You can invoke **make** with a particular way setting yourself, in order to build files related to a particular way in the current directory. eg.

```
$ make way=p
```

will build files for the profiling way only in the current directory.

## 7.10. When the canned rule isn't right

Sometimes the canned rule just doesn't do the right thing. For example, in the `nofib` suite we want the link step to print out timing information. The thing to do here is *not* to define `HS_PROG` or `C_PROG`, and instead define a special purpose rule in your own `Makefile`. By using different variable names you will avoid the canned rules being included, and conflicting with yours.

# 8. Building the documentation

## 8.1. Tools for building the Documentation

The following additional tools are required if you want to format the documentation that comes with GHC:

DocBook	Much of our documentation is written in DocBook XML, instructions on installing and configuring the DocBook tools are below.
TeX	A decent TeX distribution is required if you want to produce printable documentation. We recommend <code>teTeX</code> , which includes just about

everything you need.

Haddock

Haddock is a Haskell documentation tool that we use for automatically generating documentation from the library source code. To build documentation for the libraries (`$(GHC_TOP)/libraries`) you should build and install Haddock. Haddock requires GHC to build.

## 8.2. Installing the DocBook tools

### 8.2.1. Installing the DocBook tools on Linux

If you're on a recent RedHat (7.0+) or SuSE (8.1+) system, you probably have working DocBook tools already installed. The configure script should detect your setup and you're away.

If you don't have DocBook tools installed, and you are using a system that can handle RPM packages, you can use Rpmfind.net [<http://rpmfind.net/>] to find suitable packages for your system. Search for the packages `docbook-dtd`, `docbook-xsl-stylesheets`, `libxslt`, `libxml2`, `fop`, `xmltex`, and `dvips`.

### 8.2.2. Installing DocBook on FreeBSD

On FreeBSD systems, the easiest way to get DocBook up and running is to install it from the ports tree or a pre-compiled package (packages are available from your local FreeBSD mirror site).

To use the ports tree, do this:

```
$ cd /usr/ports/textproc/docproj
$ make install
```

This installs the FreeBSD documentation project tools, which includes everything needed to format the GHC documentation.

### 8.2.3. Installing from binaries on Windows

Probably the fastest route to a working DocBook environment on Windows is to install Cygwin [<http://www.cygwin.com/>] with the complete `Doc` category. If you are using MinGW [<http://www.mingw.org/>] for compilation, you have to help **configure** a little bit: Set the environment variables `XmllintCmd` and `XsltprocCmd` to the paths of the Cygwin executables **xmllint** and **xsltproc**, respectively, and set `fp_cv_dir_docbook_xsl` to the path of the directory where the XSL stylesheets are installed, e.g. `c:/cygwin/usr/share/docbook-xsl`.

If you want to build HTML Help, you have to install the HTML Help SDK [<http://msdn.microsoft.com/library/default.asp?url=/library/en-us/htmlhelp/html/hworiHTMLHelpStartPage.asp>], too, and make sure that **hhc** is in your `PATH`.

## 8.3. Configuring the DocBook tools

Once the DocBook tools are installed, the configure script will detect them and set up the build system accordingly. If you have a system that isn't supported, let us know, and we'll try to help.

## 8.4. Building the documentation

To build documentation in a certain format, you can say, for example,

```
$ make html
```

to build HTML documentation below the current directory. The available formats are: `dvi`, `ps`, `pdf`, `html`, and `rtf`. Note that not all documentation can be built in all of these formats: HTML documentation is generally supported everywhere, and DocBook documentation might support the other formats (depending on what other tools you have installed).

All of these targets are recursive; that is, saying `make html` will make HTML docs for all the documents recursively below the current directory.

Because there are many different formats that the DocBook documentation can be generated in, you have to select which ones you want by setting the `XMLDocWays` variable to a list of them. For example, in `build.mk` you might have a line:

```
XMLDocWays = html ps
```

This will cause the documentation to be built in the requested formats as part of the main build (the default is not to build any documentation at all).

## 8.5. Installing the documentation

To install the documentation, use:

```
$ make install-docs
```

This will install the documentation into `$(datadir)` (which defaults to `$(prefix)/share`). The exception is HTML documentation, which goes into `$(datadir)/html`, to keep things tidy.

Note that unless you set `$(XMLDocWays)` to a list of formats, the `install-docs` target won't do anything for DocBook XML documentation.

## 9. Porting GHC

This section describes how to port GHC to a currently unsupported platform. To avoid confusion, when we say “architecture” we are referring to the processor, and we use the term “platform” to refer to the combination of architecture and operating system.

There are two distinct porting scenarios:

- Your platform is already supported, but you want to compile up GHC using just a C compiler. This is a straightforward bootstrap from HC files, and is described in Section 9.1, “Bootstrapping/porting from C (.hc) files”.
- Your platform isn't supported by GHC. You will need to do an *unregisterised bootstrap*, proceed to Section 9.2, “Porting GHC to a new platform”.

### 9.1. Bootstrapping/porting from C (.hc) files

Bootstrapping GHC on a system without GHC already installed is achieved by taking the intermediate C files (known as HC files) from another GHC compilation, compiling them using `gcc` to get a working GHC.

*NOTE: GHC versions 5.xx were hard to bootstrap from C. We recommend using GHC 6.0.1 or later.*

HC files are platform-dependent, so you have to get a set that were generated on *the same platform*.

There may be some supplied on the GHC download page, otherwise you'll have to compile some up yourself.

The following steps should result in a working GHC build with full libraries:

- Make a set of HC files. On an identical system with GHC already installed, get a GHC source tree and put the following in `mk/build.mk`:

```
SRC_HC_OPTS      = -H32m -O -fasm -Rghc-timing -keep-hc-files
GhcLibHCOpts    = -O
GhcLibWays      =
SplitObjs       = NO
```

Build GHC as normal, and then make `hc-file-bundle Project=ghc` to create the tar file containing the hc files.

