Linux/IRAF Site Manager's Guide

Doug Tody
Mike Fitzpatrick

IRAF Group
National Optical Astronomy Observatories†
September 1995

ABSTRACT

An IRAF *site manager* is anyone who is responsible for installing and maintaining IRAF at a site. This document describes a variety of site management activities, including configuring the device and environment tables to provide reasonable defaults for the local site, adding interfaces for new devices, configuring and using IRAF networking, the installation and maintenance of layered software products (external packages), and configuring a custom site LOCAL package so that local software may be added to the system. Background information on multiple architecture support, compressed executable options, and the software management tools provided with the system is presented. The procedures for rebooting IRAF and performing a sysgen are described. An introduction to graphics and image display on the PC is provided. The host system resources required to run IRAF are discussed.

September 6, 1995

†Operated by the Association of Universities for Research in Astronomy, Inc. under cooperative agreement with the National Science Foundation.

Contents

1. **Introduction** 1

2. **System Setup** 2
   - 2.1. Installing the System 2
   - 2.2. Configuring the Device and Environment Tables 2
     - 2.2.1. Environment definitions 2
     - 2.2.2. The template LOGIN.CL 3
     - 2.2.3. The TAPECAP file 3
     - 2.2.4. The DEVICES.HLP file 3
     - 2.2.5. The TERMCAP file 4
     - 2.2.6. The GRAPHCAP file 4
     - 2.2.7. Configuring IRAF networking 4
     - 2.2.8. Configuring the IRAF account 6
     - 2.2.9. Configuring user accounts for IRAF 6
   - 2.3. Tuning Considerations 6
     - 2.3.2. Stripping the system to reduce disk usage 6

3. **Software Management** 7
   - 3.1. Multiple architecture support 7
   - 3.2. Shared libraries 9
   - 3.3. Layered software support 9
   - 3.4. Software management tools 9
   - 3.5. Modifying and updating a package 10
   - 3.6. Installing and maintaining layered software 11
   - 3.7. Configuring a custom LOCAL package 12
   - 3.8. Updating the full IRAF system 13
     - 3.8.1. The BOOTSTRAP 13
     - 3.8.2. The SYSGEN 14
     - 3.8.3. Localized software changes 15

4. **Graphics and Image Display under X11** 15
   - 4.1. The XFree86 environment 16
   - 4.2. Vector graphics capabilities 16
   - 4.3. Image Display capabilities 17
   - 4.4. Using the workstation with a remote compute server 17

5. **Interfacing New Graphics Devices** 18
   - 5.1. Graphics terminals 18
   - 5.2. Graphics plotters 18
   - 5.3. Image display devices 19

6. **Host System Requirements** 19
   - 6.1. Memory requirements 19
   - 6.2. Disk requirements 19

Appendix A. The IRAF Directory Structure 20

Linux/IRAF Site Manager's Guide

Doug Tody
Mike Fitzpatrick

IRAF Group
National Optical Astronomy Observatories†
September 1995

1. Introduction

Once the IRAF system has been installed it will run, but there remain many things one might want to do to tailor the system to the local site. Examples of the kinds of customizations one might want to make are the following.

- Edit the default IRAF environment definitions to provide reasonable defaults for your site.
- Make entries in the device descriptor tables for the devices in use at your site.
- Code and install new device interfaces.
- Enable and configure IRAF networking, e.g., to permit remote image display, tape drive, or file access.
- Perform various optimizations, e.g., stripping the system to reduce disk usage.
- Extend the system by installing layered software products.
- Configure a custom LOCAL package so that locally developed software may be installed in the system.

This document provides sufficient background information and instructions to guide the IRAF site manager in performing such customizations. Additional help is available via the *adass.iraf* newsgroups on USENET, by sending mail to `iraf@noao.edu`, or via the IRAF HOTLINE (520 318-8160). Contributions of interfaces developed for new devices, or any other software of general interest, are always welcome.

The IRAF software is organized in a way which attempts to isolate, so far as possible, the files or directories which must be modified to tailor the system for the local site. Most or all changes should affect only files in the local, dev, and hlib (unix/hlib) directories. Layered software products, including locally added software, reside outside of the IRAF core system directory tree and are maintained independently of the core system.

A summary of all modifications made to the IRAF system for a given IRAF release is given in the *Revisions Summary* distributed with the system. Additional information will be found in the system notes files (notes.v29, notes.v210, etc.) in the iraf/local and iraf/doc directories. This is the primary source of technical documentation for each release and should be consulted if questions arise regarding any of the system level features added in a new release of the core system.

---

†Operated by the Association of Universities for Research in Astronomy, Inc. under cooperative agreement with the National Science Foundation.

2. System Setup

2.1. Installing the System

The procedure for installing Linux/IRAF, plus other information specific to the Linux version of IRAF is documented in the *README* file in the Linux/IRAF distribution directory on the iraf.noao.edu ftp server (i.e. ftp://iraf.noao.edu/iraf/v210/LNUX/README). More detailed information on general installation procedures for IRAF on UNIX systems can be found in the *Linux/IRAF Installation Guide*, file linuxiraf.ps.Z in the distribution directory.

In short, an IRAF tape, network, or CD-ROM distribution is obtained and installed according to the instructions. The result is a full IRAF system, including both sources and executable binaries for the architectures to be supported. The system will have been modified to reflect the new IRAF root directory and should run, but will otherwise be a generic IRAF distribution. To get the most out of an IRAF installation it will be necessary to perform some of the additional steps outlined in the remainder of this document.

2.2. Configuring the Device and Environment Tables

Teaching IRAF about the devices, network nodes, external programs, and other special resources available at a site is largely a matter of editing a standard set of device descriptor and environment setup files, all of which are simple text files. The versions of these files provided with the distribution are simply those in use on the NOAO system from which the tapes were made, at the time the tapes were generated. Hence while these files may be useful as examples of properly configured descriptor files, the defaults, and many specific device entries, will in many cases be meaningless for a different site. This is harmless but it may be confusing to the user if, for example, the default printer doesn't exist at your site.

The device and environment files also contain much material which any site will need, however, so care must be taken when editing the files. Important changes may be made to the global portions of these files as part of any IRAF release. To facilitate future updates, it is wise where possible to isolate any local changes or additions so that they may be copied into the new (distributed) version of the file in a future update.

2.2.1. Environment definitions

Since IRAF is a machine and operating system independent, distributed system it has its own environment facility apart from that of the host system. Host system environment variables may be accessed as if they are part of the IRAF environment (which is sometimes useful but which can also be dangerous), but if the same variable is defined in the IRAF environment it is the IRAF variable which will be used. The IRAF environment definitions, as defined at CL startup time, are defined in a number of files in the unix/hlib directory. Chief among these is the **zzsetenv.def** file. Additional user modifiable definitions may be given in the template **login.cl** file (see §2.2.2).

