The Linux 3Dfx HOWTO

Bernd Kreimeier (bk@gamers.org) v1.03, 12 July 1997

This document describes 3Dfx graphics accelerator chip support for Linux. It lists the supported hardware, describes how to configure the drivers, and answers frequently asked questions. The intent is to bring new users up to speed more quickly and reduce the amount of traffic in the Usenet news groups and mailing lists.

Contents

1 Introduction **4**

1.1 Acknowledgments . . . 4

1.2 Revision History . . . 4

1.3 New versions of this document . . . 4

1.4 Feedback . . . 5

1.5 Distribution Policy . . . 5

2 Graphics Accelerator Technology **5**

2.1 Basics . . . 5

2.2 Hardware configurations (add-on) . . . 6

2.3 Performance limitations . . . 6

2.3.1 Fill bound . . . 6

2.3.2 Missing refresh . . . 6

2.3.3 Primitive bound . . . 7

2.4 Hardware accelerated features . . . 7

2.5 A bit of Voodoo Graphics (tm) architecture . . . 7

3 Installation **8**

3.1 Installing the board . . . 8

3.1.1 Troubleshooting the hardware installation . . . 8

3.1.2 Configuring the kernel . . . 9

3.1.3 Configuring devices . . . 9

3.2 Setting up the Displays . . . 9

3.2.1 Single screen display solution . . . 9

3.2.2 Dual screen display solution . . . 10

3.3 Installing the Glide distribution . . . 10
3.3.1 Using the detect program . . . 10
3.3.2 Using the test programs . . . 11

4 Answers To Frequently Asked Questions 11

5 FAQ: Requirements? 12
5.1 What are the system requirements? . . . 12
5.2 Does it work with Linux-Alpha? . . . 12
5.3 Which chipsets are supported? . . . 12
5.4 Which boards are supported? . . . 12
5.5 Is the Hercules Stingray 128/3D supported? . . . 14
5.6 Is the Intergraph Intense Rush supported? . . . 14

6 FAQ: Voodoo Graphics (tm)? 3Dfx? 14
6.1 Who is 3Dfx? . . . 14
6.2 What is the Voodoo Graphics (tm)? . . . 14
6.3 Where could I get additional info on Voodoo Graphics (tm)? . . . 15

7 FAQ: Glide? TexUS? 15
7.1 What is Glide anyway? . . . 15
7.2 What is TexUS? . . . 15
7.3 Is Glide freeware? . . . 15
7.4 Is the Glide source available? . . . 16
7.5 Is Linux Glide supported? . . . 16
7.6 Where could I post Glide questions? . . . 16
7.7 Where to send bug reports? . . . 16
7.8 Who is maintaining it? . . . 17
7.9 How can I contribute to Linux Glide? . . . 17
7.10 Do I have to use Glide? . . . 17
7.11 Should I program using the Glide API? . . . 17
7.12 What is the current version? . . . 17
7.13 Is Linux Glide identical to DOS/Windows Glide? . . . 18
7.14 Where to I get information on Glide? . . . 18

7.15 Where to get some Glide demos? 18

8 FAQ: Glide and SVGA? 19

9 FAQ: Glide and XFree86? 19

9.1 Does it run with XFree86? 19

9.2 Does it only run full screen? 19

9.3 What about GLX for XFree86? 19

10 FAQ: Glide versus OpenGL/Mesa? 20

10.1 Is Glide OpenGL? 20

10.2 Does Mesa work with 3Dfx? 20

10.3 Where to get additional information on OpenGL? 20

10.4 Where to get info on Mesa? 20

10.5 Where to get information on Mesa Voodoo? 20

10.6 Is there a commercial OpenGL for Linux and 3Dfx? 20

10.7 How about GLUT? 21

11 FAQ: But Quake? 21

11.1 What about the Quake GL? 21

12 FAQ: Troubleshooting? 21

12.1 Has this hardware been tested? 21

12.2 Failed to change I/O privilege? 21

12.3 Does it work without root privilege? 22

12.4 Displayed images looks awful (single screen)? 22

12.5 The last frame is still there (single or dual screen)? 22

12.6 Powersave kicks in (dual screen)? 22

12.7 My machine seem to lock (X11, single screen)? 22

12.8 My machine locks (single or dual screen)? 23

12.9 My machine locks (used with S3 VGA board)? 23

12.10No address conflict, but locks anyway? 23

12.11Compile/link error: grSstWinOpen()? 23

12.12Compile/link error: grSstOpen()? 23

12.13Cannot open shared object file? 23
12.14Mesa compilation problems 24
12.15Mesa runs, but does not access the board? 24

1 Introduction

This is the Linux 3Dfx HOWTO document. It is intended as a quick reference covering everything you need to know to install and configure 3Dfx support under Linux. Frequently asked questions regarding the 3Dfx support specifically, as well as questions about 3D graphics under Linux in general are answered, and references are given to some other sources of information on a variety of topics related to computer generated, hardware accelerated 3D graphics.

This information is only valid for Linux on the Intel platform. Some information may be applicable to other processor architectures, but I have no first hand experience or information. It is only applicable to boards based on 3Dfx technology, any other graphics accelerator hardware is beyond the scope of this document.

1.1 Acknowledgments

Much of this information found in this document has been provided by the people involved in the Linux Glide port and the beta testing process. Daryll Strauss did the port, Paul J. Metzger modified the Mesa Voodoo driver (written by David Bucciarelli) for Linux, Brian Paul integrated it with his famous Mesa library. With respect to Voodoo Graphics (tm) accelerated Mesa, additional thanks has to go to Henri Fousse and Charlie Wallace. The folks at 3Dfx, notably Gary Sanders and Gary McTaggart, provided valuable input, as did Ross Q. Smith of Quantum3D.

