



Sun™ OpenGL 1.3 for Solaris™ Implementation and Performance Guide

Beta Draft

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 805-4445-16
August 2002, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, SunSoft, SunDocs, SunExpress, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements. OpenGL is a registered trademark of Silicon Graphics, Inc

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, SunSoft, SunDocs, SunExpress, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun. OpenGL est une marque déposée de Silicon Graphics, Inc.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISÉE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.

Contents

Preface xi

1. Introduction to Sun OpenGL for Solaris Software 1

Product Functionality 1

Compatibility Issues 3

Upgrading and Motif Versions 3

MT-Safe 4

Supported Platforms 5

Where to Look for Information on OpenGL Programming 5

2. Architecture 7

A Quick Review of the Architecture 8

Software Architecture 9

 Vertex Processing Architecture 12

 Rasterization and Fragment Processing Architecture 12

3. Performance 13

Acceleration Versus Optimization 13

Multiscreen Environment Performance 14

Multisampling 15

View Frustum Culling on Display Lists 15

Display List Compression	16
General Tips on Vertex Processing	17
Vertex Arrays	17
MultiDrawArrays	17
Triangle List and Mesh Primitives	18
Consistent Data Types	18
Low Batching	20
Optimized Data Types	21
Hardware Specific Acceleration	21
Sun XVR-1000 Graphics Accelerator Performance	22
OpenGL Acceleration on the Sun XVR-1000 Graphics Accelerator	22
Sun Expert3D, Expert3D-Lite, and Sun XVR-500 Graphics Accelerator Performance	23
OpenGL Acceleration on the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 Graphics Accelerators	24
Sun Elite3D Graphics Accelerator Performance	25
Attributes	25
Pixel Operations	27
Consistent Data	27
Sun Creator and Creator3D Graphics Accelerator Performance	28
Attributes Affecting Creator3D Performance	28
Attributes Affecting Software Rendering Performance	36
Pixel Operations	37
Pixel Transfer Pipeline (ARB) Imaging Extensions and the Pixel Transform	41
Implementation	42
How To Use the Pixel Transfer Pipeline and Pixel Transform	43
Software Performance	56

4. X Visuals for Sun OpenGL for Solaris	57
Programming With X Visuals for Sun OpenGL for Solaris Software	57
Colormap Flashing for OpenGL Indexed Applications	58
GL Rendering Model and X Visual Class	59
Stereo	59
Rendering to DirectColor Visuals	61
Overlays	61
Server Overlay Visual (SOV) Convention	61
Compatibility of SOV with other Overlay Models	62
Gamma Correction	63
5. Tips and Techniques	65
Avoiding Overlay Colormap Flashing	65
Changing the Limitation on the Number of Simultaneous GLX Windows	66
Hardware Window ID Allocation Failure Message	66
Getting Peak Frame Rate	67
Identifying the Release Version	67
Determining Visuals Supported by a Specific Frame Buffer	68
Developing Applications for 64-bit	68
Common 64-bit Application Development Errors	68
Index	71

Figures

FIGURE 2-1	Basic Architecture	8
FIGURE 2-2	Sun OpenGL for Solaris Software Architecture	11
FIGURE 3-1	Hardware Rasterizer Path for Creator3D	32
FIGURE 3-2	Text Load Processing Flow	34
FIGURE 3-3	Sun OpenGL for Solaris Architecture for Drawing Pixels	38
FIGURE 3-4	Pixel Transfer Pipeline Functions and Order of Execution	42

Tables

TABLE P-1	Sun OpenGL 1.3 for Solaris Features and Benefits	xii
TABLE 2-1	Data Paths Through the OpenGL for Solaris System	10
TABLE 3-1	3D Optimized Cases	35

Preface

The *Sun OpenGL 1.3 for Solaris Implementation and Performance Guide* provides information on Sun's implementation of the OpenGL™ graphics library for the Solaris™ operating environment from Sun Microsystems, Inc.

Who Should Use This Book

This book is intended for application developers who are using Sun's OpenGL for Solaris software to port or develop OpenGL applications on Solaris. It assumes familiarity with OpenGL functionality and with the principles of 2D and 3D computer graphics.

What's New In This Release

This release of Sun™ OpenGL® for Solaris™ contains all of the functionality included in the previous release of OpenGL (version 1.2.3, shipped on 03/2002). TABLE P-1 lists the features and benefits of Sun OpenGL 1.3 for Solaris.

TABLE P-1 Sun OpenGL 1.3 for Solaris Features and Benefits

Feature	Description
Cube Map Textures	Improves the visual effects of reflection mapping using a set of six two-dimensional texture images representing the faces of a cube.
Texture Dot3 Environment Mode	Enables per-pixel lighting computations for better lighting effects.
Texture Add Environment Mode	Provides a texture function to add incoming pixel and texture source colors.
Point Parameters extension	Attenuates point size based on the distance to eye point (for example, better simulation of lights on an airport runway, and so on).
Dynamic Video Resolution extension	Scales the frame rate by dynamically adjusting frame buffer video resolution.
Gradient Background Clear	Faster clearing of background with gradient colors.
Display List Tuning	Faster display list compilation and rendering.
Bug fixes	Improved robustness and stability.

How This Book Is Organized

This book is organized as follows:

Chapter 1 “Introduction to Sun OpenGL for Solaris Software,” provides a description of the Sun OpenGL for Solaris software.

Chapter 2 “Architecture,” presents information on the Sun OpenGL for Solaris architecture.

Chapter 3 “Performance,” presents specific information on using Sun’s OpenGL library for specific hardware platforms.

Chapter 4 “X Visuals for Sun OpenGL for Solaris,” presents information on visuals for the OpenGL for Solaris product.

Chapter 5 “Tips and Techniques,” contains information that may make using the Sun OpenGL for Solaris library easier.

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Solaris™ operating environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface ¹	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

¹ The settings on your browser might differ from these settings.

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

<http://www.sun.com/documentation>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

docfeedback@sun.com

Please include the part number (805-4445-16) of your document in the subject line of your email.

Introduction to Sun OpenGL for Solaris Software

Sun OpenGL for Solaris is Sun's native implementation of the OpenGL application programming interface (API). The OpenGL API is an industry-standard, vendor-neutral graphics library. It provides a small set of low-level geometric primitives and many basic and advanced 3D rendering features, such as modeling transformations, shading, lighting, anti-aliasing, texture mapping, fog, and alpha blending.

Product Functionality

Sun OpenGL 1.3 for Solaris is an implementation based on the OpenGL 1.3 standard specifications. It also includes additional extensions that are not available in the base OpenGL 1.3 specification. These extensions are listed below in alphabetical order:

- `GL_ARB_imaging` – ARB imaging extension
- `GL_ARB_multitexture` – Multitexture extension
- `GL_ARB_texture_border_clamp` – Texture border clamp extension
- `GL_ARB_transpose_matrix` – Transpose matrix extension
- `GL_EXT_abgr` – ABGR reverse-order color format extension
- `GL_EXT_anisotropic_texture` – Anisotropic texture extension
- `GL_EXT_blend_color` – Blend color extension
- `GL_EXT_blend_logic_op` – Blend logic op extension
- `GL_EXT_blend_minmax` – Blend minmax extension
- `GL_EXT_blend_subtract` – Blend subtract extension
- `GL_EXT_compiled_vertex_array` – Compiled vertex array extension
- `GL_EXT_convolution` – Convolution extension

- `GL_EXT_histogram` – Histogram extension
- `GL_EXT_multi_draw_arrays` – Multidraw array and element extensions
- `GL_EXT_pixel_transform` – Pixel transform extension
- `GL_EXT_polygon_offset` – Polygon offset extension
- `GL_EXT_rescale_normal` – Rescale normal extension
- `GL_EXT_texture_env_combine` – Texture environment combine extension
- `GL_EXT_texture3D` – 3D texture mapping extension
- `GL_HP_convolution_border_modes` – Convolution border mode extension
- `GL_HP_occlusion_test` – HP occlusion test extension
- `GL_HP_occlusion_test` – Occlusion culling extension
- `GL_SGI_color_table` – SGI color table extension
- `GL_SGI_texture_color_table` – Texture color table extension
- `GL_SGIS_detail_texture` – Detail texture extension
- `GL_SGIS_sharpen_texture` – Sharpen texture extension
- `GL_SGIS_texture_border_clamp` – Texture border clamp extension
- `GL_SGIS_texture_filter4` – Texture filter4 extension
- `GL_SGIX_texture_lod_bias` – Texture LOD bias extension
- `GL_SGIX_texture_scale_bias` – Texture scale and bias extension
- `GL_SUN_convoution_border_modes` – Convolution border mode extension
- `GL_SUN_global_alpha` – Global alpha extension
- `GL_SUN_mesh_array` – Mesh primitive extension
- `GL_SUN_multi_draw_arrays` – Multidraw array and element extension
- `GL_SUN_slice_accum` – Slice Accumulation extension
- `GL_SUN_triangle_list` – Triangle list extension
- `GL_SUN_vertex` – Vertex extension
- `GL_SUNX_constant_data` – Constant data extension
- `GL_SUNX_geometry_compression` – Sun geometry compression extension
- `GL_SUNX_surface_hint` – Surface size hint extension
- `GLX_SGI_make_current_read` – Make current read extension
- `GLX_SGIX_fbconfig` – fbconfig extension
- `GLX_SGIX_pbuffer` – pbuffer extension
- `GLX_SUN_init_threads` – Multithread support extension
- `GLX_SUN_get_transparent_index` – Transparent pixel index extension

Note – If your application uses extensions, you will need to ensure that it also handles the functionality in an OpenGL compliant manner. Application can use the `glXQueryExtensionString(3gl)` function to programmatically determine the extensions supported by an OpenGL implementation.