The zzsetenv.def file contains a number of environment definitions. Many of these define IRAF logical directories and should be left alone. Only those definitions in the header area of the file should need to be edited to customize the file for a site. It is the default editor, default device, etc. definitions in this file which are most likely to require modification for a site.

If the name of a default device is modified, the named device must also have an entry in the **termcap** file (terminals and printers) or the **graphcap** file (graphics terminals and image displays) in iraf/dev. There must also be an *editor*`.ed` file in dev for the default editor; *edt*, *emacs*, and *vi* are examples of currently supported editors.

Sample values of those variables most likely to require modification for a site are shown below.

```
set editor      = "vi"
set printer     = "lp"
set stdplot     = "lp"
set stdimage    = "imt512"
```

For example, you may wish to change the default editor to "emacs", the default printer to "lw5", or the default image display to "imt800". Note that the values of terminal and stdgraph, which also appear in the zzsetenv.def file, have little meaning except for debugging processes run standalone, as the values of the environment variables are reset automatically by *stty* at login time. The issues of interfacing new graphics and image display devices are discussed further in §5.

2.2.2. The template LOGIN.CL

The template login.cl file hlib$login.cl, is the file used by *mkiraf* to produce the user login.cl file. The user login.cl file, after having possibly been edited by the user, is read by the CL every time a new CL is started, with the CL processing all environment and task definitions, package loads, etc., in the login file. Hence this file plays an important role in establishing the IRAF environment seen by the user.

Examples of things one might want to change in the template login.cl are the commented out environment definitions, the commented out CL parameter assignments, the foreign task definitions making up the default `user` package, and the list of packages to be loaded at startup time. For example, if there are host tasks or local packages which should be part of the default IRAF operating environment at your site, the template login.cl is the place to make the necessary changes.

2.2.3. The TAPECAP file

Beginning with V2.10 IRAF magtape devices are described by the tapecap file, dev$tapecap. This replaces the "devices" file used in earlier versions of IRAF. The tapecap file describes each local magtape device and controls all i/o to the device, as well as device allocation.

The tapecap file included in the distributed system includes some generic device entries such as "mtdat0" (HP DAT unit 0), "mtexb0" (Exabyte 8mm drive, unit 0), and so on which you may be able to use as-is to access your local magtape devices. Most likely you will want to add some device aliases, and you may need to prepare custom device entries for local devices. There must be an entry in the tapecap file for a magtape device in order to be able to access the device from within IRAF.

Instructions for adding devices to the tapecap file are given in the document *IRAF Version 2.10 Revisions Summary*, in the discussion of the new magtape system.

2.2.4. The DEVICES.HLP file

All physical devices that the user might need to access by name should be documented in the file dev$devices.hlp. Typing

```
cl> help devices
```

or just

```
cl> devices
```

in the CL will format and output the contents of this file. It is the IRAF name of the device, as given in files such as termcap, graphcap, and tapecap, which should appear in this help file.

2.2.5. The TERMCAP file

There must be entries in this file for all local terminal and printer devices you wish to access from IRAF (there is currently no "printcap" file in IRAF). The entry for a printer contains one special device-specific entry, called `DD`. This consists of three fields: the device name, e.g. "node!device", the template for the temporary spoolfile, and the UNIX command to be used to dispose of the file to the printer. On most UNIX systems it is not necessary to make use of the node name and IRAF networking to access a remote device since UNIX *lpr* already provides this capability, however it might still be useful if the desired device does not have a local *lpr* entry for some reason.

If you have a local terminal which has no entry in the IRAF termcap file, you probably already have an entry in the UNIX termcap file. Simply copy it into the IRAF file; both systems use the same termcap database format and terminal device capabilities. However, if the terminal in question is a graphics terminal with a device entry in the graphcap file, you should add a '`:gd`' capability to the termcap entry. If the graphcap entry has a different name from the termcap entry, make it '`:gd=`*gname*'.

2.2.6. The GRAPHCAP file

There must be entries in the graphcap file for all graphics terminals, batch plotters, and image displays accessed by IRAF programs. New graphics terminals will need a new entry. The IRAF file gio$doc/gio.hlp contains documentation describing how to prepare graphcap device entries. A printed copy of this document is available from the iraf/docs directory in the IRAF network archive. However, once IRAF is up you may find it easier to generate your own copy using *help*, as follows:

```
cl> help gio$doc/gio.hlp fi+ | lprint
```

which will print the document on the default IRAF printer device (use the "device=" hidden parameter to specify a different device). Alternatively, to view the file on the terminal,

```
cl> phelp gio$doc/gio.hlp fi+
```

The help pages for the IRAF tasks *showcap* and *stty* should also be reviewed as these utilities are useful for generating new graphcap entries. The i/o logging feature of *stty* is useful for determining exactly what characters your graphcap device entry is generating. The *gdevices* task is useful for printing summary information about the available graphics devices.

Help preparing new graphcap device entries is available if needed. We ask that new graphcap entries be sent back to us so that we may include them in the master graphcap file for all to benefit.

2.2.7. Configuring IRAF networking

The dev directory contains several files (`hosts`, `irafhosts`, and `uhosts`) used by the IRAF network interface. IRAF networking is used to access remote image displays, printers, magtape devices, files, images, etc. via the network. Nodes do not necessarily have to have the same architecture, or even run the same operating system, so long as they can run IRAF.

To enable IRAF networking for a Linux/IRAF system, all that is necessary is to edit the "hosts" file. Make an entry for each logical node, in the format

nodename [*aliases*] ":" *irafks.e-pathname*

following the examples given in the hosts file supplied with the distribution (which is the NOAO/Tucson hosts file). Note that there may be multiple logical entries for a single physical node.

The "uhosts" file is not used by UNIX/IRAF systems hence does not need to be modified (it used by VMS/IRAF). The "irafhosts" file is the template file used to create user .irafhosts files. It does not have to be modified, although you can do so if you wish to change the default

parameter values given in the file.

To enable IRAF networking on a particular IRAF host, the Linux **hostname** must appear as a primary name or alias somewhere in the IRAF hosts table. During process startup, the IRAF VOS looks for the system name for the current host and automatically disables networking if this name is not found. Hence IRAF networking is automatically disabled when the distributed system is first installed - unless you are unlucky enough to have installed the system on a host with the same name as one of the nodes in the NOAO host table. Note that it may be best to simply delete the NOAO host table entries since any duplicate with a local host entry will will cause the IRAF "cd" command to fail and may have other consequences.

Once IRAF networking is configured, the following command may be typed in the CL to verify that all is well:

```
cl> netstatus
```

This will print the host table and state the name of the local host. Read the output carefully to see if any problems are reported.

For IRAF networking to be of any use, it is necessary that IRAF be installed on at least two systems. In that case either system can serve as the server for an IRAF client (IRAF program) running on the other node. It is not necessary to have a separate copy of IRAF on each node, i.e., a single copy of IRAF may be NFS mounted on all nodes (you will need to run the IRAF *install* script on each client node). If it is not possible to install IRAF on a node for some reason (either directly or using NFS) it is possible to manage by installing only enough of IRAF to run the IRAF kernel server. Contact IRAF site support if you need to configure things in this manner.