Thanks to the SGML-Tools package (formerly known as Linuxdoc-SGML), this HOWTO is available in several formats, all generated from a common source file. For information on SGML-Tools see its homepage at web.inter.NL.net/users/C.deGroot/sgmltools/.

3Dfx, the 3Dfx Interactive logo, Voodoo Graphics (tm), and Voodoo Rush (tm) are registered trademarks of 3Dfx Interactive, Inc. Glide, TexUS, Pixelfx and Texelfx are trademarks of 3Dfx Interactive, Inc. OpenGL is a registered trademark of Silicon Graphics. Obsidian is a trademark of Quantum3D. Other product names are trademarks of the respective holders, and are hereby considered properly acknowledged.

1.2 Revision History

Version 1.03

First version for public release.

1.3 New versions of this document

You will find the most recent version of this document at www.gamers.org/dEngine/xf3D/.

New versions of this document will be periodically posted to the comp.os.linux.answers newsgroup. They will also be uploaded to various anonymous ftp sites that archive such information including ftp://sunsite.unc.edu/pub/Linux/docs/HOWTO/.

Hypertext versions of this and other Linux HOWTOs are available on many World-Wide-Web sites, including sunsite.unc.edu/mdw/mdw.html. Most Linux CD-ROM distributions include the HOWTOs, often under the `/usr/doc/`directory, and you can also buy printed copies from several vendors.

If you make a translation of this document into another language, let me know and I'll include a reference to it here.

1.4 Feedback

I rely on you, the reader, to make this HOWTO useful. If you have any suggestions, corrections, or comments, please send them to me (bk@gamers.org), and I will try to incorporate them in the next revision. Please add `HOWTO 3Dfx` to the Subject-line of the mail, so procmail will dump it in the appropriate folder.

Before sending bug reports or questions, *please read all of the information in this HOWTO*, and *send detailed information about the problem.*

If you publish this document on a CD-ROM or in hardcopy form, a complimentary copy would be appreciated. Mail me for my postal address. Also consider making a donation to the Linux Documentation Project to help support free documentation for Linux. Contact the Linux HOWTO coordinator, Greg Hankins (gregh@sunsite.unc.edu), for more information.

1.5 Distribution Policy

Copyright (C) 1997 Bernd Kreimeier.

This HOWTO is free documentation; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This document is distributed in the hope that it will be useful, but **without any warranty**; without even the implied warranty of **merchantability** or **fitness for a particular purpose**. See the GNU General Public License for more details.

You can obtain a copy of the GNU General Public License by writing to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

2 Graphics Accelerator Technology

2.1 Basics

This section gives a *very* cursory overview of computer graphics accelerator technology, in order to help you understand the concepts used later in the document. You should consult e.g. a book on OpenGL in

order to learn more.

Basically, 3D computer graphics often requires a lot of calculations for each single pixel on the screen. This is especially true if the application has to render a polygon world for many frames of an interactive animation. Even with low resolutions like 320x200, this consumes more processing power than even the most powerful PC could deliver.

To overcome that bottleneck, several companies have designed, manufactured and sold processors dedicated to operations needed for 3D computer graphics. So far, virtually none of the boards manufactured so far offered any Linux support. Fortunately, the manufacturer of the Voodoo Graphics (tm) and Voodoo Rush (tm) chipsets, 3Dfx, decided to support use of Voodoo Graphics (tm) based boards with Linux. The purpose of this document is to describe the support currently available.

2.2 Hardware configurations (add-on)

Graphics accelerators come in different flavors: either as a separate PCI board that is able to pass through the video signal of a (possibly 2D or video accelerated) VGA board, or as a PCI board that does both VGA and 3D graphics (effectively replacing older VGA controllers). The 3Dfx boards based on the Voodoo Graphics (tm) belong to the former category. We will get into this again later.

If there is no address conflict, any 3D accelerator board could be present under Linux without interfering, but in order to access the accelerator, you will need a driver.

2.3 Performance limitations

2.3.1 Fill bound

Hardware accelerated graphics is performance bound for several reasons. A typical bottleneck is fill rate: the total number of pixels that the hardware could possibly do under optimal conditions, within a given time interval - e.g. about 40 Mpixels/second. Given a 640x480 screen resolution and zero overdraw, the hardware won't do more than 130 frames/second.

The amount of overdraw depends on the actual depth complexity of the scene (how many polygons would a ray through a pixel intersect) and the efficiency of the visible surface determination algorithm used by the application. Drawing each pixel twice means 65 frames/second, an overdraw of 2 (drawing each pixel thrice) gets you down to about 43 frames/second.

2.3.2 Missing refresh

Next, you will probably render with double buffering, swapping back and front buffer as soon as the frame is completed. Here the refresh rate of the display comes into play: you will only switch buffers during refresh. If your application misses a 60Hz refresh on every single frame, your effective frame rate will drop to 30Hz (every second refresh). Missing two refreshes gets you down to 20Hz.

2.3.3 Primitive bound

If the scene is not very detailed (only a few polygons, but those very large, with lots of overdraw), your application will probably be fill bound - it is possible to throw more primitives (lines, triangles, polygons) at the hardware, but the pixel pipeline can't go any faster anyway.