- On the target system, unpack the HC files on top of a fresh source tree (make sure the source tree version matches the version of the HC files *exactly!*). This will place matching `.hc` files next to the corresponding Haskell source (`.hs` or `.lhs`) in the compiler subdirectory `ghc/compiler` and in the libraries (subdirectories of `libraries`).
- The actual build process is fully automated by the `hc-build` script located in the `distrib` directory. If you eventually want to install GHC into the directory `dir`, the following command will execute the whole build process (it won't install yet):

```
$ distrib/hc-build --prefix=dir
```

By default, the installation directory is `/usr/local`. If that is what you want, you may omit the argument to `hc-build`. Generally, any option given to `hc-build` is passed through to the configuration script `configure`. If `hc-build` successfully completes the build process, you can install the resulting system, as normal, with

```
$ make install
```

## 9.2. Porting GHC to a new platform

The first step in porting to a new platform is to get an *unregisterised* build working. An unregisterised build is one that compiles via vanilla C only. By contrast, a registerised build uses the following architecture-specific hacks for speed:

- Global register variables: certain abstract machine “registers” are mapped to real machine registers, depending on how many machine registers are available (see `ghc/includes/MachRegs.h`).
- Assembly-mangling: when compiling via C, we feed the assembly generated by `gcc` through a Perl script known as the *mangler* (see `ghc/driver/mangler/ghc-asm.lpr1`). The mangler rearranges the assembly to support tail-calls and various other optimisations.

In an unregisterised build, neither of these hacks are used — the idea is that the C code generated by the compiler should compile using `gcc` only. The lack of these optimisations costs about a factor of two in performance, but since unregisterised compilation is usually just a step on the way to a full registerised port, we don't mind too much.

You should go through this process even if your architecture is already has registered support in GHC, but your OS currently isn't supported. In this case you probably won't need to port any of the architecture-specific parts of the code, and you can proceed straight from the unregistered build to build a registered compiler.

Notes on GHC portability in general: we've tried to stick to writing portable code in most parts of the system, so it should compile on any POSIXish system with gcc, but in our experience most systems differ from the standards in one way or another. Deal with any problems as they arise - if you get stuck, ask the experts on <glasgow-haskell-users@haskell.org>.

Lots of useful information about the innards of GHC is available in the GHC Commentary [<http://www.cse.unsw.edu.au/~chak/haskell/ghc/comm/>], which might be helpful if you run into some code which needs tweaking for your system.

## 9.2.1. Cross-compiling to produce an unregistered GHC

NOTE! These instructions apply to GHC 6.4 and (hopefully) later. If you need instructions for an earlier version of GHC, try to get hold of the version of this document that was current at the time. It should be available from the appropriate download page on the GHC homepage [<http://www.haskell.org/ghc/>].

In this section, we explain how to bootstrap GHC on a new platform, using unregistered intermediate C files. We haven't put a great deal of effort into automating this process, for two reasons: it is done very rarely, and the process usually requires human intervention to cope with minor porting issues anyway.

The following step-by-step instructions should result in a fully working, albeit unregistered, GHC. Firstly, you need a machine that already has a working GHC (we'll call this the *host* machine), in order to cross-compile the intermediate C files that we will use to bootstrap the compiler on the *target* machine.

- On the target machine:
  - Unpack a source tree (preferably a released version). We will call the path to the root of this tree *T*.

```
•  
$ cd T  
$ ./configure --enable-hc-boot --enable-hc-boot-unregistered
```

You might need to update `configure.in` to recognise the new platform, and re-generate `configure` with `autoreconf`.

- ```
$ cd T/ghc/includes  
$ make
```
- On the host machine:
  - Unpack a source tree (same released version). Call this directory *H*.

```
•  
$ cd H  
$ ./configure
```

- Create `H/mk/build.mk`, with the following contents:

```
GhcUnregistered = YES
```

```
GhcLibHcOpts = -O -fvia-C -keep-hc-files
GhcRtsHcOpts = -keep-hc-files
GhcLibWays =
SplitObjs = NO
GhcWithNativeCodeGen = NO
GhcWithInterpreter = NO
GhcStage1HcOpts = -O
GhcStage2HcOpts = -O -fvia-C -keep-hc-files
SRC_HC_OPTS += -H32m
GhcBootLibs = YES
```

- Edit `H/mk/config.mk`:
  - change `TARGETPLATFORM` appropriately, and set the variables involving `TARGET` or `Target` to the correct values for the target platform. This step is necessary because currently `configure` doesn't cope with specifying different values for the `--host` and `--target` flags.
  - copy `LeadingUnderscore` setting from target.
- Copy `T/ghc/includes/ghcautoconf.h`, `T/ghc/includes/DerivedConstants.h`, and `T/ghc/includes/GHCConstants.h` to `H/ghc/includes`. Note that we are building on the host machine, using the target machine's configuration files. This is so that the intermediate C files generated here will be suitable for compiling on the target system.
- Touch the generated configuration files, just to make sure they don't get replaced during the build:

```
$ cd H/ghc/includes
$ touch ghcautoconf.h DerivedConstants.h GHCConstants.h mkDerivedConstants.c
$ touch mkDerivedConstantsHdr mkDerivedConstants.o mkGHCConstants mkGHCConsta
```

Note: it has been reported that these files still get overwritten during the next stage. We have installed a fix for this in GHC 6.4.2, but if you are building a version before that you need to watch out for these files getting overwritten by the Makefile in `ghc/includes`. If your system supports it, you might be able to prevent it by making them immutable:

```
$ chflags uchg ghc/includes/{ghcautoconf.h,DerivedConstants.h,GHCConstants.h}
```

- Now build the compiler:

```
$ cd H/glafp-utils && make boot && make
$ cd H/ghc && make boot && make
```