Linux/IRAF currently supports only TCP/IP based networking. Networking between any heterogeneous collection of systems is possible provided they support TCP/IP based networking (virtually all UNIX-based systems do). The situation with networking between UNIX and VMS systems is more complex. V2.9 and earlier versions of VMS/IRAF support client-side only TCP/IP using the third party Wollongong software. For V2.10 we plan to drop support for the Wollongong software and switch to the more fully-featured Multinet instead (another third party product). Contact the IRAF project for further information on networking between UNIX and VMS systems.

Once IRAF networking is enabled, objects resident on the server node may be accessed from within IRAF merely by specifying the node name in the object name, with a "*node!*" prefix. For example, if *foo* is a network node,

```
cl> page foo!hlib$motd
cl> allocate foo!mta
cl> devstatus foo!mta
```

In a network of "trusted hosts" the network connection will be made automatically, without a password prompt. A password prompt will be generated if the user does not have permission to access the remote node with UNIX commands such as *rsh*. Each user has a .irafhosts file in their UNIX login directory which can be used to exercise more control over how the system connect to remote hosts. See the discussion of IRAF networking in the *IRAF Version 2.10 Revisions Summary*, or in the V2.10 system notes file, for a more in-depth discussion of how IRAF networking works.

To keep track of where files are in a distributed file system, IRAF uses **network pathnames**. A network pathname is a name such as "foo!/tmp3/images/m51.pix", i.e., a host or IRAF filename with the node name prepended. The network pathname allows an IRAF process running on any node to access an object regardless of where it is located on the network.

Inefficiencies can result when image pixel files are stored on disks which are cross-mounted using NFS. The typical problem arises when imdir (the pixel file storage directory) is set to a path such as "/data/iraf/user/", where /data is a NFS mounted directory. Since NFS is transparent to applications like IRAF, IRAF thinks that /data is a local disk and the network

pathame for a pixel file will be something like "foo!/data/iraf" where "foo" is the hostname of the machine on which the file is written. If the image is then accessed from a different network node the image data will be accessed via an IRAF networking connection to node "foo", followed by an NFS connection to the node on which the disk is physically mounted, causing the data to traverse the network twice, slowing access and unnecessarily loading the network.

A simple way to avoid this sort of problem is to include the server name in the imdir, e.g.,

```
cl> set imdir = "server!/data/iraf/user/"
```

This also has the advantage of avoiding NFS for pixel file access - NFS is fine for small files but can load the server excessively when used to access bulk image data.

Alternatively, one can set imdir to a value such as "HDR$pixels/", or disable IRAF networking for disk file access. In both cases NFS will be used for image file access.

2.2.8. Configuring the IRAF account

The IRAF account, i.e., what one gets when one logs into Linux as "iraf", is the account used by the IRAF site manager to work on the IRAF system. Anyone who uses this account is in effect a site manager, since they have permission to modify, delete, or rebuild any part of IRAF. For these and other reasons (e.g., concurrency problems) it is recommended that all routine use of IRAF be performed from other accounts (user accounts).

If the system has been installed according to the instructions given in the installation guide the login directory for the IRAF account will be iraf/local. This directory contains both a `.login` file defining the environment for the IRAF account, and a number of other "dot" files used to setup a sample XFree86 screen of the type that an IRAF user will want, i.e., with the IRAF graphics and image display windows.

Most site managers will probably want to customize these files according to their personal preferences. In doing this please use caution to avoid losing environment definitions, etc., which are essential to the correct operation of IRAF, including IRAF software development.

The default login.cl file supplied in the IRAF login directory uses machine independent pathnames and should work as-is (no need to do a *mkiraf* - in fact *mkiraf* has safeguards against inadvertent use within the IRAF directories and may not work in iraf/local). It may be necessary to edit the .login file to modify the way the environment variable `IRAFARCH` is defined. This variable, required for software development but optional for merely using IRAF, must be set to the name of the desired machine architecture, e.g., ssun for Sun Solaris, linux for PC Linux, etc. If it is set to the name of an architecture for which there are no binaries, e.g., generic, the CL may not run, so be careful. The alias *setarch*, defined in the iraf account .login, is convenient for setting the desired architecture for IRAF execution and software development.

2.2.9. Configuring user accounts for IRAF

User accounts should be loosely modeled after the IRAF account. All that is required for a user to run IRAF is that they run *mkiraf* in their desired IRAF login directory before starting up the CL. Defining `IRAFARCH` in the user environment is not required unless the user will be doing any IRAF based software development (including IMFORT).

2.3. Tuning Considerations

2.3.1. Stripping the system to reduce disk usage

If the system is to be installed on multiple CPUs, or if a production version is to be installed on a workstation, it may be necessary or desirable to strip the system of all non-runtime files to save disk space. Linux/IRAF is unique in that a distribution that is already stripped of sources may optionally be installed, the remainder of this section deals with stripping a system that has the sources installed and how to strip source from external packages.

This equates to deleting all the sources and all the reference manuals and other documentation, excluding the online manual pages. A special utility called *rmfiles* (in the SOFTOOLS package) is provided for this purpose. It is not necessary to run *rmfiles* directly to strip the system. The preferred technique is to use "mkpkg strip" as in the following example (this may be executed from either the host system or from within IRAF).

```
% cd $iraf
% mkpkg strip
```

This will preserve all runtime files, permitting use of the standard system as well as user software development. Note that only the IRAF core system is stripped, i.e., if you want to strip any external layered software products, such as the NOAO package, a *mkpkg strip* must be executed separately for each - *cd* to the root directory of the external package first. A tape backup of a system should always be made before the system is stripped; keep the backup indefinitely as it may be necessary to restore the sources in order to, e.g., install a bug fix or add-on software product.

3. Software Management

3.1. Multiple architecture support

Often the computing facilities at a site consist of a heterogeneous network of workstations and servers. These machines will often have quite different architectures. Considering only a single vendor like Sun, as of 1992 one sees the three major architectures SPARC, Motorola 68020, and Intel 80386, and several minor variations on these architectures, i.e., the floating point options for the Sun-3, namely the Motorola 68881 coprocessor, the Sun floating point accelerator (FPA), and software floating point (Sun is trying to phase some of these out but the need for multiple architecture support is not likely to go away).

In the case of Linux/IRAF the initial distribution contains two architectures: linux and linuz. The latter option are the same linux binaries but compressed using the GNU *gzexe* task. This can result in a significant disk space savings at some cost in runtime performance. Compressed executables are automatically and transparently uncompressed when executed. Compressed executables take about half the disk space of uncompressed executables, but start up slower as the file must be uncompressed before it can be executed. In the case of IRAF the start up penalty is not as great it might otherwise be due to the CL process cache. A disadvantage of compressed executables is that they require *more* memory to execute if multiple copies of the same executable are simultaneously executed, e.g. if there are multiple IRAF users or sessions, or background IRAF jobs, because no sharing of the text segment takes place.