However, if your application insists on rendering a lot of small triangles or polygons, you might end up primitive bound. Given a PCI bandwidth of 33MHz times 32bit, or 132 MB/sec, and a per-triangle dataset of 3 vertices (9 coordinates, 16bit each, plus 3 colors, 24bit each), and a frame rate of 20Hz, you will transfer about 240K triangles/frame - not counting texture data, disk access, and other operations.

2.4 Hardware accelerated features

The rendering operations usually supported by usefull hardware accelerators are:

- Perspective correct texture mapping
- Alpha-blending, Fog
- Anti-aliasing
- Bi-linear and advanced texture filtering
- Level of detail (LOD) MIP mapping
- Sub-pixel correction
- Polygonal-based Gouraud shading and texture modulation
- Double buffering
- Depth buffering, stencil buffer

Usually, hardware allows increased screen resolution (software-only rendering being limited to 320x200 pixels for interactive frame rates), advanced filtering, true alpha channel translucency, and use true color 16bpp or 24bpp frame buffers.

2.5 A bit of Voodoo Graphics (tm) architecture

Usually, accessing texture memory and frame/depth buffer is a major bottleneck. For each pixel on the screen, there are at least one (nearest), four (bi-linear), or eight (tri-linear mipmapped) read accesses to texture memory, plus a read/write to the depth buffer, and a read/write to frame buffer memory.

The Voodoo Graphics (tm) architecture separates texture memory from frame/depth buffer memory by introducing two separate rendering stages, with two corresponding units (pixelfx and texelfx), each having a separate memory interface to dedicated memory. This gives an above-average fill rate, paid for restrictions in memory management (e.g. unused framebuffer memory can not be used for texture caching).

Moreover, a Voodoo Graphics (tm) could use two TMU's (texture management or texelfx units), and finally, two Voodoo Graphics (tm) could be combined accessing the same RAMDAC with a mechanism called Scan-Line Interleaving (SLI). SLI essentially means that each pixelfx unit effectively provides only every second scanline, which decreases bandwidth impact on each pixelfx' framebuffer memory.

3 Installation

Configuring Linux to support 3Dfx accelerators involves the following steps:

1. Installing the board.
2. Installing the Glide distribution.
3. Compiling, linking and/or running the application.

The next sections will cover each of these steps in detail.

3.1 Installing the board

Follow the manufacturer's instructions for installing the hardware or have your dealer perform the installation. It should not be necessary to select settings for IRQ, DMA channel, either Plug&Pray (tm) or factory defaults should work. The add-on boards described here are memory mapped devices and do not use IRQ's. The only kind of conflict to avoid is memory overlap with other devices.

As 3Dfx does not develop or sell any boards, do not contact them on any problems.

3.1.1 Troubleshooting the hardware installation

To check the installation and the memory mapping, do `cat /proc/pci`. The output should contain something like

```
Bus  0, device  12, function  0:
  VGA compatible controller: S3 Inc. Vision 968 (rev 0).
    Medium devsel.  IRQ 11.
    Non-prefetchable 32 bit memory at 0xf4000000.

Bus  0, device   9, function  0:
  Multimedia video controller: Unknown vendor Unknown device (rev 2).
    Vendor id=121a. Device id=1.
    Fast devsel.  Fast back-to-back capable.
    Prefetchable 32 bit memory at 0xfb000000.
```

for a Diamond Monster 3D used with a Diamond Stealth-64. Additionally a `cat /proc/cpuinfo /proc/meminfo` might be helpfull for tracking down conflicts and/or submitting a bug report.

With current kernels, you will probably get a boot warning like

```
Jun 12 12:31:52 hal kernel: Warning : Unknown PCI device (121a:1).
Please read include/linux/pci.h
```

which could be safely ignored. If you happen to have a board not very common, or have encountered a new revision, you should take the time to follow the advice in `/usr/include/linux/pci.h` and send all necessary information to linux-pcisupport@cao-vlsi.ibp.fr.

If you experience any problems with the board, you should try to verify that DOS and/or Win95 or NT support works. You will probably not receive any useful response from a board manufacturer on a bug report or request regarding Linux. Having dealt with the Diamond support e-mail system, I would not expect useful responses for other operating systems either.

3.1.2 Configuring the kernel

There is no kernel configuration necessary, as long as PCI support is enabled. The *Linux Kernel HOWTO* `<http://sunsite.unc.edu/mdw/HOWTO/Kernel-HOWTO.html>` should be consulted for the details of building a kernel.

3.1.3 Configuring devices

The current drivers to not (yet) require any special devices. This is different from other driver developments (e.g. the sound drivers, where you will find a `/dev/dsp` and `/dev/audio`). The driver uses the `/dev/mem` device which should always be available. In consequence, you need to use `setuid` or root privileges to access the accelerator board.

3.2 Setting up the Displays

There are two possible setups with add-on boards. You could either pass-through the video signal from your regular VGA board via the accelerator board to the display, or you could use two displays at the same time. Rely to the manual provided by the board manufacturer for details. Both configurations have been tried with the Monster 3D board.

3.2.1 Single screen display solution

This configuration allows you to check basic operations of the accelerator board - if the video signal is not transmitted to the display, hardware failure is possible.

Beware that the video output signal might deteoriate significantly if passed through the video board. To a degree, this is inevitable. However, reviews have complained about below-average of the cables provided e.g. with the Monster 3D, and judging from the one I tested, this has not changed.

There are other pitfalls in single screen configurations. Switching from the VGA display mode to the accelerated display mode will change resolution and refresh rate as well, even if you are using 640x480 e.g.

with X11, too. Moreover, if you are running X11, your application is responsible for demanding all keyboard and mouse events, or you might get stuck because of changed scope and exposure on the X11 display (that is effectively invisible when the accelerated mode is used) You could use SVGA console mode instead of X11.

