

# A beginner's guide to the Computing Facilities of Bristol Astrophysics Group

(version of October 6, 2009)

This document briefly lists the features of what is known as the Bristol Starlink<sup>1</sup> node and gives some basic UNIX commands to allow you to get started using it. Those with a knowledge of UNIX will probably want to read the first section and skim the rest: real beginners will probably want to skim the first section and then come back to it. Beginners are also advised to consult the Starlink User Guide and Starlink User notes (particularly SUN/145, an introduction to UNIX). New users should read the local documentation (<http://www.star.bris.ac.uk/local/>) which is updated regularly and contains this document.

## 1 Hardware

### 1.1 Computers

The node is made up of various servers, the most important of which is called sol and a number of desktop computers connected by a fast network.

The workstations all run Unix and are mostly PCs running Linux although we still have some workstations running the Solaris operating system and MacOS is beginning to appear. The table below lists them:

---

<sup>1</sup>The Starlink organisation which provided computing support and software for UK astronomy was disbanded by PPARC in June 2005

Machine	Type	Main user	RAM (Mb)	OS
sol	Sun Enterprise 250	4.13(cubicle 4.13)	1024	Solaris
centaurus	Sun Ultra 10	Mark Birkinshaw's office (3.15)	128	Solaris
lacerta	Sun Blade 100	Diana Worrall's office (3.11)	896	Solaris
virgo	Sun Ultra 10	Steve Phillipps' office (4.12)	128	Solaris
andromeda	Xeon 3GHz	Mark Taylor (4.37)	2gb	Linux
apus	P4 2.8GHz	Poemwai	512	Linux
aquarius	P4 2.8GHz	Observatory	2gb	Linux
aquila	P4 2.8GHz	login server	512	Linux
astrofft	Athlon 1GHz PC	Observatory	256	Linux
auriga	Pentium D dual core	server	2gb	Linux
bootes	Athlon64	Caz Winkworth 5.13	2gb	Linux
cetus	P4 2.8GHz	server (4.14)	2048	Linux
coldrick	P4 ?	observatory	?	Windows
coma	8 core 2.33GHz	Paul Giles	2048	Linux
crux	Athlon64 x2 2200 PC	Luke Davies 4.13	2gb	Linux
cygnus	P4 3.2GHz	Lars Koenigs 4.13	2gb	Linux
delphinus	P4 2.8GHz	Mike Masheder	2gb	Linux
dorado	P4 3.2GHz	Avon Huxor (5.13)	4gb	Linux
eridanus	P4 3.2GHz	Katy Lancaster (5.13)	2048	Linux
gemini	P4 2.8GHz	general use 4.13	1.5gb	Linux
grus	P4 3GHz	Diana Worrall	2048	Linux
hydra	Intel Core2duo	Kian Momtahan	2gb	Linux
libra	P4 2.4GHz PC	general use	512	Linux
lupus	P4 2.8GHz	Elizabeth Stanway	3.5gb	Linux
lyra	P4 2.8GHz	Amelia Bliss (4.13)	512	Linux
monoceros	Athlon 3000+ Athlon64	general use	1024	Linux
orion	P4 3GHz	Aziz	2gb	Linux
pavo	P4 2.8GHz	Rhys Morris (4.33)	1.5gb	Linux
perseus	8 core 2.33GHz	Mark Birkinshaw	4gb	Linux
pisces	P4 2.4GHz	general use	640	Linux
procyon	P3 450MHz	general use	384mb	Windows
scutum	P4 2.8GHz	General Use (4.13)	1.5gb	Linux
shaula	P4 2.4 GHz	General use	512	Linux
taurus	P4 3GHz	James Price	1gb	Linux
vela	8 core 2.0GHz	Ben Maughan	4096	Linux
vulpecula	P4 2.4GHz	visitor machine	512	Linux
aries	4 core Mac Pro	Malcolm Bremer	4gb	MacOS
Sagittarius	8 core Mac Pro	Andy Young	8gb	MacOS

Many users have Mac or dual boot Linux/windows laptops, these are not listed here.

## 1.2 Disks

The `/export/home` disk is called the home disk. It is backed up daily to several different places. This disk is for important documents, programs, theses etc, and your usage should not exceed 2gb. Once you have finished a project, you should archive the files to CD or DVD and move the files to another disk, to avoid frequent backing up of files that do not change. Data should not be kept on the home disk, you should use a data disk instead.

