

Bourne Shell Scripting

en.wikibooks.org

July 12, 2015

On the 28th of April 2012 the contents of the English as well as German Wikibooks and Wikipedia projects were licensed under Creative Commons Attribution-ShareAlike 3.0 Unported license. A URI to this license is given in the list of figures on page 123. If this document is a derived work from the contents of one of these projects and the content was still licensed by the project under this license at the time of derivation this document has to be licensed under the same, a similar or a compatible license, as stated in section 4b of the license. The list of contributors is included in chapter Contributors on page 121. The licenses GPL, LGPL and GFDL are included in chapter Licenses on page 127, since this book and/or parts of it may or may not be licensed under one or more of these licenses, and thus require inclusion of these licenses. The licenses of the figures are given in the list of figures on page 123. This PDF was generated by the \LaTeX typesetting software. The \LaTeX source code is included as an attachment (`source.7z.txt`) in this PDF file. To extract the source from the PDF file, you can use the `pdfdetach` tool including in the `poppler` suite, or the <http://www.pdfplabs.com/tools/pdftk-the-pdf-toolkit/> utility. Some PDF viewers may also let you save the attachment to a file. After extracting it from the PDF file you have to rename it to `source.7z`. To uncompress the resulting archive we recommend the use of <http://www.7-zip.org/>. The \LaTeX source itself was generated by a program written by Dirk Hünninger, which is freely available under an open source license from http://de.wikibooks.org/wiki/Benutzer:Dirk_Huenniger/wb2pdf.

Contents

1	Comparing Shells	3
1.1	Bourne shell and other Unix command shells	3
1.2	Why Bourne Shell	5
2	Running Commands	7
2.1	The easy way: the interactive session	7
2.2	The only slightly less easy way: the script	8
2.3	A little bit about Unix and multiprocessing	12
3	Environment	15
3.1	The Environments	15
3.2	Multitasking and job control	26
4	Variable Expansion	39
4.1	Substitution	39
4.2	Substitution forms	40
5	Control flow	45
5.1	Control Flow	45
5.2	Command execution	56
6	Files and streams	67
6.1	The Unix world: one file after another	67
6.2	Streams: what goes between files	68
6.3	Redirecting: using streams in the shell	69
6.4	Here documents	79
7	Modularization	81
7.1	Named functions	81
7.2	Creating a named function	82
8	Debugging and signal handling	87
8.1	Debugging Flags	87
8.2	Breaking out of a script	93
8.3	Signal trapping	95
9	Cookbook	103
9.1	Branch on extensions	103
9.2	Rename several files	103
9.3	Long command line options	104
9.4	Process certain files through xargs	105

9.5	Simple playlist frontend for GStreamer	106
10	Quick Reference	109
10.1	Useful commands	109
10.2	Elementary shell capabilities	110
10.3	IF statement	110
10.4	CASE statement	113
10.5	Loop statements	113
10.6	Credit	114
11	Command Reference	115
12	Environment reference	119
13	Contributors	121
	List of Figures	123
14	Licenses	127
14.1	GNU GENERAL PUBLIC LICENSE	127
14.2	GNU Free Documentation License	128
14.3	GNU Lesser General Public License	129

1 Comparing Shells

Almost all books like this one have a section on (or very similar to) "why you should use the shell/program flavor/language/etc. discussed in this book and not any of the others that perform the same tasks in a slightly different way". It seems to be pretty well mandatory.

However, this book will not do that. We'll talk a bit about "why Bourne Shell" of course. But you'll soon see that doesn't preclude other shells at all. And there's no good reason not to use another shell either, as we will explain a little further down.

1.1 Bourne shell and other Unix command shells

There are many Unix command shells available today. Bourne Shell is just one drop in a very large ocean. How do all these shells relate? Do they do the same things? Is one better than the other? Let's take a look at what makes a shell and what that means for all the different shells out there.

1.1.1 How it all got started...

The Unix operating system has had a unique outlook on the world ever since it was created back in the 1970s. It stands apart from most other operating systems in that its focus has always been towards *power users* : people who want to squeeze every drop of performance out of their system and have the technical knowledge to do so. Unix was designed to be programmed and modified to the desires of the user. At its core, Unix does not have a user interface; instead it is comprised of a stable OS kernel and a versatile C library. If you're not trying to do actual hard-core programming but rather are trying to do day-to-day tasks (or even just want to put a little program together quickly), pure Unix is a tremendous pain in the backside.

In other words, it was clear from the start that a tool would be needed that would allow a user to make use of the functions offered him by the coding library and kernel without actually doing serious programming work. A tool in other words that could pass small commands on to the lower-level system quickly and easily. Without the need for a compiler or other fancy editor, but with the ability to tap into the enormous power of the underlying system. Stephen Bourne set himself to the task and came up with what he called a *shell* : a small, on-the-fly compiler that could take one command at a time, translate it into the sequence of bits understood by the machine and have that command carried out. We now call this type of program an interpreter¹, but at the time, the term "shell" was much more common (since it was a shell over the underlying system for the user). Stephen's shell was

¹ <http://en.wiktionary.org/wiki/interpreter>

slim, fast, and though a bit unwieldy at times, its power is still the envy of many current operating system command-line interfaces today. Since it was designed by Stephen Bourne, this shell is called the Bourne Shell. The executable is simply called *sh* and use of this shell in scripting is still so ubiquitous, there isn't a Unix-based system on this earth that doesn't offer a shell whose executable can be reached under the name *sh*.

1.1.2 ...And how it ended up

Of course, everyone's a critic. The Bourne Shell saw tremendous use (indeed, it still does) and as a result, it became the de facto standard among Unix shells. But all sorts of people almost immediately (as well as with use) wanted new features in the shell, or a more familiar way of expressing commands, or something else. Many people built new shells that they felt continued where Bourne Shell ended. Some were completely compatible with Bourne Shell, others were less so. Some became famous, others flopped. But pretty much all of them look fondly upon Bourne Shell, the shell they call "Dad..."

A number of these shells can be run in *sh*-like mode, to more closely emulate that very first *sh*, though most people tend just to run their shells in the default mode, which provides more power than the minimum *sh*.

1.1.3 It's Bourne Shell, but not as we know it...

So there are a lot of shells around but you can find Bourne Shell everywhere, right? Good old *sh*, just sitting there faithfully until the end of time...

Well, no, not really. Most of the *sh* executables out there nowadays aren't really the Bourne Shell anymore. Through a bit of Unix magic called a link (which allows one file to masquerade as another) the *sh* executable you find on any Unix system is likely actually to be one of the shells that is based on the Bourne shell. One of the most frequently used shells nowadays (with the ascent of free and open-source operating systems like GNU and Linux) is a heavily extended form of the Bourne Shell produced by the Free Software Foundation, called Bash². Bash hasn't forgotten its roots, though: it stands for the **B**ourne **A**gain **SH**ell.

Another example of a descendant shell standing in for its ancestor is the Korn Shell (*ksh*). Also an extension shell, it is completely compatible with *sh* -- it simply adds some features. Much the same is true for *zsh*.

Finally, a slightly different category is formed by the C Shell (*csh*) and its descendant *tcsh*, native on BSD systems. These shells do break compatibility to some extent, using different syntax for many commands. Systems that use these shells as standard shells often provide a real Bourne Shell executable to run generic Bourne Shell scripts.

Having read the above, you will understand why this book doesn't have to convince you to use Bourne Shell instead of any other shell: in most cases, there's no noticeable difference. Bourne Shell and its legacy have become so ingrained in the heart and soul of the Unix

² <http://en.wikipedia.org/wiki/Bash%20%28Unix%20shell%29>

environment that you are using Bourne Shell when you are using practically any shell available to you.

1.2 Why Bourne Shell

So only one real question remains: now that you find yourself on your own, cozy slice of a Unix system, with your own shell and all its capabilities, is there any real reason to use Bourne Shell rather than using the whole range of your shell's capabilities?

Well, it depends. Probably, there isn't. For the most part of course, you *are* using Bourne Shell by using the whole potential of your shell -- your shell is probably *that* similar to the Bourne Shell. But there is one thing you might want to keep in mind: someday, you might want to write a script that you might want to pass around to other people. Of course you can write your script using the full range of options that your shell offers you; but then it might not work on another machine with another shell. This is where the role of Bourne Shell as the Latin of Unix command shells comes in -- and also where it is useful to know how to write scripts targeted specifically at the Bourne Shell. If you write your scripts for the Bourne Shell and nothing but the Bourne Shell, chances are far better than equal that your script will run straight out of the mail attachment (don't tell me you're still using boxes to ship things -- come on, get with the program) on any command shell out there.

2 Running Commands

Before we can start any kind of examination of the abilities of the Bourne Shell and how you can tap into its power, we have to cover some basic ground first: we have to discuss how to enter commands into the shell for execution by that shell.

2.1 The easy way: the interactive session

2.1.1 Taking another look at what you've probably already seen

If you have access to a Unix-based machine (or an emulator on another operating system), you've probably been using the Bourne Shell -- or one of its descendants -- already, possibly without realising. Surprise: you've been doing shell scripting for a while already!

In your Unix environment, go to a terminal; either a textual logon terminal, or a terminal-in-a-window if you're using the X Window System (look for something called *xterm* or *rxvt* or just *terminal* , if you have actually not ever done this yet). You'll probably end up looking at a screen looking something like this:

```
Ben_Tels:Local_Machine:~>_
```

or

```
The admin says: everybody, STOP TRYING TO CRASH THE SYSTEM  
Have a lot of fun!  
bzt:Another_Machine:~>_
```

or even something as simple as

```
$_
```

That's it. That's your shell: your direct access to everything the system has to offer.

2.1.2 Using the shell in interactive mode

Specifically, the program you accessed a moment ago is your shell, running in *interactive mode* : the shell is running in such a way that it displays a prompt and a cursor (the little, blinking line) and is waiting for you to enter a command for it to execute. You execute commands in interactive mode by typing them in, followed by a press of the **Enter** key.

The shell then translates your command to something the operating system understands and passes off control to the operating system so that it can actually carry out the task you have sent it. You'll notice that your cursor will disappear momentarily while the command is being carried out, and you cannot type anymore (at this point, the Bourne Shell program is no longer in control of your terminal -- the other program that you started by executing your command is). At some point the operating system will be finished working on your command and the shell will bring up a new prompt and the cursor as well and will then start waiting again for you to enter another command. Give it a try: type the command

```
ls enter
```

After a short time, you'll see a list of files in the working directory (the directory that your shell considers the "current" directory), a new prompt and the cursor.

This is the simplest way of executing shell commands: typing them in one at a time and waiting for each to complete in order. The shell is used in this way very often, both to execute commands that belong to the Bourne Shell programming language and simply to start running other programs (like the `ls` program from the example above).

2.1.3 A useful tidbit

Before we move on, we'll mention two useful key combinations when using the shell: the command to interrupt running programs and shell commands and the command to quit the shell (although, why you would ever want to *stop* using the shell is beyond me....).

To interrupt a running program or shell command, hit the Control and C keys at the same time. We'll get back to what this does exactly in a later chapter, but for now just remember this is the way to interrupt things.

To quit the shell session, hit Control+d. This key combination produces the Unix end-of-file character -- we'll talk more later about why this also terminates your shell session. Some modern shells have disabled the use of Control+d in favor of the "exit" command (shame on them). If you're using such a shell, just type the word "exit" (like with any other command) and press Enter (from here on in, I'll leave the "Enter" out of examples).

2.2 The only slightly less easy way: the script

As we saw in the last section, you can very easily execute shell commands for all purposes by starting an interactive shell session and typing your commands in at the prompt. However, sometimes you have a set of commands that you have to repeat regularly, even at different times and in different shell sessions. Of course, in the programming-centric environment of a Unix system, you can write a program to get the same result (in the C language for instance). But wouldn't it be a lot easier to have the convenience of the shell for this same task? Wouldn't it be more convenient to have a way to replay a set of commands? And to be able to compose that set as easily as you can write the single commands that you type into the shell's interactive sessions?

2.2.1 The shell script

Fortunately, there *is* such a way: the Bourne Shell's *non-interactive* mode. In this mode, the shell doesn't have a prompt or wait for your commands. Instead, the shell reads commands from a text file (which tells the shell what to do, kind of like an actor gets commands from a script -- hence, shell script). This file contains a sequence of commands, just as you would enter them into the interactive session at the prompt. The file is read by the shell from top to bottom and commands are executed in that order.

A shell script is very easy to write; you can use any text-editor you like (or even any wordprocessor or other editor, as long as you remember to save your script in plain text format). You write commands just as you would in the interactive shell. And you can run your script the moment you have saved it; no need to compile it or anything.

2.2.2 Running a shell script

To run a shell script (to have the shell read it and execute all the commands in the script), you enter a command at an interactive shell prompt as you would when doing anything else (if you're using a graphical user interface, you can probably also execute your scripts with a click of the mouse). In this case, the program you want to start is the shell program itself. For instance, to run a script called *MyScript*, you'd enter this command in the interactive shell (assuming the script is in your working directory):

Running a script

```
sh MyScript
```

Starting the shell program from inside the shell program may sound weird at first, but it makes perfect sense if you think about it. After all, you're typing commands in an *interactive mode* shell session. To run a script, you want to start a shell in *non-interactive mode*. That's what's happening in the above command. You'll note that the Bourne Shell executable takes a single parameter in the example above: the name of the script to execute.

If you happen to be using a POSIX 1003.1-compliant shell, you can also execute a single command in this new, non-interactive session. You have to use the `-c` command-line switch to tell the shell you're passing in a command instead of the name of a script:

Running a command in a new shell

```
sh -c ls
```

We'll get to why you would want to do this (rather than simply enter your command directly into the interactive shell) a little further down.

There is also another way to run a script from the interactive shell: you type the execute command (a single period) followed by the name of the script:

Sourcing a script

```
. MyScript
```

The difference between that and using the *sh* command is that the *sh* command starts a new process and the execute command does not. We'll look into this (and its importance) in the next section. By the way, this notation with the period is commonly referred to as *sourcing* a script.

2.2.3 Running a shell script the other way

There is also another way to execute a shell script, by making more direct use of a feature of the Unix operating system: the executable mode.

In Unix, each and every file has three different permissions (read, write and execute) that can be set for three different entities: the user who owns the file, the group that the file belongs to and "the world" (everybody else). Give the command

Code

```
-l
```

in the interactive shell to see the permissions for all files in the working directory (the column with up to nine letters, r, w and x for read write and execute, the first three for the user, the middle ones for the group, the right ones for the world). Whenever one of those entities has the "execute" permission, that entity can simply run the file as a program. To make your scripts executable by everybody, use the command

Code

```
chmod +x scriptname
```

as in

Making MyScript executable

```
chmod +x MyScript
```

You can then execute the script with a simple command like so (assuming it is in a directory that is in your *PATH*, the directories that the shell looks in for programs when you don't tell it exactly where to find the program):

Running a command in a new shell

```
MyScript
```

If this fails then the current directory is probably not in your PATH. You can force the execution of the script using

Making the shell look for your script in the current directory

```
./MyScript
```

At this command, the operating system examines the file, places it in memory and allows it to run like any other program. Of course, not every file makes *sense* as a program; a binary file is not necessarily a set of commands that the computer will recognize and a text file cannot be read by a computer at all. So to make our scripts run like this, we have to do something extra.

As we mentioned before, the Unix operating system starts by examining the program. If the program is a text file rather than a binary one (and cannot simply be executed), the operating system expects the first line of the file to name the interpreter that the operating system should start to interpret the rest of the file. The line the Unix operating system expects to find looks like this:

Code

```
#!/full path and name of interpreter
```

In our case, the following line should work pretty much everywhere:

Code

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

The Bourne Shell executable, to be found in the bin directory, which is right under the top of the filesystem tree. For example:

Bourne shell script with an explicit interpreter

Code

```
#!/bin/sh  
echo Hello World!
```

Output

```
World!
```

Executing shell scripts like this has several advantages. First it's less cumbersome than the other notations (it requires less typing). Second, it's an extra safety if you're going to pass your scripts around to others. Instead of relying on them to have the right shell, you can simply specify which shell they should use. If Bourne Shell is enough, that's what you

ask for. If you absolutely need *ksh* or *bash*, you specify that instead (mind you, it's not foolproof — other people can ignore your interpreter specification by running your script with one of the other commands that we discussed above, even if the script probably won't work if they do that).

Just as a sidenote, Unix doesn't limit this trick to shell scripts. Any script interpreter that expects its scripts to be plain-text can be specified in this way. You can use this same trick to make directly executable Perl scripts or Python, Ruby, etc. scripts as well as Bourne Shell scripts.

Note also that with the distributions using *bash* as their default shell, you can use the `#!/bin/sh` shebang and have typical *bash* syntax in your script. It will work. But for the same script to work with a distribution not using *bash* as its default shell (as example Debian), you will have to modify the script or to change its shebang to `#!/bin/bash`.

2.3 A little bit about Unix and multiprocessing

2.3.1 Why you want to know about multiprocessing

While this is not directly a book about Unix, there are some aspects of the Unix operating system that we must cover to fully understand why the Bourne Shell works the way it does from time to time.

One of the most important aspects of the Unix operating system — in fact, the main aspect that sets it apart from all other main-stream operating systems — is that the Unix Operating System is and always has been a multi-user, multi-processing operating system (this in contrast with other operating systems like MacOS and Microsoft's DOS/Windows operating systems). The Unix OS was always meant to run machines that would be used simultaneously by several users, who would all want to run at least one but possibly several programs at the same time. The ability of an operating system to divide the time of a machine's processor among several programs so that it seems to the user that they are all running at the same time is called *multiprocessing*. The Unix Operating System was designed from the core up with this possibility in mind and it has an effect on the way your shell sessions behave.

Whenever you start a new process (by running a program, for instance) on your Unix machine, the operating system provides that process with its very own operating environment. That environment includes some memory for the process to play in and it can also include certain predefined settings for all processes. Whenever you run the shell program, it is running in its own environment.

Whenever you start a new process from another process (for instance by issuing a command to your shell program in interactive mode), the new process becomes what is called a *child process* of the first process (the `ls` program runs as a child process of your shell, for instance). This is where it becomes important to know about multiprocessing and process interaction: a child process always starts with a *copy* of the environment of the parent process. This means two things:

1. A child process can *never* make changes to the operating environment of its parent—it only has access to a copy of that environment;
2. If you actually do *want* to make changes in the environment of your shell (or specifically want to avoid it), you have to know when a command runs as a child process and when it runs within your current shell; you might otherwise pick a variant that has the opposite effect of that which you want.

2.3.2 What does what

We have seen several ways of running a shell command or script. With respect to multiprocessing, they run in the following way:

Way of running	Runs as
Interactive mode command	<ul style="list-style-type: none"> • current environment for a shell command [Remark_0]) • child process for a new program
Shell non-interactive mode	child process
Dot-notation run command (<code>. MyScript</code>)	current environment
Through Unix executable permission with interpreter selection	child process

2.3.3 A useful thing to know: background processes

With the above, it may seem like multiprocessing is just a pain when doing shell scripting. But if that were so, we wouldn't *have* multiprocessing—Unix doesn't tend to keep things that aren't useful. Multiprocessing is a valuable tool in interacting with the rest of the system and one that you can use to work more efficiently. There are many books available on the benefits of multiprocessing in program development, but from the point of view of the Bourne Shell user and scripter the main one is the ability to hand off control of a process to the operating system *and still keep on working while that child process is running*. The way to do this is to run your process as a *background process*.

Running a process as a background process means telling the operating system that you want to start a process, but that it should not attach itself to any of the interactive devices (keyboard, screen, etc.) that its parent process is using. And more than that, it also tells the operating system that the request to start this child process should return immediately and that the parent process should then be allowed to continue working without having to wait for its child process to end.

This sounds complicated, but you have to keep in mind that this ability is completely ingrained in the Unix operating system and that the Bourne Shell was intended as an easy interface to the power of Unix. In other words: the Bourne Shell includes the ability to start a child process as a simple command of its own. Let's demonstrate how to do this and how useful the ability is at the same time, with an example. Give the following (rather pointless but still time consuming) command at the prompt:

```
N=0 && while [ $N -lt 10000 ]; do date >> scriptout; N=`expr $N + 1`; done
```


We'll get into what this says in later chapters; for now, it's enough to know that this command asks the system for the date and time and writes the result to a file named "scriptout". Since it then repeats this process 10000 times, it may take a little time to complete.

Now give the following command:

```
N=0 && while [ $N -lt 10000 ]; do date >> scriptout; N=`expr $N + 1`;  
done&
```

You'll notice that you can immediately resume using the shell (if you don't see this happening, hit Control+C and check that you have the extra ampersand at the end). After a while the background process will be finished and the scriptout file will contain another 10000 time reads.

The way to start a background process in Bourne Shell is to append an ampersand (&) to your command.

Remarks

[Remark_0]Actually, you can force a child process here as well -- we'll see how when we talk about command grouping

3 Environment

No program is an island unto itself. Not even the Bourne Shell. Each program executes within an *environment*, a system of resources that controls how the program executes and what external connections the program has and can make. And in which the program can itself make changes.

In this module we discuss the environment, the habitat in which each program and command lives and executes. We look at what the environment consists of, where it comes from and where it's going... And we discuss the most important mechanism that the shell has for passing data around: the *environment variable*.

3.1 The Environments

When discussing a Unix shell, you often come across the term "environment". This term is used to describe the context in which a program executes and is usually meant to mean a set of "environment variables" (we'll get to those shortly). But in fact there are two different terms that are somehow a program's environment and which often get mixed up together in "environment". The simpler one of these really is the collection of environment variables and actually is called the "environment". The second is a much wider collection of resources that influence the execution of a program and is called the *command execution environment*.

3.1.1 The command execution environment

Each running program, either started directly by the user from the shell or indirectly by another process, operates within a collection of global resources called its **command execution environment** (CEE).

A program's CEE contains important information such as the source and destination of data upon which the program can operate (also known as the standard input¹, standard output² and standard error³ handles). In addition, variables are defined that list the identity and home directory of the user or process that started the program, the hostname of the machine and the kind of terminal used to start the program. There are other variables too, but that's just some of the main ones. The environment also provides working space for the program, as well as a simple way of communicating with other, future programs, that will be run in the same environment.

1 Chapter 6 on page 67

2 Chapter 6 on page 67

3 Chapter 6 on page 67

The complete list of resources included in the shell's CEE is:

- Open files held in the parent process that started the shell. These files are inherited. This list of files includes the files accessed through redirection (such as standard input, output and error files).
- The current working directory: the "current" directory of the shell.
- The file creation mode: The default set of file permissions set when a new file is created.
- The active traps⁴.
- Shell parameters and variables set during the call to the shell or inherited from the parent process.
- Shell functions⁵ inherited from the parent process.
- Shell options set by *set* or *shopts*, or as command line options to the shell executable.
- Shell aliases (if available in your shell).
- The process id of the shell and of some processes started by the parent process.