If you are going to use a single screen configuration and switch modes often, remember that your monitor hardware might not enjoy this kind of use.

3.2.2 Dual screen display solution

The accelerator board does not need the VGA input signal. Instead of routing the common video output through the accelerator board, you could attach a second monitor to its output, and use both at the same time. This solution is more expensive, but gives best results, as your main display will still be hires and without the signal quality losses involved in a pass-through solution. In addition, you could use X11 and the accelerated full screen display in parallel, for development and debugging.

A common problem is that the accelerator board will not provide any video signal when not used. In consequence, each time the graphics application terminates, the hardware screensave/powersave might kick in depending on your monitors configuration. Again, your hardware might not enjoy being treated like this. You should use

```
setenv SST_DUALSCREEN 1
```

to force continued video output in this setup.

3.3 Installing the Glide distribution

The Glide driver and library are provided as a single compressed archive. Use `tar` and `gzip` to unpack, and follow the instructions in the README and INSTALL accompanying the distribution. Read the install script and run it. Installation puts everything in /usr/local/glide/include,lib,bin and sets the ld.conf to look there. Where it installs and setting ld.conf are independent actions. If you skip the ld.conf step then you need the LD_LIBRARY_PATH.

You will need to install the header files in a location available at compile time, if you want to compile your own graphics applications. If you do not want to use the installation as above (i.e. you insist on a different location), make sure that any application could access the shared libary at runtime, or you will get a response like `can't load library 'libglide.so'`.

3.3.1 Using the detect program

There is a `bin/detect` program in the distribution (the source is not available). You have to run it as root, and you will get something like

```
slot  vendorId   devId   baseAddr0  command  description
----  --------  ------  ----------  -------  -----------
```

```
00    0x8086  0x122d  0x00000000   0x0006  Intel:430FX (Triton)
07    0x8086  0x122e  0x00000000   0x0007  Intel:ISA bridge
09    0x121a  0x0001  0xfb000008   0x0002  3Dfx:video multimedia adapter
10    0x1000  0x0001  0x0000e401   0x0007  ???:SCSI bus controller
11    0x9004  0x8178  0x0000e001   0x0017  Adaptec:SCSI bus controller
12    0x5333  0x88f0  0xf4000000   0x0083  S3:VGA-compatible display co
```

as a result. If you do not have root privileges, the program will bail out with

```
Permission denied: Failed to change I/O privilege. Are you root?
```

output might come handy for a bug report as well.

3.3.2 Using the test programs

Within the Glide distribution, you will find a folder with test programs. Note that these test programs are under 3Dfx copyright, and are legally available for use only if you have purchased a board with a 3Dfx chipset. See the LICENSE file in the distribution, or their web site www.3dfx.com for details.

It is recommend to compile and link the test programs even if there happen to be binaries in the distribution. Note that some of the programs will requires some files like `alpha.3df` from the distribution to be available in the same folder. All test programs use the 640x480 screen resolution. Some will request a veriety of single character inputs, others will just state `Press A Key To Begin Test`. Beware of loss of input scope if running X11 on the same screen at the same time.

See the README.test for a list of programs, and other details.

4 Answers To Frequently Asked Questions

The following section answers some of the questions that (will) have been asked on the Usenet news groups and mailing lists. The FAQ has been subdivided into several parts for convenience, namely

- FAQ: Requirements?
- FAQ: Voodoo Graphics (tm)? 3Dfx?
- FAQ: Glide?
- FAQ: Glide and SVGA?
- FAQ: Glide and XFree86?
- FAQ: Glide versus OpenGL/Mesa?
- FAQ: But Quake?
- FAQ: Troubleshooting?

Each section lists several questions and answers, which will hopefully address most problems.

5 FAQ: Requirements?

5.1 What are the system requirements?

A Linux PC, PCI 2.1 compliant, a monitor capable of 640x480, and a 3D accelerator board based on the 3Dfx Voodoo Graphics (tm). It will work on a P5 or P6, with or without MMX.

5.2 Does it work with Linux-Alpha?

There is currently no Linux Glide distribution available for any platform besides i586. As the Glide sources are not available for distribution, you will have to wait for the binary. Quantum3D has DEC Alpha support announced for 2H97. Please contact Daryll Strauss if you are interested in supporting this.

5.3 Which chipsets are supported?

Currently, the most recent revision of the 3Dfx Voodoo Graphics (tm) chipset is supported under Linux. The Voodoo Rush (tm) chipset is not yet supported.

5.4 Which boards are supported?

This section lists the boards that are currently known to work under Linux. There are no officially supported boards, as 3Dfx does not sell any boards. The information here is based on the latest Linux kernels, at time of writing, and lists the boards that have been tested, plus boards that might work, but have yet to be checked.

It is important to recognize that Linux support for a given board does not only require a driver for the 3D accelerator component. If a board features its own VGA core as well, support by either Linux SVGA or XFree86 is required as well. Currently, an add-on solution is recommended, as it allows you to choose a regular graphics board well supported for Linux. There are other aspects discussed below.

The following configurations have been tested:

- Diamond Monster 3D with Diamond Stealth 64 3240XL
- Orchid Righteous 3D with S3-968 based graphics card
- Quantum3D Obsidian 50-4220

These are the existing Obsidian board configurations, most of them have not been tested yet, but should work as well.