Compatibility Issues

Applications compiled with the previous Sun OpenGL for Solaris libraries will run unchanged with the Sun OpenGL 1.3 for Solaris implementation. However, note the following backward compatibility issues:

- If your application uses the features in the Sun OpenGL for Solaris library that are not available in the previous release, it will not be backward compatible with the previous Sun OpenGL for Solaris libraries.
- To reduce function call overhead and improve performance for vertex calls in immediate mode, vertex commands such as `glVertex`, `glColor`, `glNormal`, `glTexCoord` and `glIndex` have been redefined as macros in the Sun OpenGL for Solaris software. You can use the compiler directive flag, `-DSUN_OGL_NO_VERTEX_MACROS`, to disable the vertex macro feature. See the `glVertex(3gl)` man page for further information.

Upgrading and Motif Versions

Sun OpenGL contains OpenGL Motif Drawing Widget libraries for Motif 2.x and 1.x libraries. Since Motif 2.x, available in the Solaris 7 system software or later compatible releases, is incompatible with Motif 1.x, when installing Sun OpenGL for Solaris software you need to make certain the OpenGL Motif Drawing Widget library symbolic links point to the appropriate version of OpenGL Motif Drawing Widget library.

If you are using pre-Solaris 7 system software, `libGLw.so` should point to `libGLw.so.1`; and if you are using Solaris 7 system software or subsequent compatible release, `libGLw.so` should point to `libGLw.so.2`.

When upgrading from Solaris 2.6 system software to Solaris 7 system software (or subsequent release), it is recommended that you reinstall Sun OpenGL software as well.

If you are using Solaris 7 system software (or subsequent release) and are developing with Motif 1.2 linking with `-lXm12`, you must link the OpenGL Motif Drawing Widgets with `-lGLw12` instead of `-lGLw`.

MT-Safe

The Sun OpenGL for Solaris library is multithread safe (MT-safe). Multiple rendering threads are allowed in a single process.

To initialize MT-safe mode for Xlib/Xt and OpenGL, applications need to call `XInitThreads()`, `XtToolkitThreadInitialize()` and `glXInitThreadsSUN()` before multiple threads are created. See man pages `glXInitThreadsSUN(3gl)`, `XInitThreads(3X11)`, and `XtToolkitThreadInitialize(3Xt)` for more detail.

If your multithreaded application uses only one OpenGL rendering thread, you can obtain better performance by avoiding calling the `glXInitThreads()` function.

Multithread safe allows OpenGL parallelism. This parallelism supports single to multiple CPUs as well as single to multiple screens.

In a non-Xinerama environment, the maximum number of supported OpenGL rendering threads is 512. In a Xinerama environment, the maximum number is $512/(N+1)$ where N is the number of screens in the Xinerama environment.



Caution – When the OpenGL renderer (see `glGetString(GL_RENDERER)`) is a graphics accelerator (not a software renderer), multiple rendering threads to the same screen might perform slower than single threaded rendering due to the overhead of context switching. If possible, avoid multithreaded rendering to a single graphics accelerated screen.

Supported Platforms

Sun OpenGL 1.3 for Solaris supports the following devices:

- All SPARC™ systems equipped with the following graphics boards are supported: PGX, PGX24, PGX32, PGX64, Creator, Creator3D, Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators. This includes Sun Blade desktops, Sun Fire servers, Ultra™ desktops, and Ultra Enterprise™ servers.
- For systems equipped with Creator, Creator3D, Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators, OpenGL functionality is accelerated in hardware.
- For systems equipped with PGX, PGX24, PGX32, and PGX64 graphics boards, OpenGL functionality is performed in software.
- For systems equipped with PGX64 graphics boards, OpenGL index rendering is not supported in 32-bit depth mode. You must configure the frame buffer to 8-bit depth for OpenGL index rendering.

Where to Look for Information on OpenGL Programming

For information on how to write an OpenGL application, see the following books:

- *OpenGL Programming Guide* by Neider, Davis, and Woo
- *OpenGL Reference Manual* by the OpenGL Architecture Review Board
- *OpenGL Programming for X Windows System* by Mark Kilgard

These books are published by Addison-Wesley and are available through your local bookstore.

For more information on OpenGL, you may want to refer to “The Design of the OpenGL Interface” written by Mark Segal and Kurt Akeley. A PostScript copy of this document is included in the `SUNWgldoc` package.

<http://www.opengl.org> is also a good source for information on OpenGL.

Architecture

The purpose of designing a graphics system architecture is to enable performance within the constraints of cost and functionality goals. Hardware design places various stages of the graphics pipeline into hardware accelerators. Software design uses the hardware features and complements the hardware by providing complete coverage of functionality.

Understanding the hardware and software architecture of a particular system will help you determine whether a feature is accelerated in the graphics hardware or implemented in software. This will enable you to identify which path through the system your application uses for the feature. With this information, you can project your application's performance. Given knowledge of performance versus functionality trade-offs, you can make informed choices about how to use the system to maximize your application's interactivity.

This chapter describes the Sun OpenGL for Solaris architecture.

A Quick Review of the Architecture

As a first step in examining the Sun OpenGL for Solaris architecture, FIGURE 2-1 shows the basic architecture of the Sun OpenGL for Solaris library.

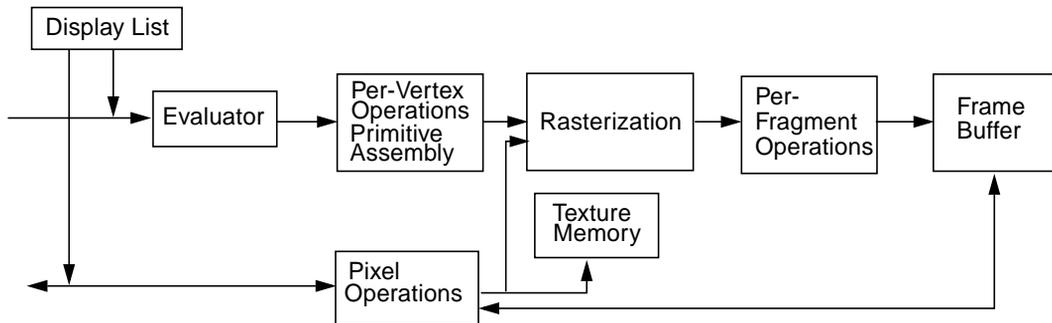


FIGURE 2-1 Basic Architecture

In the first stage of the OpenGL pipeline, vertex data enters the pipeline, and curve and surface geometry is evaluated. Next, colors, normals, and texture coordinates are associated with vertices, and vertices are transformed and lit. Vertices are then assembled into geometric primitives.

The rasterization stage converts geometric primitives into frame buffer addresses and values, or fragments. Each fragment may be altered by per-fragment operations, such as blending. Per-fragment operations store updates into the frame buffer based on incoming and previously-stored Z values (for Z buffering), blending of incoming fragment colors with stored colors, as well as masking and other logical operations.

Pixel data is processed in the pixel operation stage. The resulting data is stored as texture memory, or rasterized and processed as fragments before being written to the frame buffer.

The task of the hardware and software implementors at Sun was to implement the OpenGL functionality. The remainder of this chapter describes this implementation.

Software Architecture

Once the hardware designers have determined what the hardware will accelerate, all other decisions regarding performance fall to the software implementors. Software implementors need to consider the following questions:

What hardware features will be used?

1. What features that are not accelerated in hardware can the software optimize?
2. How will the software implement all functionality?