Don't worry if the build falls over in the RTS, we don't need the RTS yet.

- ```
$ cd H/libraries
$ make boot && make
```
- ```
$ cd H/ghc/compiler
$ make boot stage=2 && make stage=2
```

- ```
$ cd H/compat
$ make clean
$ rm .depend
$ make boot UseStagel=YES EXTRA_HC_OPTS='-O -fvia-C -keep-hc-files'
$ cd H/utils
$ make clean
$ make -k UseStagel=YES EXTRA_HC_OPTS='-O -fvia-C -keep-hc-files'
```
- ```
$ cd H
$ make hc-file-bundle Project=Ghc
```
- copy `H/*-hc.tar.gz` to `T/...`
- On the target machine:

At this stage we simply need to bootstrap a compiler from the intermediate C files we generated above. The process of bootstrapping from C files is automated by the script in `distrib/hc-build`, and is described in Section 9.1, “Booting/porting from C (.hc) files”.

```
$ ./distrib/hc-build --enable-hc-boot-unregisterised
```

However, since this is a bootstrap on a new machine, the automated process might not run to completion the first time. For that reason, you might want to treat the `hc-build` script as a list of instructions to follow, rather than as a fully automated script. This way you'll be able to restart the process part-way through if you need to fix anything on the way.

Don't bother with running `make install` in the newly bootstrapped tree; just use the compiler in that tree to build a fresh compiler from scratch, this time without booting from C files. Before doing this, you might want to check that the bootstrapped compiler is generating working binaries:

```
$ cat >hello.hs
main = putStrLn "Hello World!\n"
^D
$ T/ghc/compiler/ghc-inplace hello.hs -o hello
$ ./hello
Hello World!
```

Once you have the unregistered compiler up and running, you can use it to start a registered port. The following sections describe the various parts of the system that will need architecture-specific tweaks in order to get a registered build going.

## 9.2.2. Porting the RTS

The following files need architecture-specific code for a registered build:

|                                       |                                                                                                                                                                                                                                                      |
|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ghc/includes/MachRegs.h</code>  | Defines the STG-register to machine-register mapping. You need to know your platform's C calling convention, and which registers are generally available for mapping to global register variables. There are plenty of useful comments in this file. |
| <code>ghc/includes/TailCalls.h</code> | Macros that cooperate with the mangler (see Section 9.2.3, “The mangler”) to make proper tail-calls work.                                                                                                                                            |

|                                                               |                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ghc/rts/Adjustor.c</code>                               | Support for foreign import "wrapper" (aka foreign export dynamic). Not essential for getting GHC bootstrapped, so this file can be deferred until later if necessary.                                                                                                                                                                                                        |
| <code>ghc/rts/StgCRun.c</code>                                | The little assembly layer between the C world and the Haskell world. See the comments and code for the other architectures in this file for pointers.                                                                                                                                                                                                                        |
| <code>ghc/rts/MBlock.h</code> , <code>ghc/rts/MBlock.c</code> | These files are really OS-specific rather than architecture-specific. In <code>MBlock.h</code> is specified the absolute location at which the RTS should try to allocate memory on your platform (try to find an area which doesn't conflict with code or dynamic libraries). In <code>Mblock.c</code> you might need to tweak the call to <code>mmap()</code> for your OS. |

### 9.2.3. The mangler

The mangler is an evil Perl-script (`ghc/driver/mangler/ghc-asm.lprl`) that rearranges the assembly code output from gcc to do two main things:

- Remove function prologues and epilogues, and all movement of the C stack pointer. This is to support tail-calls: every code block in Haskell code ends in an explicit jump, so we don't want the C-stack overflowing while we're jumping around between code blocks.
- Move the *info table* for a closure next to the entry code for that closure. In unregistered code, info tables contain a pointer to the entry code, but in registered compilation we arrange that the info table is shoved right up against the entry code, and addressed backwards from the entry code pointer (this saves a word in the info table and an extra indirection when jumping to the closure entry code).

The mangler is abstracted to a certain extent over some architecture-specific things such as the particular assembler directives used to herald symbols. Take a look at the definitions for other architectures and use these as a starting point.

### 9.2.4. The splitter

The splitter is another evil Perl script (`ghc/driver/split/ghc-split.lprl`). It cooperates with the mangler to support object splitting. Object splitting is what happens when the `-split-objs` option is passed to GHC: the object file is split into many smaller objects. This feature is used when building libraries, so that a program statically linked against the library will pull in less of the library.

The splitter has some platform-specific stuff; take a look and tweak it for your system.

### 9.2.5. The native code generator

The native code generator isn't essential to getting a registered build going, but it's a desirable thing to have because it can cut compilation times in half. The native code generator is described in some detail in the GHC commentary [<http://www.cse.unsw.edu.au/~chak/haskell/ghc/comm/>].

### 9.2.6. GHCi

To support GHCi, you need to port the dynamic linker (`$(GHC_TOP)/rts/Linker.c`). The linker currently supports the ELF and PEi386 object file formats - if your platform uses one of these then things will be significantly easier. The majority of Unix platforms use the ELF format these days. Even so, there are some machine-specific parts of the ELF linker: for example, the code for resolving particu-

lar relocation types is machine-specific, so some porting of this code to your architecture and/or OS will probably be necessary.

If your system uses a different object file format, then you have to write a linker — good luck!

## 10. Known pitfalls in building Glasgow Haskell

WARNINGS about pitfalls and known “problems”:

1. One difficulty that comes up from time to time is running out of space in `TMPDIR`. (It is impossible for the configuration stuff to compensate for the vagaries of different sysadmin approaches to temp space.) The quickest way around it is **setenv** `TMPDIR /usr/tmp` or even **setenv** `TMPDIR .` (or the equivalent incantation with your shell of choice). The best way around it is to say

```
export TMPDIR=<dir>
```

in your `build.mk` file. Then GHC and the other tools will use the appropriate directory in all cases.