Obsidian 50-2200

1 pixelfx with 2MB frame buffer memory, 1 texelfx with 2MB texture memory

Obsidian 50-2400

1 pixelfx with 2MB frame buffer memory, 1 texelfx with 4MB texture memory

Obsidian 50-4400

1 pixelfx with 4MB frame buffer memory, 1 texelfx with 4MB texture memory

Obsidian 50-2220

1 pixelfx with 2MB frame buffer memory, 2 texelfx chips with 2MB texture memory, each, for a total of 4MB texture memory

Obsidian 50-4220

1 pixelfx with 4MB frame buffer memory, 2 texelfx chips with 2MB texture memory, each, for a total of 4MB texture memory. This configuration was the original "Obsidian Pro" which has been used for the 3DS Plug-in Project (now done with Datapath Realistorm). Datapath used to call this "Pro VR".

Obsidian 50-4440

1 pixelfx with 4MB frame buffer memory, 2 texelfx chips with 4MB texture memory, each, for a total of 8MB texture memory. This configuration is now the intended target for the 3DS Plug-in Project (now done with Datapath Realistorm).

Obsidian 50-2440

1 pixelfx with 2MB frame buffer memory, 2 texelfx chips with 4MB texture memory, each, for a total of 8MB texture memory.

Obsidian 100-2440

aka 2440-SLI, aka XS-100, or just "SLI".

Two PCI boards, each with 1 pixelfx with 2MB frame buffer memory and 2 texelfx chips each with 4MB texture memory, each, for a total of 8MB texture memory per board. Texture have to be stored on both boards, so this does not equal 16MB texture memory. Video output is scan line interleaved for two times the standard fill rate.

The bundling deal with additional software for Autodesk 3DS MAX is dubbed Obsidian 3DS, which originally used a 50-4220 and now comes with a 50-4440 board.

The following boards have not yet been tested:

- Deltron RealVision Flash 3D

With the current Glide 2.4, the following Voodoo Rush (tm) based boards are not expected to work with Linux:

- Hercules Stingray 128/3D
- Intergraph Intense 3D Rush

As the Voodoo Rush (tm) chipset supports operations within a window, it is meant for use on accelerated VGA boards, which in turn require XFree86 oder Linux SVGA support not yet available.

Boards that are not based on 3Dfx chipsets (e.g. manufactured by S3, Matrox, 3Dlabs, Videologic) do *not* work with the 3Dfx drivers and are beyond the scope of this document.

5.5 Is the Hercules Stingray 128/3D supported?

In this board, the 2D accelerator is mounted on a PCI card, and the Voodoo Rush (tm) chipset is mounted on a daughterboard. Currently, this board is neither supported by Linux Glide, nor by XFree86 accelerated servers. Reportedly, the XFree86 SVGA server works, according to a posting on the Mesa mailing list. It supports 8, 16 and 32 bpp.

```
# device section settings
Chipset "AT24"
Videoram 4032

# videomodes tested by Oliver Schaertel
#  25.18  28.32  for 640 x 480   (70hz)
#  61.60         for 1024 x 786  (60hz)
#  120           for 1280 x 1024 (66hz)
```

There is currently no Voodoo Rush (tm) support. It might be worth a try, but as no test boards have been provided by the manufacturers, you are in your own.

Regarding the VGA component tied to the Voodoo Rush (tm) on this board, it is a Alliance Semiconductor's ProMotion-AT3D multimedia accelerator. XFree86's support for AT3D/AT24 will not be accelerated prior to XFree86 4.0, which is quite some time away.

5.6 Is the Intergraph Intense Rush supported?

Despite the fact that this board will be a single board integrated solution, it is essentially the same chipset combo (AT3D, Voodoo Rush (tm)), thus all disclaimers made above regarding the Hercules Stringray apply here as well. According to David E. Anderson from Intergraph, they will not be providing support for Linux at this time.

6 FAQ: Voodoo Graphics (tm)? 3Dfx?

6.1 Who is 3Dfx?

3Dfx is a San Jose based manufacturer of 3D graphics accelerator hardware for arcade games, game consoles, and PC boards. Their official website is www.3dfx.com. 3Dfx does not sell any boards, but there is an associcated company, Quantum3D. See their home page at www.quantum3d.com for additional information.

6.2 What is the Voodoo Graphics (tm)?

The Voodoo Graphics (tm) is a chipset manufactured by 3Dfx. It is used in hardware acceleration boards for the PC. See the HOWTO section on supported hardware.

There is a newer chipset, the Voodoo Rush (tm), that is currently not supported with Linux.

6.3 Where could I get additional info on Voodoo Graphics (tm)?

There is a FAQ by 3Dfx, which should be available at their web site. You will find retail information at the following locations:

- www.3dfx.com
- www.quantum3d.com
- www.orchid.com
- www.diamondmm.com
- www.deltrontech.com

7 FAQ: Glide? TexUS?

7.1 What is Glide anyway?

Glide is a proprietary API plus drivers to access 3D graphics accelerator hardware based on chipsets manufactured by 3Dfx. Glide has been developed and implemented for DOS, Windows, and Macintosh, and has been ported to Linux by Daryll Strauss.

7.2 What is TexUS?

In the distribution is a `libtexus.so`, which is the 3Dfx Interactive Texture Utility Software. It is an image processing libary and utility program for preparing images for use with the 3Dfx Interactive Glide library. Features of TexUS include file format conversion, MIPmap creation, and support for 3Dfx Interactive Narrow Channel Compression textures.

The TexUS utility program `texus` reads images in several popular formats (TGA, PPM, RGT), generates MIPmaps, and writes the images as 3Dfx Interactive textures files (see e.g. alpha.3df, as found in the distribution) or as an image file for inspection. For details on the parameters for `texus`, and the API, see the TexUS documentation.