Whenever the shell executes a command that starts a child process, that command is executed in its own CEE. This CEE inherits a copy of part of the CEE of its parent, but not the entire parent CEE. The inherited copy includes:

- Open files.
- The working directory.
- The file creation mode mask.
- Any shell variables and functions that are marked to be exported to child processes.
- Traps set by the shell.

The 'set' command

The 'set' command allows you to set or disable a number of options that are part of the CEE and influence the behavior of the shell. To set an option, set is called with a command line argument of '-' followed by one or more flags. To disable the option, set is called with '+' and then the same flag. You probably won't use these options very often; the most common use of 'set' is the call without any arguments, which produces a list of all defined names in the environment (variables and functions). Here are some of the options you might get some use out of:

+/-a

When set, automatically mark all newly created or redefined variables for export.

+/-f

When set, ignore filename metacharacters⁶.

+/-n

When set, only read commands but do not execute them.

+/-v

4 Chapter 8 on page 87

5 Chapter 7 on page 81

6 Chapter 5.2.4 on page 59

When set, causes the shell to print commands as they are read from input (verbose debugging flag).

`+/-x`

When set, causes the shell to print commands as they will be executed (debugging flag).

Again, you'll probably mostly use `set` without arguments, to inspect the list of defined variables.

3.1.2 The environment and environment variables

Part of the CEE is something that is simply called the *environment*. The environment is a collection of name/value pairs called *environment variables*. These variables technically also contain the shell functions⁷, but we'll discuss those in a separate module⁸.

An environment variable is a piece of labelled storage in the environment, where you can store anything you like as long as it fits. These spaces are called variables because you can vary what you put in them. All you need to know is the name (the label) that you used for storing the content. The Bourne shell also makes use of these "environment variables". You can make scripts that examine these variables, and those scripts can make decisions based on the values stored in the variables.

An environment variable is a name/value pair of the form

Code

```
name =variable
```

which is also the way of creating a variable. There are several ways of using a variable which we will discuss in the module on substitution⁹, but for now we will limit ourselves to the simple way: if you prepend a variable name with a `$`-character, the shell will substitute the value for the variable. So, for example:

Simple use of a variable

Code

```
$ VAR=Hello  
$ echo $VAR
```

Output

As you can see from the example above, an environment variable is sort of like a bulletin board: anybody can post any kind of value there for everybody to read (as long as they have access to the board). And whatever is posted there can be interpreted by any reader

7 Chapter 7 on page 81

8 Chapter 7 on page 81

9 <http://en.wikibooks.org/wiki/Bourne%20Shell%20Scripting%2FSubstitution>

in whatever way they like. This makes the environment variable a very general mechanism for passing data along from one place to another. And as a result environment variables are used for all sorts of things. For instance, for setting global parameters that a program can use in its execution. Or for setting a value from one shell script to be picked up by another. There are even a number of environment variables that the shell itself uses in its configuration. Some typical examples:

IFS

This variable lists the characters that the shell considers to be whitespace characters.

PATH

This variable is interpreted as a list of directories (separated by colons on a Unix system). Whenever you type the name of an executable for the shell to execute but do not include the full path of that executable, the shell will look in all of these directories *in order* to find the executable.

PS1

This variable lists a set of codes. These codes instruct your shell about what the command-line prompt in the interactive shell should look like.

PWD

The value of this variable is always the path of the working directory.

The absolute beauty of environment variables, as mentioned above, is that they just contain random strings of characters without an immediate meaning. The meaning of any variable is to be interpreted by whatever program or process reads the variable. So a variable can hold literally any kind of information and be used practically anywhere. For instance, consider the following example:

Environment variables are more flexible than you thought...

```
$ echo $CMD
$ CMD=ls
$ echo $CMD
ls
$ $CMD
bin booktemp Documents Mail mbox public_html sent
```

There's nothing wrong with setting a variable to the name of an executable, then executing that executable by calling the variable as a command.

3.1.3 Different kinds of environment variables

Although you use all environment variables the same way, there are a couple of different kinds of variables. In this section we discuss the differences between them and their uses.

Named variables

The simplest and most straightforward environment variable is the named variable. We saw it earlier: it's just a name with a value, which can be retrieved by prepending a '\$' to the name. You create and define a named variable in one go, by typing the name, an equals sign and then something that results in a string of characters.

Earlier we saw the following, simple example:

Assigning a simple value to a variable

```
$ VAR=Hello
```

This just assigns a simple value. Once a variable has been defined, we can also redefine it:

Assigning a simple value to a variable

```
$ VAR=Goodbye
```

We aren't limited to straightforward strings either. We can just as easily assign the value of one variable to another:

Assigning a simple value to a variable

```
$ VAR=$PATH
```

We can even go all-out and combine several commands to come up with a value:

Assigning a combined value to a variable

```
$ PS1= "`whoami`@`hostname -s` `pwd` \u0026gt;$ "
```

In this case, we're taking the output of the three commands 'whoami', 'hostname', and 'pwd', then we add the '\$' symbol, and some spacing and other formatting just to pad things out a bit. Whew. All that, just in one piece of labeled space. As you can see environment variables can hold quite a bit, including the output of entire commands.

There are usually lots of named variables defined in your environment, even if you are not aware of them. Try the 'set' command and have a look.

Positional variables

Most of the environment variables in the shell are named variables, but there are also a couple of "special" variables. Variables that you don't set, but whose values are automatically arranged and maintained by the shell. Variables which exist to help you out, to discover information about the shell or from the environment.

The most common of these are the positional or argument variables. Any command you execute in the shell (in interactive mode or in a script) can have command-line arguments. Even if the command doesn't actually use them, they can still be there. You pass command-line arguments to a command simply by typing them after the command, like so:

Code

```
command arg0 arg1 ...
```

This is allowed for any command. Even your own shell scripts. But say that you do this (create a shell script, then execute it with arguments); how do you access the command-line arguments from your script? This is where the positional variables come in. When the shell executes a command, it automatically assigns any command-line arguments, in order, to a set of positional variables. And these variables have numbers for names: 1 through 9, accessed through \$1 through \$9. Well, actually zero through nine; \$0 is the name of the command that was executed. For example, consider a script like this:

WithArgs.sh: A script that uses command-line arguments

```
#!/bin/sh

echo $0
echo $1
echo $2
```

And a call to this script like this:

Calling the script

Code

```
$ WithArgs.sh Hello World
```

Output

```
WithArgs.sh
```

```
Hello
```

```
World
```

As you can see, the shell automatically assigned the values 'Hello' and 'World' to \$1 and \$2 (okay, technically to the variables called 1 and 2, but it's less confusing in written text to call them \$1 and \$2). What happens if we call this script with more than two arguments?

Calling the script with more arguments

Code

```
$ WithArgs.sh Hello World Mouse Cheese
```

Output

```
WithArgs.sh
```

```
Hello
```

```
World
```

This is no problem whatsoever — the extra arguments get assigned to \$3 and \$4. But we didn't use those variables in the script, so those command-line arguments are ignored. What about the opposite case (too few arguments)?

Calling the script with too few arguments...

Code

```
$ WithArgs.sh Hello
```

Output

```
WithArgs.sh
```

```
Hello
```

Again, no problem. When the script accesses \$2, the shell simply substitutes the value of \$2 for \$2. That value is nothing in this case, so we print exactly that. In this case it's not a problem, but if your script has mandatory arguments you should check whether or not they are actually there.

What about if we want 'Hello' and 'World' to be treated as one command-line argument to be passed to the script? I.e. 'Hello World' rather than 'Hello' and 'World'? We'll get deeply into that when we start talking about quoting¹⁰, but for now just surround the words with single quotes:

Calling the script with multi-word arguments

Code

```
$ WithArgs.sh 'Hello World' 'Mouse Cheese'
```

Output

¹⁰ Chapter 5.2.5 on page 62


```
WithArgs.sh
```

```
Hello World
```

```
Mouse Cheese
```

Shifting

So what happens if you have more than nine command line arguments? Then your script is too complicated. No, but seriously: then you have a little problem. It's allowed to pass more than nine arguments, but there are only nine positional variables (in Bourne Shell at least). To deal with this situation the shell includes the **shift** command:

Code

```
shift [n] n is optional and a positive integer (default 1)
```

Shift causes the positional arguments to shift left. That is, the value of \$1 becomes the old value of \$2, the value of \$2 becomes the old value of \$3 and so on. Using **shift**, you can access all the command-line arguments (even if there are more than nine). The optional integer argument to shift is the number of positions to shift (so you can shift as many positions in one go as you like). There are a couple of things to keep in mind though:

- No matter how often you shift, \$0 always remains the original command.
- If you shift *n* positions, *n* must be lower than the number of arguments. If *n* is greater than the number of arguments, no shifting occurs.
- If you shift *n* positions, the first *n* arguments are lost. So make sure you have them stored elsewhere or you don't need them anymore!
- You cannot shift back to the right.

In the module on Control flow¹¹ we'll see how you can go through all the arguments without knowing exactly how many there are.

Other, special variables

In addition to the positional variables the Bourne Shell includes a number of other, special variables with special information about the shell. You'll probably not use these as often, but it's good to know they're there. These variables are

\$#

The number of command-line arguments to the current command (changes after a use of the **shift** command!).

\$-

The shell options currently in effect (see the 'set' command¹²).

\$?

¹¹ Chapter 5.1.3 on page 53

¹² Chapter 3.1.5 on page 25

The exit status of the last command executed (0 if it succeeded, non-zero if there was an error).

\$\$

The process id of the current process.

#!

The process id of the last background command.

\$*

All the command-line arguments. When quoted¹³, expands to all command-line arguments as a single word (i.e. "\$*" = "\$1 \$2 \$3 ...").

\$@

All the command-line arguments. When quoted¹⁴, expands to all command-line arguments quoted individually (i.e. "\$@" = "\$1" "\$2" "\$3" ...).

3.1.4 Exporting variables to a subprocess

We've mentioned it a couple of times before: Unix is a multi-user, multiprocessing operating system. And that fact is very much supported by the Bourne Shell, which allows you to start up new processes right from inside a running shell. In fact, you can even run multiple processes simultaneously next to each other (but we'll get to that a little later). Here's a simple example of starting a subprocess:

Starting a new shell from the shell

```
$ sh
```

We've also talked about the Command Execution Environment and the Environment (the latter being a collection of variables). These environments can affect how programs run, so it's very important that they cannot inadvertently affect one another. After all, you wouldn't want the screen in your shell to go blue with yellow letters simply because somebody started Midnight Commander in another process, right?

One of the things that the shell does to avoid processes inadvertently affecting one another, is environment separation. Basically this means that whenever a new (sub)process is started, it has its own CEE and environment. Of course it would be damned inconvenient if the environment of a subprocess of your shell were *completely* empty; your subprocess wouldn't have a PATH variable or the settings you chose for the format of your prompt. On the other hand there is usually a good reason NOT to have certain variables in the environment of your subprocess, and it usually has something to do with not handing off too much environment data to a process if it doesn't need that data. This was particularly true

¹³ Chapter 5.2.5 on page 62

¹⁴ Chapter 5.2.5 on page 62

when running copies of MS-DOS and versions of DOS under Windows. You only HAD a limited amount of environment space, so you had to use it carefully, or ask for more space on startup. These days in a UNIX environment the space issues aren't the same, but if all your existing variables ended up in the environment of your subprocess you might still adversely affect the running of the program that you started in that subprocess (there's really something to be said for keeping your environment lean and clean in the case of subprocesses).

The compromise between the two extremes that Stephen Bourne and others came up with is this: a subprocess has an environment which contains *copies* of the variables in the environment of its parent process — but only those variables that are marked to be *exported* (i.e. copied to subprocesses). In other words, you can have any variable copied into the environment of your subprocesses, but you have to let the shell know that's what you want first. Here's an example of the distinction:

Exported and non-exported variables

```
$ echo $PATH
/usr/local/bin:/usr/bin:/bin
$ VAR=value
$ echo $VAR
value
$ sh
$ echo $PATH
/usr/local/bin:/usr/bin:/bin
$ echo $VAR

$
```

In the example above, the `PATH` variable (which is marked for export by default) gets copied into the environment of the shell that is started within the shell. But the `VAR` variable is *not* marked for export, so the environment of the second shell doesn't get a copy.

In order to mark a variable for export you use the **export** command, like so:

Code

```
export VAR0 [VAR1 VAR2 ...]
```

As you can see, you can export as many variables as you like in one go. You can also issue the **export** command without any arguments, which will print a list of variables in the environment marked for export. Here's an example of exporting a variable:

Exporting a variable

```
$ VAR=value
$ echo $VAR
value
$ sh
$ echo $VAR

$ exit Quitting the inner shell
$ export VAR This is back in the outer shell
$ sh
$ echo $VAR
value
```

More modern shells like Korn Shell and Bash have more extended forms of **export** . A common extension is to allow for definition and export of a variable in one single command. Another is to allow you to *remove* the export marking from a variable. However, Bourne Shell only supports exporting as explained above.

3.1.5 Your profile

In the previous sections we've discussed the runtime environment of every program and command you run using the shell. We've talked about the command execution environment and at some length about the piece of it simply called "the environment", which contains environment variables. We've seen that you can define your own variables and that the system usually already has quite a lot of variables to start out with.

Here's a question about those variables that the system starts out with: where do they come from? Do they descend like manna from heaven? And on a related note: what do you do if you want to create some variables automatically every time your shell starts? Or run a program every time you log in?

Those readers who have done some digging around on other operating systems will know what I'm getting at: there's usually some way of having a set of commands executed every time you log in (or every time the system starts at least). In MS-DOS for instance there is a file called `autoexec.bat`, which is executed every time the system boots. In older versions of MS-Windows there was `system.ini`. The Bourne Shell has something similar: a file in every user's home directory called `.profile` . The `$HOME/.profile` (`HOME` is a default variable whose value is your home directory) file is a shell script like any other, which is executed automatically right after you login to a new shell session. You can edit the script to have it execute any login-commands that you like.

Each specific Unix system has its own default implementation of the `.profile` script (including none — it's allowed not to have a `.profile` script). But all of them start with some variation of this:

A basic (but typical) \$HOME/.profile

```
#!/bin/sh

if [ -f /etc/profile ]; then
  . /etc/profile
fi
PS1= "`whoami`@`hostname -s` `pwd` \$_ "
export PS1
```

This .profile might surprise you a bit: where are all those variables that get set? Most of the variables that get set for you on a typical Unix system, also get set for all other users. In order to make that possible and easily maintainable, the common solution is to have each \$HOME/.profile script start by executing another shell script: */etc/profile*. This script is a systemwide script whose contents are maintained by the system administrator (the user who logs in with username *root*). This script sets all sorts of variables and calls scripts that set even more variables and generally does everything that is necessary to provide each user with a comfortable working environment.

As you can see from the example above, you can add any personal configuration you want or need to the .profile script in your directory. The call to execute the system profile script doesn't have to be first, but you probably don't want to remove it altogether.

3.2 Multitasking and job control

With the arrival of fast computers, CPUs that can switch between multiple tasks in a very small amount of time, CPUs that can actually do multiple things at the same time and networks of multiple CPUs, having the computer perform multiple tasks at the same time has become common. Fast task switching provides the illusion that the computer really is running multiple tasks simultaneously, making it possible to effectively serve multiple users at once. And the ability to switch to a new CPU task while an old task is waiting for a peripheral device makes CPU use vastly more efficient.

In order to make use of multitasking abilities as a user, you need a command environment that supports multitasking. For example, the ability to set one program to a task, then move on and start a new program *while the old one is still running*. This kind of ability allows you as a user to do multiple things at once on the same machine, as long as those programs do not interfere. Of course, you cannot always treat each program as a "fire and forget" affair; you might have to input a password, or the program might be finished and want to tell you its results. A multitasking environment must allow you to switch between the multiple programs you have running and allow those programs to send you some sort of message if your attention is needed.

To make things a little more tangible think of something like downloading files. Usually, while you're downloading files, you want to do other stuff as well — otherwise you're going to be sitting at the keyboard twiddling your thumbs a really long time when you want to

download a whole CD worth of data. So, you start up your file downloader and feed it a list of files you want to grab. Once you've entered them, you can then tell it "Go!" and it will start off by downloading the first file and continue until it finishes the last one, or until there's a problem. The smarter ones will even try to work through common problems themselves, such as files not being available. Once it starts you get the standard shell prompt back, letting you know that you can start another program.

If you want to see how far the file downloader has gotten, simply checking the files in your system against what you have on your list will tell you. But another way to notify you is via the environment. The environment can include the files that you work with, and this can help provide information about the progress of currently running programs like that file downloader. Did it download all the files? If you check the status file, you'll see that it's downloaded 65% of the files and is just working on the last three now.

Other examples of programs that don't need their hand held are programs that play music. Quite often, once you start a program that plays music tracks, you don't WANT to tell the program "Okay, now play the next track". It should be able to do that for itself, given a list of songs to play. In fact, it should not even have to hold on to the monitor; it should allow you to start running other software right after you hit the "play" button.

In this section we will explore multitasking support within the Unix shell. We will look at enabling support, at working with multiple tasks and at the utilities that a shell has available to help you.

3.2.1 Some terminology

Before we discuss the mechanics of multitasking in the shell, let's cover some terminology. This will help us discuss the subject clearly and you'll also know what is meant when you run across these terms elsewhere.

First of all, when we start a program running on a system in a process of its own, that process with that one running instance of the program is called a *job*. You'll also come across terms like process, task, instance or similar. But the term used in Unix shells is job. Second, the ability of the shell to influence and use multitasking (starting jobs and so on) is referred to as *job control*.

Job

A process that is executing an instance of a computer program.

Job control

The ability to selectively stop (suspend) the execution of jobs and continue (resume) their execution at a later point.

Note that these terms are used this way for Unix shells. Other circumstances and other contexts might allow for different definitions. Here are some more terms you'll come across:

Job ID

An ID (usually an integer) that uniquely identifies a job. Can be used to refer to jobs for different tools and commands.

Process ID (or PID)

An ID (usually an integer) that uniquely identifies a process. Can be used to refer to processes for different tools and commands. Not the same as a Job ID.

Foreground job (or foreground process)

A job that has access to the terminal (i.e. can read from the keyboard and write to the monitor).

Background job (or background process)

A job that does not have access to the terminal (i.e. cannot read from the keyboard or write to the monitor).

Stop (or suspend)

Stop the execution of a job and return terminal control to the shell. A stopped job is not a *terminated* job.

Terminate

Unload a program from memory and destroy the job that was running the program.

3.2.2 Job control in the shell: what does it mean?

A job is a program you start within the shell. By default a new job will suspend the shell and take control over the input and output: every stroke you type at the keyboard will go to the job, as will every mouse movement. Nothing but the job will be able to write to the monitor. This is what we call a *foreground* job: it's in the foreground, clearly visible to you as a user and obscuring all other jobs in the system from view.

But sometimes that way of working is very clumsy and irritating. What if you start a long-running job that doesn't need your input (like a backup of your harddrive)? If this is a foreground process you have to wait until it's done before you can do anything else. In this situation you'd much rather start the program as a *background process*: a process that is running, but that doesn't listen to the input devices and doesn't write to the monitor. Unix supports them and the shell (with job control) allows you to start any job as a background job.

But what about a middle ground? Like that file downloader? You have to start it, log into a remote server, pick your files and start the download. Only after all that does it make sense for the job to be in the background. But how do you accomplish that if you've already started the program as a foreground job? Or how about this: you're busily writing a document in your favorite editor and you just want to step out to check your mail for a moment. Do you have to shut down the editor for that? And then, after you're done with your mail, restart it, re-open your file and find where you'd left off? That's inconvenient. No, a much better idea in both cases is simply to *suspend* the program: just stop it from running any further and return to the shell. Once you're back in the shell, you can start another program (mail) and then resume the suspended program (editor) when you're done with that — and return to the program exactly where you left it. Conversely, you can

also decide to let the suspended process (downloader) continue running, but now in the background.

When we talk about job control in the shell, we are talking about the abilities described above: to start programs in the background, to suspend running programs and to resume suspended programs, either in the foreground or in the background.

3.2.3 Enabling job control

In order to do all the things we talked about in the previous section, you need two things:

- An operating system that supports job control.
- A shell that supports job control and has job control enabled.

Unix systems support multitasking and job control. Unix was designed from the ground up to support multitasking. If you come across a person claiming to be a Unix vendor but whose software doesn't support job control, call him a fraud. Then throw his install CDs away. Then throw *him* away.

Of course you've already guessed what comes next, right? I'm going to tell you Bourne Shell supports job control. And that you can rely on the same mechanisms to work in all compatible shells. Guess what: you're not correct. The original Bourne Shell has no job control support; it was a single-tasking shell. There was an extended version of the Bourne Shell though, called *jsh* (guess what the 'j' stands for...) which had job control support. To have job control in the original Bourne Shell, you had to start this extended shell in interactive mode like this:

Code

```
jsh -i
```

Within that shell you had the job control tools we will discuss in the following sections.

Pretty much every other shell written since incorporated job control straight into the basic shell and the POSIX 1003 standard has standardized the job control utilities. So you can pretty much rely on job control being available nowadays and usually also enabled by default in interactive mode (some older shells like Korn shell had support but required you to enable that support specifically). But just in case, remember that you might have to do some extra stuff on your system to use job control. There is one gotcha though: in shell scripts, you usually include an interpreter hint that calls for a Bourne Shell (i.e. `#!/bin/sh`). Since the original Bourne Shell doesn't have job control, several modern shells turn off job control by default in non-interactive mode as a compatibility feature.

3.2.4 Creating a job and moving it around

We've already talked at length about how to create a foreground job: type a command or executable name at the prompt, hit enter, there's your job. Been there, done that, bought the T-shirt.

We've also already mentioned¹⁵ how to start a background job: by adding an ampersand at the end of the command.