Since IRAF is a large system it is undesirable to have to maintain a separate copy of IRAF for each machine architecture on a network. For this reason IRAF provides support for multiple architectures within a single copy of IRAF. To be accessible by multiple network clients, this central IRAF system will typically be NFS mounted on each client. It should be noted however that it is not always possible to use the multiple architecture support within the core system to maintain a single IRAF source tree for the entire heterogeneous network. The Host System Interface (HSI) for IRAF ports is different for platforms as diverse as Linux and Sun so there should be a separate installation for each system to minimize difficulties (the update schedules usually differ as well so maintaining the same version is also more difficult). In the case of supporting Linux with both the normal and compressed executable options installed, one installation is sufficient since it’s generally only the binaries that will differ (i.e. the OS is still the same). Almost any combination of architectures may be supported by a single copy of an external package.

Multiple architecture support is implemented by separating the IRAF sources and binaries into different directory trees. The sources are architecture independent and hence sharable by machines of any architecture. All of the architecture dependence is concentrated into the

binaries, which are collected together into the so-called BIN directories, one for each architecture. The BIN directory contains all the object files, object libraries, and executables, supporting both IRAF execution and software development for that architecture. A given system can support any number of BIN directories, and therefore any number of architectures.

In IRAF terminology, when we refer to an "architecture" what we really mean is a type of BIN. The correspondence between BINs and hardware architectures is not necessarily one-to-one, i.e., multiple BINs can exist for a single compiler architecture by compiling the system with different compilation flags, as different versions of the software, and so on. Examples of some currently supported software architectures are shown below.

Architecture	*System*	*Description*
generic	any	no binaries
linux	PC	PC platforms running Linux
linuz	PC	PC platforms running Linux - compressed binaries
sx86	PC	PC platforms running Solaris x86
ssun	Sun-4	Sun SPARC under Solaris (RISC) architecture, integral fpu
sparc	Sun-4	Sun SPARC (RISC) architecture, integral fpu
f68881	Sun-3	mc68020, 68881 floating point coprocessor
pg	Sun-4	Sun/IRAF compiled for profiling
ddec	Decstation	DEC Fortran version of DSUX/IRAF
dmip	Decstation	MIPS Risc Fortran version of DSUX/IRAF
rs6000	IBM	IBM RS/6000 running AIX
irix	SGI	SGI IRIX, MIPS cpu
f2c	Macintosh	A/UX, using Fortran-to-C translation and GCC

Most of these correspond to hardware architectures or floating point hardware options. The exceptions are the generic architecture, which is what the distributed system is configured to by default (to avoid having any architecture dependent binary files mingled with the sources), and the "pg" architecture, which is not normally distributed to user sites, but is a good example of a custom software architecture used for software development.

When running IRAF on a system configured for multiple architectures, selection of the BIN (architecture) to be used is controlled by the UNIX environment variable `IRAFARCH`, e.g.,

```
% setenv IRAFARCH linux
```

would cause IRAF to run using the linux architecture, corresponding to the BIN directory bin.linux. Once inside the CL one can check the current architecture by entering one of the following commands (the output in each case is shown as well).

```
cl> show IRAFARCH
linux
```

or

```
cl> show arch
.linux
```

If IRAFARCH is undefined at CL startup time a default architecture will be selected based on the current machine architecture, the available floating point hardware, and the available BINs. The IRAFARCH variable controls not only the architecture of the executables used to run IRAF, but the libraries used to link IRAF programs, when doing software development from within the IRAF or host environment.

Additional information on multiple architecture support is provided in the system notes file for V2.8, file doc$notes.v28.

3.2. Shared libraries

The IRAF system executables (CORE and NOAO executables and HSI executables) are statically linked non-Elf executables hence should be usable on any recent (e.g. Slackware 2.3, Linux 1.2.11 kernel) system. The current distribution does not include an IRAF shareable image, this may appear in a future release.

3.3. Layered software support

An IRAF installation consists of the core IRAF system and any number of external packages, or "layered software products". As the name suggests, layered software products are layered upon the core IRAF system. Layered software requires the facilities of the core system to run, and is portable to any computer which already runs IRAF. Any number of layered products can be installed in IRAF to produce the IRAF system seen by the user at a given site.

The support provided by IRAF for layered software is essentially the same as that provided for maintaining the core IRAF system itself (the core system is a special case of a layered package). Each layered package (usually this refers to a suite of subpackages) is a system in itself, similar in structure to the core IRAF system. Hence, there is a LIB, one or more BINs, a help database, and all the sources and runtime files. A good example of an external package is the NOAO package. Except for the fact that NOAO is rooted in the IRAF directories, NOAO is equivalent to any other layered product, e.g., STSDAS, TABLES, XRAY, CTIO, NLOCAL, ICE, and so on. In general, layered products should be rooted somewhere outside the IRAF directory tree to simplify updates.

3.4. Software management tools

IRAF software management is performed with a standard set of tools, consisting of the tasks in the SOFTOOLS package, plus the host system editors and debuggers. Some of the most important and often used tools for IRAF software development and software maintenance are the following.

`mkhelpdb` — Updates the HELP database of the core IRAF system or an external package. The core system, and each external package, has its own help database. The help database is the machine independent file `helpdb.mip` in the package library (LIB directory). The help database file is generated with *mkhelpdb* by compiling the `root.hd` file in the same directory.

`mkpkg` — The "make-package" utility. Used to make or update package trees. Will update the contents of the current directory tree. When run at the root iraf directory, updates the full IRAF system; when run at the root directory of an external package, updates the external package. Note that updating the core IRAF system does not update any external packages (including NOAO). When updating an external package, the package name must be specified, e.g., "*mkpkg -p noao*".

`rmbin` — Descends a directory tree or trees, finding and optionally listing or deleting all binary files therein. This is used, for example, to strip the binaries from a directory tree to leave only sources, to force *mkpkg* to do a full recompile of a package, or to locate all the binaries files for some reason. IRAF has its own notion of what a binary file is. By default, files with the "known" file extensions (.[aoe], .[xfh] etc.) are classified as binary or text (machine independent) files immediately, while a heuristic involving

examination of the file data is used to classify other files. Alternatively, a list of file extensions to be searched for may optionally be given.

`rtar,wtar` These are the portable IRAF tarfile writer (*wtar*) and reader (*rtar*). About the only reasons to use these with Linux/IRAF are if one wants to move only the machine independent or source files (*wtar*, like *rmbin*, can discriminate between machine generated and machine independent files), or if one is importing files written to a tarfile on a VMS/IRAF system, where the files are blank padded and the trailing blanks need to be stripped with *rtar*.