Each of our machines should have a large disk or disks attached to it, use these as data disks to store your data.

A large amount of disk space is available on the node. The table below lists some of the disks available, and who has priority on those disks.

Disk name	Local to	Size (Gb)	Availability
/export/home	sol	45	All
/data/project	sol	36	project students
/usr/local	sol	20	software
/data/centaurus1,2,3,4,5	centaurus	70	Primarily Mark Birkinshaw
/data/perseus1,2,5	perseus	1500	Primarily Mark Birkinshaw
/data/dmw1,2,3	lacerta	10	Primarily Diana Worrall
/data/grus1	grus	360	Primarily Diana Worrall
/data/lacerta1,2,3,4,5	lacerta	100	Primarily Diana Worrall
/data/delphinus1	delphinus	500	Mike Masheder
/data/scratch1	sol	4	All
/data/astrofft	astrofft	30	Observatory
/data/aquarius1	aquarius	500	Observatory
/data/aries1	aries	2000	Malcolm Bremer
/data/vela1,2	vela	2000	Ben Maughan
/data/lupus1,2,3,2	vela	120	Bremer/Stamway/Davies
/data/eridanus1	eridanus	200	Katy Lancaster
/data/pavo1	pavo	300	Mainly Rhys Morris

All other disks should be usable by everyone. But please check with the main user of the machine before using up vast quantities of their disk space.

The first column above gives the UNIX path to the disk. Your home directory (the directory you start in when you log in) is on /export/home, in /export/home/<your user name>. On other general usage data disks you may create yourself a working directory, usually named after you. E.g.

```
> cd /data/crux2
> mkdir fred
```

sets up work space on disk crux2 for the user fred. You are encouraged to keep data on one of the disks other than /export/home. Consult an expert to see which disks are currently suitable.

Each disk is physically attached to (local to) a particular workstation, although all are available across the network. When doing disk-intensive work (e.g. data reduction) it is very important that you should use a disk local to the machine you are running the disk-intensive processing on; accessing a disk across the network is very much slower than accessing it locally, and the load put on the network by intensive remote use will slow everybody else down too.

You can find out the current usage of all the disks with the commands `df -k` or `diskuse`. Individual disks have a usage list on the root directory (e.g. `cat /export/home/usage`).

Some disks are regularly backed up while others are not. Please consult the local web page about backups for more information.

### 1.3 Printers

Two printers are attached to the node. Astroprint is a black and white laser printer and astrocolour is colour laser printer. Both print Postscript files and plain text.

On Linux, the command to print is `lpr` (on Suns it is `lp`). E.g. to send the file `test.ps` to the laser printer,

```
> lpr -P astroprint test.ps
```

Astroprint is the default printer, so to print to this you may omit the `-P` option:

```
lpr test.ps
```

You can check the status of the print queues with `lpq -P astroprint`.

Both printers can print on transparencies for production of overheads. The laser printer uses standard photocopier printable transparencies.

Take great care not to mix up transparencies, **the laser printer can be seriously damaged by the wrong kind of transparency**. Always ask if you are unsure.

Loading the printers with paper or changing toner or ink cartridges is easy, but if you haven't done it before ask someone experienced to show you how first.

As a backup printing service we can use the Physics departmental printer queues `document` (double sided) and `laserprint` (single sided). Use the `lpr -P document testfile` and `lpr -P laserprint` commands respectively. Output is produced in the photocopier room on the 3rd floor. The same printer can also produce A3 output using the `lpr -P a3print testfile` command

For colour laser printing use either `astrocolour` or the departmental colour

laser printer in room 1.14, the queue name is colour114. Printing to `kprinter` instead of `lpr` from `acoread`, `OpenOffice` or even from the command line gives you a great deal of flexibility regarding how your printout looks. You can save paper by putting two pages of output on one page of A4 for example. There is a global `print` command which does this using the `enscript` program (see local web pages for more info).