In response to these questions, the Sun OpenGL for Solaris software developers implemented OpenGL as follows:

- Accelerated OpenGL by using all features of the Creator, Creator3D, Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators.
- For the Creator and Creator3D systems, optimized line and point transformation and clip test, and a subset of texture lookup and filtering.
- System hardware acceleration.
- Implemented OpenGL to its complete specification by writing code for primitive assembly and vertex processing, including:
 - Coordinate transformations
 - Texture coordinate generation
 - Clipping
- Implemented two forms of software rasterization for OpenGL features not rasterized in hardware:
 - Optimized software rasterizer for many texturing functions and pixel operations. Software rasterization is done by the CPU using an optimized implementation. On an UltraSPARC CPU, some features, such as texturing rasterization, may be handled using software code employing the VIS instruction set.
 - A software rasterizer for all features not handled by the hardware or by the VIS software.

This implementation of the Sun OpenGL for Solaris library allows devices with varying capabilities to run efficiently. It enables Sun OpenGL for Solaris applications to run on the following types of devices:

- Model coordinate device (Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 or Sun XVR-1000 graphics accelerator system) – Handles most OpenGL functionality in hardware, including vertex processing, primitive assembly, rasterization, and fragment operations.

- Device coordinate device (Creator or Creator3D graphics accelerator system) – Performs vertex processing. Rasterization and fragment processing is handled in hardware.
- Memory mappable devices (PGX graphics, software renderer) – Vertex processing, primitive assembly, rasterization, and fragment processing are performed in software, and the results are written to the memory-mapped frame buffer.

FIGURE 2-2 on page 11 illustrates the graphics software architecture of the Sun OpenGL for Solaris product. This figure shows the paths that application data can take through the OpenGL system, depending on the type of hardware device the application is running on. TABLE 2-1 summarizes the data paths with reference to several hardware platforms.

TABLE 2-1 Data Paths Through the OpenGL for Solaris System

Platform	Vertex Processing	Rasterization	Performance
MC device (Elite3D, Expert3D, Expert3D-Lite XVR-500 or XVR-1000)	Hardware vertex processing	Hardware rasterizer	Fastest path
	Software vertex processing	Hardware rasterizer	Fast path
	Software vertex processing	Software rasterizer	Slow path
DC device (Creator3D or Creator)	Software vertex processing	Hardware rasterizer	Fast path
	Software vertex processing	Software rasterizer	Slow path
Memory map (PGX, software renderer)	Software vertex processing	Software rasterizer	Only path

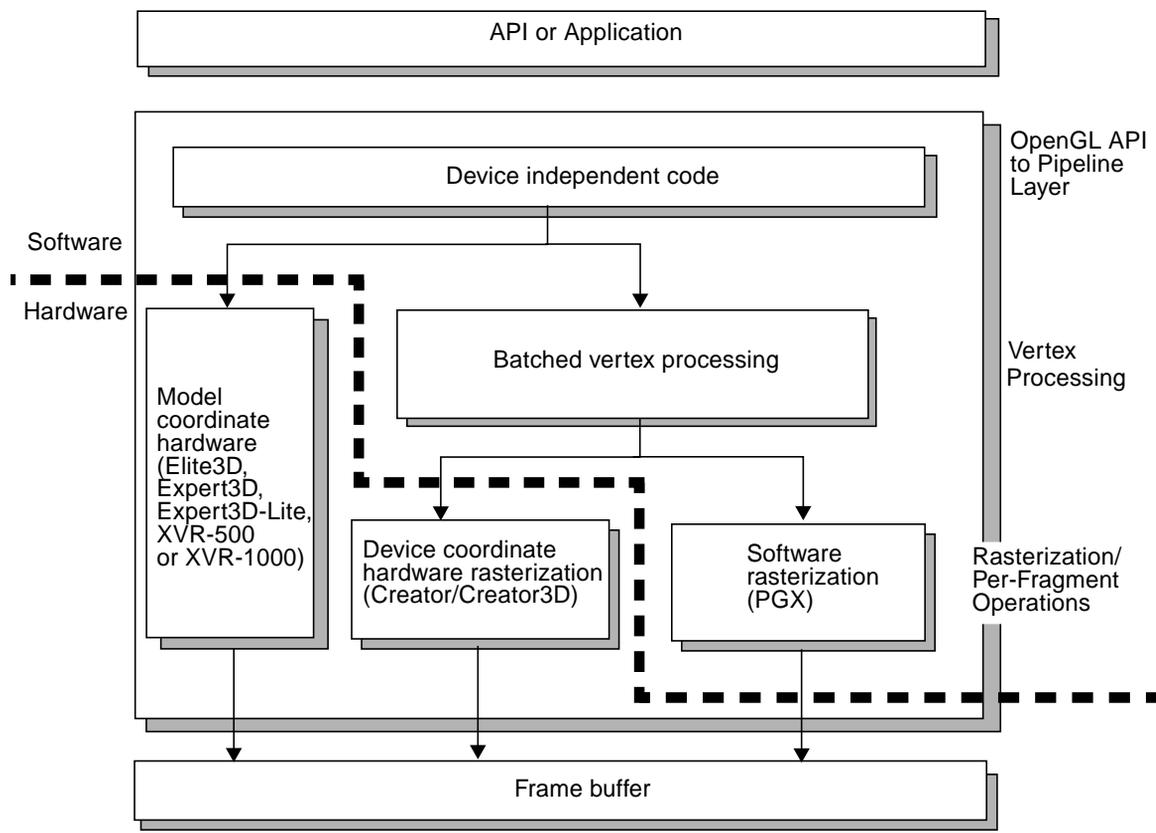


FIGURE 2-2 Sun OpenGL for Solaris Software Architecture

Vertex Processing Architecture

As FIGURE 2-2 on page 11 shows, Sun's OpenGL implementation handles vertex processing in several ways:

- **Hardware vertex processing** – On model coordinate devices, vertex processing is done via the hardware. In addition to hardware acceleration, the model coordinate (MC) pipeline is optimized for vertex arrays and display list mode. The model coordinate pipeline also recognizes consistent data pattern within `glBegin/glEnd` pairs. If the data is consistent, the software is able to use hardware resources efficiently.
- **Software vertex processor** – This is the fully optimized path from the software implementor's point of view. The principal optimization is that the model coordinate software pipeline recognizes consistent data types within `glBegin/glEnd` pairs: if the data is consistent, the software pipeline is able to use CPU resources efficiently.

The OpenGL vertex array commands result in the best performance for vertex processing on all hardware platforms. For repeated rendering of the same geometry, display lists provide significant performance benefits over immediate mode rendering.

Rasterization and Fragment Processing Architecture

Rasterization and fragment processing is handled in one of the following ways:

- **Hardware rasterizer** – The graphics subsystem handles lines, points, and triangles, and does simple fragment processing, such as blending and the depth-buffer test.
- **Optimized software rasterizer** – The CPU does software rasterization using an optimized implementation. On an UltraSPARC CPU, some features, such as texturing rasterization, may be handled by the UltraSPARC CPU using software code employing the VIS instruction set.
- **Software rasterizer** – The CPU does software rasterization using a generic, unoptimized implementation. The generic software rasterizer is approximately one-sixth the speed of the optimized software rasterizer.

Performance

This chapter provides performance information that you can use to tune your application to make the best use of Sun hardware graphics accelerators. The first section defines two terms commonly used when discussing hardware and software performance. The second section provides general advice on how to optimize vertex processing performance for a variety of platforms. The third section describes the graphics hardware architecture. The subsequent sections provide specific techniques to ensure maximum performance on the Sun Elite3D, Creator, Creator3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators.

Acceleration Versus Optimization

When discussing performance, understanding how the hardware implementor, software implementor, and application programmer define and differentiate the terms *hardware acceleration* and *software optimization* is helpful.

- To the hardware designer, hardware accelerating OpenGL means implementing logic in the form of gates and data paths for OpenGL functions.
- To the OpenGL software implementor, accelerating OpenGL functions means writing software to use the graphics hardware features. In addition, the software implementor can *optimize* OpenGL features that are not accelerated in hardware by writing highly tuned code to make the performance of those features as efficient as possible.
- To the OpenGL application programmer, acceleration typically means the speed at which various combinations of geometry and OpenGL state render, with the goal generally being interactive performance.

With these definitions in mind, the next sections describe the OpenGL architecture and the implementation of this architecture in the Sun OpenGL for Solaris software.