7.3 Is Glide freeware?

Nope. Glide is neither GPL'ed nor subject to any other public license. See LICENSE in the distribution for any details. Glide is provided as binary only, and you should neither use nor distribute any files but the ones released to the public, if you have not signed an NDA. The Glide distribution including the test program sources are copyrighted by 3Dfx.

The same is true for all the sources in the Glide distribution. In the words of 3Dfx: These are not public domain, but they can be freely distributed to owners of 3Dfx products only. No card, No code!

7.4 Is the Glide source available?

Nope. The Glide source is made available only based on a special agreement and NDA with 3Dfx.

7.5 Is Linux Glide supported?

Currently, Linux Glide is unsupported. Basically, it is provided under the same disclaimers as the GLQuake DLL.

However, 3Dfx definitely wants to provide as much support as possible, and is in the process of setting up some prerequisites. For the time being, you will have to rely on the 3Dfx newsgroup (see below).

In addition, the Quantum3D web page claims that Linux support (for Obsidian) is planned for both Intel and AXP architecture systems in 2H97.

7.6 Where could I post Glide questions?

There are newsgroups currently available only on the NNTP server news.3dfx.com run by 3Dfx. This USENET groups are dedicated to 3Dfx and Glide in general, and will mainly provide assistance for DOS, Win95, and NT. The current list is:

```
3dfx.d3d.drivers
3dfx.events
3dfx.game.titles
3dfx.games.glquake
3dfx.glide
3dfx.glide.linux
3dfx.oem.products.diamond.monster3d
3dfx.oem.products.hercules.stingray128-3d
3dfx.oem.products.orchid.righteous3d
3dfx.oem.products.quantum3d.obsidian
3dfx.oem.products.realvision.flash3d
3dfx.products
3dfx.test
```

Please use news.3dfx.com/3dfx.glide.linux for all Lnux Glide related questions.

A mailing list dedicated to Linux Glide is in preparation (ETA in late august). Send mail to majordomo@gamers.org, no subject, body of the message `info linux-3dfx` to get information about the posting guidelines, the hypermail archive and how to subscribe to the list or the digest when available.

7.7 Where to send bug reports?

Currently, you should rely on the newsgroup (see above), that is news.3dfx.com/3dfx.glide.linux. There is no official support e-mail set up yet. For questions not specific to Linux Glide, make sure to use the other newsgroups.

7.8 Who is maintaining it?

3Dfx will appoint an official maintainer soon. Currently, inofficial maintainer of the Linux Glide port is Daryll Strauss. Please post bug reports in the newsgroup (above). If you are confident that you found a bug not previously reported, please mail to Daryll at daryll@harlot.rb.ca.us

7.9 How can I contribute to Linux Glide?

You could submit precise bug reports. Providing sample programs to be included in the distribution is another possibility. A major contribution would be adding code to the Glide based Mesa Voodoo driver source. See section on Mesa Voodoo below.

7.10 Do I have to use Glide?

Yes. As of now, there is no other Voodoo Graphics (tm) driver available for Linux.

7.11 Should I program using the Glide API?

That depends on the application you are heading for. Glide is a proprietary API that is partly similar to OpenGL or Mesa, partly contains features only available as EXTensions to some OpenGL implementations, and partly contains features not available anywhere but within Glide.

If you want to use the OpenGL API, you will need Mesa (see below). Mesa, namely the Mesa Voodoo driver, offers an API resembling the well documented and widely used OpenGL API. However, the Mesa Voodoo driver is in early alpha, and you will have to accept performance losses and lack of support for some features.

In summary, the decision is up to you - if you are heading for maximum performance while accepting potential problems with porting to non-3Dfx hardware, Glide is not a bad choice. If you care about maintenance, OpenGL might be the best bet in the long run.

7.12 What is the current version?

The version of Linux Glide to be released to the public is 2.4, as is the next release of Glide for DOS/Windows.

Note that this HOWTO has been written based on Linux Glide 2.3.1, as Glide 2.4 has not been released, and the Linux Glide 2.4 port as not been finished yet. As the API did not change, and as there are no changes planned for the Linux Glide distribution, this document should still address the most common problems.

7.13 Is Linux Glide identical to DOS/Windows Glide?

The version of Linux Glide to be publicly released will be 2.4, following the release of DOS/Windows Glide 2.4. API and implementation are supposed to be identical.

Glide 2.2 has been ported to Linux in April 1997. The port of Glide 2.3.1 has been done in June 1997. Both lacked a key optimization for triangle setup, which will be included in the public Linux Glide 2.4 release. The previous ports have never been publicly available, and have been used for beta tests only.

7.14 Where to I get information on Glide?

There is exhaustive information available from 3Dfx. You could download it from their home page at www.3dfx.com/software/download_glide.html. These are for free, presuming you bought a 3Dfx hardware based board. Please read the licensing regulations.

Basically, you should look for some of the following:

- *Glide Release Notes*
- *Glide Programming Guide*
- *Glide Reference Manual*
- *Glide Porting Guide*
- *TexUs Texture Utility Software*
- *ATB Release Notes*
- *Installing and Using the Obsidian*

These are available as Microsoft Word documents, and part of the Windows Glide distribution, i.e. the self-extracting archive file. Postscript copies for separate download should be available at www.3dfx.com as well. Note that the release numbers are not always in sync with those of Glide.

7.15 Where to get some Glide demos?

You will find demo sources for Glide within the distribution (test programs), and on the 3Dfx home page. The problem with the latter is that some require ATB. To port these demos to Linux, the event handling has to be completely rewritten.