## 1.4 CDs, DVDs and USB Memory Sticks

A machine with a DVD writer can write DVDs and CDs, most of our machines have this capability. Instructions for writing CDs and DVDs can be found on our local documentation pages. As regards reading optical media, on a Linux machine, a window should pop up containing the files on the DVD. If this does not happen, you may need to manually mount the disk, again see the local documentation

## 1.5 Tape

A number of machines have tape drives, as follows:

Machine	Location	Drive type	Device name
sol	4.13	Exabyte Mammoth	/dev/rmt/2
pavo	4.33	DDS-2 DAT drive	/dev/rmt/0
pavo	4.33	DLT drive	/dev/rmt/1
ask	4.s13	Exabyte drive	/dev/st0
ask	4.13	DDS-3 DAT drive	/dev/st1
lacerta	3.11	DDS-2 DAT drive	/dev/rmt/0
lacerta	3.11	DDS-4 DAT drive	/dev/rmt/1

You must be logged in to the machine to which a tape drive is physically attached in order to use it. The UNIX tape archiving command is `tar`. E.g.

```
> tar -cf /dev/rmt/0 wombats
```

on sol, archives the directory wombats to the tape in the DAT drive.

```
> tar -cf /dev/st0 wingnuts
```

on bootes, archives the directory wingnuts to the tape in the DAT drive.

```
> tar -xf /dev/rmt/1
```

on sol extracts the files from the tape whose device name is /dev/rmt/1.

```
> tar -tvf /dev/st1
```

lists the files on the tape in bootes' Exabyte drive. For more help on using `tar`, see `man tar`.

By default the tape drives automatically rewind at the end of every operation. This can be a nuisance if, for example, you want to archive several files on the tape sequentially. To prevent it refer to the tape devices suffixed with an `n` or, on Linux, prefixed by an `n`;

```
> tar -cf /dev/rmt/0n wombats
```

```
> tar -cf /dev/nst0 wombats
```

both commands archive the directory `wombats` to tape and do not rewind afterwards on Solaris and Linux.

The UNIX command `mt` allows you to manipulate tape devices. To eject the tape in the DAT drive, for example, you would use

```
mt -f /dev/rmt/0 offline
```

For more possible commands see `man mt`.

## 1.6 Floppy disks

On a Linux machine machine, you can use the commands `mdir`, `mcopy`, `mren` (`rename`), `mdelete` etc to manipulate files on an MS DOS format floppy. Floppy disks can also be formatted for the Linux ext2 file system, users can mount disks using the `mount /mnt/floppy` and use the usual Linux commands, before typing `umount /mnt/floppy` to flush the data to disk before removing the floppy from the drive.

Some Suns (lacerta, scorpius, leo, virgo and centaurus) have floppy disk drives. To mount a floppy disk for reading or writing, place it in the drive and type `volcheck`. The disk should then appear at `/floppy/floppy0`. Before manually ejecting the disk, type `eject` to ensure that all data has been written out to it.

## 2 Basic UNIX/Linux commands

### 2.1 Files

In UNIX, files store data and directories contain files. Examples of some basic file- and directory-manipulating commands are listed below.

Command	Effect
<code>ls</code>	List the files in the current working directory
<code>ls -al</code>	Get a long listing of all files
<code>ls /data/scratch</code>	List files in the directory /data/scratch
<code>rm file1 [file2 file3...]</code>	Delete one or more files
<code>cp file1 file2</code>	Copy file1 to file2
<code>cp file1 file2 ... newdir</code>	Copy files 1, 2... to directory newdir
<code>mv file1 file2</code>	Rename file1 to file2
<code>mv file1 file ... newdir</code>	Move files 1, 2... to directory newdir
<code>cat file1</code>	Print file1 on the terminal
<code>cat file1 file2... &gt; newfile</code>	Join files 1, 2... together and put the result in newfile.
<code>cd dir</code>	Change the current working directory to be dir. <code>cd ..</code> to go back up one level.
<code>mkdir dir</code>	Make new directory dir
<code>rmdir dir</code>	Delete the (empty) directory dir
<code>less file</code>	View file on the terminal, moving around with cursor keys

The `man` command gives access to on-line documentation. E.g. `man ls` would tell you about all the options for the `ls` command. `man -k wombats` searches for all commands relating to wombats.