Creating a background job

```
$ ls * > /dev/null &
[1] 4808
$
```

But that suddenly looks different than when we issued commands previously; there's a "[1]" and some number there. The "[1]" is the *job ID* and the number is the *process ID*. We can use these numbers to refer to the process and the job that we just created, which is useful for using tools that work with jobs. When the task finishes, you will receive a notice similar to the following:

Job done

```
[1]+  Done      ls * > /dev/null &
```

One of the tools that you use to manage jobs is the 'fg' command. This command takes a background job and places it in the foreground. For instance, consider a background job that actually takes some time to complete:

A heftier job

```
while [ $CNT -lt 200000 ]; do echo $CNT >> outp.txt; CNT=$(expr $CNT + 1); done
&
```

We haven't gotten into flow control¹⁶ yet, but this writes 200,000 integers to a file and takes some time. It also runs in the background. Say that we start this job:

Starting the job

```
$ CNT=0
$ while [ $CNT -lt 200000 ]; do echo $CNT >> outp.txt; CNT=$(expr $CNT + 1);
  done &
[1] 11246
```

The job is given job ID 1 and process ID 11246. Let's move the process to the foreground:

15 Chapter 2.3.3 on page 13

16 Chapter 5 on page 45

Moving the job to the front

```
$ fg %1
while [ $CNT -lt 200000 ]; do
    echo $CNT >> outp.txt; CNT=$(expr $CNT + 1);
done
```

The job is now running in the foreground, as you can tell from the fact that we are not returned a prompt. Now type the **CTRL+Z** keyboard combination:

Stopping the job

```
'CTRL+Z'
[1]+  Stopped                  while [ $CNT -lt 200000 ]; do
    echo $CNT >> outp.txt; CNT=$(expr $CNT + 1);
done
$
```

Did you notice the shell reports the job as stopped? Try using the 'cat' command to inspect the outp.txt file. Try it a couple of times; the contents won't change. The job is not a background job; it's not running at all! The job is suspended. Many programs recognize the **CTRL+Z** combination to suspend. And even those that don't usually have some way of suspending themselves.

3.2.5 Moving to the background and stopping in the background

Once a job is suspended, you can resume it either in the foreground or the background. To resume in the foreground you use the 'fg' command discussed earlier. You use 'bg' for the background:

Code

```
bg jobId
```

To resume our long-lasting job that writes numbers, we do the following:

Resuming the job in the background

```
$ bg %1
[1]+ while [ $CNT -lt 200000 ] do
    echo $CNT >> outp.txt; CNT=`expr $CNT + 1`;
done &
$
```

The output indicates that the job is running again. In the background this time, since we are also returned a prompt.

Can we also stop a process in the background? Sure, we can move it to the foreground and hit 'CTRL+Z'. But can we also do it directly? Well, there is no utility or command to do it. Mostly, you wouldn't *want* to do it — the whole point of putting it in the background was to let it run without bothering anybody or requiring attention. But if you really want to, you can do it like this:

Code

```
kill -SIGSTOP jobId
```

or

Code

```
kill -SIGSTOP processId
```

We'll get back to what this does exactly later¹⁷, when we talk about signals.

3.2.6 Job control tools and job status

We mentioned before that the POSIX 1003.1 standard has standardized a number of the job control tools that were included for job control in the jsh shell and its successors. We've already looked at a couple of these tools; in this section we will cover the complete list.

The standard list of job control tools consists of the following:

bg

Moves a job to the background.

fg

Moves a job to the foreground.

jobs

Lists the active jobs.

kill

Terminate a job or send a signal to a process.

CTRL+C

Terminate a job (same as 'kill' using the SIGTERM signal).

CTRL+Z

Suspend a foreground job.

wait

Wait for background jobs to terminate.

¹⁷ Chapter 8.3.2 on page 99

All of these commands can take a job specification as an argument. A job specification starts with a percent sign and can be any of the following:

%n

A job ID (n is number).

%s

The job whose command-line started with the string s.

%?s

The jobs whose command-lines *contained* the string s.

%%

The current job (i.e. the most recent one that you managed using job control).

%+

The current job (i.e. the most recent one that you managed using job control).

%-

The previous job.

We've already looked at 'bg', 'fg', and CTRL+Z and we'll cover 'kill' in a later section¹⁸. That leaves us with 'jobs' and 'wait'. Let's start with the simplest one:

Code

```
wait [job spec] ...
```

job spec is a specification as listed above.

'Wait' is what you call a *synchronization mechanism*: it causes the invoking process to suspend until all background jobs terminate. Or, if you include one or more job specifications, until the jobs you list have terminated. You use 'wait' if you have fired off multiple jobs (simply to make use of a system's parallel processing capabilities) and you cannot proceed safely until they're all done.

The 'wait' command is used in quite advanced scripting. In other words, you might not use it all that often. Here's a command that you probably will use regularly though:

Code

```
jobs [-lnprs] [job spec] ...
```

- -l lists the process IDs as well as normal output
- -n limits the output to information about jobs whose status has changed since the last status report
- -p lists only the process ID of the jobs' process group leader
- -r limits output to data on running jobs
- -s limits output to data on stopped jobs

¹⁸ Chapter 8.3.1 on page 95

- *job spec* is a specification as listed above

The *jobs* command reports information and status about active jobs (don't confuse active with running!). It is important to remember though, that this command reports on jobs and not processes. Since a job is local to a shell, the 'jobs' command cannot see across shells. The 'jobs' command is a primary source of information on jobs that you can apply job control to; for starters, you'll use this command to retrieve job IDs if you don't remember them. For example, consider the following:

Using 'jobs' to report on jobs

Code

```
$ CNT0=0
$ while [ $CNT0 -lt 200000 ]; do echo $CNT0 >> outtemp0.txt; CNT0=`expr $CNT0 +
  1`; done&
[1] 26859
$ CNT1=0
$ while [ $CNT1 -lt 200000 ]; do echo $CNT1 >> outtemp1.txt; CNT1=`expr $CNT1 +
  1`; done&
[2] 31331
$ jobs
```

Output

```
[1]-  Running                  while [ $CNT0 -lt 200000 ]; do
     echo $CNT0 >> outtemp0.txt; CNT0=`expr $CNT0 + 1`;
done &
[2]+  Running                  while [ $CNT1 -lt 200000 ]; do
     echo $CNT1 >> outtemp1.txt; CNT1=`expr $CNT1 + 1`;
done &
```

Speaking of state (which is reported by the 'jobs' command), this is a good time to talk about the different states we have. Jobs can be in any of several states, sometimes even in more than one state at the same time. The 'jobs' command reports on state directly after the job id and order. We recognize the following states:

Running

This is where the job is doing what it's supposed to do. You probably don't need to interrupt it unless you really want to give the program your personal attention (for example, to stop the program, or to find out how far through a file download has proceeded). You'll generally find that anything in the foreground that's not waiting for your attention is in this state, unless it's been put to sleep.

Sleeping

When programs need to retrieve input that's not yet available, there is no need for them to continue using CPU resources. As such, they will enter a sleep mode until another batch of input arrives. You will see more sleeping processes, since they are not as likely to be processing data at an exact moment of time.

Stopped

The **stopped** state indicates that the program was stopped by the operating system. This usually occurs when the user suspends a foreground job (e.g. pressing CTRL-Z) or if it receives SIGSTOP. At that point, the job cannot actively consume CPU resources and aside from still being loaded in memory, won't impact the rest of the system. It will resume at the point where it left off once it receives the SIGCONT signal or is otherwise resumed from the shell. The difference between sleeping and stopped is that "sleep" is a form of waiting until a planned event happens, whereas "stop" can be user-initiated and indefinite.

Zombie

A **zombie** process appears if the parent's program terminated before the child could provide its return value to the parent. These processes will get cleaned up by the *init* process but sometimes a reboot will be required to get rid of them.

3.2.7 Other job control related tools

In the last section¹⁹ we discussed the standard facilities that are available for job control in the Unix shell. However, there are also a number of non-standard tools that you might come across. And even though the focus of this book is Bourne Shell scripting (particularly as the lingua franca of Unix shell scripting) these tools are so common that we would be remiss if we did not at least mention them.

Shell commands you might come across

In addition to the tools previously discussed, there are two shell commands that are quite common: 'stop' and 'suspend'.

Code

```
stop job ID
```

The 'stop' command is a command that occurs in the shells of many System V-compatible Unix systems. It is used to suspend background processes — in other words, it is the equivalent of 'CTRL+Z' for background processes. It usually takes a job ID, like most of these commands. On systems that do not have a 'stop' command, you should be able to stop background processes by using the 'kill' command to send a SIGSTOP signal to the background process.

Code

```
suspend job ID  
suspend [-f]
```

The other command you might come across is the 'suspend' command. The 'suspend' command is a little tricky though, since it doesn't always mean the same thing on all systems and all shells. There are two variations known to the authors at this time, both of which are shown above. The first, obvious one takes a job ID argument and suspends the indicated job; really it's just the same as 'CTRL+Z'.

¹⁹ Chapter 3.2.6 on page 32

The second variant of 'suspend' doesn't take a job ID at all, which is because it doesn't suspend any random job. Rather, it suspends the execution of the shell in which the command was issued. In this variant the -f argument indicates the shell should be suspended even if it is a login shell. To resume the shell execution, send it a SIGCONT signal using the 'kill' command.

The process snapshot utility

The last tool we will discuss is the process snapshot utility, 'ps'. This utility is not a shell tool at all, but it occurs in some variant on pretty much every system and you will want to use it often. Possibly more often even than the 'jobs' tool.

The 'ps' utility is meant to report on running processes in the system. Processes, not jobs — meaning it can see across shell instances. Here's an example of the 'ps' utility:

Using the 'ps' utility

Code

```
$ ps x
```

Output

```
PID   TTY    STAT  TIME  COMMAND
32094 tty5   R     3:37:21 /bin/sh
37759 tty5   S     0:00:00 /bin/ps
```

Typical process output includes the process ID, the ID of the terminal the process is connected to (or running on), the CPU time the process has taken and the command issued to start the process. Possibly you also get a process state. The process state is indicated by a letter code, but by-and-large the same states are reported as for job reports: **R**unning, **S**leeping, **sT**opped and **Z**ombie. Different 'ps' implementations may use different or more codes though.

The main problem with writing about 'ps' is that it is not exactly standardized, so there are different command-line option sets available. You'll have to check the documentation on your system for specific details. Some options are quite common though, so we will list them here:

-a

List all processes except group leader processes.

-d

List all processes except session leaders.

-e

List all processes, without taking into account user id and other access limits.

-f

Produce a full listing as output (i.e. all reporting options).

-g *list*

Limit output to processes whose group leader process IDs are mentioned in *list* .

-l

Produce a long listing.

-p *list*

Limit output to processes whose process IDs are mentioned in *list* .

-s *list*

Limit output to processes whose session leader process IDs are mentioned in *list* .

-t *list*

Limit output to processes running on terminals mentioned in *list* .

-u *list*

Limit output to processes owned by user accounts mentioned in *list* .

The 'ps' tool is useful for monitoring jobs across shell instances and for discovering process IDs for signal transmission.

4 Variable Expansion

In the Environment¹ module we introduced the idea of an environment variable as a general way of storing small pieces of data. In this module we take an in-depth look at using those variables: 'variable expansion', 'parameter substitution' or just 'substitution'.

4.1 Substitution

The reason that using a variable is called substitution is that the shell literally replaces each reference to any variable with its value. This is done while evaluating the command-line, which means that the variable substitution is made *before* the command is actually executed.

The simplest way of using a variable is the way we've already seen, prepending the variable name with a '\$'. So for instance:

Simple use of a variable
Code

```
$ USER=JoeSixpack  
$ echo $USER
```

Output

Of course, once the substitution is made the result is still just the text that was in the variable. The interpretation of that text is still done by whatever program is run. So for example:

Variables do not make magic
Code

```
$ USER=JoeSixpack  
$ ls $USER
```

Output

¹ Chapter 3 on page 15

```
: cannot access JoeSixpack: No such file or directory
```

Basic variable expansion is already quite flexible. You can use it as described above, but you can also use variables to create longer strings. For instance, if you want to set the log directory for an application to the "log" directory in your home directory, you might fill in the setting like this:

```
$HOME/log
```

And if you're going to use that setting more often, you might want to create your own variable like this:

```
LOGDIR=$HOME/log
```

And, of course, if you want specific subdirectories for logs for different programs, then the logs for the Wyxliz application go into directory

```
$LOGDIR/Wyxliz/
```

4.2 Substitution forms

The Bourne Shell has a number of different syntaxes for variable substitution, each with its own meaning and use. In this section we examine these syntaxes.

4.2.1 Basic variable substitution

We've already talked at length about basic variable substitution: you define a variable, stick a '\$' in front of it, the shell substitutes the value for the variable. By now you're probably bored of hearing about it.

But we've not talked about one situation that you might run into with basic variable substitution. Consider the following:

Adding some text to a variable's value
Code

```
$ ANIMAL=duck  
$ echo One $ANIMAL, two $ANIMALs
```

Output

```
duck, two
```

So what went wrong here? Well, obviously the shell substituted nothing for the ANIMAL variable, but why? Because with the extra 's' the shell thought we were asking for the non-existent ANIMALs variable. But what gives there? We've used variables in the middle of strings before (as in '/home/ANIMAL/logs'). But an 's' is not a '/': an 's' can be a valid

part of a variable name, so the shell cannot tell the difference. In cases where you explicitly have to separate the variable from other text, you can use braces:

Adding some text to a variable's value, take II
Code

```
$ ANIMAL=duck
$ echo One $ANIMAL, two ${ANIMAL}s
```

Output

```
duck, two ducks
```

Both cases (with and without the braces) count as basic variable substitution and the rules are exactly the same. Just remember not to leave any spaces between the braces and the variable name.

4.2.2 Substitution with a default value

Since a variable can be empty, you'll often write code in your scripts to check that mandatory variables actually have a value. But in the case of optional variables it is usually more convenient not to check, but to use a default value for the case that a variable is not defined. This case is actually so common that the Bourne Shell defines a special syntax for it: the dash. Since a dash can mean other things to the shell as well, you have to combine it with braces — the final result looks like this:

Code

```
${ <nowiki/>varname [-default <nowiki/>]}
  varname is the name of the variable
```

and *default* is the value used if *varname* is not defined

Again, don't leave any spaces between the braces and the rest of the text. The way to use this syntax is as follows:

Default values

Code

```
$ THIS_ONE_SET=Hello
$ echo $THIS_ONE_SET ${THIS_ONE_NOT:-World}
```

Output

```
World
```

Compare that to this:

Default not needed

Code

```
$ TEXT=aaaaahhhhhhh  
$ echo Say ${TEXT:-bbbbbbbb}
```

Output

```
aaaaahhhhhhh
```

Interestingly, the colon is optional; so `${VAR:-default}` has the same result as `${VAR-default}`.

4.2.3 Substitution with default assignment

As an extension to default values, there's a syntax that not only supplies a default value but assigns it to the unset variable at the same time. It looks like this:

Code

```
`${ <nowiki/>varname [:]=default <nowiki/>}  
  varname is the name of the variable
```

and *default* is the value used **and assigned** if *varname* is not defined

As usual, avoid spaces in between the braces. Here's an example that demonstrates how this syntax works:

Default value assignment

```
$ echo $NEWVAR  
  
$ echo ${NEWVAR:=newval}  
newval  
$ echo $NEWVAR  
newval
```

As with the default value syntax, the colon is optional.

4.2.4 Substitution for actual value

This substitution is sort of a quick test to see if a variable is defined (and that's usually what it's used for). It's sort of the reverse of the default value syntax and looks like this:

Code

```
`${ <nowiki/>varname [:+substitute <nowiki/>}  
  varname is the name of the variable
```

and *substitute* is the value used if *varname* is defined

This syntax returns the substitute value if the variable **is** defined. That sounds counterintuitive at first, especially if you ask what is returned if the variable is **not** defined — and learn that the value is nothing. Here’s an example:

Actual value substitution

```
$ echo ${NEWVAR:+newval}

$ NEWVAR=oldval
$ echo ${NEWVAR:+newval}
newval
```

So what could possibly be the use of this notation? Well, it’s used often in scripts that have to check whether lots of variables are set or not. In this case the fact that a variable has a value means that a certain option has been activated, so you’re interested in knowing that the variable *has* a value, not what that value is. It looks sort of like this (pseudocode, this won’t actually work in the shell):

Default value assignment

```
if ${SPECIFIC_OPTION_VAR:+optionset} == optionset then ...
```

Of course, in this notation the colon is optional as well.

4.2.5 Substitution with value check

This final syntax is sort of a debug check to check whether or not a variable is set. It looks like this:

Code

```
${ <nowiki/>varname [:]?message <nowiki/>}
  varname is the name of the variable
```

and *message* is the printed if *varname* is not defined

With this syntax, if the variable is defined everything is okay. Otherwise, the message is printed and the command or script exits with a non-zero exit status. Or, if there is no message, the text "parameter null or not set" is printed. As usual the colon is optional and you may not have a space between the colon and the variable name.

You can use this syntax to check that the mandatory variables for your scripts have been set and to print an error message if they are not.

Default value assignment

```
$ echo ${SOMEVAR:?has not been set}
-sh: SOMEVAR: has not been set
$ echo ${SOMEVAR:?}
-sh: SOMEVAR: parameter null or not set
```

5 Control flow

So far we've talked about basics and theory. We've covered the different shells available and how to get shell scripts running in the Bourne Shell. We've talked about the Unix environment and we've seen that you have variables that control the environment and that you can use to store values for your own use. What we haven't done yet, though, is actually *done* anything. We haven't made the system act, jump through hoops, fetch the newspaper or do the dishes.

In this chapter it's time to get serious. In this chapter we talk programming — how to write programs that make decisions and execute commands. In this chapter we talk about control flow and command execution.

5.1 Control Flow

What is the difference between a program launcher and a command shell? Why is Bourne Shell a tool that has commanded power and respect the world over for decades and not just a stupid little tool you use to start *real* programs? Because Bourne Shell is not just an environment that launches programs: Bourne Shell is a fully programmable environment with the power of a full programming language at its command. We've already seen in Environment¹ that Bourne Shell has variables in memory. But Bourne Shell can do more than that: it can make decisions and repeat commands. Like any real programming language, Bourne Shell has *control flow*, the ability to steer the computer.

5.1.1 Test: evaluating conditions

Before we can make decisions in shell scripts, we need a way of evaluating conditions. We have to be able to check the state of certain affairs so that we can base our decisions on what we find.

Strangely enough the actual shell doesn't include any mechanism for this. There is a tool for exactly this purpose called **test** (and it was literally created for use in shell scripts), but nevertheless it is not strictly part of the shell. The 'test' tool evaluates conditions and returns either **true** or **false**, depending on what it finds. It returns these values in the form of an exit status (in the \$? shell variable): a zero for **true** and something else for **false**. The general form of the test command is

Code

```
test condition
```

¹ Chapter 3 on page 15

as in

A test for string equality

```
test "Hello World" = "Hello World"
```

This test for the equality of two strings returns an exit status of zero. There is also a shorthand notation for 'test' which is usually more readable in scripts, namely square brackets:

Code

```
[ condition ]
```

Note the spaces between the brackets and the actual condition – don't forget them in your own scripts. The equivalent of the example above in shorthand is

A shorter test for string equality

```
[ "Hello World" = "Hello World" ]
```

'Test' can evaluate a number of different kinds of conditions, to fit with the different kinds of tests that you're likely to want to carry out in a shell script. Most specific shells have added on to the basic set of available conditions, but Bourne Shell recognizes the following:

File conditions

-b *file*

file exists and is a block special file

-c *file*

file exists and is a character special file

-d *file*

file exists and is a directory

-f *file*

file exists and is a regular data file

-g *file*

file exists and has its set-group-id bit set

-k *file*

file exists and has its sticky bit set

-p *file*

file exists and is a named pipe

-r *file*

file exists and is readable

-s *file*

file exists and its size is greater than zero

-t [**n**]

The open file descriptor with number *n* is a terminal device; *n* is optional, default 1

-u *file*

file exists and has its set-user-id bit set

-w *file*

file exists and is writable

-x *file*

file exists and is executable

String conditions

-n *s*

s has non-zero length

-z *s*

s has zero length

s0 = *s1*

s0 and *s1* are identical

s0 != *s1*

s0 and *s1* are different

s

s is not null (often used to check that an environment variable has a value)

Integer conditions

n0 -eq *n1*

n0 is equal to *n1*

n0 -ge *n1*

n0 is greater than or equal to *n1*

n0 -gt *n1*

n0 is strictly greater than *n1*

n0 -le *n1*

n0 is less than or equal to *n1*

n0 -lt *n1*

$n0$ is strictly less than $n1$

n0 -ne $n1$

$n0$ is not equal to $n1$

Finally, conditions can be combined and grouped:

$\backslash(B \backslash)$

Parentheses are used for grouping conditions (don't forget the backslashes). A grouped condition (B) is true if B is true.

! B

Negation; is true if B is false.

$B0$ -a $B1$

And; is true if $B0$ and $B1$ are both true.

$B0$ -o $B1$

Or; is true if either $B0$ or $B1$ is true.

5.1.2 Conditional execution

Okay, so now we know how to evaluate some conditions. Let's see how we can make use of this ability to do some programming.

All programming languages need two things: a form of decision making or conditional execution and a form of repetition or looping. We'll get to looping later, for now let's concentrate on conditional execution. Bourne Shell supports two forms of conditional execution, the **if**-statement and the **case**-statement.

The **if**-statement is the most general of the two. Its general form is

Code

```
if command-list

then command-list

elif command-list

then command-list

...

else command-list

fi
```

This command is to be interpreted as follows:

1. The command list following the **if** is executed.

2. If the last command returns a status zero, the command list following the first **then** is executed and the statement terminates after completion of the last command in this list.
3. If the last command returns a non-zero status, the command list following the first **elif** (if there is one) is executed.
4. If the last command returns a status zero, the command list following the next **then** is executed and the statement terminates after completion of the last command in this list.
5. If the last command returns a non-zero status, the command list following the next **elif** (if there is one) is executed and so on.
6. If no command list following the **if** or an **elif** terminates in a zero status, the command list following the **else** (if there is one) is executed.
7. The statement terminates. If the statement terminated without an error, the return status is zero.

It is interesting to note that the **if**-statement allows command lists everywhere, including in places where conditions are evaluated. This means that you can execute as many compound commands as you like before reaching a decision point. The only command that affects the outcome of the decision is the last one executed in the list.

In most cases though, for the sake of readability and maintainability, you will want to limit yourself to one command for a condition. In most cases this command will be a use of the 'test' tool.

Example of a simple if statement

Code

```
if [ 1 -gt 0 ]
then
  echo YES
fi
```

Output

Example of an if statement with an else clause

Code

```
if [ 1 -le 0 ]
then
  echo YES
else
  echo NO
fi
```

Output

Example of a full if statement with an else clause and two elifs

Code

```
rank=captain

if [ "$rank" = colonel ]
then
    echo Hannibal Smith
elif [ "$rank" = captain ]
then
    echo Howling Mad Murdock
elif [ "$rank" = lieutenant ]
then
    echo Templeton Peck
else
    echo B.A. Baracus
fi
```

Output

Mad Murdock

The **case** -statement is sort of a special form of the **if** -statement, specialized in the kind of test demonstrated in the last example: taking a value and comparing it to a fixed set of expected values or patterns. The case statement is used very frequently to evaluate command line arguments to scripts. For example, if you write a script that uses switches to identify command line arguments, you know that there are only a limited number of legal switches. The **case** -statement is an elegant alternative to a potentially messy **if** -statement in such a case.