2. In compiling some support-code bits, e.g., in `ghc/rts/gmp` and even in `ghc/lib`, you may get a few C-compiler warnings. We think these are OK.
3. When compiling via C, you'll sometimes get “warning: assignment from incompatible pointer type” out of GCC. Harmless.
4. Similarly, **archiving** warning messages like the following are not a problem:

```
ar: filename GlaIOMonad__1_2s.o truncated to GlaIOMonad_
ar: filename GlaIOMonad__2_2s.o truncated to GlaIOMonad_
...
```

5. In compiling the compiler proper (in `compiler/`), you *may* get an “Out of heap space” error message. These can vary with the vagaries of different systems, it seems. The solution is simple:
  - If you're compiling with GHC 4.00 or later, then the *maximum* heap size must have been reached. This is somewhat unlikely, since the maximum is set to 64M by default. Anyway, you can raise it with the `-optCrts-M<size>` flag (add this flag to `<module>_HC_OPTS` **make** variable in the appropriate Makefile).
  - For GHC < 4.00, add a suitable `-H` flag to the Makefile, as above.
 and try again: **make**. (see Section 7.6, “Pattern rules and options” for information about `<module>_HC_OPTS`.) Alternatively, just cut to the chase:

```
$ cd ghc/compiler
$ make EXTRA_HC_OPTS=-optCrts-M128M
```

6. If you try to compile some Haskell, and you get errors from GCC about lots of things from `/usr/include/math.h`, then your GCC was mis-installed. **fixincludes** wasn't run when it should've been. As **fixincludes** is now automatically run as part of GCC installation, this bug also suggests that you have an old GCC.
7. You *may* need to re-**ranlib** your libraries (on Sun4s).

```
$ cd $(libdir)/ghc-x.xx/sparc-sun-sunos4
$ foreach i ( `find . -name '*.a' -print` ) # or other-shell equiv...
?   ranlib $i
```

```
? # or, on some machines: ar s $i  
? end
```

We'd be interested to know if this is still necessary.

8. GHC's sources go through **cpp** before being compiled, and **cpp** varies a bit from one Unix to another. One particular gotcha is macro calls like this:

```
SLIT("Hello, world")
```

Some **cpps** treat the comma inside the string as separating two macro arguments, so you get

```
:731: macro `SLIT' used with too many (2) args
```

Alas, **cpp** doesn't tell you the offending file! Workaround: don't put weird things in string args to **cpp** macros.

## 11. Platforms, scripts, and file names

GHC is designed both to be built, and to run, on both Unix and Windows. This flexibility gives rise to a good deal of brain-bending detail, which we have tried to collect in this chapter.

### 11.1. Windows platforms: Cygwin, MSYS, and MinGW

The build system is built around Unix-y makefiles. Because it's not native, the Windows situation for building GHC is particularly confusing. This section tries to clarify, and to establish terminology.

#### 11.1.1. MinGW

MinGW (Minimalist GNU for Windows) [<http://www.mingw.org>] is a collection of header files and import libraries that allow one to use **gcc** and produce native Win32 programs that do not rely on any third-party DLLs. The current set of tools include GNU Compiler Collection (**gcc**), GNU Binary Utilities (Binutils), GNU debugger (Gdb), GNU make, and a assorted other utilities.

The down-side of MinGW is that the MinGW libraries do not support anything like the full Posix interface.

#### 11.1.2. Cygwin and MSYS

You can't use the MinGW to *build* GHC, because MinGW doesn't have a shell, or the standard Unix commands such as **mv**, **rm**, **ls**, nor build-system stuff such as **make** and **darcs**. For that, there are two choices: Cygwin [<http://www.cygwin.com>] and MSYS [<http://www.mingw.org/msys.shtml>]:

- Cygwin comes with compilation tools (**gcc**, **ld** and so on), which compile code that has access to all of Posix. The price is that the executables must be dynamically linked with the Cygwin DLL, so that *you cannot run a Cywin-compiled program on a machine that doesn't have Cygwin*. Worse, Cygwin is a moving target. The name of the main DLL, `cygwin1.dll` does not change, but the implementation certainly does. Even the interfaces to functions it exports seem to change occasionally.
- MSYS is a fork of the Cygwin tree, so they are fundamentally similar. However, MSYS is by design much smaller and simpler. Access to the file system goes through fewer layers, so MSYS is quite a bit faster too.

Furthermore, MSYS provides no compilation tools; it relies instead on the MinGW tools. These

compile binaries that run with no DLL support, on any Win32 system. However, MSYS does come with all the make-system tools, such as **make**, **autoconf**, **darcs**, **ssh** etc. To get these, you have to download the MsysDTK (Developer Tool Kit) package, as well as the base MSYS package.

MSYS does have a DLL, but it's only used by MSYS commands (**sh**, **rm**, **ssh** and so on), not by programs compiled under MSYS.

### 11.1.3. Targeting MinGW

We want GHC to compile programs that work on any Win32 system. Hence:

- GHC does invoke a C compiler, assembler, linker and so on, but we ensure that it only invokes the MinGW tools, not the Cygwin ones. That means that the programs GHC compiles will work on any system, but it also means that the programs GHC compiles do not have access to all of Posix. In particular, they cannot import the (Haskell) Posix library; they have to do their input output using standard Haskell I/O libraries, or native Win32 bindings.

We will call a GHC that targets MinGW in this way *GHC-mingw*.

- To make the GHC distribution self-contained, the GHC distribution includes the MinGW **gcc**, **as**, **ld**, and a bunch of input/output libraries.

So *GHC targets MinGW*, not Cygwin. It is in principle possible to build a version of GHC, *GHC-cygwin*, that targets Cygwin instead. The up-side of *GHC-cygwin* is that Haskell programs compiled by *GHC-cygwin* can import the (Haskell) Posix library. *We do not support GHC-cygwin, however; it is beyond our resources.*

While *GHC targets MinGW*, that says nothing about how GHC is *built*. We use both MSYS and Cygwin as build environments for GHC; both work fine, though MSYS is rather lighter weight.