### 2.2 Networking

The `ssh` command and its synonym the `ssh` command allow you to log in to other machines in the local cluster. You do not need to specify the domain name (.star.bris.ac.uk) part of the Internet address.

```
orion> slogin sol
sol>
```

Log out with `logout` or `exit` as usual.

`ssh` allows you to execute a single command or series of commands on a remote machine without having to log in; so

```
orion> ssh sol ps -e -u fred
```

allows a user on orion to list fred's processes on sol.

It is recommended that you use secure shell programs whenever possible, because of their numerous security features.

In ancient times, other tools, which did not use encryption were used to access systems. You may come across the `telnet` and `ftp` commands, anonymous ftp is still used, but non-anonymous ftp and telnet should never be used.

On Linux; mozilla, firefox and Konqueror can all be used to browse the web, on Suns mozilla is the recommended browser. You will need X-windows running to use a web browser.

E-mail clients available on the system include `pine`, `elm`, `mutt`, `evolution`, `thunderbird` and `mozilla`. The choice is a matter of preference. (Note that you will not be able to send e-mail from the Starlink system until you are registered with the University Computing Service.) Please do **not** use e-mail to transfer large files between machines.

Transferring files can again be done in two ways. `scp` is easiest for transferring files between two machines in the Starlink cluster (though it's not usually necessary to do this since almost all disks are visible on other machines). E.g.

```
scp sol:/some/disk/file1 file2
```

copies the remote file from sol to file2 on your local machine, while

```
scp localfile1 leo:/some/other/disk
```

copies a local file to leo. For file transfer outside the cluster, try scp or you can make your files available via our anonymous ftp server or from your web pages, see below for more details on these options.

If you are transferring a large number of files, consider using `rsync`. It can



compare two directories on separate machines and copy only the files that have changed, more details and examples can be found on our local pages.

You can also use a remote desktop to access our systems, see the local documentation on VNC.

### 3 Anonymous ftp

The group runs an anonymous ftp server. If you make a directory for yourself in `/data/ftp/pub/` using `mkdir /data/ftp/pub/<your username>` then your files will be available by anonymous ftp from `ftp.star.bris.ac.uk` in the `/pub/<your username>` directory or on the web at `ftp://ftp.star.bris.ac.uk/pub/<your username>`.

When retrieving files from a remote ftp server, you connect using the command `ftp ftp.some.server.com`. When prompted for your username, you should reply 'anonymous' and use your e-mail address as the password. Once you have logged in, you can use a limited set of file management commands: `dir` or `ls` to get a directory, `cd` to change remote directory, and `get` and `put` to transfer files from and to the remote machine respectively. `mget` and `mput` allow you to use wild-card characters (e.g. `mget *` to get all files in a remote directory). Before you transfer binary files (anything other than plain text) be sure to use the command `bin` to specify that the program should not try to translate any characters, otherwise the files will be corrupted.

### 4 Web pages

You may set up your own web page on the Starlink node's web server (`http://www.star.bris.ac.uk`). Go to the directory `/export/home/WWW` and create a new directory with the same name as your user id:

```
> cd /export/home/WWW
> mkdir fred
```

Create a file named `index.html`; this will be your home page. The URL of this page is `http://www.star.bris.ac.uk/<your user id>`. Other files you put in the newly created directory will be available as `http://www.star.bris.ac.uk/<your user id>/<filename>`.

Some local computing information, including an HTML version of this guide, is available at <http://www.star.bris.ac.uk/local/>.

You may export information to others on the Web. If you wish to make large files available to the world (e.g. sending data to a collaborator) then it is not advisable to copy them to your area in `/export/home/WWW`. Instead set up a symbolic link to the file you wish to export using the `ln -s` command:

```
> cd /export/home/WWW/fred
> ln -s /data/crux2/fred/a-very-large-file.dat large-file.dat
```

The first filename you give here is the name of an existing file, the second is the name of a link that will be created. The link is simply a pointer to the existing file and takes up no space. In the example above, someone from outside the group would then be able to access Fred's very large file as <http://www.star.bris.ac.uk/fred/large-file.dat>.

Files must be world-readable to be exported in this way. See the man pages or other documentation for information on `chmod` and file protection.

(Symbolic links may be used anywhere, not just in `/export/home/WWW`. You may find them useful in organising your filesystem efficiently.)