The general form of the case statement is

Code

```
case value in

    pattern0 ) command-list-0 ;;

    pattern1 ) command-list-1 ;;

    ...

esac
```

The value can be any value, including an environment variable. Each pattern is a regular expression and the command list executed is the one for the first pattern that matches the value (so make sure you don't have overlapping patterns). Each command list must end with a double semicolon. The return status is zero if the statement terminates without syntax errors.

The last 'if'-example revisited

Code

```
rank=captain

case $rank in
```

```

colonel) echo Hannibal Smith;;
captain) echo Howling Mad Murdock;;
lieutenant) echo Templeton Peck;;
sergeant) echo B.A. Baracus;;
*) echo OOPS;;
esac

```

Output

Mad Murdock

If versus case: what is the difference?

So what exactly is the difference between the **if** - and **case** -statements? And what is the point of having two statements that are so similar? Well, the technical difference is this: the **case** -statement works off of data available to the shell (like an environment variable), whereas the **if** -statement works off the exit status of a program or command. Since fixed values and environment variables depend on the shell but the exit status is a concept general to the Unix system, this means that the **if** -statement is more general than the **case** -statement.

Let's look at a slightly larger example, just to put the two together and compare:

```

#!/bin/sh

if [ "$2" ]
then
    sentence="$1 is a"
else
    echo Not enough command line arguments! >&2
    exit 1
fi

case $2 in
    fruit|veg*) sentence="$sentence vegetarian!";;
    meat) sentence="$sentence meat eater!";;
    *) sentence="${sentence}n omnivore!";;
esac

echo $sentence

```

Note that this is a shell script and that it uses positional variables to capture command-line arguments. The script starts with an **if** -statement to check that we have the right number of arguments – note the use of 'test' to see if the value of variable \$2 is not null and the exit status of 'test' to determine how the **if** -statement proceeds. If there are enough arguments, we assume the first argument is a name and start building the sentence that is the result of our script. Otherwise we write an error message (to stderr, the place to write errors; read

all about it in Files and streams²) and exit the script with a non-zero return value. Note that this `else` statement has a command list with more than one command in it.

Assuming we got through the `if`-statement without trouble, we get to the `case`-statement. Here we check the value of variable `$2`, which should be a food preference. If that value is either fruit or something starting with veg, we add a claim to the script result that some person is a vegetarian. If the value was exactly meat, the person is a meat eater. Anything else, he is an omnivore. Note that in that last case pattern clause we have to use curly braces in the variable substitution; that's because we want to add a letter `n` directly onto the existing value of sentence, without a space in between.

Let's put the script in a file called `'preferences.sh'` and look at the effect of some calls of this script:

Calling the script with different effects

```
$ sh preferences.sh
Not enough command line arguments!
$ sh preferences.sh Joe
Not enough command line arguments!
$ sh preferences.sh Joe fruit
Joe is a vegetarian!
$ sh preferences.sh Joe veg
Joe is a vegetarian!
$ sh preferences.sh Joe vegetables
Joe is a vegetarian!
$ sh preferences.sh Joe meat
Joe is a meat eater!
$ sh preferences.sh Joe meat potatoes
Joe is a meat eater!
$ sh preferences.sh Joe potatoes
Joe is an omnivore!
```

5.1.3 Repetition

In addition to conditional execution mechanisms every programming language needs a means of repetition, repeated execution of a set of commands. The Bourne Shell has several mechanisms for exactly that: the `while`-statement, the `until`-statement and the `for`-statement.

² Chapter 6 on page 67

The while loop

The **while** -statement is the simplest and most straightforward form of repetition statement in Bourne shell. It is also the most general. Its general form is this:

Code

```
while command-list1
do command-list2
done
```

The **while** -statement is interpreted as follows:

1. Execute the commands in command list 1.
2. If the exit status of the last command is non-zero, the statement terminates.
3. Otherwise execute the commands in command list 2 and go back to step 1.
4. If the statement does not contain a syntax error and it ever terminates, it terminates with exit status zero.

Much like the **if** -statement, you can use a full command list to control the **while** -statement and only the last command in that list actually controls the statement. But in reality you will probably want to limit yourself to one command and, as with the **if** -statement, you will usually use the 'test' program for that command.

A while loop that prints all the values between 0 and 10

Code

```
counter=0

while [ $counter -lt 10 ]
do
    echo $counter
    counter=`expr $counter + 1`
done
```

Output

```
0
1
2
3
4
5
6
7
8
9
```

The **while** -statement is commonly used to deal with situations where a script can have an indeterminate number of command-line arguments, by using the **shift** command and the special '\$#' variable that indicates the number of command-line arguments:

Printing all the command-line arguments

```
#!/bin/sh

while [ $ -gt 0 ]
do
    echo $1
    shift
done
```

The until loop

The **until** -statement is also a repetition statement, but it is sort of the semantic opposite of the **while** -statement. The general form of the **until** -statement is

Code

```
until command-list1

do command-list2

done
```

The interpretation of this statement is almost the same as that of the **while** -statement. The only difference is that the commands in command list 2 are executed as long as the last command of command list 1 returns a non-zero status. Or, to put it more simply: command list 2 is executed as long as the condition of the loop is *not* met.

Whereas **while** -statements are mostly used to establish some effect ("repeat until done"), **until** -statements are more commonly used to poll for the existence of some condition or to wait until some condition is met. For instance, assume some process is running that will write 10000 lines to a certain file. The following **until** -statement waits for the file to have grown to 10000 lines:

Waiting for myfile.txt to grow to 10000 lines

```
until [ $lines -eq 10000 ]
do
    lines=`wc -l dates | awk '{print $1}'`
    sleep 5
done
```

The for loop

In the section on Control flow³, we discussed that the difference between **if** and **case** was that the first depended on command exit statuses whereas the second was closely linked to data available in the shell. That kind of pairing also exists for repetition statements: **while** and **until** use command exit statuses and **for** uses data explicitly available in the shell.

The **for** -statement loops over a fixed, finite set of values. Its general form is

Code

```
for name in w1 w2 ...
```

```
do command-list
```

```
done
```

This statement executes the command list for each value named after the 'in'. Within the command list, the "current" value w_1 is available through the variable *name*. The value list must be separated from the 'do' by a semicolon or a newline. And the command list must be separated from the 'done' by a semicolon or a newline. So, for example:

A for loop that prints some values

Code

```
for myval in Abel Bertha Charlie Delta Easy Fox Gumbo Henry India
do
  echo $myval Company
done
```

Output

```
Company
Bertha Company
Charlie Company
Delta Company
Easy Company
Fox Company
Gumbo Company
Henry Company
India Company
```

The **for** statement is used a lot to loop over command line arguments. For that reason the shell even has a shorthand notation for this use: if you leave off the 'in' and the values part, the command assumes $*$ as the list of values. For example:

Using for to loop over command line arguments

Code

³ Chapter 5.1.3 on page 53

```
#!/bin/sh

for arg
do
  echo $arg
done
```

Output

```
$ sh loop_args.sh A B C D
```

```
A
```

```
B
```

```
C
```

```
D
```

5.2 Command execution

In the last section on Control Flow⁴ we discussed the major programming constructs and control flow statements offered by the Bourne Shell. However, there are lots of other syntactic constructs in the shell that allow you to control the way commands are executed and to embed commands in other commands. In this section we discuss some of the more important ones.

5.2.1 Command joining

Earlier, we looked at the `if`-statement as a method of conditional execution. In addition to this expansive statement the Bourne Shell also offers a method of directly linking two commands together and making the execution of one of them conditional on the outcome (the exit status) of the other. This is useful for making quick, inline decisions on command execution. But you probably wouldn't want to use these constructs in a shell script or for longer command sequences, because they aren't the most readable.

You can join commands together using the `&&` and `||` operators. These operators (which you might recognize as borrowed from the C programming language) are *short circuiting* operators: they make the execution of the second command dependent on the exit status of the first and so can allow you to avoid unnecessary command executions.

The `&&` operator joins two commands together and only executes the second if the exit status of the first is zero (i.e. the first command "succeeds"). Consider the following example:

4 Chapter 5.1 on page 45

Attempt to create a file and delete it again *if the creation succeeds*

```
echo Hello World > tempfile.txt && rm tempfile.txt
```

In this example the deletion would be pointless if the file creation fails (because the file system is read-only, say). Using the `&&` operator prevents the deletion from being attempted if the file creation fails. A similar – and possibly more useful – example is this:

Check if a file exists and make a backup copy if it does

```
test -f my_important_file && cp my_important_file backup
```

In contrast to the `&&` operator, the `||` operator executes the second command only if the exit status of the first command is *not* zero (i.e. it fails). Consider the following example:

Make sure we do not overwrite a file; create a new file only if it doesn't exist yet

```
test -f my_file || echo Hello World > my_file
```

For both these operators the exit status of the joined commands is the exit status of the last command that actually got executed.

5.2.2 Command grouping

You can join multiple commands into one command list by joining them using the `;` operator, like so:

Create a directory and change into it all in one go

```
mkdir newdir;cd newdir
```

There is no conditional execution here; all commands are executed, even if one of them fails.

When joining commands into a command list, you can group the commands together for clarity and some special handling. There are two ways of grouping commands: using curly braces and using parentheses.

Grouping using curly braces is just for clarity; it doesn't add any semantics to joining using semicolons. The only differences between with braces and without are that if you use braces you *must* insert an extra semicolon after your command list and you have to remember to

put spaces between the braces and your command list or the shell won't understand what you mean. Here's an example:

Create a directory and change into it all in one go, grouped with curly braces

```
{ mkdir newdir;cd newdir; }
```

The parentheses are far more interesting. When you group a command list with parentheses, it is executed... in a separate process. This means that whatever you do in the command list doesn't affect the environment in which you gave the command. Consider the example above again, with braces and parentheses:

Create a directory and change into it all in one go, grouped with curly braces

Code

```
/home$ { mkdir newdir;cd newdir; }
```

Output

```
/home/newdir$
```

Create a directory and change into it all in one go, grouped with parentheses

Code

```
/home$ (mkdir newdir;cd newdir)
```

Output

```
/home$
```

Here's another one:

Creating shell variables in the current and in a new environment

Code

```
$ VAR0=A  
$ (VAR1=B)  
$ echo \"$VAR0\" \"$VAR1\"
```

Output

```
"A" ""
```

5.2.3 Command substitution

In the chapter on Environment⁵ we talked about variable substitution. The Bourne Shell also supports *command substitution*. This is sort of like variable substitution, but instead of a variable being replaced by its value a command is replaced by its output. We saw an example of this earlier when discussing the **while** -statement, where we assigned the outcome of an arithmetic expression evaluation to an environment variable.

Command substitution is accomplished using either of two notations. The original Bourne Shell used grave accents (*'command '*), which is still generally supported by most shells. Later on the POSIX 1003.1 standard added the **\$(command)** notation. Consider the following examples:

Making a daily backup (old-skool)

```
cp myfile backup/myfile-`date`
```

Making a daily backup (POSIX 1003.1)

```
cp myfile backup/myfile-$(date)
```

5.2.4 Regular expressions and metacharacters

Usually, in the day-to-day tasks that you do with your shell, you will want to be explicit and exact about which files you want to operate on. After all, you want to delete a specific file and not a random one. And you want to send your network communications to the network device file and not to the keyboard.

But there are times, especially in scripting, when you will want to be able to operate on more than one file at a time. For instance, if you have written a script that makes a regular backup of all the files in your home directory whose names end in ".dat". If there are a lot of those files, or there are more being made each day with new names each time, then you do not want to have to name all those files explicitly in your backup script.

We have also seen another example of not wanting to be too explicit: in the section on the **case** -statement, there is an example where we claim that somebody is a vegetarian if he likes fruit or anything starting with "veg". We could have included all sorts of options there and been explicit (although there are an infinite number of words you can make that start with "veg"). But we used a *pattern* instead and saved ourselves a lot of time.

For exactly these cases the shell supports a (limited) form of *regular expressions*: patterns that allow you to say something like "I mean every string, every sequence of characters,

⁵ Chapter 3 on page 15

that looks sort of like this”. The shell allows you to use these regular expressions anywhere (although they don’t always make sense — for example, it makes no sense to use a regular expression to say *where* you want to copy a file). That means in shell scripts, in the interactive shell, as part of the **case** -statement, to select files, general strings, anything.

In order to create regular expressions you use one or more *metacharacters* . Metacharacters are characters that have special meaning to the shell and are automatically recognized as part of regular expressions. The Bourne shell recognizes the following metacharacters:

Matches any string.

?

Matches any single character.

[characters]

Matches any character enclosed in the angle brackets.

[!characters]

Matches any character *not* enclosed in the angle brackets.

pat0 |pat1

Matches any string that matches *pat0* **or** *pat1* (only in **case** -statement patterns!)

Here are some examples of how you might use regular expressions in the shell:

List all files whose names end in “.dat”

```
ls *.dat
```

List all files whose names are “file-” followed by two characters followed by “.txt”

```
ls file-?.txt
```

Make a backup copy of all text files, with a datestamp

```
for i in *.txt; do cp $i backup/${i}`date +%Y%m%d`; done
```

List all files in the directories Backup0 and Backup1

```
ls Backup[01]
```

List all files in the *other* backup directories

```
ls Backup[!01]
```

Execute all shell scripts whose names start with "myscript" and end in ".sh"

```
myscript*.sh
```

Regular expressions and hidden files

When selecting files, the metacharacters match all files except files whose names start with a period ("."). Files that start with a period are either special files or are assumed to be configuration files. For that reason these files are semi-protected, in the sense that you cannot just pick them up with the metacharacters. In order to include these files when selecting with regular expressions, you must include the leading period explicitly. For example:

Listing all files whose names start with a period

Code

```
/home$ ls .*
```

Output

```
.
..
.profile
```

The example above shows a listing of period files. In this example the listing includes '.profile', which is the user configuration file for the Bourne Shell. It also includes the special directories '.' (which means "the current directory") and '..' (which is the parent directory of the current directory). You can address these special directories like any other. So for instance

Code

```
.
```


is the same semantically as just 'ls' and

Code

```
..
```

changes your working directory to the parent directory of the directory that was your working directory before.

5.2.5 Quoting

When you introduce special characters like the metacharacters discussed in the previous section, you automatically get into situations when you really don't want those special characters evaluated. For example, assume that you have a file whose name includes an asterisk (*). How would you address that file? For example:

Metacharacters in file names can cause problems

Code

```
echo Test0 > asterisk*.file
echo Test1 > asteriskSTAR.file
cat asterisk*.file
```

Output

```
0
```

```
Test1
```

Clearly what is needed is a way of temporarily turning metacharacters off. The Bourne Shell built-in quoting mechanisms do exactly that. In fact, they do a lot more than that. For instance, if you have a file name with spaces in it (so that the shell cannot tell the different words in the file name belong together) the quoting mechanisms will help you deal with that problem as well.

There are three quoting mechanisms in the Bourne Shell:

\

backslash, for single character quoting.

”

single quotes, to quote entire strings.

””

double quotes, to quote entire strings but still allow for some special characters.

The simplest of these is the backslash, which quotes the character that immediately follows it. So, for example:

Echo with an asterisk

Code

```
echo *
```

Output

```
asterisk*.file asterisking.file backup biép.txt BLAAT.txt conditional1.sh condit
ional1.sh~ conditional.sh conditional.sh~ dates error_test.sh error_test.sh~ fil
e with spaces.txt looping0.sh looping1.sh out_nok out_ok preferences.sh pre
ferences.sh~ test.txt Echoing an asterisk
```

Code

```
echo \*
```

Output

```
*
```

So the backslash basically disables special character interpretation for the duration of one character. Interestingly, the newline character is also considered a special character in this context, so you can use the backslash to split commands to the interpreter over multiple lines. Like so:

A multiline command

Code

```
echo This is a \  
>very long command!
```

Output

```
is a very long command!
```

The backslash escape also works for file names with spaces:

Difficult file to list...

Code

```
ls file with spaces.txt
```

Output

```
ls: cannot access file: No such file or directory
ls: cannot access with: No such file or directory
ls: cannot access spaces.txt: No such file or directory Listing the file using escapes
```

Code

```
ls file\_with\_spaces.txt
```

Output

```
with spaces.txt
```

But what if you want to pass a backslash to the shell? Well, think about it. Backslash disables interpretation of one character, so if you want to use a backslash for anything else... then `'\'` will do it!

So we've seen that a backslash allows you to disable special character interpretation for a single character by quoting it. But what if you want to quote a lot of special characters all at once? As you've seen above with the filename with spaces, you *can* quote each special character separately, but that gets to be a drag really quickly. Usually it's quicker, easier and less error-prone simply to quote an entire string of characters in one go. To do exactly that you use single quotes. Two single quotes quote the entire string they surround, disabling interpretation of all special characters in that string — with the exception of the single quote (so that you can *stop* quoting as well). For example:

Quoting to use lots of asterisks

Code

```
echo '*****'
```

Output

```
*****
```

So let's try something. Let's assume that for some strange reason we would like to print three asterisks (`'''***'''`), then a space, then the current working directory, a space and three more asterisks. We know we can disable metacharacter interpretation with single quotes so this should be no biggy, right? And to make life easy, the built-in command `'pwd'` prints the working directory, so this is really easy:

Printing the working directory with decorations

Code

```
echo '*** `pwd` ***'
```

Output

```
*** `pwd` ***
```

So what went wrong? Well, the single quotes disable interpretation of all special characters. So the grave accents we used for the command substitution didn't work! Can we make it work a different way? Like by using the Path of Working Directory environmentvariable (`$PWD`)? Nope, the `$`-character won't work either.

This is a typical Goldilocks problem. We want to quote some special characters, but not all. We could use backslashes, but that doesn't do enough to be convenient (it's too cold). We can use single quotes, but that kills too many special characters (it's too hot). What we need is quoting that's *juuuust riiiiight*. More to the point, what we want (and more often than you think) is to disable all special character interpretation except variable and command substitution. Because this is a common desire the shell supports it through a separate quoting mechanism: the double quote. The double quote disables all special character interpretation except the grave accent (command substitution), the `$` (variable substitution) and the double quote (so you can *stop* quoting). So the solution to our problem above is:

Printing the working directory with decorations, take II
Code

```
echo "*** `pwd` ***"
```

Output

```
*** /home/user/examples ***
```

By the way, we actually cheated a little bit above for educational purposes (hey, you try coming up with these examples); we could also have solved the problem like this:

Printing the working directory with decorations, alternative
Code

```
echo '***' `pwd` '***'
```

Output

```
*** /home/user/examples ***
```


6 Files and streams

6.1 The Unix world: one file after another

When you think of a computer and everything that goes with it, you usually come up with a mental list of all sorts of different things:

- The computer itself
- The monitor
- The keyboard
- The mouse
- Your hard drive with your files and directories on it
- The network connection leading to the Internet
- The printer
- The DVD player
- et cetera

Here's a surprise for you: Unix doesn't have any of these things. Well, almost. Unix certainly has *files*. Unix has endless reams of files. And since Unix has files, it also has a concept of "between files" (think of it this way: if your universe consists only of boxes, you automatically know about spaces where there are no boxes as well). But Unix knows nothing else than that. Everything in the whole (Unix) universe is a file.

Everything is a file. Even things that are really weird things to think of as files, are files. Your (data) files are files. Your directories are files. Your hard drive is a file. Your *keyboard*, *monitor* and *printer* are files. Yes, really: your keyboard is a read-only file of infinite size. Your monitor and printer are infinitely sized write-only files. Your network connection is a read/write file.

At this point you're probably asking: *Why?* Why would the designers of the Unix system have come up with this madness? Why is everything a file? The answer is: because if everything is a file, you can **treat** everything like a file. Or, put a different way, you can treat everything in the Unix world the same way. And, as we will see shortly, that means you can also *combine* virtually everything using file operations.

Before we move on, here's an extra level of weirdness for you: everything in Unix is a file. Including the *processes* that run programs. Effectively this means that running programs are also files. Including the interactive shell session that you've been running to practice scripting in. Yes, really, that text screen with the blinking cursor is also a file. And we can prove it too. You might recall that in the chapter on Running Commands¹ we mentioned you can exit the shell using the Ctrl+d key combination. Because that combination produces the Unix character for... that's right, end-of-file!

¹ Chapter 2.1.3 on page 8

6.2 Streams: what goes between files

As we mentioned in the previous section, everything in Unix is a file -- except that which sits *between* files. Between files Unix defines a mechanism that allows data to move, bit by bit, from one file to another: the **stream**. A stream is literally what it sounds like: a little river of bits pouring from one file into another. Although actually a bridge would probably have been a better name because unlike a stream (which is a constant flow of water) the flow of bits between files need not be constant, or even used at all.

6.2.1 The standard streams

Within the Unix world it is a general convention that each file is connected to at least three streams (that's because that turned out to be the most useful number for those files that are processes, or running programs). There can be more and in fact each file can cause itself to be connected to any number of streams (a program can print and open a network connection, for instance). But there are three basic streams available to all files, even though they may not always be useful or used. These streams are called the "standard" streams:

Standard in (stdin)

the standard stream for input into a file.

Standard out (stdout)

the standard stream for output out of a file.

Standard error (stderr)

the standard stream for error output from a file.

As you can probably tell, these streams are very geared towards those files that are actually processes of the system. In fact many programming languages (like C, C++, Java and Pascal) use exactly these conventions for their standard I/O operations. And since the Unix operating system family includes them in the core of the system definition, these streams are also central to the Bourne Shell.

6.2.2 Getting hold of the standard streams in your scripts

So now we know that there's a general mechanism for basic input and output in Unix; but how do you get hold of these streams in a script? What do you have to do to hook your script up to the standard out, or read from the standard in? Well, the happy answer is: nothing. Your scripts are automatically connected to the standard in, out and error stream of the process that is running them. When you read input, it automatically comes from the standard in. Your output goes straight to the standard out. And program errors go right to the standard error. In fact you've already used these streams: every example so far that has printed anything has done so to the standard output stream of your script.

And what about the shell in interactive mode? Does that use those standard streams as well? Yes, it does. In interactive mode, the standard in stream is connected to the keyboard file. And the standard output and standard error are connected to the monitor file.

6.2.3 Okay... But what good is it?

This discussion on files and streams has been very interesting so far and a nice insight into the depths of Unix. But what good does it do you to know all this? Ah, glad you asked!

The Bourne Shell has some built-in features that allow you to do neat tricks involving files and their streams. You see, files don't just have streams -- you can also cross-connect the streams of two files. At the end of the previous section we said that the standard input of the interactive session is connected to the keyboard file. In fact it is connected to the *standard output* stream of the keyboard file. And the standard output and error of the interactive session are connected to the *standard input* of the monitor file. So you can connect the streams of the interactive session to the streams of devices.