`xc` The X (SPP) compiler. This is analogous to the UNIX *cc* except that it can compile ".x" or SPP source files, knows how to link with the IRAF system libraries, knows how to read the environment of external packages, and so on.

The SOFTOOLS package contains other tasks of interest, e.g., a program *mktags* for making a tags file for the *vi* editor, a help database examine tool, and other tasks. Further information on these tasks is available in the online help pages.

3.5. Modifying and updating a package

IRAF applications development is most conveniently performed from within the IRAF environment, since testing must be done from within the environment. The usual edit-compile-test development cycle is illustrated below. This takes place within the *package directory* containing all the files specific to a given package.

- Edit one or more source files.
- Use *mkpkg* to compile any modified files, or files which include a modified file, and relink the package executable.
- Test the new executable.

The mkpkg file for a package can be written to do anything, but by convention the following commands are usually provided.

`mkpkg` The *mkpkg* command with no arguments does the default mkpkg operation; for a subpackage this is usually the same as *mkpkg relink* below. For the root mkpkg in a layered package it udpates the entire layered package.

`mkpkg libpkg.a` Updates the package library, compiling any files which have been modified or which reference include files which have been modified. Private package libraries are intentionally given the generic name libpkg.a to symbolize that they are private to the package.

`mkpkg relink` Rebuilds the package executable, i.e., updates the package library and relinks the package executable. By convention, this is the file xx_*pkgname*.e in the package directory, where *pkgname* is the package name.

`mkpkg install` Installs the package executable, i.e., renames the xx_foo.e file to x_foo.e in the global BIN directory for the layered package to which the subpackage *foo* belongs.

`mkpkg update` Does everything, i.e., a *relink* followed by an *install*.

If one wishes to test the new program before installing it one should do a *relink* (i.e., merely type "mkpkg" since that defaults to relink), then run the host system debugger on the resultant executable. The process is debugged standalone, running the task by giving its name to the standalone process interpreter. The CL task *dparam* is useful for dumping a task's parameters to a text file to avoid having to answer parameter queries during process execution. The LOGIPC debugging facility introduced in V2.10 is also useful for debugging subprocesses. If the new program is to be tested under the CL before installation, a *task* statement can be interactively typed into the CL to cause the CL to run the "xx_" version of the package executable, rather than old installed version.

When updating a package other than in the core IRAF system, the `-p` flag, or the equivalent `PKGENV` environment variable, must be used to indicate the system or layered product being updated. For example, "mkpkg -p noao update" would be used to update one of the subpackages of the NOAO layered package. If the package being updated references any libraries or include files in *other* layered packages, those packages must be indicated with a "-p pkgname" flag as well, to cause the external package to be searched.

The CL process cache can complicate debugging and testing if one forgets that it is there. When a task is run under the CL, the executing process remains idle in the CL process cache following task termination. If a new executable is installed while the old one is still in the process cache, the CL will automatically run the new executable (the CL checks the modify date on the executable file every time a task is run). If however an executable is currently running, either in the process cache or because some other user is using the program, it may not be possible to set debugger breakpoints.

A full description of these techniques is beyond the scope of this manual, but one need not be an expert at IRAF software development techniques to perform simple updates. Most simple revisions, e.g., bug fixes or updates, can be made by merely editing or replacing the affected files and typing

```
cl> mkpkg
```

or

```
cl> mkpkg update
```

to update the package.

3.6. Installing and maintaining layered software

The procedures for installing layered software products are similar to those used to install the core IRAF system, or update a package. Layered software may be distributed in source only form, or with binaries; it may be configured for a single architecture, or may be preconfigured to support multiple architectures. The exact procedures to be followed to install a layered product will in general be product dependent, and should be documented in the installation guide for the product.

In brief, the procedure to be followed should resemble the following:

- Create the root directory for the new software, somewhere outside the IRAF directories.
- Restore the files to disk from a tape or network archive distribution file.
- Edit the core system file hlib$extern.pkg to "install" the new package in IRAF. This file is the sole link between the IRAF core system and the external package.

- Configure the package BIN directory or directories, either by restoring the BIN to disk from an archive file, or by recompiling and relinking the package with *mkpkg*.

As always, there are some little things to watch out for. When using *mkpkg* on a layered product, you must give the name of the system being operated upon, e.g.,

```
cl> mkpkg -p foo update
```

where *foo* is the system or package name, e.g., "noao", "local", etc. The `-p` flag can be omitted by defining `PKGENV` in your UNIX environment, but this only works for updates to a single package.

An external system of packages may be configured for multiple architecture support by repeating what was done for the core system. One sets up several BIN directories, one for each architecture, named `bin.`*arch*, where *arch* is "linux", "ssun", etc. These directories, or symbolic links to the actual directories, go into the root directory of the external system. A symbolic link `bin` pointing to an empty directory bin.generic, and the directory itself, are added to the system's root directory. The system is then stripped of its binaries with *rmbin*, if it is not already a source only system. Examine the file zzsetenv.def in the layered package LIB directory to verify that the definition for the system BIN (which may be called anything) includes the string "(arch)", e.g.,

```
set noaobin = "noao$bin(arch)/"
```

The binaries for each architecture may then be generated by configuring the system for the desired architecture and running *mkpkg* to update the binaries, for example,

```
cl> cd foo
cl> mkpkg linux
cl> mkpkg -p foo update >& spool &
```

where *foo* is the name of the system being updated. If any questions arise, examination of a working example of a system configured for multiple architecture support (e.g., the NOAO packages) may reveal the answers.

Once installed and configured, a layered product may be deinstalled merely by archiving the package directory tree, deleting the files, and commenting out the affected lines of hlib$extern.pkg. With the BINs already configured reinstallation is a simple matter of restoring the files to disk and editing the extern.pkg file.

3.7. Configuring a custom LOCAL package

Anyone who uses IRAF enough will eventually want to add their own software to the system, by copying and modifying the distributed versions of programs, by obtaining and installing isolated programs written elsewhere, or by writing new programs of their own. A single user can do this by developing software for their own personal use, defining the necessary *task* statements etc. to run the software in their personal login.cl or loginuser.cl file. To go one step further and install the new software in IRAF so that it can be used by everyone at a site, one must configure a custom local package.

The procedures for configuring and maintaining a custom LOCAL package are similar to those outlined in §3.5 for installing and maintaining layered software, since a custom LOCAL will in fact be a layered software product, possibly even something one might want to export to another site (although custom LOCALs may contain non-portable or site specific software).

To make a custom local you make a copy of the "template local" package (iraf$local) somewhere outside the IRAF directory tree, change the name to whatever you wish to call the new layered package, and install it as outlined in §3.5. The purpose of the template local is to provide the framework necessary for a external package; a couple of simple tasks are provided in the template local to serve as examples. Once you have configured a local copy of the template local and gotten it to compile and link, it should be a simple matter to add new tasks to

the existing framework.