Multiscreen Environment Performance

For a multiscreen Xinerama environment, OpenGL internally spawns a rendering thread for each screen, so that the rendering is done in parallel. Note that an application's OpenGL performance in multiscreen Xinerama may not be equal to its performance in single-screen non-Xinerama mode. For example:

- For true immediate mode (that is, rendering using `glVertex` calls), the OpenGL/Xinerama performance is approximately 32% to 50% of its non-Xinerama performance. This is caused by the extra overhead of copying GL commands to dispatch buffer, the extra layer of dispatch function calls and the thread synchronization cost of the multiple rendering threads.
- For applications using vertex arrays, the OpenGL/Xinerama performance is approximately 40% to 70% for `glDrawElements`, and 30% to 45% for `glDrawArrays`, of its non-Xinerama performance.
- For applications using display list mode, the OpenGL/Xinerama performance is approximately 70% to 99% of its non-Xinerama performance. And in some cases, the performance can exceed its non-Xinerama performance.
- If your application window resides entirely within one screen, OpenGL will revert back to single-screen rendering and deliver the full performance as in non-Xinerama mode.
- Note that since OpenGL spawns a separate rendering thread for each screen, it's important that you configure your system to balance the CPU and frame buffer needs. In general, we recommend allocating one processor per screen for your graphics computation.

In summary, for the best OpenGL performance in multiscreen Xinerama:

1. Use display list mode.
2. Allocate at least one CPU per screen for graphics rendering.

Multisampling

Multisampling is a technique for doing full screen antialiasing. It works by sampling each pixel multiple times and filters the samples before they are displayed on the screen. This removes jagged edges on 3D data and results in higher image quality.

Performance will be slower when multisampling is enabled due to the extra processing required to filter the samples and to generate each sample. The filtering operation normally happens when the buffers are swapped or when pixels are read from the frame buffer. Applications which have a high frame rate or do many frame buffer reads could experience a loss in performance. Also, applications which are fill rate limited could see a loss in performance due to the overhead of generating each sample during rasterization.

To enable multisampling on the XVR-1000, see the *Sun XVR-1000 Graphics Accelerator Installation and User's Guide* which can be found at <http://docs.sun.com>

To enable multisampling on the Sun Expert3D, Sun Expert3D-lite, or Sun XVR-500 graphics accelerators, see the FAQ on Sun's developer OpenGL web page at:

<http://www.sun.com/software/graphics/OpenGL>

View Frustum Culling on Display Lists

Sun OpenGL performs view frustum culling internally on display list data. Each block of geometric data in a display list is checked against the current view frustum and rendering is skipped for blocks of primitives that are completely outside of the view frustum. OpenGL display list view frustum culling can significantly improve rendering performance for certain cases.

To best use OpenGL's view frustum culling, the application should define display lists based on the spatial locality of the geometry. In other words, objects that are in close proximity to one another should be put into the same display list. This allows the view frustum culling implementation to generate a minimum-sized bounding box for the geometry. Rendering of the entire set can be skipped when the bounding box is outside the current view frustum.

Display list view frustum culling is ON by default and can be turned OFF by setting the environment variable:

```
% setenv SUN_OGL_VFCULL off
```

Display List Compression

On UltraSPARC platforms, Sun OpenGL can optionally compress display lists to reduce the total memory footprint used by the system. This can also provide a potential performance improvement when used with frame buffers capable of hardware geometry decompression.

Display list compression is OFF by default and can be turned ON by setting the variable `SUN_OGL_COMPRESSED_DISPLAY_LISTS` in the runtime environment. OpenGL will still honor this variable when running on frame buffers without hardware geometry decompression, but in those cases rendering is accomplished through a software decompression process that can significantly decrease rendering performance depending on processor speed and the size of the display list.

Display list compression can range from 6:1 to none at all over uncompressed display lists depending upon the content. In particular, primitives containing texture coordinates cannot be compressed. OpenGL will compress all display list commands that it can, leaving incompatible commands in their original form. The best compression ratios result from display lists containing vertices with consistent, supported data types within `glBegin/glEnd` pairs (see the discussion on Consistent Data Types below).

Other primitive data that cannot be compressed include those with indexed data, homogeneous coordinates, edge flags, coordinate evaluation, and `glMaterial` commands.

Display list build times are greater when they are being compressed. This again varies depending on the content of the display lists, but can range up to 10:1. Whether such a performance cost is significant depends upon how often display lists are built versus how often they are rendered, and should be weighed against the memory savings and rendering performance improvement or degradation experienced by individual applications.

Applications interested in using display list compression should simply set the `SUN_OGL_COMPRESSED_DISPLAY_LISTS` variable at runtime and examine the resulting performance to see if it is useful for their specific needs.

```
% setenv SUN_OGL_COMPRESSED_DISPLAY_LISTS on
```

General Tips on Vertex Processing

To achieve the best vertex processing performance on all Sun platforms, follow these guidelines:

1. Use vertex arrays or display list mode rather than immediate mode whenever rendering data repeatedly.
2. Use consistent patterns of data types between `glBegin(3gl)` and `glEnd(3gl)`. Consistent data types are described in “Consistent Data Types” on page 18.
3. If you must use immediate mode, try to include as many primitives of the same type as possible between one `glBegin` and the corresponding `glEnd`.
4. Define at most eight primitive states per vertex when they are defined between `glBegin` and `glEnd`.
5. If vertex array is used, try to stay in vertex array mode, rather than switching between vertex array and immediate mode.

These guidelines are discussed in the sections that follow.

Vertex Arrays

Vertex array commands provide the best performance for vertex processing of big primitives because they avoid the function call overhead of passing one vertex, color, and normal at a time. Instead of calling an OpenGL command for each vertex, you can pre-specify arrays of vertices, colors, and normals, and use them to define a primitive or set of primitives of the same type with a single command. Interleaved vertex arrays may enable even faster performance, since the application passes the data packed in a single array.

MultiDrawArrays

Sun OpenGL for Solaris contains the extension `glMultiDrawArraysSUN()`. This function allows multiple strips of primitives to be rendered with one call to OpenGL. Because of reduced function call and setup overhead, this function can provide significant speed improvement when an object contains many short strips. For some implementations of this function, there may be additional performance gains if the strips are contiguous in the vertex array. As with the standard `glDrawArrays()`, using interleaved vertex arrays gives even better performance.

Triangle List and Mesh Primitives

Sun OpenGL for Solaris contains extensions for rendering `TriangleList`, `TriangleMesh`, and `QuadMesh` primitives. A single `TriangleList` can be used to substitute a combination of triangles, triangle strips, and triangle fan; a single `TriangleMesh` can be used to substitute a set of triangle strips; and a single `QuadMesh` can be used to substitute a set of quadrilateral strips. These extended primitives can provide significant speed improvement over standard simple primitives because of reduced function call overhead.

See `glBegin`, `glDrawMeshArraysSUN`, and `glReplacementCodeSUN` man pages for more details.

Consistent Data Types

For the Sun OpenGL for Solaris implementation on all Sun platforms, vertex processing is optimized if the application provides consistent, supported data types within a `glBegin/glEnd` pair. Data types are consistent when the commands between one vertex call, such as `glVertex3fv`, and the next vertex call include identical patterns of data types in the identical order. In other words, consistent data is data for which the pattern is the same for each vertex, except when `glCallList` or `glEval*` is included. For example, the following set of commands is consistent because the primitive is defined by the repeating set of calls `glColor3fv(3gl); glVertex3fv(3gl)`.

```
glBegin(GL_LINES);
    glColor3fv(...);
    glVertex3fv(...);
    glColor3fv(...);
    glVertex3fv(...);
    glColor3fv(...);
    glVertex3fv(...);
glEnd();
```

As another example, the following set of commands is consistent since each vertex contains the same data – a color, normal, and vertex in repeating order.

```
glBegin(GL_LINES);
    glColor3f(...);
    glNormal3f(...);
    glVertex3f(...);
    glColor3f(...);
    glNormal3f(...);
    glVertex3f(...);
glEnd();
```

Note – The *f versions of the calls may be used interchangeably with the *fv versions.

Inconsistent data types do not follow a repeating, supported pattern. In the first example below, the data is inconsistent because the first vertex has a normal, but the second vertex doesn't. In the second example, the order is reversed in the second set of commands, although both vertices have a color and a normal.

```
glBegin(GL_LINES);
    glNormal3fv(...);
    glColor3fv(...);
    glVertex3fv(...);
    glColor3fv(...);
    glVertex3fv(...);
glEnd();
```

```
glBegin(GL_LINES);
    glColor3fv(...);
    glNormal3fv(...);
    glVertex3fv(...);
    glColor3fv(...);
    glVertex3fv(...);
glEnd();
```

For general information on the vertex data that can be specified between `glBegin(3gl)` and `glEnd(3gl)` calls, see the `glBegin(3gl)` reference page.

Low Batching

Sun OpenGL for Solaris performs best when given big primitives. If small primitives are sent to the library, the library will try to batch these primitives together, providing that the primitives are of the same primitive type, with the same consistent data pattern, and there are no attribute state changes outside the `glBegin` call.