An alternative to using the Web to export files is to use our own anonymous ftp server. For more information, see <http://www.star.bris.ac.uk/local/ftp.html>.

## 5 X-Windows

The Windowing environment on Linux and Suns is generically called X-Windows, or X11. There are many desktop systems that use X-windows, you may encounter Gnome or KDE on the Linux machines for example.

Some useful programs you could have on your desktop could be; a clock, calendar, web browser and mail client. Some of these can be started from icons on the desktop or the taskbar while some application programs you will have to start from the command line.

If you type a command such as `mozilla` into a terminal window, your command prompt won't return until the application has finished. It's normally sensible to specify that applications of this sort should run in the background. Do this by putting an `&` after the command.

```
sol> mozilla http://www.star.bris.ac.uk/ &  
sol>
```

## 6 Important configuration files

There are several files in your home directory that you may need to edit in order to customise the way you access the system. Filenames of configuration files in UNIX start with a dot (.); such files are hidden from `ls` unless you give it the argument `-a`:

```
> ls -a  
. .Xauthority .logged-in .gnome .kde  
.. .cshrc .login
```

The most important of these files is Your `.cshrc` file. It contains commands which are run whenever you start a new shell (when you log in for the first time, open a new window, `slogin` to a new machine and so on. This is the place to put configuration commands. For example, you may want to add aliases.

```
alias lsall ls -a
```

Setting environment variables is best done in this file. The command to do this is `setenv`:

```
> setenv EDITOR emacs
```

sets the variable `EDITOR` to 'emacs'. You can see what environment variables you currently have set with `printenv`.

Important environment variables include:

Name	Effect
DISPLAY	The name of the X display that windowing commands will use by default. If you get ‘Can’t open display’ errors, check this. DISPLAY should be a fully qualified host name followed by :0.0; <pre>&gt; echo \$DISPLAY scorpius.star.bris.ac.uk:0.0</pre>
EDITOR	The name of your default editor.
PATH	A colon-separated list of directories to search for commands. The default path is quite minimal.
LD_LIBRARY_PATH	As PATH, but for libraries.
MANPATH	As PATH, but for manual pages.

The `.login` file is similar to the `.cshrc` but is executed only when you start a ‘login shell’. It’s probably easiest to put all important configuration commands in the `.cshrc`.

## 7 Text and document processing

### 7.1 Editors

Various applications are available to edit text on the Starlink system, including (in order of sophistication) `pico`, `vi`, `nedit`, `gedit`, `emacs` and `xemacs`. Emacs can be used in its own X-window or you can run it in a terminal using the `-nw` flags (useful when logging on to remote systems).

Emacs is the recommended text editor; it is powerful and widely used. You can start it with the name of a file to edit:

```
sol> emacs guide.tex &
sol>
```

Emacs is a full-screen editor. You can move around with the cursor keys, page up/page down keys, mouse etc and insert and delete text in the obvious way. Emacs is partially menu-driven; use the mouse to select items from the drop-down menus at the top of the window.

Some useful emacs keystrokes are:

Key combination	Effect
control-X control-C	Quit emacs, prompting to save any changed files
control-X control-S	Save the file currently being edited to disk
control-X control-F	Open a new file
control-A	Move to start of line
control-E	Move to end of line
control-K	Delete from current cursor position to end of line, or (if at end of line already) close up.
control-S	Search forward
control-S control-S	Repeat previous forward search
control-R	Search backwards
control-x u	Undo previous action
control-space	Set mark
control-W	Cut text between cursor and mark
control-Y	Paste previously cut text at cursor
ESC <	Go to top of document
ESC >	Go to end of document
ESC q	Word-wrap current paragraph

There is a lot of emacs information on the web and emacs has its own extensive on-line documentation; see the ‘Help’ menu.

Note that it’s not efficient to start a new version of emacs whenever you want to edit a new file. Instead use control-X control-F to open a new file. Use the ‘Buffers’ menu or control-x o to switch between buffers.

Emacs is configurable; you may want to copy someone else’s .emacs file and customise it.

## 7.2 Document processing

The typesetting programs T<sub>E</sub>X and L<sup>A</sup>T<sub>E</sub>X are installed on the node. L<sup>A</sup>T<sub>E</sub>X is now the standard for scientific document processing.