3.8. Updating the full IRAF system

This section will describe how to recompile or relink IRAF. Before we get into this however, it should be emphasized that *most users will never need to recompile or relink IRAF*. In fact, this is not something that one should attempt lightly - don't do it unless you have some special circumstance which requires a custom build of the system (such as a port). Even then you might want to set up a second copy of IRAF to be used for the experiment, keeping the production system around as the standard system. If you change the system it is a good idea to make sure that you can undo the change.

While the procedure for building IRAF is straightforward, it is easy to make a mistake and without considerable knowledge of IRAF it may be difficult to recover from such a a mistake (for example, running out of disk space during a build, or an architecture mismatch resulting in a corrupted library). More seriously, the software - Linux, the GCC and F2C compilers, the local system configuration, and IRAF - is changing constantly. A build of IRAF brings all these things together at one time, and every build needs to be independently and carefully tested. An OS upgrade or a new version of the GCC/F2C compilers may not yet be supported by the version of IRAF you have locally. Any problems with the Linux configuration can cause a build to fail, or introduce bugs.

The precompiled binaries we ship with IRAF have been carefully prepared and tested, usually over a period of months. They are the same as are used at NOAO and at most IRAF sites, so even if there are bugs they will likely have already been seen elsewhere and a workaround determined. If the bugs are new then since we have the exact same IRAF system we are more likely to be able to reproduce and fix the bug. Often the bug is not in the IRAF software at all but in the host system or IRAF configuration. As soon as an executable is rebuilt (even something as simple as a relink) you have new, untested, software.

3.8.1. The BOOTSTRAP

To fully build IRAF from the sources is a three step process. First the system is "bootstrapped", which builds the host system interface (HSI) executables. A "sysgen" of the core system is then performed; this compiles all the system libraries and builds the core system applications. Finally, the bootstrap is then repeated, to make use of some of the functions from the IRAF libraries compiled in step two.

To bootstrap Linux/IRAF, login as IRAF and enter the commands shown below. This takes a while and generates a lot of output, so the output should be spooled in a file.

For a PC Linux (using uncompressed binaries) system:

```
% cd $iraf
% mkpkg linux
% cd $iraf/unix
% reboot >& spool &
```

There are two types of bootstrap, the initial bootstrap starting from a source only system, called the NOVOS bootstrap, and the final or VOS bootstrap, performed once the IRAF system libraries `libsys.a` and `libvops.a` exist. The bootstrap script *reboot* will automatically determine whether or not the VOS libraries are available and will perform a NOVOS bootstrap if the libraries cannot be found. It is important to restore the linux architecture before attempting a bootstrap, as otherwise a NOVOS bootstrap will be performed.

3.8.2. The SYSGEN

By sysgen we refer to an update of the core IRAF system - all of the files comprising the runtime system, excluding the HSI which is generated by the bootstrap. On a source only system, the sysgen will fully recompile the core system, build all libraries and applications, and link and install the executables. On an already built system, the sysgen scans the full IRAF directory tree to see if anything is out of date, recompiles any files that need it, then relinks and installs new executables.

To do a full sysgen of IRAF one merely runs *mkpkg* at the IRAF root. If the system is configured for multiple architecture support one must repeat the sysgen for each architecture. Each sysgen builds or updates a single BIN directory. Since a full sysgen takes a long time and generates a lot of output which later has to be reviewed, it is best to run the job in batch mode with the output redirected. For example, to update the linux binaries:

```
% cd $iraf
% mkpkg linux
% mkpkg >& spool &
```

To watch what is going on after this command has been submitted and while it is running, try

```
% tail -f spool
```

Sysgens are restartable, so if the sysgen aborts for any reason, simply fix the problem and start it up again. Modules that have already been compiled should not need to be recompiled. How long the sysgen takes depends upon how much work it has to do. The worst case is if the system and applications libraries have to be fully recompiled. If the system libraries already exist they will merely be updated. Once the system libraries are up to date the sysgen will relink the core system executables.

A full sysgen generates a lot of output, too much to be safely reviewed for errors by simply paging the spool file. Enter the following command to review the output (this assumes that the output has been saved in a file named "spool").

```
% mkpkg summary
```

It is normal for a number of compiler messages warning about assigning character data to an integer variable to appear in the spooled output if the full system has been compiled. There should be no serious error messages if a supported and tested system is being recompiled.

The above procedure only updates the core IRAF system. To update a layered product one must repeat the sysgen process for the layered system. For example, to update the linux binaries for the NOAO package:

```
% cd $iraf/noao
% mkpkg linux
% mkpkg -p noao >& spool &
```

This must be repeated for each supported architecture. Layered systems are independent of one another and hence must be updated separately.

To force a full recompile of the core system or a layered package, one can use *rmbin* to delete the objects, libraries, etc. scattered throughout the system, or do a "mkpkg generic" and then delete the `OBJS.arc.Z` file in the BIN one wishes to regenerate (the latter approach is probably safest).

A full IRAF core system sysgen currently takes anywhere from 3 to 30 hours or more, depending upon the system (e.g. from 30 hours on a VAX 11/750, to 3 hours on a big modern server). On most systems a full sysgen is a good job to run overnight.

3.8.3. Localized software changes

The bootstrap and the sysgen are unusual in that they update the entire HSI, core IRAF system, or layered package. Many software changes are more localized. If only a few files are changed a sysgen will pick up the changes and update whatever needs to be updated, but for localized changes a sysgen really does more than it needs to (if the changes are scattered all over the system an incremental sysgen-relink will still be best).

To make a localized change to a core system VOS library and update the linked applications to reflect the change all one really needs to do is change the desired source files, run *mkpkg* in the library source directory to compile the modules and update the affected libraries, and then relink the system binaries.

For example, assume we want to make a minor change to some files in the VOS interface IMIO, compiling for the linux architecture. We could do this as follows (this assumes that one is logged in as IRAF and that the usual IRAF environment is defined).

```
% whoami
iraf
% cd $iraf
% mkpkg linux
% cd imio
    (edit the files)
% mkpkg                          # update IMIO libraries (libex)
%
% cd $iraf/bin.linux             # save copy of old libex.a
% cp libex.a libex.a.V210
%
% mkpkg relink                   # relink and install new binaries
```

Changing applications is even easier. Ensure that the system architecture is set correctly (i.e. "mkpkg linux" at the iraf or layered package root), edit the affected files in the package source directory, and type "mkpkg -p <pkgname> update" in the root directory of the package being edited. This will compile any modified files, and link and install a new executable. You can do this from within the CL and immediately run the revised program.

We should emphasize again that, although we document the procedures for making changes to the software here, to avoid introducing bugs we do not recommend changing any of the IRAF software except in unusual (or at least carefully controlled) circumstances. To make custom changes to an application, it is best to make a local copy of the full package somewhere outside the standard IRAF system. If changes are made to the IRAF system software it is best to set up an entire new copy of IRAF on a machine separate from the normal production installation, so that one can experiment at will without affecting the standard system. An alternative which does not require duplicating the full system is to use the `IRAFULIB` environment variable. This can be used to safely experiment with custom changes to the IRAF system software outside the main system; IRAFULIB lets you define a private directory to be searched for IRAF global include files, libraries, executables, etc., allowing you to have your own private versions of any of these. See the system notes files for further information on how to use IRAFULIB.