But wait. Do you remember the remark above that the point of Unix considering everything to be a file was that everything gets treated like a file? This is why that was important: you can connect a stream from *any* file to a stream of *any* other file. You can connect your interactive shell session to the printer or the network rather than to the monitor (or in addition to the monitor) using streams. You can run a program and have its output go directly to the printer by reconnecting the standard output stream of the program. You can connect the standard output stream of one program directly to the standard input stream of **another program** and make chains of programs. And the Bourne Shell makes it really simple to do all that.

Do you suddenly feel like you've stuck your fingers in the electrical socket? That's the feeling of the raw power of the shell flowing through your body...

6.3 Redirecting: using streams in the shell

As explained in the previous section, the shell process is connected by standard streams to (by default) the keyboard and the monitor. But very often you will want to change this linking. Connecting a file to a stream is a very common operation, so would expect it to be called something like "connecting" or "linking". But since the Bourne Shell has default connections and everything you do is always a *change* in the default connections, connecting a file to a (different) stream using the shell is actually called **redirecting** .

There are several operators built in to the Bourne Shell that relate to redirecting. The most basic and general one is the pipe operator, which we will examine in some detail further on. The others are related to redirecting to file.

6.3.1 Redirecting to file

As we explained (or rather: hinted at) in the previous section, one of the enormously powerful features of the Bourne Shell on top of a Unix operating system is the ability to

chain programs together. Execute a program, have it produce output, then automatically send that output to another program as input. The possible combinations are endless, as is the power of what you can achieve.

One of the most common places where you might want to send a program's output is to a file in the file system. And this time by file we mean a regular, classic data file and not a Unix "everything is a file including your hardware" file. In order to achieve this you can imagine that we can use the chaining mechanism described above: let a program generate output through the standard output stream, then connect that stream (i.e. *redirect the output*) to the standard input stream of a program that creates a data file in the file system. And this would absolutely work. However, redirecting to a data file is such a common operation that you don't need a separate end-of-chain program for it. Redirecting to file is built straight into the Bourne Shell, through the following operators:

process > data file

redirect the output of *process* to the data file; create the file if necessary, overwrite its existing contents otherwise.

process >> data file

redirect the output of *process* to the data file; create the file if necessary, append to its existing contents otherwise.

process < data file

read the contents of the data file and redirect that contents to *process* as input.

Redirecting output

Let's take a closer look at these operators through some examples. Take the simple Bourne shell script below called 'hello.sh':

A simple shell script that generates some output

```
#!/bin/sh
echo Hello
```

This code may be run in any of the ways described in the chapter Running Commands². When we run the script, it simply outputs the string "Hello" to the screen and then returns us to our prompt. But let's say we want to redirect the output to a file instead. We can use the redirect operators to do that easily:

Redirecting the output to a data file

```
$ hello.sh > myfile.txt
$
```

This time, we don't see the string 'Hello' on the screen. Where's it gone? Well, exactly where we wanted it to: into the (new) data file called 'myfile.txt'. Let's examine this file using the 'cat' command:

Examining the results of redirecting some output

```
$ cat myfile.txt
Hello
$
```

Let's run the program again, this time using the '>>' operator instead, and then examine 'myfile.txt' again using the 'cat' command:

Redirecting using the append redirect

```
$ hello.sh >> myfile.txt
$ cat myfile.txt
Hello
Hello
$
```

You can see that 'myfile.txt' now consists of two lines — the output has been added to the end of the file (or concatenated); this is due to the use of the '>>' operator. If we run the script again, this time with the single greater-than operator, we get:

Redirecting using the overwrite redirect

```
$ hello.sh > myfile.txt
$ cat myfile.txt
Hello
$
```

Just one 'Hello' again, because the '>' will always overwrite the contents of an existing file if there is one.

Redirecting input

Okay, so it's clear we can redirect output to a data file. But what about reading from a data file? That's also pretty common. The Bourne Shell helps us here as well: the entire process of reading a file and pumping its data into a stream is captured by the '<' operator.

By default 'stdin' is fed from your keyboard; run the 'cat' command without any arguments and it will just sit there, waiting for you to type something:

```
cat ???  
  
$ cat  
  
I can type all day here, and I never seem to get my prompt back from  
  
this stupid machine.  
  
I have even pressed RETURN a few times !!!  
  
.....etc.....etc
```

In fact 'cat' will sit there all day until you type a 'Ctrl+D' (the 'End of File Character' or 'EOF' for short). To redirect our standard input from somewhere else use the '<' (less-than operator):

Redirecting into the standard input

```
$ cat < myfile.txt  
Hello  
$
```

So 'cat' will now read from the text file 'myfile.txt'; the 'EOF' character is also generated at the end of file, so 'cat' will exit as before.

Note that we previously used 'cat' in this format:

Code

```
$ cat myfile.txt
```

Which is functionally identical to

Code

```
$ cat < myfile.txt
```

However, these are two fundamentally different mechanisms: one uses an argument to the command, the other is more general and redirects 'stdin' — which is what we're concerned

with here. It's more convenient to use 'cat' with a filename as argument, which is why the inventors of 'cat' put this in. However, not all programs and scripts are going to take arguments so this is just an easy example.

Combining file redirects

It's possible to redirect 'stdin' and 'stdout' in one line:

Redirecting input to and output from cat at the same time

```
$ cat < myfile.txt > mynewfile.txt
```

The command above will copy the contents of 'myfile.txt' to 'mynewfile.txt' (and will overwrite any previous contents of 'mynewfile.txt'). Once again this is just a convenient example as we normally would have achieved this effect using 'cp myfile.txt mynewfile.txt'.

Redirecting standard error (and other streams)

So far we have looked at redirecting the "normal" standard streams associated with files, i.e. the files that you use if everything goes correctly and as planned. But what about that other stream? The one meant for errors? How do we go about redirecting that? For example, if we wanted to redirect error data into a log file.

As an example, consider the ls command. If you run the command 'ls myfile.txt', it simply lists the filename 'myfile.txt' – if that file exists. If the file 'myfile.txt' does NOT exist, 'ls' will return an error to the 'stderr' stream, which by default in Bourne Shell is also connected to your monitor.

So, lets run 'ls' a couple of times, first on a file which does exist and then on one that doesn't:

Listing an existing file

Code

```
$ ls myfile.txt
```

Output

```
.txt
```

```
$
```

and then:

Listing a non-existent file

Code

```
$ ls nosuchfile.txt
```

Output

```
: no such file or directory
```

```
$
```

And again, this time with 'stdout' redirected only:

Trying to redirect...

Code

```
$ ls nosuchfile.txt > logfile.txt
```

Output

```
: no such file or directory
```

```
$
```

We still see the error message; 'logfile.txt' will be created but will be empty. This is because we have now redirected the stdout stream, while the error message was written to the error stream. So how do we tell the shell that we want to redirect the error stream?

In order to understand the answer, we have to cover a little more theory about Unix files and streams. You see, deep down the reason that we can redirect stdin and stdout with simple operators is that redirecting those streams is so common that the shell lets us use a shorthand notation for those streams. But actually, to be completely correct, we should have told the shell in every case *which* stream we wanted to redirect. In general you see, the shell cannot know: there could be tons of streams connected to any file. And in order to distinguish one from the other each stream connected to a file has a number associated with it: by convention 0 is the standard in, 1 is the standard out, 2 is standard error and any other streams have numbers counting on from there. To redirect any particular stream you prepend the redirect operator with the stream number (called the *file descriptor* . So to redirect the error message in our example, we prepend the redirect operator with a 2, for the stderr stream:

Redirecting the stderr stream

Code

```
$ ls nosuchfile.txt 2> logfile.txt
```

Output

```
$
```

No output to the screen, but if we examine 'logfile.txt':

Checking the logfile

Code

```
$ cat logfile.txt
```

Output

```
: no such file or directory
$
```

As we mentioned before, the operator without a number is a shorthand notation. In other words, this:

Code

```
$ cat < inputfile.txt > outputfile.txt
```

is actually short for

Code

```
$ cat 0< inputfile.txt 1> outputfile.txt
```

We can also redirect both 'stdout' and 'stderr' independently like this:

Code

```
$ ls nosuchfile.txt > stdout.txt 2>logfile.txt
```

```
$
```

'stdout.txt' will be blank , 'logfile.txt' will contain the error as before.

If we want to redirect stdout and stderr to the same file, we can use the file descriptor as well:

Code

```
$ ls nosuchfile.txt > alloutput.txt 2>&1
```

Here '2>&1' means something like 'redirect stderr to the same file stdout has been redirected to'. Be careful with the ordering! If you do it this way:

Code

```
$ ls nosuchfile.txt 2>&1 > alloutput.txt
```

you will redirect stderr to the file that stdout points to, then send stdout somewhere else — and both streams will end up being redirected to different locations.

Special files

We said earlier that the redirect operators discussed so far all redirect to data files. While this is technically true, Unix magic still means that there's more to it than just that. You see, the Unix file system tends to contain a number of special files called "devices", by convention collected in the /dev directory. These device files include the files that represent your hard drive, DVD player, USB stick and so on. They also include some special files, like /dev/null (also known as the bit bucket; anything you write to this file is discarded). You can redirect to device files as well as to regular data files. Be careful here; you really

don't want to redirect raw text data to the boot sector of your hard drive (and you can!). But if you know what you're doing, you can use the device files by redirecting to them (this is how DVDs are burned in Linux, for instance).

As an example of how you might actually use a device file, in the 'Solaris' flavour of Unix the loudspeaker and its microphone can be accessed by the file '/dev/audio'. So:

Code

```
# cat /tmp/audio.au > /dev/audio
```

Will play a sound, whereas:

Code

```
# cat < /dev/audio > /tmp/mysound.au
```

Will record a sound.(you will need to CTRL-C this to finish...)

This is fun:

Code

```
# cat < /dev/audio > /dev/audio
```

Now wave the microphone around whilst shouting - Jimi Hendrix style feedback. Great stuff. You will probably need to be logged in as 'root' to try this by the way.

Some redirect warnings

The astute reader will have noticed one or two things in the discussion above. First of all, a file can have more than just the standard streams associated with it. Is it legal to redirect those? Is it even *possible*? The answer is, technically, yes. You can redirect stream 4 or 5 of a file (if they exist). Don't try it though. If there's more than a few streams in any direction, you won't know which stream you're redirecting. Plus, if a program needs more than the standard streams it's a good bet that program also needs its extra streams going to a specific location.

Second, you might have noticed that file descriptor 0 is, by convention, the standard input stream. Does that mean you can redirect a program's standard input **away** from the program? Could you do the following?

Code

```
$ cat 0> somewhere_else
```

The answer is, yes you can. And yes, things will break if you do.

6.3.2 Pipes, Tees and Named Pipes

So, after all this talk about redirecting to file, we finally get to it: general redirecting by cross-connecting streams. The most general form of redirecting and the most powerful one to boot. It's called a pipe and is performed using the pipe operator '|'. Pipes allow you to join two processes together through a "pipeline", which directly connects the stdout of one file to the stdin of another.

As an example let's consider the 'grep' command which returns a matching string, given a keyword and some text to search. And let's also use the ps command, which lists running processes on the machine. If you give the command

Code

```
$ ps -eaf
```

it will generally list pagefuls of running processes on your machine, which you would have to sift through manually to find what you want. Let's say you are looking for a process which you know contains the word 'oracle'; use the output of 'ps' to pipe into grep, which will only return the matching lines:

Code

```
$ ps -eaf | grep oracle
```

Now you will only get back the lines you need. What happens if there's still loads of these? No problem, pipe the output to the command 'more' (or 'pg'), which will pause your screen if it fills up:

Code

```
$ ps -ef | grep oracle | more
```

What about if you want to kill all those processes? You need the 'kill' program, plus the process number for each process (the second column returned by the ps command). Easy:

Code

```
$ ps -ef | grep oracle | awk '{print $2}' | xargs kill -9
```

In this command, 'ps' lists the processes and 'grep' narrows the results down to oracle. The 'awk' tool pulls out the second column of each line. And 'xargs' feeds each line, one at a time, to 'kill' as a command line argument.

Pipes can be used to link as many programs as you wish within reasonable limits (and we don't know what these limits are!)

Don't forget you can still use the redirectors in combination:

Code

```
$ ps -ef | grep oracle > /tmp/myprocesses.txt
```

There is another useful mechanism that can be used with pipes: the 'tee'. To understand tee, imagine a pipe shaped like a 'T' - one input, two outputs:

Code

```
$ ps -ef | grep oracle | tee /tmp/myprocesses.txt
```

The 'tee' will copy whatever is given to its stdin and redirect this to the argument given (a file); it will also then send a further copy to its stdout - which means you can effectively intercept the pipe, take a copy at this stage, and carry on piping up other commands; useful maybe for outputting to a logfile, and copying to the screen.

A note on piped commands: piped processes run in parallel on the Unix environment. Sometimes one process will be blocked, waiting for input from another process. But each process in a pipeline is, in principle, running simultaneously with all the others.

Named pipes

There is a variation on the in-line pipe which we have been discussing called the 'named pipe'. A named pipe is actually a file with its own 'stdin' and 'stdout' - which you attach processes to. This is useful for allowing programs to talk to each other, especially when you don't know exactly when one program will try and talk to the other (waiting for a backup to finish etc.) and when you don't want to write a complicated network-based listener or do a clumsy polling loop.

To create a 'named pipe', you use the 'mkfifo' command (fifo=first in, first out; so data is read out in the same order as it is written into).

Code

```
$ mkfifo mypipe
$
```

This creates a named pipe called 'mypipe'; next we can start using it.

This test is best run with two terminals logged in:

1. From 'terminal a'

Code

```
$ cat < mypipe
```

The 'cat' will sit there waiting for an input.

2. From 'terminal b'

Code

```
$ cat myfile.txt > mypipe
$
```

This should finish immediately. Flick back to 'terminal a'; this will now have read from the pipe and received an 'EOF', and you will see the data on the screen; the command will have finished, and you are back at the command prompt.

Now try the other way round:

1. From terminal 'b'

Code

```
$ cat myfile.txt > mypipe
```

This will now sit there, as there isn't another process on the other end to 'drain' the pipe - it's blocked.

2. From terminal 'a'

Code

```
$ cat < mypipe
```

As before, both processes will now finish, the output showing on terminal 'a'.

6.4 Here documents

So far we have looked at redirecting from and to data files and cross-connecting data streams. All of these shell mechanisms are based on having a "physical" source for data — a process or a data file. Sometimes though, you want to feed some data into a target without having a source for it. In these cases you can use an "on the fly" document called a *here document*. A here document means that you open a virtual text document (in memory), type into it as usual, close it and then treat it like any normal file.

Creating a here document is done using a variation on the input redirect operator: the '<<'' operator. Like the input redirect operator, the here document operator takes an argument. For the input redirect operator this operand is the name of the file to be streamed in. For the here document operator it is the string that will terminate the here document. So using the here document operator looks like this:

Code

```
target << terminator string
```

```
here document contents
```

```
terminator string
```

For example:

Using a here document

Code

```
cat << %%
> This is a test.
> This test uses a here document.
> Hello world.
> This here document will end upon the occurrence of the string "%%" on a
  separate line.
> So this document is still open now.
> But now it will end....
> %%
```

Output

```
is a test.
This test uses a here document.
Hello world.
This here document will end upon the occurrence of the string "%%" on a separate line.
So this document is still open now.
But now it will end....
```

When using here documents in combination with variable or command substitution, it is important to realize that substitutions are carried out *before* the here document is passed on. So for example:

Using a here document with substitutions

Code

```
$ COMMAND=cat
$ PARAM='Hello World!!'
$ $COMMAND <<%
> `echo $PARAM`
> %
```

Output

```
World!!
```

7 Modularization

If you've ever done any programming in a different environment than the shell, you're probably familiar with the following scenario: you're writing your program, happily typing away at the keyboard, until you notice that

- you have to repeat some code you typed earlier because your program has to perform exactly the same actions in two different locations; or
- your program is just too *long* to understand anymore.

In other words, you've reached the point where it becomes necessary to divide your program up into modules that can be run as separate subprograms and called as often as you like. Working in the Bourne Shell is no different than working in any other language in this respect. Sooner or later you're going to find yourself writing a shell script that's just too long to be practical anymore. And the time will have come to divide your script up into modules.

7.1 Named functions

Of course, the easy and obvious way to divide a script into modules is just to create a couple of different shell scripts — just a few separate text files with executable permissions. But using separate files isn't always the most practical solution either. Spreading your script over multiple files can make it hard to maintain. Especially if you end up with shell scripts that aren't really meaningful unless they are called specifically from one other, particular shell script.

Especially for this situation the Bourne Shell includes the concept of a *named function*: the possibility to associate a name with a command list and execute the command list by using the name as a command. This is what it looks like:

Code

```
name () command group  
name is a text string
```

and *command group* is any grouped command list (either with curly braces or parentheses)

This functionality is available throughout the shell and is useful in several situations. First of all, you can use it to break a long shell script up into multiple modules. But second, you can use it to define your own little macros in your own environment that you don't want to create a full script for. Many modern shells include a built-in command for this called *alias*, but old-fashioned shells like the original Bourne Shell did not; you can use named functions to accomplish the same result.

7.2 Creating a named function

7.2.1 Functions with a simple command group

Let's start off simply by creating a function that prints "Hello World!". And let's call it "hw". This is what it looks like:

Hello world as a named function

```
hw() {  
> echo 'Hello World!!!';  
>}
```

We can use exactly the same code in a shell script or in the interactive shell — the example above is from the interactive shell. There are several things to notice about this example. First of all, we didn't need a separate keyword to define a function, just the parentheses did it. To the shell, function definitions are like extended variable definitions. They're part of the environment; you set them just by defining a name and a meaning.

The second thing to note is that, once you're past the parentheses, all the normal rules hold for the command group. In our case we used a command group with braces, so we needed the semicolon after the echo command. The string we want to print contains exclamation points, so we have to quote it (as usual). And we were allowed to break the command group across multiple lines, even in interactive mode, just like normal.

Here's how you use the new function: Calling our function

Code

```
$ hw
```

Output

```
World!!
```

7.2.2 Functions that execute in a separate process

The definition of a function takes a command group. *Any* command group. Including the command group with parentheses rather than braces. So if we want, we can define a function that runs as a subprocess in its own environment as well. Here's hello world again, in a subprocess:

Hello world as a named function

```
hw() ( echo 'Hello World!!!' )
```

It's all on one line this time to keep it short, but the same rules apply as before. And of course the same environment rules apply as well, so any variables defined in the function will not be available anymore once the function ends.

7.2.3 Functions with parameters

If you've done any programming in a different programming language you know that the most useful functions are those that take parameters. In other words, ones that don't always rigidly do the same thing but can be influenced by values passed in when the function is called. So here's an interesting question: can we pass parameters to a function? Can we create a definition like

Code

```
(ARG0, ARG1) { do something with ARG0 and ARG1 }
```

And then make a call like 'functionWithParams(Hello, World)'? Well, the answer is simple: no. The parentheses are just there as a flag for the shell to let it know that the previous name is the name of a function rather than a variable and there is no room for parameters.

Or actually, it's more a case of the above being the simple answer rather than the answer being simple. You see, when you execute a function you are executing a command. To the shell there's really very little difference between executing a named function and executing 'ls'. It's a command like any other. And it may not be able to have parameters, but like any other command it can certainly have *command line arguments*. So we may not be able to define a function with parameters like above, but we can certainly do this:

Functions with command-line arguments

Code

```
$ repeatOne () { echo $1; }  
$ repeatOne 'Hello World!'
```

Output

```
World!
```

And you can use any other variable from the environment as well. Of course, that's a nice trick for when you're calling a function from the command line in the interactive shell. But what about in a shell script? The positional variables for command-line arguments are already taken by the arguments to the shell script, right? Ah, but wait! Each command executed in the shell (no matter how it was executed) has its **own** set of command-line arguments! So there's no interference and you can use the same mechanism. For example, if we define a script like this:

function.sh: A function in a shell script

```
#!/bin/sh

myFunction() {
    echo $1
}

echo $1
myFunction
myFunction "Hello World"
echo $1
```

Then it executes exactly the way we want:

Executing the function.sh script

Code

```
$ . function.sh 'Goodbye World!!'
```

Output

```
World!
```

```
Hello World
```

```
Goodbye World!
```

7.2.4 Functions in the environment

We've mentioned it before, but let's delve a little deeper into it now: what are functions exactly? We've hinted that they're an alias for a command list or a macro and that they're part of the environment. But what *is* a function exactly?

A function, as far as the shell is concerned, is just a very verbose variable definition. And that's really all it is: a name (a text string) that is associated with a value (some more text) and can be replaced by that value when the name is used. Just like a shell variable. And we can prove it too: just define a function in the interactive shell, then give the 'set' command (to list all the variable definitions in your current environment). Your function will be in the list.

Because functions are really a special kind of shell variable definition, they behave exactly the same way "normal" variables do:

- Functions are defined by listing their name, a definition operator and then the value of the function. Functions use a different definition operator though: `()` instead of `=`. This tells the shell to add some special considerations to the function (like not needing the `$` character when using the function).
- Functions are part of the environment. That means that when commands are issued from the shell, functions are also copied into the copy of the environment that is given to the issued command.
- Functions can also be passed to new subprocesses if they are marked for export, using the `export` command. Some shells will require a special command-line argument to `export` for functions (bash, for instance, requires you to do an `export -f` to export functions).
- You can drop function definitions by using the `unset` command.

Of course, when you use them functions behave just like commands (they are expanded into a command list, after all). We've already seen that you can use command-line arguments with functions and the positional variables to match. But you can also redirect¹ input and output to and from commands and pipe commands together as well.

¹ Chapter 6 on page 67

8 Debugging and signal handling

In the previous sections we've told you all about the Bourne Shell and how to write scripts using the shell's language. We've been careful to include all the details we could think of so you will be able to write the best scripts you can. But no matter how carefully you've paid attention and no matter how carefully you write your scripts, the time will come to pass when something you've written simply will not work — no matter how sure you are it should. So how do you proceed from here?

In this module we cover the tools the Bourne Shell provides to deal with the unexpected. Unexpected behavior of your script (for which you must debug the script) and unexpected behavior *around* your script (caused by signals being delivered to your script by the operating system).

8.1 Debugging Flags

So here you are, in the middle of the night having just finished a long and complicated shell script, just poured your heart and soul into it for three days straight while living on nothing but coffee, cola and pizza... and it just won't work. Somewhere in there is a bug that is just eluding you. Something is going wrong, some unexpected behavior has popped up or something else is driving you crazy. So how are you going to debug this script? Sure, you can pump it full of 'echo' commands and debug that way, but isn't there an easier way?