In your build tree, you build a compiler called **ghc-inplace**. It uses the **gcc** that you specify using the `-with-gcc` flag when you run **configure** (see below). The makefiles are careful to use **ghc-inplace** (not **gcc**) to compile any C files, so that it will in turn invoke the correct **gcc** rather than whatever one happens to be in your path. However, the makefiles do use whatever **ld** and **ar** happen to be in your path. This is a bit naughty, but (a) they are only used to glom together `.o` files into a bigger `.o` file, or a `.a` file, so they don't ever get libraries (which would be bogus; they might be the wrong libraries), and (b) Cygwin and MinGW use the same `.o` file format. So it's ok.

### 11.1.4. File names

Cygwin, MSYS, and the underlying Windows file system all understand file paths of form `c:/tmp/foo`. However:

- MSYS programs understand `/bin`, `/usr/bin`, and map Windows's lettered drives as `/c/tmp/foo` etc. The exact mount table is given in the `doc` subdirectory of the MSYS distribution.

When it invokes a command, the MSYS shell sees whether the invoked binary lives in the MSYS `/bin` directory. If so, it just invokes it. If not, it assumes the program is not an MSYS program, and walks over the command-line arguments changing MSYS paths into native-compatible paths. It does this inside sub-arguments and inside quotes. For example, if you invoke

```
foogle -B/c/tmp/baz
```

the MSYS shell will actually call `foogle` with argument `-Bc:/tmp/baz`.

- Cygwin programs have a more complicated mount table, and map the lettered drives as `/cygdrive/c/tmp/foo`.

The Cygwin shell does no argument processing when invoking non-Cygwin programs.

### 11.1.5. Crippled ld

It turns out that on both Cygwin and MSYS, the `ld` has a limit of 32kbytes on its command line. Especially when using split object files, the make system can emit calls to `ld` with thousands of files on it. Then you may see something like this:

```
(cd Graphics/Rendering/OpenGL/GL/QueryUtils_split && /mingw/bin/ld -r -x -o ../Que
/bin/sh: /mingw/bin/ld: Invalid argument
```

The solution is either to switch off object file splitting (set `SplitObjs` to `NO` in your `build.mk`), or to make the module smaller.

### 11.1.6. Host System vs Target System

In the source code you'll find various `ifdefs` looking like:

```
#ifdef mingw32_HOST_OS
...blah blah...
#endif

and

#ifdef mingw32_TARGET_OS
...blah blah...
#endif
```

These macros are set by the configure script (via the file `config.h`). Which is which? The criterion is this. In the `ifdefs` in GHC's source code:

- The "host" system is the one on which GHC itself will be run.
- The "target" system is the one for which the program compiled by GHC will be run.

For a stage-2 compiler, in which `GHCi` is available, the "host" and "target" systems must be the same. So then it doesn't really matter whether you use the `HOST_OS` or `TARGET_OS` `cpp` macros.

## 11.2. Wrapper scripts

Many programs, including GHC itself and `hsc2hs`, need to find associated binaries and libraries. For *installed* programs, the strategy depends on the platform. We'll use GHC itself as an example:

- On Unix, the command `ghc` is a shell script, generated by adding installation paths to the front of the source file `ghc.sh`, that invokes the real binary, passing `-Bpath` as an argument to tell `ghc` where to find its supporting files.
- On vanilla Windows, it turns out to be much harder to make reliable script to be run by the native Windows shell `cmd` (e.g. limits on the length of the command line). So instead we invoke the GHC binary directly, with no `-B` flag. GHC uses the Windows `getExecDir` function to find where the executable is, and from that figures out where the supporting files are.

(You can find the layout of GHC's supporting files in the section "Layout of installed files" of Section 2 of the GHC user guide.)

Things work differently for *in-place* execution, where you want to execute a program that has just been built in a build tree. The difference is that the layout of the supporting files is different. In this case, whether on Windows or Unix, we always use a shell script. This works OK on Windows because the script is executed by MSYS or Cygwin, which don't have the shortcomings of the native Windows **cmd** shell.

## 12. Instructions for building under Windows

This section gives detailed instructions for how to build GHC from source on your Windows machine. Similar instructions for installing and running GHC may be found in the user guide. In general, Win95/Win98 behave the same, and WinNT/Win2k behave the same.

Make sure you read the preceding section on platforms (Section 11, "Platforms, scripts, and file names") before reading section. You don't need Cygwin or MSYS to *use* GHC, but you do need one or the other to *build* GHC.

### 12.1. Installing and configuring MSYS

MSYS is a lightweight alternative to Cygwin. You don't need MSYS to *use* GHC, but you do need it or Cygwin to *build* GHC. Here's how to install MSYS.

- Go to <http://www.mingw.org/download.shtml> and download the following (of course, the version numbers will differ):
  - The main MSYS package (binary is sufficient): `MSYS-1.0.9.exe`
  - The MSYS developer's toolkit (binary is sufficient): `msysDTK-1.0.1.exe`. This provides **make**, **autoconf**, **ssh** and probably more besides.Run both executables (in the order given above) to install them. I put them in `c:/msys`
- Set the following environment variables
  - `PATH`: add `c:/msys/1.0/bin` and `c:/msys/1.0/local/bin` to your path. (Of course, the version number may differ.) MSYS mounts the former as both `/bin` and `/usr/bin` and the latter as `/usr/local/bin`.
  - `HOME`: set to your home directory (e.g. `c:/userid`). This is where, among other things, **ssh** will look for your `.ssh` directory.
  - `SHELL`: set to `c:/msys/1.0/bin/sh.exe`
  - `CVS_RSH`: set to `c:/msys/1.0/bin/ssh.exe`. Only necessary if you are using CVS.
  - `MAKE_MODE`: set to `UNIX`. (I'm not certain this is necessary for MSYS.)
- Check that the `CYGWIN` environment variable is *not* set. It's a bad bug that MSYS is affected by this, but if you have `CYGWIN` set to `"ntsec ntea"`, which is right for Cygwin, it causes the MSYS **ssh** to bogusly fail complaining that your `.ssh/identity` file has too-liberal permissions.