For example, the following primitives will be batched together by the library.

```
glBegin(GL_TRIANGLES);
    glNormal3fv(...);
    glVertex3fv(...);
    glNormal3fv(...);
    glVertex3fv(...);
    glNormal3fv(...);
    glVertex3fv(...);
glEnd();

glBegin(GL_TRIANGLES);
    glNormal3fv(...);
    glVertex3fv(...);
    glNormal3fv(...);
    glVertex3fv(...);
    glNormal3fv(...);
    glVertex3fv(...);
glEnd();
```

The following example shows that the primitives are not batched together because the `glColor3fv` call outside the `glBegin` call breaks the batching of the primitives.

```
glBegin(GL_LINES);
    glVertex3fv(...);
    glVertex3fv(...);
glEnd();

glColorfv(...);
glBegin(GL_LINES);
    glVertex3fv(...);
    glVertex3fv(...);
glEnd();
```

Optimized Data Types

On any platform that uses the software pipeline for model coordinate rendering, your application will get better performance if it can pass vertex data in patterns for which the software pipeline has optimized code. Optimized data patterns are consistent data patterns that contain none of the following:

- `glEdgeFlag*()`
- `glMaterial*()`
- `glEvalCoord*()`
- both `glColor*()` and `glIndex*()`
- both `glTexCoord*()` and `glIndex*()`

Hardware Specific Acceleration

Graphics hardware architectures can be designed to meet varying constraints of cost and CPU performance. High-performance model coordinate (MC) devices, such as Sun Expert3D, Sun XVR-500 and Sun XVR-1000 graphics accelerators, typically implement vertex processing and transformations in hardware. A model coordinate device may perform lighting, coordinate transformations, clipping, and culling as well as rasterization and fragment processing in hardware, thereby providing very fast performance.

At a different performance level, rasterization devices typically use the host CPU to perform vertex processing and use the rasterization hardware to convert device coordinate geometry into pixel values. The Creator3D systems are examples of device coordinate (DC) devices. The graphics hardware architecture of the Creator3D graphics system is designed as follows:

- Primitive assembly and vertex processing are performed on the UltraSPARC™ CPU. Texturing operations are also performed on the CPU.
- Rasterization and fragment processing are performed in the Creator3D graphics hardware subsystem. The Creator3D graphics system accelerates rasterization of lines, points, and triangles, and also accelerates per-fragment operations such as the pixel ownership test, scissor test, depth buffer test, blending, logical operations, line anti-aliasing, line stippling, and polygon stippling.

The benefit of building custom hardware for graphics is that when operations are parallelized in hardware circuits, turning on features (like both Z-buffering and blending) has a very small performance cost. If a feature is provided in hardware, the hardware is usually designed to allow sustained throughput for that feature. Thus, you can make full use of features that have been implemented in hardware without experiencing performance degradation.

The benefit of putting graphics functions in software is that since the CPU is a required and shared computing resource, using it for graphics operations imposes no additional financial cost. The disadvantage is that each additional graphics operation requires CPU cycle time. When an application asks more of the CPU, the CPU may perform more slowly.

Sun XVR-1000 Graphics Accelerator Performance

The Sun XVR-1000 graphics accelerator is an UltraSPARC™ port architecture (UPA) bus high-resolution, high-performance graphics frame buffer that provides 30-bit color and 3D acceleration. The graphics board has full hardware support for 2D and 3D texture mapping, as well as multisampling and an S-video port.

OpenGL Acceleration on the Sun XVR-1000 Graphics Accelerator

The Sun XVR-1000 graphics accelerator provides a complete acceleration of the OpenGL API, including 2D and 3D texture-mapping and image processing. The Sun XVR-1000 graphics accelerator will accelerate the entire 3D OpenGL graphics pipeline in hardware, including all geometry operations, triangle setup, texturing and pixel operations.

In addition, Sun XVR-1000 graphics hardware acceleration is used for matrix transformations, perspective, and viewport transformations, area fills, block moves, puts and gets, 2D/3D and antialiased points vectors and polygons, alpha operations, window clipping, fog, stencil, depth buffering, texture mapping, Pbuffers, and accumulation buffers.

The Sun XVR-1000 graphics accelerator supports both 8-bit and 24-bit OpenGL visuals.

The General Tips on vertex processing, noted earlier in this chapter also apply for the Sun XVR-1000 graphics accelerator. Applications should strive to use Display List mode, rather than Immediate Mode. Use of consistent Data Types is preferred, as is high batching of primitive data. Vertex Arrays and MultiDraw Arrays will provide for the best immediate mode rendering performance.

The following are attributes that cause slower rendering by the Sun XVR-1000 graphics accelerator:

- Rendering polygons with `GL_POLYGON_SMOOTH` enabled
- Rendering lines with `GL_LINE_SMOOTH` enabled, or points with `GL_POINT_SMOOTH` on an indexed color raster
- `glEnable` of `GL_LIGHTING` on an indexed color raster
- Setting `glRenderMode` to anything other than `GL_RENDER`
- Setting `glPointSize` to larger than 13.0
- Rendering a line primitive (`GL_LINES`, `GL_LINE_STRIP`, `GL_LINE_LOOP`), or a polygon with `glPolygonMode` set to `GL_LINE` or `GL_POINT` with texture mapping enabled
- Multi-texturing with alpha blending enabled
- Multi-texturing with `GL_COMBINE_EXT` set to `GL_ADD_SIGNED_EXT` or `scale != 1.0` (applies to both `rgb` and `alpha`)
- Setting `glPolygonOffset` `scale` to a value outside the range `(-8 ... 8)`

In general, you can obtain the best performance from batching vertices together in a single primitive, or a large group of like primitives. Display list rendering is generally faster than vertex arrays, which is faster than immediate mode.

Enabling multisampling will yield better rendering quality at the expense of performance.

Sun Expert3D, Expert3D-Lite, and Sun XVR-500 Graphics Accelerator Performance

The Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators are high resolution, high performance PCI graphics frame buffers that provide hardware texture mapping. It is designed to fill the needs of Suns customers who use texture mapping extensively, notably in defense, geophysical and digital content creation applications.

OpenGL Acceleration on the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 Graphics Accelerators

The Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators provide a complete acceleration of the OpenGL API, including 2D and 3D texture-mapping and image processing. The Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators will accelerate the entire 3D OpenGL graphics pipeline in hardware, including all geometry operations, triangle setup, texturing and pixel operations.

In addition, Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerator hardware acceleration is used for matrix transformations, perspective, and viewport transformations, area fills, block moves, puts and gets, 2D/3D and antialiased points vectors and polygons, alpha operations, window clipping, fog, stencil, depth buffering, texture mapping, Pbuffers, and accumulation buffers.

The Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators supports both 8-bit and 24-bit OpenGL visuals.

The General Tips on vertex processing, noted earlier in this chapter, also apply to the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators. Applications should strive to use Display List mode, rather than Immediate Mode. Use of consistent Data Types is preferred, as is high batching of primitive data. Vertex Arrays and MultiDraw Arrays will provide for the best immediate mode rendering performance.

Immediate-Mode OpenGL graphics operations are likely to scale with CPU performance on the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators.

When clearing the OpenGL color and depth buffers, there is a performance advantage to clearing both buffers from the same `glClear()` call. The performance difference on the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators may be an order of magnitude.

Virtual texture maps are supported. The Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators do not limit the number of textures that are defined. As much texture is supported as can be defined in Virtual Memory on Sun Expert3D, Expert3D-Lite or Sun XVR-500 graphics accelerator systems.

The largest single texture map supported by the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators is 16 megabytes. All texture maps should be this size or smaller. The smallest texel format uses 16 bits per texel.

The list of unaccelerated features for the Sun Expert3D, Expert3D-Lite, and Sun XVR-500 graphics accelerators is small:

- Fill rates limit the performance of large, non-textured triangles.

- Wide antialiased points (Enabled `GL_POINT_SMOOTH` and `glPointSize > 3.0`)
- Wide antialiased lines (Enabled `GL_LINE_SMOOTH` and `glLineWidth > 1.0`)

Sun Elite3D Graphics Accelerator Performance

The Sun Elite3D graphics accelerator performance is affected by attributes that force the slower rendering layers to be used.

Attributes

Primitive-attribute settings affect performance; therefore, you will get a better level of performance if you can avoid setting the attributes listed below. In general, the more host processing needed, the slower the resulting rendering, so it is advantageous for an application to avoid those attributes that force the slower rendering layers to be used.