Generally speaking the most insightful way to debug any program is to follow the execution of the program along statement by statement to see what the program is doing exactly. The most advanced form of this (offered by modern IDEs) allows you to trace into a program by stopping the execution at a certain point and examining its internal state. The Bourne Shell is, unfortunately, not that advanced. But it does offer the next best thing: command tracing. The shell can print each and every command as it is being executed.

The tracing functionality (there are two of them) is activated using either the 'set' command or by passing parameters directly to the shell executable when it is called. In either case you can use the `-x` parameter, the `-v` parameter or both.

-v

Turns on verbose mode; each command is printed by the shell as it is read.

-x

This turns on command tracing; every command is printed by the shell as it is executed.

8.1.1 Debugging

Let's consider the following script:

divider.sh: Script with a potential error

```
#!/bin/sh

DIVISOR=${1:-0}
echo $DIVISOR
expr 12 / $DIVISOR
```

Let's execute this script and not pass in a command-line argument (so we use the default value 0 for the DIVISOR variable):

Running the script

Code

```
$ sh divider.sh
```

Output

```
0
```

```
expr: division by zero
```

Of course it's not too hard to figure out what went wrong in this case, but let's take a closer look anyway. Let's see what the shell executed, using the **-x** parameter:

Running the script with tracing on

Code

```
$ sh -x divider.sh
```

Output

```
+ DIVISOR=0
```

```
+ echo 0
```

```
0
```

```
+ expr 12 / 0
```

```
expr: division by zero
```

So indeed, clearly the shell tried to have a division by zero evaluated. Just in case we're confused about where the zero came from, let's see which commands the shell actually read:

Running the script in verbose mode

Code

```
$ sh -v divider.sh
```

Output

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

```
DIVISOR=${1:-0}
```

```
echo $DIVISOR
```

```
0
```

```
expr 12 / $DIVISOR
```

```
expr: division by zero
```

So obviously, the script read a command with a variable substitution that didn't work out very well. If we combine these two parameters the resulting output tells the whole, sad story:

Running the script with maximum debugging

Code

```
$ sh -xv divider.sh
```

Output

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

```
DIVISOR=${1:-0}
```

```
+ DIVISOR=0
```

```
echo $DIVISOR
```

```
+ echo 0
```

```
0
```

```
expr 12 / $DIVISOR
```

```
+ expr 12 / 0
```

`expr: division by zero`

There is another parameter that you can use to debug your script, the `-n` parameter. This causes the shell to read the commands but not execute them. You can use this parameter to do a syntax checking run of your script.

8.1.2 Places to put your parameters

As you saw in the previous section, we used the shell parameters by passing them in as command-line parameters to the shell executable. But couldn't we have put the parameters inside the script itself? After all, there is an interpreter hint in there... And surely enough, we can do exactly that. Let's modify the script a little and try it.

The same script, but now with parameters to the interpreter hint

```
#!/bin/sh -xv

DIVISOR=${1:-0}
echo $DIVISOR
expr 12 / $DIVISOR
```

Running the script

Code

```
$ chmod +x divider.sh
$ ./divider.sh
```

Output

```
#!/bin/sh

DIVISOR=${1:-0}

+ DIVISOR=0

echo $DIVISOR

+ echo 0

0

expr 12 / $DIVISOR
```

```
+ expr 12 / 0
```

```
expr: division by zero
```

So there's no problem there. But there is a little gotcha. Let's try running the script again:

Running the script again

Code

```
$ sh divider.sh
```

Output

```
0
```

```
expr: division by zero
```

So what happened to the debugging that time? Well, you have remember that the interpreter hint is used when you try to execute the script as an executable in its own right. But in the last example, we weren't doing that. In the last example we called the shell ourselves and passed it the script as a parameter. So the shell executed without any debugging activated. It would have worked if we'd done a "sh -xv divider.sh" though.

What about sourcing the script (i.e. using the dot notation)?

Running the script again

Code

```
$ . divider.sh
```

Output

```
0
```

```
expr: division by zero
```

This time the script was executed by the same shell process that is running the interactive shell for us. And the same principle applies: no debugging there either. Because the interactive shell was not started with debugging flags. But we can fix that as well; this is where the 'set' command comes in:

Running the script again

Code

```
$ set -xv  
$ . divider.sh
```

Output

```
. divider.sh

+ . divider.sh

#!/bin/sh -vx

DIVISOR=${1:-0}

++ DIVISOR=0

echo $DIVISOR

++ echo 0

0

expr 12 / $DIVISOR

++ expr 12 / 0

expr: division by zero
```

And now we have debugging active in the interactive shell and we get a full trace of the script. In fact, we even get a trace of the interactive shell *calling* the script! But now what happens if we start a new shell process with debugging on in the interactive shell? Does it carry over?

Running the script again
Code

```
$ sh divider.sh
```

Output

```
divider.sh

+ sh divider.sh

0

expr: division by zero
```

Well, we certainly get a trace of the script being called, but no trace of the script itself. The moral of the story is: when debugging, make sure you know which shell you're activating the trace in.

By the way, to turn tracing in the interactive shell off again you can either do a 'set +xv' or simply a 'set -'.

8.2 Breaking out of a script

When writing or debugging a shell script it is sometimes useful to exit out (to stop the execution of the script) at certain points. You use the 'exit' built-in command to do this. The command looks simply like this:

Code

```
exit [n]
```

Where n (optional) is the exit status of the script.

If you leave off the optional exit status, the exit status of the script will be the exit status of the last command that executed before the call to 'exit'.

For example:

Exiting from a script

```
#!/bin/sh -x
echo hello
exit 1
```

If you run this script and then test the output status, you will see (using the "\$?" built-in variable):

Checking the exit status

Code

```
echo $?
```

Output

```
1
```

There's one thing to look out for when using 'exit': 'exit' actually terminates the executing process. So if you're executing an executable script with an interpreter hint or have called the shell explicitly and passed in your script as an argument that is fine. But if you've sourced the script (used the dot-notation), then your script is being run by the process running your interactive shell. So you may inadvertently terminate your shell session and log yourself out by using the 'exit' command!

There's a variation on 'exit' meant specifically for blocks of code that are not processes of their own. This is the 'return' command, which is very similar to the 'exit' command:

Code

`return [n]`

Where `n` (optional) is the exit status of the block.

Return has exactly the same semantics as 'exit', but is primarily intended for use in shell functions (it makes a function return without terminating the script). Here's an example:

exit_and_return.sh: A script with a function and an explicit return

```
#!/bin/sh

sayHello() {
    echo 'Hi there!!'
    return 2
}

echo 'Hello World!!'
sayHello
echo $?
echo 'Goodbye!!'
exit
```

If we run this script, we see the following:

Running the script

Code

```
./exit_and_return.sh
```

Output

```
World!!
```

```
Hi there!!
```

```
2
```

```
Goodbye!!
```

The function returned with a testable exit status of 2. The overall exit status of the script is zero though, since the last command executed by the script ('echo Goodbye!!') exited without errors.

You can also use a 'return' statement to exit a shell script that you executed by sourcing it (the script will be run by the process that runs the interactive shell, so that's not a subprocess). But it's usually not a good idea, since this will limit your script to being sourced: if you try to run it any other way, the fact that you used a 'return' statement alone will cause an error.

8.3 Signal trapping

A syntax, command error or call to 'exit' is not the only thing that can stop your script from executing. The process that runs your script might also suddenly receive a signal from the operating system. Signals are a simple form of event notification: think of a signal as a little light suddenly popping on in your room, to let you know that somebody outside the room wants your attention. Only it's not just one light. The Unix system usually allows for lots of different signals so it's more like having a wall full of little lamps, each of which could suddenly start blinking.

On a single-process operating system like MS-DOS, life was simple. The environment was single-process, meaning your code (once running) had complete machine control. Any signal arriving was always a hardware interrupt (e.g. the computer signalling that the floppy disk was ready to read) and you could safely ignore all those signals if you didn't need external hardware; either it was some device event you weren't interested in, or something was really wrong — in which case the computer was crashing anyway and there was nothing you could do.

On a Unix system, life is not so easy. On Unix, signals can come from all over the place (including other programs). And you never have complete control of the system either. A signal may be a hardware interrupt, or another program signalling, or the user who got fed up with waiting, logged in to a second shell session and is now ordering your process to die. On the bright side, life is still not too complicated. Most Unix systems (and certainly the Bourne Shell) come with default handling for most signals. Usually you can still safely ignore signals and let the shell or the OS deal with them. In fact, if the signal in question is number 9 (loosely translated: **KILL!! KILL!! DIE!! DIE, RIGHT NOW!!**), you probably *should* ignore it and let the OS kill your process.

But sometimes you just have to do your own signal handling. That might be because you've been working with files and want to do some cleanup before your process dies. Or because the signal is part of your multi-process program design (e.g. listening for signal 16, which is "user-defined signal 1"). Which is why the Bourne Shell gives us the 'trap' command.

8.3.1 Trap

The trap command is actually quite simple (especially if you've ever done event-driven programming of any kind). Essentially the trap command says "if one of the following signals is received by this process, do this". It looks like this:

Code

```
trap [command string] signal0 [signal1] ...
```

command string is a string containing the commands to execute if a signal is trapped

and *signal_n* is a signal to be trapped.

For example, to trap user-defined signal 1 (commonly referred to as SIGUSR1) and print "Hello World" whenever it comes along, you would do this:

Trapping SIGUSR1

```
$ trap "echo Hello World" 16
```

Most Unix systems also allow you to use symbolic names (we'll get back to these a little later). So you can probably also do this:

Trapping SIGUSR1 (little easier)

```
$ trap "echo Hello World" SIGUSR1
```

And if you can do that, you can usually also do this:

Trapping SIGUSR1 (even easier)

```
$ trap "echo Hello World" USR1
```

The command string passed to 'trap' is a string that contains a command list. It's not treated as a command list though; it's just a string and it isn't interpreted until a signal is caught. The command string can be any of the following:

A string

A string containing a command list. Any and all commands are allowed and you can use multiple commands separated by semicolons as well (i.e. a command list).

'<nowiki/>'

The empty string. Actually this is the same as the previous case, since this is the empty command string. This causes the shell to execute nothing when a signal is trapped — in other words, to ignore the signal.

Nothing, the null string. This resets the signal handling to the default signal action (which is usually "kill process").

Following the command list you can list as many signals as you want to be associated with that command list. The traps that you set up in this manner are valid for every command that *follows* the 'trap' command.

Right about now it's probably a good idea to look at an example to clear things up a bit. You can use 'trap' anywhere (as usual) including the interactive shell. But most of the time you will want to introduce traps into a script rather than into your interactive shell process. Let's create a simple script that uses the 'trap' command:

A simple signal trap

```
!/bin/sh

trap 'echo Hello World' SIGUSR1

while [ 1 -gt 0 ]
do
    echo Running....
    sleep 5
done
```

This script in and of itself is an endless loop, which prints "Running..." and then sleeps for five seconds. But we've added a 'trap' command (**before** the loop, otherwise the trap would never be executed and it wouldn't affect the loop) that prints "Hello World" whenever the process receives the SIGUSR1 signal. So let's start the process by running the script:

Infinite loop...

Code

```
$ ./trap_signal.sh
```

Output

```
....
Running....
Running....
Running....
Running....
Running....
...
```

To spring the trap, we must send the running process a signal. To do that, log into a new shell session and use a process tool (like 'ps') to find the correct process id (PID):

Finding the process ID

Code

```
$ ps -ef | grep signal
```

Output

```
      10865  7067  0 15:08 pts/0    00:00:00 /bin/sh ./trap_signal.sh
bzt      10808 10415  0 15:12 pts/1    00:00:00 fgrep signal
```

Now, to send a signal to that process, we use the 'kill' command which is built into the Bourne Shell:

Code

```
kill [-signal] ID [ID] ...
```

`-signal` is the signal to send (optional; default is 15, or SIGTERM)

and `ID` are the PIDs of the processes to send the signal to (at least one of them)

As the name suggests, 'kill' was actually intended to kill processes (this fits with the default signal being SIGTERM and the default signal handler being terminate). But in fact what it does is nothing more than send a signal to a process. So for example, we can send a SIGUSR1 to our process like this:

Let's trip the trap...

Code

```
kill -SIGUSR1 10865
```

Output

```
...  
Running....  
Running....  
Running....  
Running....  
Running....  
Running....  
Hello World  
Running....  
Running....  
...
```

You might notice that there's a short pause before "Hello World!" appears; it won't happen until the running 'sleep' command is done. But after that, there it is. But you might be a little surprised: the signal didn't kill the process. That's because 'trap' *completely* replaces the signal handler with the commands you set up. And an 'echo Hello World' alone won't kill a process... The lesson here is a simple one: if you want your signal trap to terminate your process, make sure you include an 'exit' command.

Between having multiple commands in your command list and potentially trapping lots of signals, you might be worried that a 'trap' statement can become messy. Fortunately, you can also use shell functions as commands in a 'trap'. The following example illustrates that and the difference between an exiting event handler and a non-exiting event handler:

A trap with a shell function as a handler

```

!/bin/sh

exit_with_grace() {
    echo Goodbye World
    exit
}

trap "exit_with_grace" USR1 TERM QUIT
trap "echo Hello World" USR2

while [ 1 -gt 0 ]
do
    echo Running....
    sleep 5
done

```

8.3.2 System signals

Here's the official definition of a signal from the POSIX-1003 2001 edition standard:

A mechanism by which a process or thread may be notified of, or affected by, an event occurring in the system.

Examples of such events include hardware exceptions and specific actions by processes.

The term signal is also used to refer to the event itself.

In other words, a signal is some sort of short message that is sent from one process (possibly a system process) to another. But what does that mean exactly? What does a signal look like? The definition given above is kind of vague...

If you have any feel for what happens in computing when you give a vague definition, you already know the answer to the questions above: every Unix flavor that was developed came up with its own definition of "signal". They pretty much all settled on a message that consists of an integer (because that's simple), but not exactly the *same* list everywhere. Then there was some standardization and Unix systems organized themselves into the System V and BSD flavors and at last everybody agreed on the following definition:

The system signals are the signals listed in `/usr/include/sys/signal.h`.

God, that's helpful...

Since then a lot has happened, including the definition of the POSIX-1003 standard. This standard, which standardizes most of the Unix interfaces (including the shell in part 1 (1003.1)) finally came up with a standard list of symbolic signal names and default handlers. So usually, nowadays, you can make use of that list and expect your script to work on most systems. Just be aware that it's not *completely* fool-proof...

POSIX-1003 defines the signals listed in the table below. The values given are the *typical* numeric values, but they aren't mandatory and you shouldn't rely on them (but then again, you use symbolic values in order not to use actual values).

POSIX system signals			
Signal	Default action	Description	Typical value(s)
SIGABRT	Abort with core dump	Abort process and generate a core dump	6
SIGALRM	Terminate	Alarm clock.	14
SIGBUS	Abort with core dump	Access to an undefined portion of a memory object.	7, 10
SIGCHLD	Ignore	Child process terminated, stopped	20, 17, 18
SIGCONT	Continue process (if stopped)	Continue executing, if stopped.	19,18,25
SIGFPE	Abort with core dump	Erroneous arithmetic operation.	8
SIGHUP	Terminate	Hangup.	1
SIGILL	Abort with core dump	Illegal instruction.	4
SIGINT	Terminate	Terminal interrupt signal.	2
SIGKILL	Terminate	Kill (cannot be caught or ignored).	9
SIGPIPE	Terminate	Write on a pipe with no one to read it (i.e. broken pipe).	13
SIGQUIT	Terminate	Terminal quit signal.	3
SIGSEGV	Abort with core dump	Invalid memory reference.	11
SIGSTOP	Stop process	Stop executing (cannot be caught or ignored).	17,19,23
SIGTERM	Terminate	Termination signal.	15
SIGTSTP	Stop process	Terminal stop signal.	18,20,24
SIGTTIN	Stop process	Background process attempting read.	21,21,26
SIGTTOU	Stop process	Background process attempting write.	22,22,27

POSIX system signals			
Signal	Default action	Description	Typical value(s)
SIGUSR1	Terminate	User-defined signal 1.	30,10,16
SIGUSR2	Terminate	User-defined signal 2.	31,12,17
SIGPOLL	Terminate	Pollable event.	-
SIGPROF	Terminate	Profiling timer expired.	27,27,29
SIGSYS	Abort with core dump	Bad system call.	12
SIGTRAP	Abort with core dump	Trace/breakpoint trap	5
SIGURG	Ignore	High bandwidth data is available at a socket.	16,23,21
SIGVTALRM	Terminate	Virtual timer expired.	26,28
SIGXCPU	Abort with core dump	CPU time limit exceeded.	24,30
SIGXFSZ	Abort with core dump	File size limit exceeded.	25,31

Earlier on¹ we talked about job control and suspending and resuming jobs. Job suspension and resuming is actually completely based on sending signals to processes, so you can in fact control job stopping and starting completely using 'kill' and the signal list. To suspend a process, send it the SIGSTOP signal. To resume, send it the SIGCONT signal.

8.3.3 Err... ERR?

If you go online and read about 'trap', you might come across another kind of "signal" which is called **ERR**. It's used with 'trap' the same way regular signals are, but it isn't really a signal at all. It's used to trap command errors (i.e. non-zero exit statuses), like this:

Error trapping

Code

```
$ trap 'echo HELLO WORLD' ERR
$ expr 1 / 0
```

Output

```
: division by zero
```

¹ Chapter 3.2.5 on page 31

HELLO WORLD

So why didn't we cover this "signal" earlier, when we were discussing 'trap'²? Well, we saved it until the discussion on system and non-system signals for a reason: ERR isn't standard at all. It was added by the Korn Shell to make life easier, but not adopted by the POSIX standard and it certainly isn't part of the original Bourne Shell. So if you use it, remember that your script may not be portable anymore.

² Chapter 8.3.1 on page 95

9 Cookbook

```
action=edit&section=new}} Post a new Cookbook entry]
```

If you use the title box, then you do not need to put a title in the body.

9.1 Branch on extensions

When writing a bash script which should do different things based on the extension of a file, the following pattern is helpful.

```
#filepath should be set to the name(with optional path) of the file in question
ext=${filepath##*.}
if [[ "$ext" == txt ]] ; then
    #do something with text files
fi
```

(Source: slike.com Bash FAQ¹).

9.2 Rename several files

This recipe shows how to rename several files following a pattern.

In this example, the user has huge collection of screenshots. This user wants to rename the files using a Bourne-compatible shell. Here is an "ls" at the shell prompt to show you the filenames. The goal is to rename images like "snapshot1.png" to "nethack-kernigh-22oct2005-01.png"².

```
$ ls
snapshot1.png  snapshot25.png  snapshot40.png  snapshot56.png  snapshot71.png
snapshot10.png snapshot26.png  snapshot41.png  snapshot57.png  snapshot72.png
snapshot11.png snapshot27.png  snapshot42.png  snapshot58.png  snapshot73.png
snapshot12.png snapshot28.png  snapshot43.png  snapshot59.png  snapshot74.png
snapshot13.png snapshot29.png  snapshot44.png  snapshot6.png   snapshot75.png
snapshot14.png snapshot3.png   snapshot45.png  snapshot60.png  snapshot76.png
snapshot15.png snapshot30.png  snapshot46.png  snapshot61.png  snapshot77.png
snapshot16.png snapshot31.png  snapshot47.png  snapshot62.png  snapshot78.png
snapshot17.png snapshot32.png  snapshot48.png  snapshot63.png  snapshot79.png
snapshot18.png snapshot33.png  snapshot49.png  snapshot64.png  snapshot8.png
snapshot19.png snapshot34.png  snapshot5.png   snapshot65.png  snapshot80.png
snapshot2.png  snapshot35.png  snapshot50.png  snapshot66.png  snapshot81.png
```

1 http://www.slike.com/howtos/bash_faq.html#Get+a+file%27s+basename%2C+dirname%2C+extension%2C+etc%3F

2 <http://en.commons.org/wiki/Image%3Anethack-kernigh-22oct2005-01.png>

```

snapshot20.png snapshot36.png snapshot51.png snapshot67.png snapshot82.png
snapshot21.png snapshot37.png snapshot52.png snapshot68.png snapshot83.png
snapshot22.png snapshot38.png snapshot53.png snapshot69.png snapshot9.png
snapshot23.png snapshot39.png snapshot54.png snapshot7.png
snapshot24.png snapshot4.png snapshot55.png snapshot70.png

```

First, to add a "0" (zero) before snapshots 1 through 9, write a for loop (in effect, a short shell script).

- Use `?` which is a filename pattern for a single character. Using it, I can match snapshots 1 through 9 but miss 10 through 83 by saying `snapshot?.png`.
- Use `${parameter#pattern}` to substitute the value of `parameter` with the `pattern` removed from the beginning. This is to get rid of "snapshot" so I can put in "snapshot0".
- Before actually running the loop, insert an "echo" to test that the commands will be correct.

```

$ for i in snapshot?.png; do echo mv "$i" "snapshot0${i#snapshot}"; done
mv snapshot1.png snapshot01.png
mv snapshot2.png snapshot02.png
mv snapshot3.png snapshot03.png
mv snapshot4.png snapshot04.png
mv snapshot5.png snapshot05.png
mv snapshot6.png snapshot06.png
mv snapshot7.png snapshot07.png
mv snapshot8.png snapshot08.png
mv snapshot9.png snapshot09.png

```

That seems good, so run it by removing the "echo".

```

$ for i in snapshot?.png; do mv "$i" "snapshot0${i#snapshot}"; done

```

An `ls` confirms that this was effective.

Now change prefix "snapshot" to "nethack-kernigh-22oct2005-". Run a loop similar to the previous one:

```

$ for i in snapshot*.png; do
> mv "$i" "nethack-kernigh-22oct2005-${i#snapshot}"
> done

```

This saves the user from typing 83 "mv" commands.

9.3 Long command line options

The builtin `getopts` does not support long options so the external `getopt` is required. (On some systems, `getopt` *also* does not support long options, so the next example will not work.)

```

eval set -- $(getopt -l install-opts: "" "$@")
while true; do
  case "$1" in

```

```

--install-opts)
    INSTALL_OPTS=$2
    shift 2
    ;;
--)
    shift
    break
    ;;
esac
done

echo $INSTALL_OPTS

```

The call to `getopt` quotes and reorders the command line arguments found in `$@`. `set` then makes replaces `$@` with the output from `getopt`

Another example of `getopt` use can also be found in the [Advanced Bash Script Guide](#)³

9.4 Process certain files through xargs

In this recipe, we want to process a large list of files, but we must run one command for each file. In this example, we want to convert the sampling rates of some sound files to 44100 hertz. The command is `sox file.ogg -r 44100 conv/file.ogg`, which converts `file.ogg` to a new file `conv/file.ogg`. We also want to skip files that are already 44100 hertz.