Here are some points to bear in mind when using MSYS:

- MSYS does some kind of special magic to binaries stored in `/bin` and `/usr/bin`, which are by

default both mapped to `c:/msys/1.0/bin` (assuming you installed MSYS in `c:/msys`). Do not put any other binaries (such as GHC or Alex) in this directory or its sub-directories: they fail in mysterious ways. However, it's fine to put other binaries in `/usr/local/bin`, which maps to `c:/msys/1.0/local/bin`.

- MSYS seems to implement symbolic links by copying, so sharing is lost.
- Win32 has a **find** command which is not the same as MSYS's find. You will probably discover that the Win32 **find** appears in your PATH before the MSYS one, because it's in the *system* PATH environment variable, whereas you have probably modified the *user* PATH variable. You can always invoke **find** with an absolute path, or rename it.
- MSYS comes with **bzip**, and MSYS's **tar**'s `-j` will bunzip an archive (e.g. `tar xvjf foo.tar.bz2`). Useful when you get a bzip'd dump.

## 12.2. Installing and configuring Cygwin

Install Cygwin from <http://www.cygwin.com/>. The installation process is straightforward; we install it in `c:/cygwin`.

You must install enough Cygwin *packages* to support building GHC. If you miss out any of these, strange things will happen to you. There are two ways to do this:

- The direct, but laborious way is to select all of the following packages in the installation dialogue: **cvs**, **openssh**, **autoconf**, **binutils** (includes `ld` and (I think) `ar`), **gcc**, **flex**, **make**. To see these packages, click on the "View" button in the "Select Packages" stage of Cygwin's installation dialogue, until the view says "Full". The default view, which is "Category" isn't very helpful, and the "View" button is rather unobtrusive.
- The clever way is to point the Cygwin installer at the **ghc-depends** package, which is kept at <http://haskell.org/ghc/cygwin>. When the Cygwin installer asks you to "Choose a Download Site", choose one of the offered mirror sites; and then type "<http://haskell.org/ghc/cygwin>" into the "User URL" box and click "Add"; now two sites are selected. (The Cygwin installer remembers this for next time.) Click "Next".

In the "Select Packages" dialogue box that follows, click the "+" sign by "Devel", scroll down to the end of the "Devel" packages, and choose **ghc-depends**. The package **ghc-depends** will not actually install anything itself, but forces additional packages to be added by the Cygwin installer.

Now set the following user environment variables:

- Add `c:/cygwin/bin` and `c:/cygwin/usr/bin` to your PATH
- Set `MAKE_MODE` to `UNIX`. If you don't do this you get very weird messages when you type **make**, such as:  

```
/c: /c: No such file or directory
```
- Set `SHELL` to `c:/cygwin/bin/bash`. When you invoke a shell in Emacs, this `SHELL` is what you get.
- Set `HOME` to point to your home directory. This is where, for example, **bash** will look for your `.bashrc` file. Ditto **emacs** looking for `.emacsrc`

Here are some things to be aware of when using Cygwin:

- Cygwin doesn't deal well with filenames that include spaces. "Program Files" and "Local files" are common gotchas.
- Cygwin implements a symbolic link as a text file with some magical text in it. So other programs that don't use Cygwin's I/O libraries won't recognise such files as symlinks. In particular, programs compiled by GHC are meant to be runnable without having Cygwin, so they don't use the Cygwin library, so they don't recognise symlinks.
- See the notes in Section 12.1, "Installing and configuring MSYS" about **find** and **bzip**, which apply to Cygwin too.
- Some script files used in the make system start with "**#!/bin/perl**", (and similarly for **sh**). Notice the hardwired path! So you need to ensure that your `/bin` directory has at least **sh**, **perl**, and **cat** in it. All these come in Cygwin's `bin` directory, which you probably have installed as `c:/cygwin/bin`. By default Cygwin mounts "/" as `c:/cygwin`, so if you just take the defaults it'll all work ok. (You can discover where your Cygwin root directory / is by typing **mount**.) Provided `/bin` points to the Cygwin `bin` directory, there's no need to copy anything. If not, copy these binaries from the `cygwin/bin` directory (after fixing the `sh.exe` stuff mentioned in the previous bullet).
- By default, cygwin provides the command shell `ash` as `sh.exe`. It seems to be fine now, but in the past we saw build-system problems that turned out to be due to bugs in `ash` (to do with quoting and length of command lines). On the other hand `bash` seems to be rock solid. If this happens to you (which it shouldn't), in `cygwin/bin` remove the supplied `sh.exe` (or rename it as `ash.exe`), and copy `bash.exe` to `sh.exe`. You'll need to do this in Windows Explorer or the Windows **cmd** shell, because you can't rename a running program!

## 12.3. Configuring SSH

**ssh** comes with both Cygwin and MSYS. (Cygwin note: you need to ask for package **openssh** (not **ssh**) in the Cygwin list of packages; or use the **ghc-depends** package -- see Section 12.2, "Installing and configuring Cygwin".)

There are several strange things about **ssh** on Windows that you need to know.

- The programs **ssh-keygen1**, **ssh1**, and **cvs**, seem to lock up **bash** entirely if they try to get user input (e.g. if they ask for a password). To solve this, start up `cmd.exe` and run it as follows:

```
c:\tmp> set CYGWIN32=TTY
c:\tmp> c:/user/local/bin/ssh-keygen1
```

- (Cygwin-only problem, I think.) **ssh** needs to access your directory `.ssh`, in your home directory. To determine your home directory **ssh** first looks in `c:/cygwin/etc/passwd` (or wherever you have Cygwin installed). If there's an entry there with your userid, it'll use that entry to determine your home directory, *ignoring the setting of the environment variable \$HOME*. If the home directory is bogus, **ssh** fails horribly. The best way to see what is going on is to say

```
ssh -v cvs.haskell.org
```

which makes **ssh** print out information about its activity.