In addition, you might find useful some of the OpenGL demo sources accompanying Mesa and GLUT. While the Glide API is different from the OpenGL API, they target the same hardware rendering pipeline.

8 FAQ: Glide and SVGA?

You should have no problems running Glide based applications either single or dual screen using VGA modes. It might be a good idea to set up the 640x480 resolution in the SVGA modes, too, if you are using a single screen setup.

9 FAQ: Glide and XFree86?

9.1 Does it run with XFree86?

Basically, the Voodoo Graphics (tm) hardware does not care about X. The X server will not even notice that the video signal generated by the VGA hardware does not reach the display in single screen configurations. If your application is not written X aware, Glide switching to full screen mode might cause problems (see troubleshooting section). If you do not want the overhead of writing an X11-aware application, you might want to use SVGA console mode instead.

So yes, it does run with XFree86, but no, it is not cooperating if you don't write your application accordingly.

9.2 Does it only run full screen?

See above. The Voodoo Graphics (tm) hardware is not window environment aware, neither is Linux Glide.

9.3 What about GLX for XFree86?

There are a couple of problems.

The currently supported Voodoo Graphics (tm) hardware and the available revision of Linux Glide are full screen only, and not set up to share a framebuffer with a window environment. Thus GLX or other integration with X11 is not yet possible.

The Voodoo Rush (tm) might be capable of cooperating with XFree86 (that is, an SVGA compliant board will work with the XFree86 SVGA server), but it is not yet supported by Linux Glide, nor do S3 or other XFree86 servers support these boards yet.

In addition, GLX is tied to OpenGL or, in the Linux case, to Mesa. The XFree86 team is currently working on integrating Mesa with their X Server. GLX is in beta, XFree86 3.3 has the hooks for GLX. See Steve Parker's GLX pages at www.cs.utah.edu/ sparker/xfree86-3d/ for the most recent information. Currently, Mesa still uses its GLX emulation with Linux.

10 FAQ: Glide versus OpenGL/Mesa?

10.1 Is Glide OpenGL?

No, Glide is a proprietary 3Dfx API which several features specific to the Voodoo Graphics (tm) and Voodoo Rush (tm). A 3Dfx OpenGL is in preparation (see below). Several Glide features would require EXTensions to OpenGL, some of which already found in other implementations (e.g. paletted textures).

The closest thing to a hardware accelerated Linux OpenGL you could currently get is Brian Paul's Mesa along with David Bucciarelli's Mesa Voodoo driver (see below).

10.2 Does Mesa work with 3Dfx?

As of Mesa 2.3 Beta3, Mesa works with Linux Glide 2.2, similar to Mesa with Glide for DOS/Windows. There are patches to Mesa 2.3b3 for Linux Glide 2.3.1. Later versions of Mesa will work with Linux Glide 2.4; as the API did not change, the patches to Mesa-2.3b3 might be sufficient. The Glide distribution is not part of the Mesa distribution.

You will need to get the Mesa library archive from the iris.ssec.wisc.edu FTP site.

10.3 Where to get additional information on OpenGL?

Use Mark Kilgard's Gateway to OpenGL Info at reality.sgi.com/mjk_asd/opengl-links.html, and proceed from there.

10.4 Where to get info on Mesa?

The Mesa home page is at www.ssec.wisc.edu/ brianp/Mesa.html. There is an archive of the Mesa mailing list. at www.iqm.unicamp.br/mesa/. This list is not specific to 3Dfx and Glide, but if you are interested in using 3Dfx hardware to accelerate Mesa, it is a good place to start.

10.5 Where to get information on Mesa Voodoo?

For latest information on the Mesa Voodoo driver maintained by David Bucciarelli tech.hmw@plus.it see the home page at www-hmw.caribel.pisa.it/fxmesa/.

10.6 Is there a commercial OpenGL for Linux and 3Dfx?

3Dfx has publicly announced an OpenGL implementation for Windows for this year (2H97). It is not known whether this will be available for Linux as well.

As for third party commercial OpenGL, I am aware of three products:

- MetroLink MetroOpenGL

- XInside OpenGL
- Evans & Sutherland OpenGL

The latter was distributed by Portable Graphics, and was a straight Linux port of the OpenGL reference software implementation, with a link kit for an older revision of the XFree86 X servers. Portable Graphics never promised hardware support. To my knowledge, this product is no longer available.

The other two promised support for hardware accelerators, but both are tied to proprietary ports of the X server, and both do not support any 3D acceleration, as far as I know.

10.7 How about GLUT?

Mark Kilgard's GLUT distribution is a very good place to get sample applications plus a lot of useful utilities. You will find it at reality.sgi.com/mjk_asd/glut3/, and you should get it anyway. The current release is GLUT 3.4.

However, as GLUT handles double buffers, windows, events, and other operations closely tied to hardware and operating system, a Voodoo-GLUT requires several specific modifications. There is an alpha release available as part of the most recent Mesa distribution (David Bucciarelli, Henri Fousse).

11 FAQ: But Quake?

11.1 What about the Quake GL?

The DLL released by 3Dfx is only available for Windows. It supports a Quake-specific a subset of OpenGL only (see http://www.cs.unc.edu/ martin/3dfx.html for an inofficial list of supported code paths). No, this DLL is not ported to Linux. No, there is no Glide based Quake version, not even for Windows. I have no information about a Linux glQuake.

12 FAQ: Troubleshooting?

12.1 Has this hardware been tested?

See hardware requirements above, there is a list of hardware that has been found to work.