Attributes that Force the Use of the CPU for Vertex Processing

The Elite3D graphics accelerator performs most all vertex processing operations. If an attribute is set and it cannot be handled directly by the hardware, the host CPU must handle it, as well as the other model coordinate functions. Rasterization of the resulting device coordinate geometry is still performed by the hardware, although at a reduced speed. The drop off in performance may be anywhere from 20% to 50%.

- Using vertices with homogeneous coordinates – `glVertex4()`
- Calling (`glMaterial()`) between `glBegin()` and `glEnd()`.
- Enabling indexed color linear fog – `glIndex()` and `GL_FOG_MODE, GL_LINEAR`

Environment Variables Affecting Read Performance

- `unsetenv SUN_OGL_ABGR_READPIX_NOCONFORM` (default)

The alpha value read back from the frame buffer during `glReadPixels` with the `GL_ABGR_EXT` format is always 1.0. This is conformant but slower than the following variable:

- `setenv SUN_OGL_ABGR_READPIX_NOCONFORM`

The alpha value read back from the frame buffer during `glReadPixels` with the `GL_ABGR_EXT` format is undefined. This up to 30% faster than the conformant version. For Elite3D, the alpha value is not stored in the frame buffer anyway. Consequently, if the application does not use the alpha value, this version is a significantly faster way to read pixels back from the frame buffer.

Attributes that Force the Use of the Generic Software Rasterizer

Elite3D hardware rasterizers cannot be used under some attribute setting combinations. Consequently, software rasterization must be used to scan convert primitives. This is the slowest data path, with a noticeable performance drop from the DC layer rendering mentioned above. If your application requires any of the following attributes for performance critical functionality, you may want to determine whether this performance is acceptable. If not, you can evaluate whether the use of these attributes is advisable.

1. Rasterization attributes

- Enabling immediate mode texture mapping – To take advantage of Elite3D’s hardware texture mapping, use display list rendering. In immediate mode, a considerable slowdown in rendering results. Except when `glHint(GL_SURFACE_SIZE_HINT_SUNX, GL_LARGE_SUNX)` is set, hardware texture mapping acceleration will be turned on. The performance of certain applications with large textured surfaces will improve.
- Environment mapping – `glTexGen(Coord, pname, GL_EYE_LINEAR` or `GL_SPHERE_MAP)`
- 3D texture mapping – `GL_TEXTURE_3D`. Note that 3D texturing refers to using a 3D cube of textures, not 1D or 2D texturing of 3D geometry, which can be accelerated by Elite3D.
- Enabling surface antialiasing (`GL_POLYGON_SMOOTH`) for any of the triangle, quad, or polygon primitives rendered with `glPolygonMode == GL_FILL`
- Enabling indexed color antialiasing
- Stippled lines (`GL_LINE_STIPPLE`) where the line stipple scale factor is larger than 15
- `glColorTable(GL_TEXTURE_COLOR_TABLE_SGI(), ...size bigger than 256 entries` – Elite3D hardware texture mapping can only support up to 256 table entries.
- BGRA pixel formats
- Packed pixel formats
- Separate specular color
- Texture level of detail control

2. Texturing attributes

- Texture-mapped lines and dots – Elite3D hardware texture mapping is only supported for filled primitives.

3. Fragment attributes

- Blending or alpha test requiring destination alpha. There is no alpha buffer in an Elite3D accelerator; therefore, any blending operation or alpha test that requires destination alpha is rasterized by software.
- Indexed color exponential fog. Elite3D supports fog for RGB colors only. Exponential fog must be applied on a per-pixel basis; therefore, rendering must be done at the software raster level.

Pixel Operations

For Elite3D, pixel operations are handled in the same way as the Creator Graphics accelerator. The optimized attributes are the same for both accelerators. For information on which attribute settings optimize performance on the Creator Graphics, see “Pixel Operations” on page 37.

Consistent Data

In most cases, other than texture mapping, the Elite3D can handle any data, inconsistent or consistent, without calling the software pipeline for vertex processing. If the vertex data is consistent, the Elite3D can handle it more efficiently, resulting in better performance.

If the data is in a display list, is consistent, and doesn't violate the guidelines listed on “Attributes” on page 25, it can be block copied to the hardware. For most applications, this is the fastest way to get data to the Elite3D hardware.

If the data is inconsistent, or if the application uses immediate mode rendering the data cannot be block copied. Instead, it is written to the hardware one word at a time. This is slower than block copying, but still much faster than the performance on a DC device or performance via the software pipeline. The OpenGL for Solaris implementation for the Elite3D calls the software pipeline as a last resort, for a small set of attributes or types of geometry, for example using homogeneous coordinates or calling `glMaterial` inside a primitive. For examples of consistent data patterns, see “Consistent Data Types” on page 18.

Sun Creator and Creator3D Graphics Accelerator Performance

The Ultra Creator and Creator 3D Graphics systems accelerate rasterization of lines, points, and triangles as well as most per-fragment operations. Vertex processing and texturing operations are performed on the UltraSPARC CPU. The Sun OpenGL for Solaris implementation for the Creator and Creator3D frame buffers uses all features of the Creator graphics subsystem.

Rasterization and fragment processing is handled in one of three ways:

- Creator3D hardware rasterizer – Handles lines, points, and triangles, and does simple fragment processing.
- Optimized software rasterizer – UltraSPARC VIS (Visual Instruction Set) handles many texturing functions and pixel operations.
- Generic software rasterizer – Performs rasterization for all features not handled by the hardware or by the VIS software.

The following sections provide specific information on attribute use and pixel operations on these platforms.

Attributes Affecting Creator3D Performance

Primitive-attribute settings affect performance; therefore, you will get a better level of performance if you can avoid setting the attributes listed below. In some cases, the listed attributes simply increase the amount of processing in the hardware or optimized software data paths. In other cases, setting these attributes forces the use of the software rasterizer, resulting in slow performance.

Attributes that Increase Vertex Processing Overhead

Attributes that result in more vertex processing overhead include:

- Enabling lighting.
- Turning on user specified clip planes (`GL_CLIP_PLANE[i]`).
- Enabling color material (`GL_COLOR_MATERIAL`).
- Enabling non-linear fog (`glFog(GL_FOG_MODE, GL_EXP{2})`). An exception to this is using `RGBA` mode on Creator3D Series 2.
- Enabling `GL_NORMALIZE`.

- Turning on polygon offset. However, polygon offset is optimized for the case when the factor parameter of the `glPolygonOffset` call is set to 0.0. Users may have to adjust the units parameter accordingly to avoid stitching for this case.

Primitive Types and Vertex Data Patterns that Increase Vertex Processing Overhead

Types and patterns that result in more vertex processing overhead are:

- Using a vertex data pattern for `GL_POINTS`, `GL_LINES`, `GL_LINE_STRIP`, and `GL_LINE_LOOP`, *other than* one of the following repeating patterns. These are the patterns that are maximally accelerated.

```
V3F:
glVertex3f(...);
...
C3F_V3F:
glColor3f(...);
glVertex3f(...);
...
C4F_V3F:
glColor4f(...);
glVertex3f(...);
...
V2F:
glVertex2f(...);
...
C3F_V2F:
glColor3f(...);
glVertex2f(...);
...
C4F_V2F:
glColor4f(...);
glVertex2f(...);
...
```

Note – All vertex data patterns, other than one of the above repeating patterns, take more memory.

- Using `glDrawElements` in immediate mode.

Attributes That Increase Hardware Rasterization Overhead

Attributes that result in slower hardware rasterization are:

- Enabling line antialiasing (`GL_LINE_SMOOTH`)
- Enabling point antialiasing (`GL_POINT_SMOOTH`)

Environment Variables Affecting Read Performance

- `unsetenv SUN_OGL_ABGR_READPIX_NOCONFORM` (default)

The alpha value read back from the frame buffer during `glReadPixels` with the `GL_ABGR_EXT` format is always 1.0. This is conformant but slower than the following variable.

- `setenv SUN_OGL_ABGR_READPIX_NOCONFORM`

The alpha value read back from the frame buffer during `glReadPixels` with the `GL_ABGR_EXT` format is undefined. This is up to 30% faster than the conformant version. For Creator, the alpha value is not stored in the frame buffer anyway. Consequently, if the application does not use the alpha value, this version is a significantly faster way to read pixels back from the frame buffer.

Attributes That Force the Use of the Software Rasterizer

Setting the following attributes forces the use of the software rasterizer. This is the slowest data path. If your application requires any of the following attributes for performance critical functionality, you may want to determine whether this performance is acceptable. If not, you can evaluate whether the use of these attributes is advisable.