You can fix this problem, either by correcting the home-directory field in `c:/cygwin/etc/passwd`, or by simply deleting the entire entry for your userid. If you do that,

`ssh` uses the `$HOME` environment variable instead.

- To protect your `.ssh` from access by anyone else, right-click your `.ssh` directory, and select `Properties`. If you are not on the access control list, add yourself, and give yourself full permissions (the second panel). Remove everyone else from the access control list. Don't leave them there but deny them access, because 'they' may be a list that includes you!
- In fact `ssh` 3.6.1 now seems to *require* you to have Unix permissions 600 (read/write for owner only) on the `.ssh/identity` file, else it bombs out. For your local C drive, it seems that `chmod 600 identity` works, but on Windows NT/XP, it doesn't work on a network drive (exact details obscure). The solution seems to be to set the `$CYGWIN` environment variable to "ntsec neta". The `$CYGWIN` environment variable is discussed in the Cygwin User's Guide [<http://cygwin.com/cygwin-ug-net/using-cygwinenv.html>], and there are more details in the Cygwin FAQ [[http://cygwin.com/faq/faq\\_4.html#SEC44](http://cygwin.com/faq/faq_4.html#SEC44)].

## 12.4. Other things you need to install

You have to install the following other things to build GHC, listed below.

On Windows you often install executables in directories with spaces, such as "Program Files". However, the make system doesn't deal with this situation (it'd have to do more quoting of binaries), so you are strongly advised to put binaries for all tools in places with no spaces in their path. On both MSYS and Cygwin, it's perfectly OK to install such programs in the standard Unixy places, `/usr/local/bin` and `/usr/local/lib`. But it doesn't matter, provided they are in your path.

- Install an executable GHC, from <http://www.haskell.org/ghc>. This is what you will use to compile GHC. Add it in your `PATH`: the installer tells you the path element you need to add upon completion.
- Install an executable Happy, from <http://www.haskell.org/happy>. Happy is a parser generator used to compile the Haskell grammar. Under MSYS or Cygwin you can easily build it from the source distribution using

```
$ ./configure
$ make
$ make install
```

This should install it in `/usr/local/bin` (which maps to `c:/msys/1.0/local/bin` on MSYS). Make sure the installation directory is in your `PATH`.

- Install an executable Alex. This can be done by building from the source distribution in the same way as Happy. Sources are available from <http://www.haskell.org/alex>.
- GHC uses the *mingw* C compiler to generate code, so you have to install that (see Section 11.1, "Windows platforms: Cygwin, MSYS, and MinGW"). Just pick up a mingw bundle at <http://www.mingw.org/>. We install it in `c:/mingw`.

*On MSYS*, add `c:/mingw/bin` to your `PATH`. MSYS does not provide `gcc`, `ld`, `ar`, and so on, because it just uses the MinGW ones. So you need them in your path.

*On Cygwin*, do not add any of the *mingw* binaries to your path. They are only going to get used by explicit access (via the `--with-gcc` flag you give to `configure` later). If you do add them to your path you are likely to get into a mess because their names overlap with Cygwin binaries. On the other hand, you *do* need `ld`, `ar` (and perhaps one or two other things) in your path. The Cygwin ones are fine, but you must have them; hence needing the Cygwin binutils package.

- We use **emacs** a lot, so we install that too. When you are in `$(GHC_TOP)/compiler`, you can use "make tags" to make a TAGS file for emacs. That uses the utility `$(GHC_TOP)/ghc/utis/hasktags/hasktags`, so you need to make that first. The most convenient way to do this is by going `make boot` in `$(GHC_TOP)/ghc`. The `make tags` command also uses **etags**, which comes with **emacs**, so you will need to add `emacs/bin` to your `PATH`.
- You might want to install GLUT in your MSYS/Cygwin installation, otherwise the GLUT package will not be built with GHC.
- Finally, check out a copy of GHC sources from the darcs repository, following the instructions at <http://hackage.haskell.org/trac/ghc/wiki/GhcDarcs>.

## 12.5. Building GHC

OK! Now go read the documentation above on building from source (Section 4, "Building from source"); the bullets below only tell you about Windows-specific wrinkles.

- If you used **autoconf** instead of **autoreconf**, you'll get an error when you run `./configure`:

```
...lots of stuff...
creating mk/config.h
mk/config.h is unchanged
configuring in ghc
running /bin/sh ./configure --cache-file=../config.cache --srcdir=.
./configure: ./configure: No such file or directory
configure: error: ./configure failed for ghc
```

- **autoreconf** seems to create the file `configure` read-only. So if you need to run `autoreconf` again (which I sometimes do for safety's sake), you get

```
/usr/bin/autoconf: cannot create configure: permission denied
```

Solution: delete `configure` first.

- After **autoreconf** run `./configure` in `$(GHC_TOP)/` thus:

```
$ ./configure --host=i386-unknown-mingw32 --with-gcc=c:/mingw/bin/gcc
```

This is the point at which you specify that you are building GHC-mingw (see Section 11.1.1, "MinGW").

Both these options are important! It's possible to get into trouble using the wrong C compiler!

Furthermore, it's *very important* that you specify a full MinGW path for **gcc**, not a Cygwin path, because GHC (which uses this path to invoke **gcc**) is a MinGW program and won't understand a Cygwin path. For example, if you say `--with-gcc=/mingw/bin/gcc`, it'll be interpreted as `/cygdrive/c/mingw/bin/gcc`, and GHC will fail the first time it tries to invoke it. Worse, the failure comes with no error message whatsoever. GHC simply fails silently when first invoked, typically leaving you with this:

```
make[4]: Leaving directory `~/cygdrive/e/ghc-stagel/ghc/rts/gmp'
../../ghc/compiler/ghc-inplace -optc-mno-cygwin -optc-O
-optc-Wall -optc-W -optc-Wstrict-prototypes -optc-Wmissing-prototypes
-optc-Wmissing-declarations -optc-Winline -optc-Waggregate-return
-optc-Wbad-function-cast -optc-Wcast-align -optc-I../includes
```

```

    -optc-I. -optc-Iparallel -optc-DCOMPILING_RTS
    -optc-fomit-frame-pointer -O2 -static
    -package-name rts -O -dcore-lint -c Adjustor.c -o Adjustor.o
make[2]: *** [Adjustor.o] Error 1
make[1]: *** [all] Error 1
make[1]: Leaving directory `/cygdrive/e/ghc-stage1/ghc'
make: *** [all] Error 1

```

Be warned!