12.2 Failed to change I/O privilege?

You need to be root, or `setuid` your application to run a Glide based application. For DMA, the driver accesses `/dev/mem`, which is not writeable for anybody but root, with good reasons. See the README in the Glide distribution for Linux.

12.3 Does it work without root privilege?

There are compelling case where the setuid requirement is a problem, obviously. There are currently solutions in preparation, which require changes to the library internals itself.

12.4 Displayed images looks awful (single screen)?

If you are using the analog pass through configuration, the common SVGA or X11 display might look pretty bad. You could try to get a better connector cable than the one provided with the accelerator board (the ones delivered with the Diamond Monster 3D are reportedly worse then the one accompanying the Orchid Righteous 3D), but up to a degree there will inevitably be signal loss with an additional transmission added.

If the 640x480 full screen image created by the accelerator board does look awful, this might indicate a real hardware problem. You will have to contact the board manufacturer, not 3Dfx for details, as the quality of the video signal has nothing to do with the accelerator - the board manufacturer chooses the RAMDAC, output drivers, and other components responsible.

12.5 The last frame is still there (single or dual screen)?

You terminated your application with Ctrl-C, or it did not exit normally. The accelerator board will dutifully provide the current content of the framebuffer as a video signal unless told otherwise.

12.6 Powersave kicks in (dual screen)?

When you application terminates in dual screen setups, the accelerator board does not provide video output any longer. Thus powersave kicks each time. To avoid this, use

```
setenv SST_DUALSCREEN 1
```

12.7 My machine seem to lock (X11, single screen)?

If you are running X when calling a Glide application, you probably moved the mouse out of the window, and the keyboard inputs do not reach the application anymore.

If you application is supposed to run concurrently with X11, it is recommend to expose a full screen window, or use the `XGrabPointer` and `XGrabServer` functions to redirect all inputs to the application while the X server cannot access the display. Note that grabbing all input with `XGrabPointer` and `XGrabServer` does not qualify as well-behaved application, and that your program might block the entire system.

If you experience this problem without running X, be sure that there is no hardware conflict (see below).

12.8 My machine locks (single or dual screen)?

If the system definitely does not respond to any inputs (you are running two displays and know about the loss of focus), you might experience a more or less subtle hardware conflict. See installation troubleshooting section for details.

If there is no obvious address conflict, there might still be other problems (below). If you are writing your own code the most common reason for locking is that you didn't snap your vertices. See the section on snapping in the Glide documentation.

12.9 My machine locks (used with S3 VGA board)?

It is possible you have a problem with memory region overlap specific to S3. There is some info and a patch to the so-called S3 problem in the 3Dfx web site, but these apply to Windows only. To my understanding, the cause of the problem is that some S3 boards (older revisions of Diamond Stealth S3 968) reserve more memory space than actually used, thus the Voodoo Graphics (tm) has to be mapped to a different location. However, this has not been reported as a problem with Linux, and might be Windows-specific.

12.10 No address conflict, but locks anyway?

If you happen to use a motherboard with non-standard or incomplete PCI support, you could try to shuffle the boards a bit. I am running an ASUS TP4XE that has that non-standard modified "Media Slot", i.e. PCI slot4 with additional connector for ASUS-manufactured SCSI/Sound combo boards, and I experienced severe problems while running a Diamond Monster 3D in that slot. The system operates flawlessly since I put the board in one of the regular slots.

12.11 Compile/link error: grSstWinOpen()?

As Linux Glide will be version 2.4, this error should not occur. This function was not available in Glide and thus Linux Glide 2.2;; the latter has never been released to the public.

12.12 Compile/link error: grSstOpen()?

Your application's source is based on Glide 2.2, and this function was removed in Glide 2.3;. It is not available anymore and so may not be used with Linux Glide 2.4. Modify your application to use the function `grSstWinOpen` instead.

As the Linux Glide integration with Mesa was based on Glide 2.2 originally, earlier versions of Mesa might produce compile time errors. The Mesa-2.3b3 release has been patched to be used with Linux Glide 2.3.1; make sure to get both the distribution and the patches, or, preferably, a later revision of Mesa.

12.13 Cannot open shared object file?

```
test25: error in loading shared libraries
```

`libglide2x.so: cannot open shared object file: No such file or directory`

If, for whatever reasons, you have a binary lying around compiled for a different revision of Linux Glide, or if there is an inconsistency in your `ldconfig` setup, the program will not find the shared library. Check the name (e.g. `libglide2x.so`) and make sure that the proper flags are used when compiling and linking - e.g. `-lglide` might not work with the default installation.

Note that the naming of Linux Glide revisions follows the conventions used in the 3Dfx Windows distribution, not the conventions common for Linux.

12.14 Mesa compilation problems

Be sure to set `USE_GLIDE_FULLSCREEN` in `fxmesa.h`. Check whether the linkage options (e.g. `-lglide`) match the name of the Linux Glide library installed (e.g. `-lglide2x` instead). Be sure to use the patches to Mesa-2.3b3 or a later release, as all Mesa releases up to 2.3b3 were based on Linux Glide 2.2. See above.

12.15 Mesa runs, but does not access the board?

Be sure that you recompiled all the libraries (including the toolkits the demo programs use - remember that GLUT does not yet support Voodoo Graphics (tm)), and that you removed the older libraries, run `ldconfig`, and/or set your `LD_LIBRARY_PATH` properly. Mesa supports several drivers in parallel (you could use X11 SHM, off screen rendering, and Mesa Voodoo at the same time), and you might have to create and switch contexts explicitely (see `MakeCurrent` function) if the Voodoo Graphics (tm) isn't chosen by default.