First, we need the sampling rates of our files. One way is to use the `file` command:

```

$ file *.ogg
audio_on.ogg:      Ogg data, Vorbis audio, mono, 44100 Hz, ~80000 bps
beep_1.ogg:       Ogg data, Vorbis audio, stereo, 44100 Hz, ~193603 bps
cannon_1.ogg:     Ogg data, Vorbis audio, mono, 48000 Hz, ~96000 bps
...

```

(The files in this example are from [Secret Maryo Chronicles](#)⁴.) We can use `grep -v` to filter out all lines that contain '44100 Hz':

```

$ file *.ogg | grep -v '44100 Hz '
cannon_1.ogg:     Ogg data, Vorbis audio, mono, 48000 Hz, ~96000 bps
...
jump_small.ogg:  Ogg data, Vorbis audio, mono, 8000 Hz, ~22400 bps
live_up.ogg:     Ogg data, Vorbis audio, mono, 22050 Hz, ~40222 bps
...

```

We finished with "grep" and "file", so now we want to remove the other info and leave only the filenames to pass to "sox". We use the text utility `cut`. The option `-d:` divides each line into fields at the colon; `-f1` selects the first field.

³ <http://www.tldp.org/LDP/abs/html/extmisc.html#EX33A>

⁴ <http://www.secretmaryo.org/>

```
$ file *.ogg | grep -v '44100 Hz' | cut -d: -f1
cannon_1.ogg
...
jump_small.ogg
live_up.ogg
...
```

We can use another pipe to supply the filenames on the standard input, but "sox" expects them as arguments. We use **xargs**, which will run a command repeatedly using arguments from the standard input. The **-n1** option specifies one argument per command. For example, we can run **echo sox** repeatedly:

```
$ file *.ogg | grep -v '44100 Hz' | cut -d: -f1 | xargs -n1 echo sox
sox cannon_1.ogg
...
sox itembox_set.ogg
sox jump_small.ogg
...
```

However, these commands are wrong. The full command for `cannon_1.ogg`, for example, is **sox cannon_1.ogg -r 44100 conv/cannon_1.ogg**. "xargs" will insert incoming data into placeholders indicated by "{}". We use this strategy in our pipeline. If we have doubt, then first we can build a test pipeline with "echo":

```
$ file *.ogg | grep -v '44100 Hz' | cut -d: -f1 | \
> xargs -i 'echo sox {} -r 44100 conv/{}'
sox cannon_1.ogg -r 44100 conv/cannon_1.ogg
...
sox itembox_set.ogg -r 44100 conv/itembox_set.ogg
sox jump_small.ogg -r 44100 conv/jump_small.ogg
...
```

It worked, so let us remove the "echo" and run the "sox" commands:

```
$ mkdir conv
$ file *.ogg | grep -v '44100 Hz' | cut -d: -f1 | \
> xargs -i 'sox {} -r 44100 conv/{}'
```

After a wait, the converted files appear in the `conv` subdirectory. The above three lines alone did the entire conversion.

9.5 Simple playlist frontend for GStreamer

If you have GStreamer, the command `gst-launch filesrc location=filename ! decodebin ! audioconvert ! esdsink` will play a sound or music file of any format for which you have a GStreamer plugin. This script will play through a list of files, optionally looping through them. (Replace "esdsink" with your favorite sink.)

```
#!/bin/sh
loop=false
if test x"$1" == x-l; then
    loop=true
    shift
fi

while true; do
    for i in "$@"; do
        if test -f "$i"; then
            echo "${0##*/}: playing $i" > /dev/stderr
            gst-launch filesrc location="$i" ! decodebin ! audioconvert ! esdsink
        else
            echo "${0##*/}: not a file: $i" > /dev/stderr
        fi
    done
    if $loop; then true; else break; fi
done
```

This script demonstrates some common Bourne shell tactics:

- "loop" is a boolean variable. It works because its values "true" and "false" are both Unix commands (and sometimes shell builtins), thus you can use them as conditions in `if` and `while` statements.
- The shell builtin "shift" removes \$1 from the argument list, thus shifting \$2 to \$1, \$3 to \$2, and so forward. This script uses it to process an "-l" option.
- The substitution `${0##*/}` gives everything in \$0 after the last slash, thus "playlist", not "/home/musicfan/bin/playlist".

10 Quick Reference

This final section provides a fast lookup reference for the materials in this document. It is a collection of thumbnail examples and rules that will be cryptic if you haven't read through the text.

10.1 Useful commands

Command	Effect
cat	Lists a file or files sequentially.
cd	Change directories.
chmod ugo+rwx	Set read, write and execute permissions for user, group and others.
chmod a-rwx	Remove read, write and execute permissions from all.
chmod 755	Set user write and universal read-execute permissions
chmod 644	set user write and universal read permissions.
cp	Copy files.
expr 2 + 2	Add 2 + 2.
fgrep	Search for string match.
grep	Search for string pattern matches.
grep -v	Search for no match.
grep -n	List line numbers of matches.
grep -i	Ignore case.
grep -l	Only list file names for a match.
head -n5 source.txt	List first 5 lines.
less	View a text file one screen at a time; can scroll both ways.
ll	Give a listing of files with file details.
ls	Give a simple listing of files.
mkdir	Make a directory.
more	Displays a file a screenfull at a time.
mv	Move or rename files.
paste f1 f2	Paste files by columns.
pg	Variant on "more".
pwd	Print working directory.
rm	Remove files.
rm -r	Remove entire directory subtree.
rmdir	Remove an empty directory.
sed 's/txt/TXT/g'	Scan and replace text.
sed 's/txt/d'	Scan and delete text.

Command	Effect
sed '/txt/q'	Scan and then quit.
sort	Sort input.
sort +1	Skip first field in sorting.
sort -n	Sort numbers.
sort -r	Sort in reverse order.
sort -u	Eliminate redundant lines in output.
tail -5 source.txt	List last 5 lines.
tail +5 source.txt	List all lines after line 5.
tr '[A-Z]' '[a-z]'	Translate to lowercase.
tr '[a-z]' '[A-Z]'	Translate to uppercase.
tr -d ' _ '	Delete underscores.
uniq	Find unique lines.
wc	Word count (characters, words, lines).
wc -w	Word count only.
wc -l	Line count.

10.2 Elementary shell capabilities

Command	Effect
Initialize a shell variable.	
echo \$shvar	Display a shell variable.
export shvar	Allow subshells to use shell variable.
mv \$f \${f}2 <i>or</i> mv \${f}{,2}	Append "2" to file name in shell variable.
\$1, \$2, \$3, ...	Command-line arguments.
\$0	Shell-program name.
\$#	Number of arguments.
\$*	Complete argument list (all in one string).
@	Complete argument list (string for every argument).
?	Exit status of the last command executed.
shift 2	Shift argument variables by 2.
read v	Read input into variable "v".
. mycmds	Execute commands in file.

10.3 IF statement

The if statement executes the command between **if** and **then** . If the command returns not 0 then the commands between **then** and **else** are executed - otherwise the command between **else** and **fi** .

```

if test "${1}" = "red" ; then
    echo "Illegal code."
elif test "${1}" = "blue" ; then
    echo "Illegal code."
else

```

```

    echo "Access granted."
fi

```

```

if [ "$1" = "red" ]
then
    echo "Illegal code."
elif [ "$1" = "blue" ]
then
    echo "Illegal code."
else
    echo "Access granted."
fi

```

Test Syntax Variations

Most test commands can be written using more than one syntax. Mastering and consistently using one form may be a programming best-practice, and may be a more efficient use of overall time.

String Tests

String Tests are performed by the `test` command. See `help test` for more details. To make scripts look more like other programming languages the synonym `[...]` was defined which does exactly the same as `test`.

Command	Effect
<pre>test "\$shvar" = "red" ["\$shvar" = "red"]</pre>	String comparison, true if match.
<pre>test "\$shvar" != "red" ["\$shvar" != "red"]</pre>	String comparison, true if no match.
<pre>test -z "\${shvar}" test "\$shvar" = "" ["\$shvar" = ""]</pre>	True if null variable.
<pre>test -n "\${shvar}" test "\$shvar" != "" [-n "\$shvar"] ["\$shvar" != ""]</pre>	True if not null variable.

Arithmetic tests

simple arithmetics can be performed with the `test` for more complex arithmetics the `let` command exists. See `help let` for more details. Note that for `let` command variables don't need to be prefixed with '\$' and the statement need to be one argument, use '...' when there are spaces inside the argument. Like with `test` a synonym - `((...))` - was defined to make shell scripts look more like ordinary programs.

Command	Effect
---------	--------

Command	Effect
test "\$nval" -eq 0 let 'nval == 0' ["\$nval" -eq 0] ((nval == 0))	Integer test; true if equal to 0.
test "\$nval" -ge 0 let 'nval >= 0' ["\$nval" -ge 0] ((nval >= 0))	Integer test; true if greater than or equal to 0.
test "\$nval" -gt 0 let 'nval > 0' ["\$nval" -gt 0] ((nval > 0))	Integer test; true if greater than 0.
test "\$nval" -le 0 let 'nval <= 0' ["\$nval" -le 0] ((nval <= 0))	Integer test; true if less than or equal to 0.
test "\$nval" -lt 0 let 'nval < 0' ["\$nval" -lt 0] ((nval < 0))	Integer test; true if less than to 0.
test "\$nval" -ne 0 let 'nval != 0' ["\$nval" -ne 0] ((nval != 0))	Integer test; true if not equal to 0.
let 'y + y > 100' ((y + y >= 100))	Integer test; true when $x + y \geq 0$

File tests

Command	Effect
test -d tmp [-d tmp]	True if "tmp" is a directory.
test -f tmp [-f tmp]	True if "tmp" is an ordinary file.
test -r tmp [-r tmp]	True if "tmp" can be read.
test -s tmp [-s tmp]	True if "tmp" is nonzero length.
test -w tmp [-w tmp]	True if "tmp" can be written.
test -x tmp [-x tmp]	True if "tmp" is executable.

Boolean tests

Boolean arithmetic is performed by a set of operators. It is important to note then the operators execute programs and compare the result codes. Because boolean operators are often combined with `test` command a unifications was created in the form of `[[...]]`.

Command	Effect
<code>test -d /tmp && test -r /tmp</code> <code>[[-d /tmp && -r /tmp]]</code>	True if <code>/tmp</code> is a directory and can be read.
<code>test -r /tmp test -w /tmp[[-r /tmp -w /tmp]]</code>	True if <code>tmp</code> can be read or written.
<code>test ! -x /tmp</code> <code>[[! -x /tmp]]</code>	True if the file is not executable

10.4 CASE statement

```

case "$1"
in
  "red")      echo "Illegal code."
              exit;;
  "blue")     echo "Illegal code."
              exit;;
  "x"|"y")   echo "Illegal code."
              exit;;
  *)         echo "Access granted.";;
esac

```

10.5 Loop statements

```

for nvar in 1 2 3 4 5
do
  echo $nvar
done

```

```

for file          # Cycle through command-line arguments.
do
  echo $file
done

```

```

while [ "$n" != "Joe" ]    # Or:  until [ "$n" == "Joe" ]
do
  echo "What's your name?"
  read n
  echo $n
done

```

There are `break` and `continue` commands that allow you to exit or skip to the end of loops as the need arises.

Instead of `[[` we can use `test`. `[[` requires space after and before the brackets and there should be spaces between arguments.

10.6 Credit

This content was originally from <http://www.vectorsite.net/tsshell.html> and was originally in the public domain.

11 Command Reference

The Bourne Shell offers a large number of built-in commands that you can use in your shell scripts. The following table gives an overview:

Bourne Shell command reference

Command	Description
:	A null command that returns a 0 (true) exit value.
. <i>file</i>	Execute. The commands in the specified file are read and executed by the shell. Commonly referred to as <i>sourcing</i> a file.
#	Ignore all the text until the end of the line. Used to create comments in shell scripts.
#! <i>shell</i>	Interpreter hint. Indicates to the OS which interpreter to use to execute a script.
bg [job] ...	Run the specified jobs (or the current job if no arguments are given) in the background.
break [n]	Break out of a loop. If a number argument is specified, break out of n levels of loops.
case	See ../Control flow ¹
cd [directory]	Switch to the specified directory (default \$HOME).
continue [n]	Skip the remaining commands in a loop and continue the loop at the next iteration. If an integer argument is specified, skip n loops.
echo string	Write string to the standard output.
eval string ...	Concatenate all the arguments with spaces. Then re-parse and execute the command.
exec [command arg ...]	Execute command in the current process.
exit [exitstatus]	Terminate the shell process. If exitstatus is given it is used as the exit status of the shell; otherwise the exit status of the last completed command is used.
export name ...	Mark the named variables or functions for export to child process environments.
fg [job]	Move the specified job (or the current job if not specified) to the foreground.
for	See ../Control flow ² .

1 Chapter 5 on page 45

2 Chapter 5 on page 45

Bourne Shell command reference

Command	Description
hash -rv command ...	The shell maintains a hash table which remembers the locations of commands. With no arguments whatsoever, the hash command prints out the contents of this table. Entries which have not been looked at since the last cd command are marked with an asterisk; it is possible for these entries to be invalid. With arguments, the hash command removes the specified commands from the hash table (unless they are functions) and then locates them. The -r option causes the hash command to delete all the entries in the hash table except for functions.
if	See ../Control flow ³ .
jobs	This command lists out all the background processes which are children of the current shell process.
-signal] PID ...	Send signal to the jobs listed by ID. If no signal is specified, send SIGTERM. If the -l option is used, lists all the signal names defined on the system.
newgrp [group]	Temporarily move your user to a new group. If no group is listed, move back to your user's default group.
pwd	Print the working directory.
read variable [...]	Read a line from the input and assign each individual word to a listed variable (in order). Any leftover words are assigned to the last variable.
readonly name ...	Make the listed variables read-only.
return [n]	Return from a shell function. If an integer argument is specified it will be the exit status of the function.
set [{ -options +options -- }] arg ...	The set command performs three different functions. With no arguments, it lists the values of all shell variables. If options are given, it sets the specified option flags or clears them. The third use of the set command is to set the values of the shell's positional parameters to the specified args. To change the positional parameters without changing any options, use "--" as the first argument to set. If no args are present, the set command will clear all the positional parameters (equivalent to executing "shift \$#".)
shift [n]	Shift the positional parameters n times.
test	See ../Control flow ⁴ .

3 Chapter 5 on page 45

4 Chapter 5 on page 45

Bourne Shell command reference

Command	Description
trap [action] signal ...	Cause the shell to parse and execute action when any of the specified signals are received.
type [name ...]	Show whether a command is a UNIX command, a shell built-in command or a shell function.
ulimit	Report on or set resource limits.
umask [mask]	Set the value of umask (the mask for the default file permissions <i>not</i> assigned to newly created files). If the argument is omitted, the umask value is printed.
unset name ...	Drop the definition of the given names in the shell.
wait [job]	Wait for the specified job to complete and return the exit status of the last process in the job. If the argument is omitted, wait for all jobs to complete and the return an exit status of zero.
while	See ../Control flow ⁵ .

12 Environment reference

In the section on the environment¹ we discussed the concept of environment variables. We also mentioned that there are usually a large number of environment variables that are created centrally in `/etc/profile`. There are a number of these that have a predefined meaning in the Bourne Shell. They are not set automatically, mind, but they have meaning when they are set.

On most systems there are far more predefined variables than we list here. And some of these will mean something to your shell (most shells have more options than the Bourne Shell). Check your shell's documentation for a listing. The ones below are meaningful to the Bourne Shell and are usually also recognized by other shells.

Bourne Shell environment variables

Variable	Meanings
HOME	The user's home directory. Set automatically at login from the user's login directory in the password file
PATH	The default search path for executables.
CDPATH	The search path used with the <code>cd</code> builtin, to allow for shortcuts.
LANG	The directory for internationalization files, used by localizable programs.
MAIL	The name of a mail file, that will be checked for the arrival of new mail.
MAILCHECK	The frequency in seconds that the shell checks for the arrival of mail.
MAILPATH	A colon ":" separated list of file names, for the shell to check for incoming mail.
PS1	The control string for your prompt, which defaults to "\$ ", unless you are the superuser, in which case it defaults to "# "
PS2	The control string for your secondary prompt, which defaults to "> ". The secondary prompt is what you see when you break a command over more than one line.
PS4	The character string you see before the output of an execution trace (<code>set -x</code>); defaults to "+ "
IFS	Input Field Separators. Basically the characters the shell considers to be whitespace. Normally set to <code><space></code> , <code><tab></code> , and <code><newline></code> .
TERM	The terminal type, for use by the shell.

¹ Chapter 3 on page 15

13 Contributors

Edits	User
1	ABCD ¹
16	Adrignola ²
1	Albmont ³
5	Avicennasis ⁴
1	Aya ⁵
1	Az1568 ⁶
189	BenTels ⁷
3	Dirk Hünninger ⁸
1	Eric119 ⁹
1	Geocachernemesis ¹⁰
1	Guanabot~enwikibooks ¹¹
1	Guanaco ¹²
3	Hagindaz ¹³
1	JesseW ¹⁴
10	Jguk ¹⁵
1	Jomegat ¹⁶
20	Kernigh ¹⁷
17	Krischik ¹⁸
1	Minderaser ¹⁹
1	Mjbmrbot ²⁰
1	Monobi ²¹

-
- 1 <https://en.wikibooks.org/wiki/User:ABCD>
 - 2 <https://en.wikibooks.org/wiki/User:Adrignola>
 - 3 <https://en.wikibooks.org/wiki/User:Albmont>
 - 4 <https://en.wikibooks.org/wiki/User:Avicennasis>
 - 5 <https://en.wikibooks.org/wiki/User:Aya>
 - 6 <https://en.wikibooks.org/wiki/User:Az1568>
 - 7 <https://en.wikibooks.org/wiki/User:BenTels>
 - 8 https://en.wikibooks.org/wiki/User:Dirk_H%25C3%25BCnniger
 - 9 <https://en.wikibooks.org/wiki/User:Eric119>
 - 10 <https://en.wikibooks.org/wiki/User:Geocachernemesis>
 - 11 <https://en.wikibooks.org/wiki/User:Guanabot~enwikibooks>
 - 12 <https://en.wikibooks.org/wiki/User:Guanaco>
 - 13 <https://en.wikibooks.org/wiki/User:Hagindaz>
 - 14 <https://en.wikibooks.org/wiki/User:JesseW>
 - 15 <https://en.wikibooks.org/wiki/User:Jguk>
 - 16 <https://en.wikibooks.org/wiki/User:Jomegat>
 - 17 <https://en.wikibooks.org/wiki/User:Kernigh>
 - 18 <https://en.wikibooks.org/wiki/User:Krischik>
 - 19 <https://en.wikibooks.org/wiki/User:Minderaser>
 - 20 <https://en.wikibooks.org/wiki/User:Mjbmrbot>
 - 21 <https://en.wikibooks.org/wiki/User:Monobi>

- 3 Sigma 7²²
- 4 Thenub314²³
- 2 Webaware²⁴
- 2 Xania²⁵

²² https://en.wikibooks.org/wiki/User:Sigma_7

²³ <https://en.wikibooks.org/wiki/User:Thenub314>

²⁴ <https://en.wikibooks.org/wiki/User:Webaware>

²⁵ <https://en.wikibooks.org/wiki/User:Xania>

List of Figures

- GFDL: Gnu Free Documentation License. <http://www.gnu.org/licenses/fdl.html>
- cc-by-sa-3.0: Creative Commons Attribution ShareAlike 3.0 License. <http://creativecommons.org/licenses/by-sa/3.0/>
- cc-by-sa-2.5: Creative Commons Attribution ShareAlike 2.5 License. <http://creativecommons.org/licenses/by-sa/2.5/>
- cc-by-sa-2.0: Creative Commons Attribution ShareAlike 2.0 License. <http://creativecommons.org/licenses/by-sa/2.0/>
- cc-by-sa-1.0: Creative Commons Attribution ShareAlike 1.0 License. <http://creativecommons.org/licenses/by-sa/1.0/>
- cc-by-2.0: Creative Commons Attribution 2.0 License. <http://creativecommons.org/licenses/by/2.0/>
- cc-by-2.0: Creative Commons Attribution 2.0 License. <http://creativecommons.org/licenses/by/2.0/deed.en>
- cc-by-2.5: Creative Commons Attribution 2.5 License. <http://creativecommons.org/licenses/by/2.5/deed.en>
- cc-by-3.0: Creative Commons Attribution 3.0 License. <http://creativecommons.org/licenses/by/3.0/deed.en>
- GPL: GNU General Public License. <http://www.gnu.org/licenses/gpl-2.0.txt>
- LGPL: GNU Lesser General Public License. <http://www.gnu.org/licenses/lgpl.html>
- PD: This image is in the public domain.
- ATTR: The copyright holder of this file allows anyone to use it for any purpose, provided that the copyright holder is properly attributed. Redistribution, derivative work, commercial use, and all other use is permitted.
- EURO: This is the common (reverse) face of a euro coin. The copyright on the design of the common face of the euro coins belongs to the European Commission. Authorised is reproduction in a format without relief (drawings, paintings, films) provided they are not detrimental to the image of the euro.
- LFK: Lizenz Freie Kunst. <http://artlibre.org/licence/lal/de>
- CFR: Copyright free use.

- EPL: Eclipse Public License. <http://www.eclipse.org/org/documents/epl-v10.php>

Copies of the GPL, the LGPL as well as a GFDL are included in chapter Licenses²⁶. Please note that images in the public domain do not require attribution. You may click on the image numbers in the following table to open the webpage of the images in your webbrowser.

²⁶ Chapter 14 on page 127

14 Licenses

14.1 GNU GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <http://fsf.org/>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed. Preamble

The GNU General Public License is a free, copyleft license for software and other kinds of works.

The licenses for most software and other practical works are designed to take away your freedom to share and change the works. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change all versions of a program—to make sure it remains free software for all its users. We, the Free Software Foundation, use the GNU General Public License for most of our software; it applies also to any other work released this way by its authors. You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for them if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs, and that you know you can do these things.

To protect your rights, we need to prevent others from denying you these rights or asking you to surrender the rights. Therefore, you have certain responsibilities if you distribute copies of the software, or if you modify it: responsibilities to respect the freedom of others.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must pass on to the recipients the same freedoms that you received. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

Developers that use the GNU GPL protect your rights with two steps: (1) assert copyright on the software, and (2) offer you this License giving you legal permission to copy, distribute and/or modify it.

For the developers' and authors' protection, the GPL clearly explains that there is no warranty for this free software. For both users' and authors' sake, the GPL requires that modified versions be marked as changed, so that their problems will not be attributed erroneously to authors of previous versions.