4. Graphics and Image Display under X11

Linux/IRAF may be used from any ordinary video or graphics terminal, or from a workstation run either as a terminal or under any X Window System variant such as XFree86 for PCs. To keep things simple we will limit our discussion here to XFree86 using the *fvwm* window manager, but running under *mwm*, *twm*, *olwm* or any other X11 desktop is similar. The X11 support package, which includes the *xgterm* and *ximtool* programs for graphics and imaging, is system independent and is distributed separately from IRAF. IRAF can also be used with other graphics and image display servers, e.g. *xterm* and *saoimage*. The *x11iraf* utilities are available from the IRAF network archives or by contacting IRAF site support.

XFree86 provides standard window tools such as *xterm* and *rxvt* which can be used to run IRAF. *rxvt* is ok as a terminal but cannot do graphics. *xterm* can be used for IRAF graphics but doesn't support status line interaction very well, and lacks color graphics, a full-screen cursor, and error recovery on cursor reads. Nonetheless it has been used successfully for IRAF graphics on many platforms for years.

Most people will prefer to use *xgterm* and *ximtool* (or a similar display tool such as *saoimage*) for IRAF graphics and imaging. *xgterm* is based on *xterm*, providing an equivalent vt100 (text window) capability but a much enhanced graphics capability. *ximtool* provides a general interactive image display capbility, including support for multiple image frame buffers and frame blinking, independent zoom, pan, and color enhancement for each frame, and many other features. Both programs are implemented at the host level as general purpose window system tools, and are useful independently of IRAF. Detailed documentation on the basic operation and use of these programs is still in preparation (as of September 1995) but introductory documentation is available from the ftp archive in the same directory as the programs themselves. Our concern in this document is with the use of these programs with IRAF.

4.1. The XFree86 environment

The graphics and image display tools provided with IRAF operate within the XFree86 windowing environment much like the standard tools provided with XFree86. To help illustrate the use of these tools, Linux/IRAF is distributed with a sample XFree86 environment already configured for the IRAF account. This consists of the following files in the IRAF account login directory, iraf$local.

`.Xdefaults`	Sets up the defaults for how the window system looks, e.g., defines the X resources controlling window colors, fonts, etc.
`.fvwmrc`	An example of a simple custom rootmenu for the *fvwm* window manager, including entries for xgterm and ximtool. Other window managers will rely on a different configuration file.
`.xinitrc`	Executed at window system start up time to create all the windows.

No one screen layout will suit all users or all applications. Everyone will wish to customize the workstation screen to suit their preferences and the type of work they are doing. However, the configuration provided works and should be useful as an example of how to make things function correctly.

4.2. Vector graphics capabilities

The standard graphics terminal emulator for Linux/IRAF under XFree86 is *xgterm*, which emulates a conventional dual plane text/graphics terminal. This software terminal is driven via an ASCII datastream like a conventional hard terminal (except that the effective baud rate is much higher). The text window behaves like the PC console and the graphics window behaves like a Tektronix 4012, plus some IRAF oriented extensions. Since xgterm emulates standard text and graphics devices non-IRAF programs can easily be run as well as IRAF programs.

Configuring IRAF to use xgterm is very simple. The following command does the job. This is normally executed by the login.cl or loginuser.cl file at login time.

```
cl> stty xgterm
```

Further information on xgterm may be found in the *xgterm.info* file in the IRAF network archive with the xgterm binaries or by contacting site support.

4.3. Image Display capabilities

Image display for IRAF running in the XFree86 environment is provided by *ximtool* or a comparable IRAF-compatible display server. The current *ximtool* program provides a basic display capability, including programmed access from the IRAF environment to load images, interactive windowing of the display, pseudocolor, an interactive image cursor readback capability, zoom and pan, a variety of frame buffer sizes, independent frame buffer and display window sizing, up to four frames, each with its own state, and programmable frame blink. *ximtool* runs as a display server, meaning that it sits idle most of the time, waiting for some client, e.g., IRAF, to send it an image to be displayed via some form of interprocess communication.

To use ximtool from within IRAF one must define the logical device and enable image cursor input. For example,

```
cl> reset stdimage = imt512
```

would configure IRAF and ximtool for use with a 512 pixel square frame buffer (image display image memory). A variety of frame buffer sizes are predefined; see the `imtoolrc` file (normally in /usr/local/lib) for a complete list of possible configurations or use the IRAF *gdevices* command.

The image cursor is enabled by

```
cl> reset stdimcur = stdimage
```

This is the default for Linux/IRAF. Setting `stdimcur` to "text" disables the image cursor, allowing cursor values to be typed in interactively in the terminal window. This is useful, for example, when running image oriented programs from a simple terminal.

The standard IRAF interface to the display server is the *display* program in the TV package. Automatic determination of the optimum intensity mapping to the 200 ximtool greylevels is provided. Entire frames can be displayed, or one can write to subregions of the display. Other programs useful with the image display include *imexamine*, used to interactively examine images under image cursor control, *imedit*, used to edit images using the display, and *tvmark*, used to write color graphics into a display frame.

The display server has the capability of displaying the cursor (mouse) position and pixel value in image pixel units as the mouse is moved about in the window. In addition, text file cursor lists can be generated and displayed, or the image cursor can be read interactively from within IRAF. The image cursor may be called up at any time by typing

```
cl> =imcur
```

into the CL. Applications programs which read the interactive image cursor will do this automatically during program execution.

4.4. Using the workstation with a remote compute server

A common mode of operation with a workstation is to run Linux/IRAF under XFree86 directly on the PC which runs IRAF, accessing files either on a local disk, or on a remote disk via a network interface (NFS, IRAFKS, etc.). It is also possible, however, to run XFree86 with xgterm and ximtool on the workstation, but run IRAF on a remote node, e.g., some powerful compute server such as a large Sun server, a large VAX, or a vector minisupercomputer or supercomputer, possibly quite some distance away. This is done by logging onto the workstation, starting up X and a *xgterm* window, logging onto the remote machine with *rlogin*, *telnet*, or whatever, and starting up IRAF on the remote node.

After IRAF comes up one need only type

```
cl> stty xgterm
cl> reset node = hostname!
```

to tell the remote IRAF that it is talking to a xgterm window and that the image display is on the network node *hostname*. The trailing exclamation point is required in V2.10.4 and later

versions of IRAF to avoid interpretation of general environment variables as network logical node names.