1. Rasterization attributes

- In Indexed color mode, enabling line anti-aliasing (`GL_LINE_SMOOTH`) or point anti-aliasing (`GL_POINT_SMOOTH`)
- Enabling polygon anti-aliasing (`GL_POLYGON_SMOOTH`)
- Stippled lines (`GL_LINE_STIPPLE`) where the line stipple scale factor is larger than 15
- Antialiased line width not equal to 1.0
- Non-antialiased (“jaggy”) points with `glPointSize(3gl)` greater than 1.0

Note – The only anti-aliased point size supported by Creator3D and Creator is 1.0. `glPointSize` is ignored for anti-aliased points. Although the nominal antialiased point size is 1.0, the actual visible size is approximately 1.5.

2. Fragment Attributes

- Blending (`GL_BLEND`) forces the use of the software rasterizer unless both the source and destination blend functions are in the following set of blend functions supported by the hardware: `GL_ZERO`, `GL_ONE`, `GL_SRC_ALPHA`, `GL_ONE_MINUS_SRC_ALPHA`
- Enabling the stencil test (`GL_STENCIL_TEST`) on Creator3D or Creator3D Series 2. (Enabling the stencil test does not force the use of the software rasterizer on Creator3D Series 3 because it supports hardware stencilling).

On the UltraSPARC platform, a VIS optimized software rasterizer is used for smooth-shaded non-textured stenciled triangles whenever the `glStencilOp` parameter *fail* is anything other than `GL_INCR` or `GL_DECR` and the depth test does not affect the stencil buffer. (This is the case when depth test is disabled or the `glStencilOp` parameters *zfail* and *zpass* are identical).

- Enabling any type of fog in Indexed color mode

FIGURE 3-1 shows the data path for hardware rasterization on the Creator3D system.

3. Texturing Attributes

- Color Table—When the `GL_TEXTURE_COLOR_TABLE_SGI` extension is used, the only `glTexEnv` texture base internal formats that are accelerated are `GL_LUMINANCE`, `GL_LUMINANCE_ALPHA` and `GL_INTENSITY`.
- The texture environment mode `glTexEnv GL_TEXTURE_ENV_MODE` of `GL_BLEND` is not accelerated when it is used with the `GL_TEXTURE_COLOR_TABLE_SGI` extension.
- Fog—On Creator3D, only linear fog is accelerated. On Creator3D Series 2, all types of RGBA fog are accelerated.

Attributes That Vary Optimized Texturing Speed

Texturing makes extensive use of VIS on UltraSparc platforms and allows for large textures. Texturing speed naturally increases with faster CPUs (a 300 Mhz UltraSPARC CPU is 1.6 times faster than a 167 Mhz CPU). Though texturing fill rates are slower on a host CPU than on dedicated hardware, the system costs are lower.

Stencil and some fragment blending cases are slow. The rest are fast (done by Creator 3D hardware).

Some texturing attributes are handled by generic code and result in the slowest texturing speed when the `GL_TEXTURE_COLOR_TABLE_SGI` extension is used with texture environment color blending or base internal formats of `GL_ALPHA`, `GL_RGB`, or `GL_RGBA`.

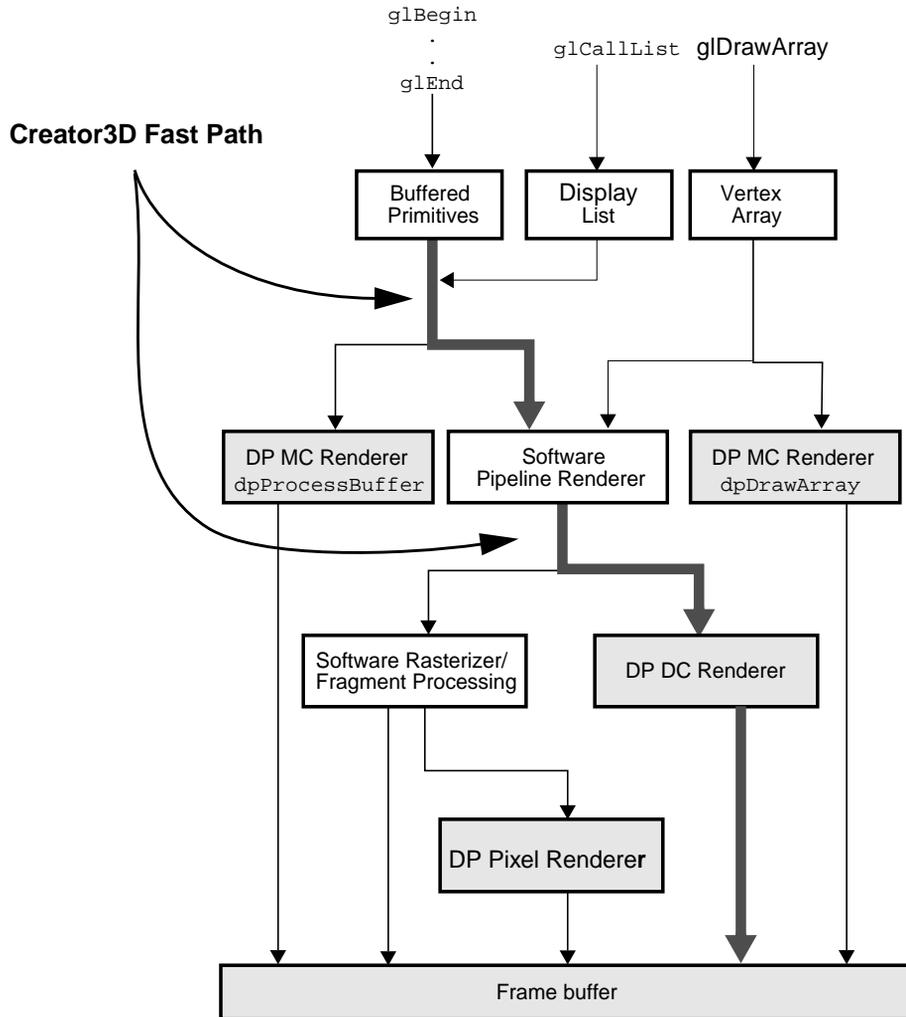


FIGURE 3-1 Hardware Rasterizer Path for Creator3D

Texturing attributes with the most impact on speed are:

- Minification filter
- Texture Coordinate Interior/Exterior Classification (per triangle)
- All wrap modes set to `GL_REPEAT`
- Texture Color Lookup Table

The VIS optimized software rasterizer will vary in texturing speed based on the texturing attributes specified. The factors affecting texturing speed are listed below. Note that this is *variance within the optimized path*, not the difference between the optimized and generic paths.

- **Projection Type**—The type of projection matrix. Orthographic is faster than perspective.
- **Wrap Mode**—Best speed is when all dimensions (`GL_TEXTURE_WRAP_x`) are set to `GL_REPEAT`. If all the texture wrap modes are `GL_REPEAT`, this case is specially optimized. If any of the texture wrap modes are `GL_CLAMP`, the standard texture wrap routine is used, but it is slower than the special case.
- **Dimension**—In general, 2D texturing is faster than 3D texturing, since there is one less texture coordinate to deal with. However, this does not mean it is better to use many 2D textures to approximate 3D texturing since the texture load time (see next section) may significantly increase the overhead.
- **Minfilter**—The fastest `GL_TEXTURE_MIN_FILTER` parameter is `GL_NEAREST`, which is approximately 4x the speed of `GL_LINEAR`. The approximate relative speed in decreasing order is: `GL_NEAREST`, `GL_NEAREST_MIPMAP_NEAREST`, `GL_NEAREST_MIPMAP_LINEAR`, `GL_LINEAR`, `GL_LINEAR_MIPMAP_NEAREST`, and `GL_LINEAR_MIPMAP_LINEAR`.
- **Magfilter**—For `GL_TEXTURE_MAG_FILTER`, the same speed ratio of 4x applies to `GL_NEAREST` vs. `GL_LINEAR`. Note, however, that `GL_TEXTURE_MAG_FILTER` is ignored when `GL_TEXTURE_MIN_FILTER` is set to `GL_NEAREST` or `GL_LINEAR`. This can be overridden with a shell environment variable but this will slow down texturing speed for `GL_NEAREST` and `GL_LINEAR`, since they now have to perform level-of-detail calculations to determine when to use `GL_TEXTURE_MAG_FILTER`. The shell environment variable that forces this slower behavior is: `setenv SUN_OGL_MAGFILTER "conformant"`
- **Texture Coordinate Classification**—If all texture coordinates of a triangle/quad/polygon are at LEAST 1/2 texel inside away from the texture map edge, the primitive is considered interior and are render faster than those whose texture coordinates touch or cross the texture map's edges. If any vertex touches or crosses the texture map edge, the primitive is considered exterior. If a primitive is interior, the texture edge related attributes such as wrap modes and texture border no longer affect the texturing speed.
- **Env Mode**—The fastest `glTexEnv()` `GL_TEXTURE_ENV_MODE` is `GL_REPLACE`, followed closely by `GL_MODULATE`. `GL_DECAL` is the same speed as `GL_REPLACE`.
- **Color Table**—The use of the extension `GL_TEXTURE_COLOR_TABLE_SGI` will reduce texturing speed.
- **Texture Color Lookup Table**—Using this table causes significant slowdown of texturing speed. Only cases of one or two channel lookups are optimized - `GL_LUMINANCE`, `GL_INTENSITY`, `GL_LUMINANCE_ALPHA`. Three or four channel lookups (`GL_RGB`, `GL_RGBA`) go to a generic code routine that is slower than the special case.