Relevant UNIX commands are

Command	Effect
<code>tex file.tex</code>	Run $\text{T}_{\text{E}}\text{X}$ on file.tex
<code>latex file.tex</code>	Run $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ on file.tex
<code>xdvi file.dvi</code>	Display the output (dvi) file on the terminal. X-session required.
<code>dvips file.dvi</code>	Translate the dvi file to PostScript for printing
<code>bibtex file</code>	insert references from the bibtex database into the document

(.tex, .dvi extensions may be omitted from these commands.)

Copies of the  $\text{T}_{\text{E}}\text{X}$ book and  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ book are available, and there are a number of local experts.

The `tth` command can be used to convert  $\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$  to HTML for export on the web — consult the documentation at [http://www.star.bris.ac.uk/local/tth\\_manual.html](http://www.star.bris.ac.uk/local/tth_manual.html).

If you wish to read documents and spreadsheets created using Microsoft products, then the `OpenOffice` suite of programs is available, see the local web pages for setup instructions. Some of the Linux machines have `AbiWord` which can do the same thing. Text can be extracted from a word file using the `catdoc` utility on Linux and Suns. eg

```
cat departmental_minutes.doc | catdoc
```

## 7.3 Pictures and Graphics

There are a number of ways of creating and manipulating images in the Unix/Linux environment. Line drawings can be created with `xfig` or any number of Linux utilities. There is a powerful tool called `gimp` available which can do all sorts of image manipulation. There is a handy tool called `xv` on all systems which can be used for taking a quick look at images.

## 7.4 Miscellaneous

Miscellaneous text utilities include `ispell` (spell checker), `wc` (word/line counter), `grep` (search text file for specified text), ...

## 8 Programming

On Linux, the FORTRAN compiler is `f77`, the C compiler is `gcc` and the C++ compiler is `c++`. There are also a FORTRAN 90 compilers available from Sun and Intel on Linux.

The only shells available are `sh`, `csh` and `tcsh`. `tcsh` is the default. You should never need to use `csh`. Extensive shell-script programming in `tcsh` is generally considered a bad idea. `perl` is available.

Use of the `make` utility is recommended for anything other than the simplest programming tasks.

## 9 Specialist software

The Starlink Software Collection (SSC) contains many programs that may be useful for data reduction. See URL: <http://www.star.rl.ac.uk>. The latest versions of IRAF (optical/X-ray data analysis) and AIPS (radio data analysis) are also installed. IDL is available on all machines, although it is better to run it on faster machines with more memory.

### 9.1 AIPS

See the system manager to be assigned an AIPS number; you will need to do this before you can start AIPS.

The AIPS startup script is `/aips/START.AIPS`. For AIPS use, it is more than usually important to make sure that you are running data reduction processes on the same machine as your data.

## 10 Jobs and processes

Each program currently executing on a given workstation is called a process, and has a numeric process id (pid). You can get a list of the processes you are currently running with `ps -e -u <your user name>` or list all current

processes with `ps -ef`. `top` lists the processes that are currently using the most resources.

Sensible use of the available CPU and memory is important in order to make sure that the system works well for everyone. Before running a program that is likely to consume large amounts of either, use `top` to check what processes are already running and how much free memory the system has. Most workstations have only one processor; if two CPU-intensive tasks are run on the same single processor machine they will both run at half speed.

Sol has two processors but it is a file server for the node, so slowing it down affects all users, not just those logged in to sol. If your CPU-intensive job is non-interactive (i.e. it will run to completion without further input from you) and is likely to last more than an hour or so, consider using `renice` to flag it as having non-interactive priority. Find the pid of your process with `top`, and then `renice 19 <pid>`. This will speed up interactive processes (xterms, shells and so on) while allowing your job to run at full speed when there are no interactions currently happening. Do not run long CPU-intensive jobs on the machines located in offices without checking with their owners that this is OK.

Memory-intensive jobs should be run with even greater care, as the system has been known to crash when very large jobs are running. Check the amount of memory on the machine you intend to use, use the `top` command or refer to the table at the beginning of this document. Consult the interactive user of a machine before starting to use it for memory-intensive work.

It may be worthwhile to run very demanding jobs at a fixed time late at night. The command to do this is `at`.