In this mode one is effectively using the PC as a sort of super terminal with powerful graphics and image display capabilities. One gets the best of both worlds, i.e., a state of the art user interface, and the compute power of a large machine. It matters little what operating system is used on the remote machine, so long as it also runs IRAF. Except for the details of the login sequence, operation is completely transparent; xgterm does not care whether the process it is talking to is on a local or remote node. Performance, e.g,. for image loads, is often *better* than when everything is run directly on the local node, due to the more powerful server unless connected by a slow modem line or network connection.

5. Interfacing New Graphics Devices

There are three types of graphics devices that concern us here. These are the graphics terminals, graphics plotters, and image displays.

5.1. Graphics terminals

The IRAF system as distributed is capable of talking to just about any conventional graphics terminal or terminal emulator, using the *stdgraph* graphics kernel supplied with the system. All one need do to interface to a new graphics terminal is add new graphcap and termcap entries for the device. This can take anywhere from a few hours to a few days, depending on one's level of expertise, and the characteristics of the device. Be sure to check the contents of the dev$graphcap file to see if the terminal is already supported, before trying to write a new entry. Useful documentation for writing graphcap entries is the GIO reference manual and the HELP pages for the *showcap* and *stty* tasks (see §2.2.6). Assistance with interfacing new graphics terminals is available via the IRAF Hotline.

5.2. Graphics plotters

The current IRAF system comes with several graphics kernels used to drive graphics plotters. The standard plotter interface the SGI graphics kernel, which is interfaced as the tasks *sgikern* and *stdplot* in the PLOT package. Further information on the SGI plotter interface is given in the paper *The IRAF Simple Graphics Interface*, a copy of which is included with the IRAF installation kit.

SGI device interfaces for most plotter devices already exist, and adding support for new devices is straightforward. Sources for the SGI device translators supplied with the distributed system are maintained in the directory iraf/unix/gdev/sgidev. NOAO serves as a clearinghouse for new SGI plotter device interfaces; contact us if you do not find support for a local plotter device in the distributed system, and if you plan to implement a new device interface let us know so that we may help other sites with the same device.

The older NCAR kernel is used to generate NCAR metacode and can be interfaced to an NCAR metacode translator at the host system level to get plots on devices supported by host-level NCAR metacode translators. The host level NCAR metacode translators are not included in the standard IRAF distribution, but public domain versions of the NCAR implementation for UNIX systems are widely available. A site which already has the NCAR software may wish to go this route, but the SGI interface will provide a more efficient and simpler solution in most cases.

The remaining possibility with the current system is the *calcomp* kernel. Many sites will have a Calcomp or Versaplot library (or Calcomp compatible library) already available locally. To make use of such a library to get plotter output on any devices supported by the interface, one may copy the library to the hlib directory and relink the Calcomp graphics kernel.

A graphcap entry for each new device will also be required. Information on preparing graphcap entries for graphics devices is given in the GIO design document, and many actual

working examples will be found in the graphcap file. The best approach is usually to copy one of these and modify it.

5.3. Image display devices

The standard image display facility for a PC running the XFree86 window system is the *ximtool* or *saoimage* display server. XImtool is available from the iraf.noao.edu (140.252.1.1) ftp archives, SAOimage was developed for IRAF by SAO; distribution kits are available from the IRAF network archive.

Some interfaces for hardware image display devices are also available, although a general display interface is not yet included in the system. Only the IIS model 70 and 75 are currently supported by NOAO. Interfaces for other devices are possible using the current datastream interface, which is based on the IIS model 70 datastream protocol with extensions for passing the WCS, image cursor readback, etc. (see the ZFIOGD driver in unix/gdev). This is how all the current displays, e.g., ximtool and saoimage, and the IIS devices, are interfaced, and there is no reason why other devices could not be interfaced to IRAF via the same interface. Eventually this prototype interface will be obsoleted and replaced by a more general interface.

6. Host System Requirements

IRAF is currently supported on all PC models capable of running Linux. A minimal system for running Linux/IRAF will be a personal machine with at least 500 Mb of disk storage and 8-16Mb of memory. For scientific use, a color screen (16 inch is fine) capable of at least 1024x768 resolution and several hundred Mb of free disk is desirable. A CD-ROM is desirable but not required. A DAT or Exabyte drive on a SCSI adapter may be required for bulk data interchange and system backups. This is a minimal system; PC hardware is so cost effective today that a system with 16-32 Mb of RAM and several Gb of disk is recommended. For optimimum performance a fast Pentium processor, a 256 Kb or larger cache and a fast local bus such as PCI is recommended.

6.1. Memory requirements

The windowing systems used in these workstations tend to be very memory intensive; trying to do image processing only makes matters worse. Even though XFree86 and FVWM try to be very efficient with memory, the typical screen with ten or so windows uses a lot of memory. Interactive performance will suffer greatly if the system pages a lot. No system should be configured with less than 8 Mb of main memory; 20 Mb or more is recommended if you plan to do a lot of image processing but 16 Mb is about the practical minimum for most uses. Get at least 32 Mb if you can afford it and you plan to do a lot of data processing with the machine. These comments refer to single user systems, servers could require even more (we have little experience with the use of Linux for server applications however).

6.2. Disk requirements

The amount of disk required by a user depends greatly on the application, so it is hard to recommend a minimum disk size. For a system with access to a central server, 200-300 Mb of local SCSI or IDE disk is fine. For a standalone system with no access to large server, 500-600 Mb is about the minimum provided you install a stripped IRAF and Linux system. If large data sets will be processed several Gb of disk may be required. IDE is fine for small disks but SCSI is recommended (in some cases required) for large disks. Modern PCs can have very respectable data transfer rates and they can be very effective for serious image processing applications that required workstations or servers not long ago.

Appendix A. The IRAF Directory Structure

A graph of the current full Linux/IRAF directory structure is given at the end of this document. The main branches of the tree are summarized below. Beneath the directories shown are some 300 subdirectories, the largest directory trees being `sys`, `pkg`, and `noao`. The entire contents of all directories other than `unix`, `local`, and `dev` are fully portable, and are identical in all installations of IRAF sharing the same version number.

bin	- the IRAF BIN directories
dev	- device tables (termcap, graphcap, etc.)
doc	- assorted IRAF manuals
lib	- the system library; global files
local	- iraf login directory; locally added software
math	- sources for the mathematical libraries
noao	- packages for NOAO data reduction
pkg	- the IRAF applications packages
sys	- the virtual operating system (VOS)
unix	- the UNIX host system interface (HSI = kernel + bootstrap utilities)

The contents of the `unix` directory (host system interface) are as follows:

as	- assembler sources
bin	- the HSI BIN directories
boot	- bootstrap utilities (mkpkg, rtar, wtar, etc.)
gdev	- graphics device interfaces (SGI device translators)
hlib	- host dependent library; global files
os	- OS interface routines (UNIX/IRAF kernel)
reboot	- executable script run to reboot the HSI
shlib	- shared library facility sources
sun	- gterm and imtool sources

If you will be working with the system much at the system level, it will be well worthwhile to spend some time exploring these directories and gaining familiarity with the system.