If you want to build GHC-cygwin (Section 11.1.2, “Cygwin and MSYS”) you'll have to do something more like:

```
$ ./configure --with-gcc=...the Cygwin gcc...
```

- If you are paranoid, delete `config.cache` if it exists. This file occasionally remembers out-of-date configuration information, which can be really confusing.
- You almost certainly want to set

```
SplitObjs = NO
```

in your `build.mk` configuration file (see Section 6.3, “Getting the build you want”). This tells the build system not to split each library into a myriad of little object files, one for each function. Doing so reduces binary sizes for statically-linked binaries, but on Windows it dramatically increases the time taken to build the libraries in the first place.

- Do not attempt to build the documentation. It needs all kinds of wierd Jade stuff that we haven't worked out for Win32.

## 12.6. A Windows build log using Cygwin

Here is a complete, from-scratch, log of all you need to build GHC using Cygwin, kindly provided by Claus Reinke. It does not discuss alternative choices, but it gives a single path that works.

- Install some editor (vim, emacs, whatever)
- Install cygwin (<http://www.cygwin.com>)
  - ; i used 1.5.16-1, installed in `c:\cygwin`
- run 'setup.exe'
  - Choose a DownloadSource:
    - select 'download from internet';
  - Select Root Install Directory:
    - root dir: `c:\cygwin`;
    - install for: all users;
    - default file type: unix
  - Select Local Package Directory
    - choose a spare temporary home
  - Select Your Internet Connection
    - Use IE5 settings
  - Choose a Download Site
    - Choose your preferred main mirror and
    - Add '<http://www.haskell.org/ghc/cygwin>'
  - Select Packages
    - In addition to 'Base' (default install),
    - select 'Devel->ghc-depends'
- Install mingw (<http://www.mingw.org/>)



## A

- Adjustor.c, 29
- Alex, 4
- ALL\_DIRS, 17
- autoconf, pre-supposed, 4
- autoreconf, 8

## B

- bash, presupposed (Parallel Haskell only), 5
- boilerplate architecture, 16
- boilerplate.mk, 10, 15, 16, 17
- booting GHC from .hc files, 24
- bugs
  - known, 3
  - mailing list, 3
  - seporing, 3
- build trees, 7
- build.mk, 9
- Building from source, 5
- building GHC from .hc files, 24
- building pitfalls, 30

## C

- config.h, 9
- config.h.in, 9
- config.mk, 8, 17
- config.mk.in, 8
- configure, 3
- C\_OBJS, 18
- C\_PROG, 21
- C\_SRCS, 17

## D

- darcs repository, 2
- dependencies, omitting, 14
- Disk space needed, 3
- DocBook, pre-supposed, 22

## E

- EXCLUDED\_SRCS, 17
- EXTRA\_CC\_OPTS, 20
- EXTRA\_SRCS, 17

## F

- FAST, makefile variable, 14
- fastmake, 14

## G

- GCC (GNU C compiler), pre-supposed, 4
- GHC, pre-supposed, 3
- GhcWithHc, 16
- GHC\_TOP, 7

## H

Haddock, 23  
happy, 2  
Happy, 4  
HAPPY\_SRCS, 18  
HC\_OPTS, 16, 18  
HSC\_SRCS, 18  
HS\_IFACES, 17  
HS\_OBJS, 17  
HS\_PROG, 15, 16, 20, 21  
HS\_SRCS, 17

## I

include, directive in Makefiles, 15  
install, 21

## L

LIBRARY, 21  
link trees, for building, 7  
lndir, 7

## M

MachRegs.h, 28  
make  
GNU, 4  
makefile architecture, 14  
Makefile inclusion, 15  
makefile targets, 12  
Makefile, minimal, 15  
Makefile, recursing into subdirectories, 21  
MBlock.c, 29  
MBlock.h, 29  
mkshadowdir, 7

## N

NoFibSubDirs, 16

## O

OBJS, 16, 18  
opts.mk, 18

## P

paths.mk, 17  
Pattern rules, 19  
Perl, pre-supposed, 3  
pitfalls, in building, 30  
platform, 8  
Platform settings, 18  
porting GHC, 24  
pre-supposed utilities, 3  
pre-supposed: autoconf, 4  
pre-supposed: DocBook, 22  
pre-supposed: GCC (GNU C compiler), 4  
pre-supposed: GHC, 3  
pre-supposed: Perl, 3  
pre-supposed: PVM3 (Parallel Virtual Machine), 5  
pre-supposed: sed, 5

pre-supposed: TeX, 22  
problems, building, 30  
PVM, 3  
PVM3 (Parallel Virtual Machine), pre-supposed, 5

## R

ranlib, 30  
recursion, in makefiles, 21

## S

SCRIPT\_OBJS, 18  
SCRIPT\_SRCS, 18  
sed, pre-supposed, 5  
Source distributions, 2  
Source, building from, 5  
SRCS, 17  
SRC\_CC\_OPTS, 19  
SRC\_HC\_OPTS, 16  
StgCRun.c, 29  
SUBDIRS, 21  
suffix.mk, 18, 19

## T

TailCalls.h, 28  
target.mk, 15, 16, 20, 21  
targets, standard makefile, 12  
TeX, pre-supposed, 22  
tmp, running out of space in, 30  
TMPDIR, 30  
TOP, 15

## U

utilities, pre-supposed, 3

## W

way management, 21