Some devices are designed to deny users access to install or run modified versions of the software inside them, although the manufacturer can do so. This is fundamentally incompatible with the aim of protecting users' freedom to change the software. The systematic pattern of such abuse occurs in the area of products for individuals to use, which is precisely where it is most unacceptable. Therefore, we have designed this version of the GPL to prohibit the practice for those products. If such problems arise substantially in other domains, we stand ready to extend this provision to those domains in future versions of the GPL, as needed to protect the freedom of users.

Finally, every program is threatened constantly by software patents. States should not allow patents to restrict development and use of software on general-purpose computers, but in those that do, we wish to avoid the special danger that patents applied to a free program could make it effectively proprietary. To prevent this, the GPL assures that patents cannot be used to render the program non-free.

The precise terms and conditions for copying, distribution and modification follow. TERMS AND CONDITIONS 0. Definitions.

“This License” refers to version 3 of the GNU General Public License.

“Copyright” also means copyright-like laws that apply to other kinds of works, such as semiconductor masks.

“The Program” refers to any copyrightable work licensed under this License. Each licensee is addressed as “you”. “Licensees” and “recipients” may be individuals or organizations.

To “modify” a work means to copy from or adapt all or part of the work in a fashion requiring copyright permission, other than the making of an exact copy. The resulting work is called a “modified version” of the earlier work or a work “based on” the earlier work.

A “covered work” means either the unmodified Program or a work based on the Program.

To “propagate” a work means to do anything with it that, without permission, would make you directly or secondarily liable for infringement under applicable copyright law, except executing it on a computer or modifying a private copy. Propagation includes copying, distribution (with or without modification), making available to the public, and in some countries other activities as well.

To “convey” a work means any kind of propagation that enables other parties to make or receive copies. Mere interaction with a user through a computer network, with no transfer of a copy, is not conveying.

An interactive user interface displays “Appropriate Legal Notices” to the extent that it includes a convenient and prominently visible feature that (1) displays an appropriate copyright notice, and (2) tells the user that there is no warranty for the work (except to the extent that warranties are provided), that licensees may convey the work under this License, and how to view a copy of this License. If the interface presents a list of user commands or options, such as a menu, a prominent item in the list meets this criterion. 1. Source Code.

The “source code” for a work means the preferred form of the work for making modifications to it. “Object code” means any non-source form of a work.

A “Standard Interface” means an interface that either is an official standard defined by a recognized standards body, or, in the case of interfaces specified for a particular programming language, one that is widely used among developers working in that language.

The “System Libraries” of an executable work include anything other than the work as a whole, that (a) is included in the normal form of packaging a Major Component, but which is not part of that Major Component, and (b) serves only to enable use of the work with that Major Component, or to implement a Standard Interface for which an implementation is available to the public in source code form. A “Major Component”, in this context, means a major essential component (kernel, window system, and so on) of the specific operating system (if any) on which the executable work runs, or a compiler used to produce the work, or an object code interpreter used to run it.

The “Corresponding Source” for a work in object code form means all the source code needed to generate, install, and (for an executable work) run the object code and to modify the work, including scripts to control those activities. However, it does not include the work’s System Libraries, or general-purpose tools or generally available free programs which are used unmodified in performing those activities but which are not part of the work. For example, Corresponding Source includes interface definition files associated with source files for the work, and the source code for shared libraries and dynamically linked subprograms that the work is specifically designed to require, such as by intimate data communication or control flow between those subprograms and other parts of the work.

The Corresponding Source need not include anything that users can regenerate automatically from other parts of the Corresponding Source.

The Corresponding Source for a work in source code form is that same work. 2. Basic Permissions.

All rights granted under this License are granted for the term of copyright on the Program, and are irrevocable provided the stated conditions are met. This License explicitly affirms your unlimited permission to run the unmodified Program. The output from running a covered work is covered by this License only if the output, given its content, constitutes a covered work. This License acknowledges your rights of fair use or other equivalent, as provided by copyright law.

You may make, run and propagate covered works that you do not convey, without conditions so long as your license otherwise remains in force. You may convey covered works to others for the sole purpose of having them make modifications exclusively for you, or provide you with facilities for running those works, provided that you comply with the terms of this License in conveying all material for which you do not control copyright. Those thus making or running the covered works for you must do so exclusively on your behalf, under your direction and control, on terms that prohibit them from making any copies of your copyrighted material outside their relationship with you.

Conveying under any other circumstances is permitted solely under the conditions stated below. Sublicensing is not allowed; section 10 makes it unnecessary. 3. Protecting Users’ Legal Rights From Anti-Circumvention Law.

No covered work shall be deemed part of an effective technological measure under any applicable law fulfilling obligations under article 11 of the WIPO copyright treaty adopted on 20 December 1996, or similar laws prohibiting or restricting circumvention of such measures.

When you convey a covered work, you waive any legal power to forbid circumvention of technological measures to the extent such circumvention is effected by exercising rights under this License with respect to the covered work, and you disclaim any intention to limit operation or modification of the work as a means of enforcing, against the work’s users, your or third parties’ legal rights to forbid circumvention of technological measures. 4. Conveying Verbatim Copies.

You may convey verbatim copies of the Program’s source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice; keep intact all notices stating that this License and any non-permissive terms added in accord with section 7 apply to the code; keep intact all notices of the absence of any warranty; and give all recipients a copy of this License along with the Program.

You may charge any price or no price for any copy that you convey, and you may offer support for a warranty protection for a fee. 5. Conveying Modified Source Versions.

You may convey a work based on the Program, or the modifications to produce it from the Program, in the form of source code under the terms of section 4, provided that you also meet all of these conditions:

* a) The work must carry prominent notices stating that you modified it, and giving a relevant date. * b) The work must carry prominent notices stating that it is released under this License and any conditions added under section 7. This requirement modifies the requirement in section 4 to “keep intact all notices”. * c) You must license the entire work, as a whole, under this License to anyone who comes into possession of a copy. This License will therefore apply, along with any applicable section 7 additional terms, to the whole of the work, and all its parts, regardless of how they are packaged. This License gives no permission to license the work in any other way, but it does not invalidate such permission if you have separately received it. * d) If the work has interactive user interfaces, each must display Appropriate Legal Notices; however, if the Program has interactive interfaces that do not display Appropriate Legal Notices, your work need not make them do so.

A compilation of a covered work with other separate and independent works, which are not by their nature extensions of the covered work, and which are not combined with it such as to form a larger program, in or on a volume of a storage or distribution medium, is called an “aggregate” if the compilation and its resulting copyright are not used to limit the access or legal rights of the compilation’s users beyond what the individual works permit. Inclusion of a covered work in an aggregate does not cause this License to apply to the other parts of the aggregate. 6. Conveying Non-Source Forms.

You may convey a covered work in object code form under the terms of sections 4 and 5, provided that you also convey the machine-readable Corresponding Source under the terms of this License, in one of these ways:

* a) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by the Corresponding Source fixed on a durable physical medium customarily used for software interchange. * b) Convey the object code in, or embodied in, a physical product (including a physical distribution medium), accompanied by a written offer, valid for at least three years and valid for as long as you offer spare parts or customer support for that product model, to give anyone who possesses the object code either (1) a copy of the Corresponding Source for all the software in the product that is covered by this License, on a durable physical medium customarily used for software interchange, for a price no more than your reasonable cost of physically performing this conveying of source, or (2) access to copy the Corresponding Source from a network server at no charge. * c) Convey individual copies of the object code with a copy of the written offer to provide the Corresponding Source. This alternative is allowed only occasionally and noncommercially, and only if you received the object code with such an offer, in accord with subsection 6b. * d) Convey the object code by offering access from a designated place (gratis or for a charge), and offer equivalent access to the Corresponding Source in the same way through the same place at no further charge. You need not require recipients to copy the Corresponding Source along with the object code. If the place to copy the object code is a network server, the Corresponding Source may be on a

different server (operated by you or a third party) that supports equivalent copying facilities, provided you maintain clear directions next to the object code saying where to find the Corresponding Source. Regardless of what server hosts the Corresponding Source, you remain obligated to ensure that it is available for as long as needed to satisfy these requirements. * e) Convey the object code using peer-to-peer transmission, provided you inform other peers where the object code and Corresponding Source of the work are being offered to the general public at no charge under subsection 6d.

A separable portion of the object code, whose source code is excluded from the Corresponding Source as a System Library, need not be included in conveying the object code work.

A “User Product” is either (1) a “consumer product”, which means any tangible personal property which is normally used for personal, family, or household purposes, or (2) anything designed or sold for incorporation into a dwelling. In determining whether a product is a consumer product, doubtful cases shall be resolved in favor of coverage. For a particular product received by a particular user, “normally used” refers to a typical or common use of that class of product, regardless of the status of the particular user or of the way in which the particular user actually uses, or expects to use, is expected to use, the product. A product is a consumer product regardless of whether the product has substantial commercial, industrial or non-consumer uses, unless such uses represent the only significant mode of use of the product.

“Installation Information” for a User Product means any methods, procedures, authorization keys, or other information required to install and execute modified versions of a covered work in that User Product from a modified version of its Corresponding Source. The information must suffice to ensure that the continued functioning of the modified object code is in no case prevented or interfered with solely because modification has been made.

If you specify an object code work under this section in, or with, or specifically for use in, a User Product, and the conveying occurs as part of a transaction in which the right of possession and use of the User Product is transferred to the recipient in perpetuity or for a fixed term (regardless of how the transaction is characterized), the Corresponding Source conveyed under this section must be accompanied by the Installation Information. But this requirement does not apply if neither you nor any third party retains the ability to install modified object code on the User Product (for example, the work has been installed in ROM).

The requirement to provide Installation Information does not include a requirement to continue to provide support service, warranty, or updates for a work that has been modified or installed by the recipient, or for the User Product in which it has been modified or installed. Access to a network may be denied when the modification itself materially and adversely affects the operation of the network or violates the rules and protocols for communication across the network.

Corresponding Source conveyed, and Installation Information provided, in accord with this section must be in a format that is publicly documented (and with an implementation available to the public in source code form), and must require no special password or key for unpacking, reading or copying. 7. Additional Terms.

“Additional permissions” are terms that supplement the terms of this License by making exceptions from one or more of its conditions. Additional permissions that are applicable to the entire Program shall be treated as though they were included in this License, to the extent that they are valid under applicable law. If additional permissions apply only to part of the Program, that part may be used separately under those permissions, but the entire Program remains governed by this License without regard to the additional permissions.

When you convey a copy of a covered work, you may at your option remove any additional permissions from that copy, or from any part of it. (Additional permissions may be written to require their own removal in certain cases when you modify the work.) You may place additional permissions on material, added by you to a covered work, for which you have or can give appropriate copyright permission.

Notwithstanding any other provision of this License, for material you add to a covered work, you may (if authorized by the copyright holders of that material) supplement the terms of this License with terms:

* a) Disclaiming warranty or limiting liability differently from the terms of sections 15 and 16 of this License; or * b) Requiring preservation of specified reasonable legal notices or author attributions in that material or in the Appropriate Legal Notices displayed by works containing it; or * c) Prohibiting misrepresentation of the origin of that material, or requiring that modified versions of such material be marked in reasonable ways as different from the original version; or * d) Limiting the use for publicity purposes of names of licensors or authors of the material; or * e) Declining to grant rights under trademark law for use of some trade names, trademarks, or service marks; or * f) Requiring indemnification of licensors and authors of that material by anyone who conveys the material (or modified versions of it) with contractual assumptions of liability to the recipient, for any liability that these contractual assumptions directly impose on those licensors and authors.

All other non-permissive additional terms are considered “further restrictions” within the meaning of section 10. If the Program as you received it, or any part of it, contains a notice stating that it is governed by this License along with a term that is a further restriction, you may remove that term. If a license document contains a further restriction but permits relicensing or conveying under this License, you may add to a covered work material governed by the terms of that license document, provided that the further restriction does not survive such relicensing or conveying.

If you add terms to a covered work in accord with this section, you must place, in the relevant source files, a statement of the additional terms that apply to those files, or a notice indicating where to find the applicable terms.

Additional terms, permissive or non-permissive, may be stated in the form of a separately written license, or stated as exceptions; the above requirements apply either way. 8. Termination.

You may not propagate or modify a covered work except as expressly provided under this License. Any attempt otherwise to propagate or modify it is void, and will automatically terminate your rights under this License (including any patent licenses granted under the third paragraph of section 11).

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates

your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License for any work (from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, you do not qualify to receive new licenses for the same material under section 10. 9. Acceptance Not Required for Having Copies.

You are not required to accept this License in order to receive or run a copy of the Program. Ancillary propagation of a covered work occurring solely as a consequence of using peer-to-peer transmission to receive a copy likewise does not require acceptance. However, nothing other than this License grants you permission to propagate or modify any covered work. These actions infringe copyright if you do not accept this License. Therefore, by modifying or propagating a covered work, you indicate your acceptance of this License to do so. 10. Automatic Licensing of Downstream Recipients.

Each time you convey a covered work, the recipient automatically receives a license from the original licensors, to run, modify and propagate that work, subject to this License. You are not responsible for enforcing compliance by third parties with this License.

An “entity transaction” is a transaction transferring control of an organization, or substantially all assets of one, or subdividing an organization, or merging organizations. If propagation of a covered work results from an entity transaction, each party to that transaction who receives a copy of the work also receives whatever licenses to the work the party’s predecessor in interest had or could give under the previous paragraph, plus a right to possession of the Corresponding Source of the work from the predecessor in interest, if the predecessor has it or can get it with reasonable efforts.

You may not impose any further restrictions on the exercise of the rights granted or affirmed under this License. For example, you may not impose a license fee, royalty, or other charge for exercise of rights granted under this License, and you may not initiate litigation (including a cross-claim or counterclaim in a lawsuit) alleging that any patent claim is infringed by making, using, selling, offering for sale, or importing the Program or any portion of it. 11. Patents.

A “contributor” is a copyright holder who authorizes use under this License of the Program or a work on which the Program is based. The work thus licensed is called the contributor’s “contributor version”.

A contributor’s “essential patent claims” are all patent claims owned or controlled by the contributor, whether already acquired or hereafter acquired, that would be infringed by some manner, permitted by this License, of making, using, or selling its contributor version, but do not include claims that would be infringed only as a consequence of further modification of the contributor version. For purposes of this definition, “control” includes the right to grant patent sublicenses in a manner consistent with the requirements of this License.

Each contributor grants you a non-exclusive, worldwide, royalty-free patent license under the contributor’s essential patent claims, to make, use, sell, offer for sale, import and otherwise run, modify and propagate the contents of its contributor version.

In the following three paragraphs, a “patent license” is any express agreement or commitment, however denominated, not to enforce a patent (such as an express promise to practice a patent or covenant not to sue for patent infringement). To “grant” such a patent license to a party means to make such an agreement or commitment not to enforce a patent against the party.

If you convey a covered work, knowingly relying on a patent license, and the Corresponding Source of the work is not available for anyone to copy, free of charge and under the terms of this License, through a publicly available network server or other readily accessible means, then you must either (1) cause the Corresponding Source to be so available, or (2) arrange to deprive yourself of the benefit of the patent license for this particular work, or (3) arrange, in a manner consistent with the requirements of this License, to extend the patent license to downstream recipients. “Knowingly relying” means you have actual knowledge that, but for the patent license, your conveying the covered work in a country, or your recipient’s use of the covered work in a country, would infringe one or more identifiable patents in that country that you have reason to believe are valid.

If, pursuant to or in connection with a single transaction or arrangement, you convey, or propagate by procuring conveyance of, a covered work, and grant a patent license to some of the parties receiving the covered work authorizing them to use, propagate, modify or convey a specific copy of the covered work, then the patent license you grant is automatically extended to all recipients of the covered work and works based on it.

A patent license is “discriminatory” if it does not include within the scope of its coverage, prohibits the exercise of, or is conditioned on the non-exercise of one or more of the rights that are specifically granted under this License. You may not convey a covered work if you are a party to an arrangement with a third party that is in the business of distributing software, under which you make payment to the third party based on the extent of your activity of conveying the work, and under which the third party grants, to any of the parties who would receive the covered work from you, a discriminatory patent license (a) in connection with copies of the covered work conveyed by you (or copies made from those copies), or (b) primarily for and in connection with specific products or compilations that contain the covered work, unless you entered into that arrangement, or that patent license was granted, prior to 28 March 2007.

Nothing in this License shall be construed as excluding or limiting any implied license or other defenses to infringement that may otherwise be available to you under applicable patent law. 12. No Surrender of Others’ Freedom.

If conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot convey a covered work so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not convey it at all. For example, if you agree to terms that obligate you to collect a royalty for further conveying from those to whom you convey the Program, the only way you could satisfy

both those terms and this License would be to refrain entirely from conveying the Program. 13. Use with the GNU Affero General Public License.

Notwithstanding any other provision of this License, you have permission to link or combine any covered work with a work licensed under version 3 of the GNU Affero General Public License into a single combined work, and to convey the resulting work. The terms of this License will continue to apply to the part which is the covered work, but the special requirements of the GNU Affero General Public License, section 13, concerning interaction through a network will apply to the combination as such. 14. Revised Versions of this License.

The Free Software Foundation may publish revised and/or new versions of the GNU General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies that a certain numbered version of the GNU General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that numbered version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the GNU General Public License, you may choose any version ever published by the Free Software Foundation.

If the Program specifies that a proxy can decide which future versions of the GNU General Public License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Program.

14.2 GNU Free Documentation License

Version 1.3, 3 November 2008

Copyright (c) 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed. 0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference. 1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “Document”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “you”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “Modified Version” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “Secondary Section” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “Invariant Sections” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “Cover Texts” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “Transparent” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “Opaque”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “Title Page” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

The “publisher” means any person or entity that distributes copies of the Document to the public.

A section “Entitled XYZ” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses

Later license versions may give you additional or different permissions. However, no additional obligations are imposed on any author or copyright holder as a result of your choosing to follow a later version. 15. Disclaimer of Warranty.

THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION. 16. Limitation of Liability.

IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MODIFIES AND/OR CONVEYS THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. 17. Interpretation of Sections 15 and 16.

following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “Acknowledgements”, “Dedications”, “Endorsements”, or “History”). To “Preserve the Title” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties; any other indication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License. 2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies. 3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document’s license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first one listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general networking-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document. 4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version precisely as the full-text version of the Modified Version filling the role of the Document, this licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

* A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission. * B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if that is fewer than five), unless they release you from this requirement. * C. State on the Title page the name of the publisher of the Modified Version, as the publisher. * D. Preserve all the copyright notices of the Document. * E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices. * F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the medium below. * G. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the “History” section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission. * K. For any section Entitled “Acknowledgements” or “Dedications”, Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein. * L. Preserve all the Invariant Sections of the Document, unaltered in their text and

if the disclaimer of warranty and limitation of liability provided above cannot be made legal local effect according to their terms, reviewing courts shall apply local law that most closely approximates an absolute waiver of all civil liability in connection with the Program, unless a warranty or assumption of liability accompanies a copy of the Program in return for a fee.

END OF TERMS AND CONDITIONS
How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively state the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

```
<one line to give the program’s name and a brief idea of what it does.>
Copyright (C) <year> <name of author>
```

This program is free software: you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation, either version 3 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

```
in their titles. Section numbers or the equivalent are not considered part of the section titles. * M. Delete any section Entitled “Endorsements”. Such a section may not be included in the Modified Version. * N. Do not retittle any existing section to be Entitled “Endorsements” or to conflict in title with any Invariant Section. * O. Preserve any Warranty Disclaimers.
```

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version’s license notice. These titles must be distinct from any other section titles.

You may add a section Entitled “Endorsements”, provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or of the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added (by or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version. 5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”. 6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document. 7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate. 8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title

You should have received a copy of the GNU General Public License along with this program. If not, see <<http://www.gnu.org/licenses/>>.

Also add information on how to contact you by electronic and paper mail.

If the program does terminal interaction, make it output a short notice like this when it starts in an interactive mode:

```
<program> Copyright (C) <year> <name of author> This program
comes with ABSOLUTELY NO WARRANTY; for details type 'show w'.
This is free software, and you are welcome to redistribute it under
certain conditions; type 'show c' for details.
```

The hypothetical commands ‘show w’ and ‘show c’ should show the appropriate parts of the General Public License. Of course, your program’s commands might be different; for a GUI interface, you would use an “about box”.

You should also get your employer (if you work as a programmer) or school, if any, to sign a “copyright disclaimer” for the program, if necessary. For more information on this, and how to apply and follow the GNU GPL, see <<http://www.gnu.org/licenses/>>.

The GNU General Public License does not permit incorporating your program into proprietary programs. If you permit it as a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License. But first, please read <<http://www.gnu.org/philosophy/why-not-lgpl.html>>.

(section 1) will typically require changing the actual title. 9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it. 10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <<http://www.gnu.org/copyleft/>>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy’s public statement of acceptance of a version permanently authorizes you to choose that version for the Document. 11. RELICENSING

“Massive Multiauthor Collaboration Site” (or “MMC Site”) means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A “Massive Multiauthor Collaboration” (or “MMC”) contained in the site means any set of copyrightable works thus published on the MMC site.

“CC-BY-SA” means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

“Incorporate” means to publish or republish a Document, in whole or in part, as part of another Document.

An MMC is “eligible for relicensing” if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing. ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C) YEAR YOUR NAME. Permission is granted to copy,
distribute and/or modify this document under the terms of the GNU
Free Documentation License, Version 1.3 or any later version pub-
lished by the Free Software Foundation; with no Invariant Sections,
no Front-Cover Texts, and no Back-Cover Texts. A copy of the License
is included in the section entitled “GNU Free Documentation License”.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the “with ... Texts.” line with this:

```
with the Invariant Sections being LIST THEIR TITLES, with the
Front-Cover Texts being LIST, and with the Back-Cover Texts being
LIST.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

14.3 GNU Lesser General Public License

GNU LESSER GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright © 2007 Free Software Foundation, Inc. <<http://fsf.org/>>

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

0. Additional Definitions.

As used herein, “this License” refers to version 3 of the GNU Lesser General Public License, and the “GNU GPL” refers to version 3 of the GNU General Public License.

“The Library” refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An “Application” is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the “Linked Version”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

1. Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

2. Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

* a) under this License, provided that you make a good faith effort to ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or * b) under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

3. Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

* a) Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License. * b) Accompany the object code with a copy of the GNU GPL and this license document.

4. Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

* a) Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License. * b) Accompany the Combined Work with a copy of the GNU GPL and this license document. * c) For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document. * d) Do one of the following: o 0) Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source. o 1) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (a) uses at run time a copy of the Library already present on the user’s computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version. * e) Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 4d0, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option 4d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

5. Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

* a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License. * b) Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

6. Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy’s public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.