Texture Memory Usage

The OpenGL library uses packed textures when possible to minimize texture space. The OpenGL library internally makes a copy of the user specified texture. Internal formats `GL_LUMINANCE`, `GL_ALPHA`, and `GL_INTENSITY` are stored as packed 8-bits per texel textures. The internal format `GL_LUMINANCE_ALPHA` is stored as 16-bit per texel textures. Note that this copy can be avoided by using constant data extension (see `GL_CONSTANT_DATA_SUNX` man page for more details). On the Elite3D, if MIP mapping is used, an additional copy is made internally (resulting in a total of two internal copies).

Attributes That Vary Texture Load Time

The time to load the texture image into a texture object or a display list will vary depending on the pixel store and pixel transfer attributes specified when the texture is specified.

FIGURE 3-2 shows the texture load processing flow.

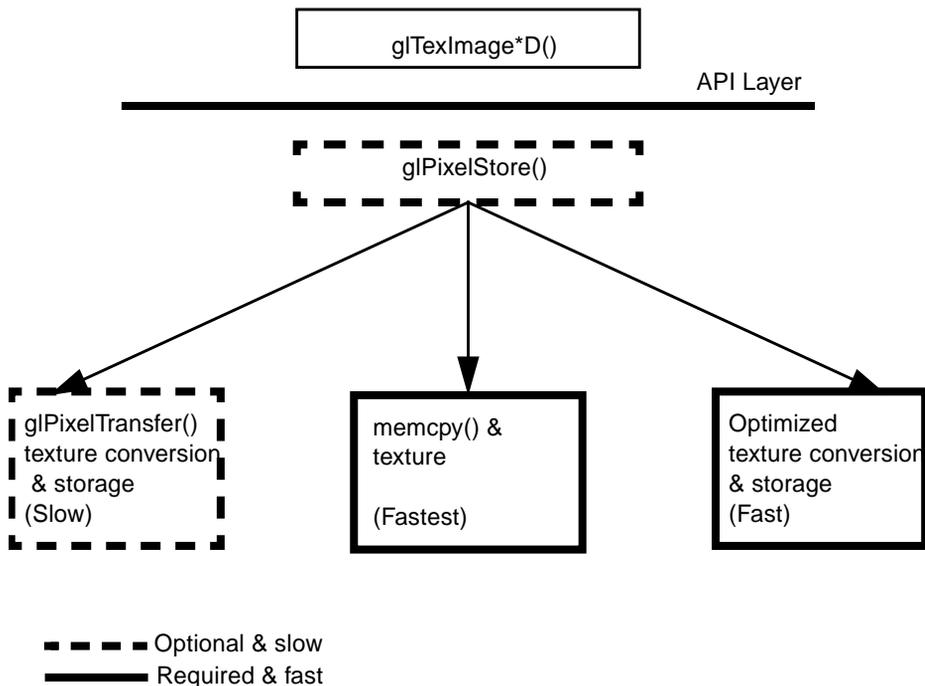


FIGURE 3-2 Text Load Processing Flow

The following recommendations should be followed where possible to reduce texture load time:

- Use Constant Data Extension – This extension eliminates the texture load time by using the application-provided texture data pointer internally. To use this extension, the application should use one of the format/internal format combinations listed in TABLE 3-1 and the texture data type should be `GL_UNSIGNED_BYTE`. Using this extension restricts when or how the texture data pointer can be deleted or changed. For a complete description refer to `GL_CONSTANT_DATA_SUNX` man page.
- Use texture objects where possible.
- If multiple textures are being used, put the textures in texture objects and use `glBindTexture` to switch among the textures. This ensures that the internal copy of texture is evaluated only once.
- Internally, 1D, 2D, and 3D textures use packed representation to minimize memory usage.
- If application uses pixel store (but not pixel transfer) and therefore constant data extension cannot be used, for textures using data type `GL_UNSIGNED_BYTE`, the following format/base internal format combination give the best loading performance.

TABLE 3-1 3D Optimized Cases

Format	Base Internal Format
<code>GL_RED</code>	<code>GL_INTENSITY</code>
<code>GL_RED</code>	<code>GL_LUMINANCE</code>
<code>GL_ALPHA</code>	<code>GL_ALPHA</code>
<code>GL_LUMINANCE</code>	<code>GL_INTENSITY</code>
<code>GL_LUMINANCE</code>	<code>GL_LUMINANCE</code>
<code>GL_ABGR_EXT</code>	<code>GL_RGBA</code>

Attributes Affecting Software Rendering Performance

This section applies when pure software rendering is being used. This happens on the single-buffered Creator platform when `glDrawBuffer(3gl)` is set to `GL_BACK` or `GL_FRONT_AND_BACK`. The data presented here is also valid for the PGX platforms. Note that for non-Ultra machines, VIS rasterization is replaced by an optimized software rasterizer.

Attributes That Increase Vertex Processing Overhead

Attributes that result in more vertex processing overhead are:

- Enabling lighting.
- Turning on user specified clip planes (`GL_CLIP_PLANE[i]`).
- Enabling color material (`GL_COLOR_MATERIAL`).
- Enabling non-linear fog (`glFog (GL_FOG_MODE, GL_EXP{2})`). An exception to this is using RGBA mode on Creator3D Series 2.
- Enabling `GL_NORMALIZE`.
- Turning on polygon offset. However, polygon offset is optimized when the factor parameter of the `glPolygonOffset` call is set to 0.0. Users may have to adjust the units parameter accordingly to avoid stitching for this case.

Attributes That Force the Use of the Generic Software Rasterizer

Setting the following attributes forces the use of the generic software rasterizer. This is the slowest data path. If your application requires any of the following attributes for performance critical functionality, you may want to determine whether this performance is acceptable. If not, you can evaluate whether the use of these attributes is advisable.

1. Texturing Attributes

- All three-dimensional texturing attributes result in the use of the generic software rasterizer.
- Two-dimensional texture mapping (`GL_TEXTURE_2D`) in the following cases:
 - i. Texture environment mode `glTexEnv GL_TEXTURE_ENV_MODE` is set to `GL_BLEND`.
 - ii. `glTexEnv` texture base internal format is `GL_ALPHA`.

- iii. Texturing of points is handled by the generic software.
- iv. Fog is enabled.
- v. Any use of the SGI Texture Color Table (`GL_SGI_texture_color_table`) extension.

2. Fragment Attributes

- Enabling any type of fog in Indexed color mode.
- Enabling blending (`glBlendFunc(3gl)`) except when the source blending factor is `GL_SRC_ALPHA` and the destination blending factor is `GL_ONE_MINUS_SRC_ALPHA`. This case is optimized.
- Enabling logical operations.
- Enabling depth test `glEnable(GL_DEPTH_TEST)` forces the use of the optimized software rasterizer. If depth test is enabled, and if `glDepthFunc(3gl)` is on, enabling any Z comparison other than `GL_LESS` or `GL_LEQUAL` forces the use of the generic software rasterizer.
- Enabling alpha test.
- Setting `glDrawBuffer(3gl)` to `GL_BACK` or `GL_FRONT_AND_BACK`, or setting `glReadBuffer(3gl)` to `GL_BACK`.

Pixel Operations

Under optimal conditions, the commands `glDrawPixels(3gl)`, `glReadPixels(3gl)`, and `glCopyPixels(3gl)` are optimized on the Creator and Creator3D systems using the VIS instruction set on the UltraSPARC CPU. Bitmap operations using the command `glBitmap(3gl)` are accelerated in the Creator3D font registers. However, some attribute settings result in the use of the software rasterizer for pixel operations.

FIGURE 3-3 shows the rasterization and fragment processing architecture for `glDrawPixels(3gl)`. The figure shows the optimized and unoptimized paths for pixel rendering. Your application will experience performance degradation for each functional box that it needs. In addition, performance degradation will occur if the data type is not unsigned byte; in this case, the data must be reformatted internally.

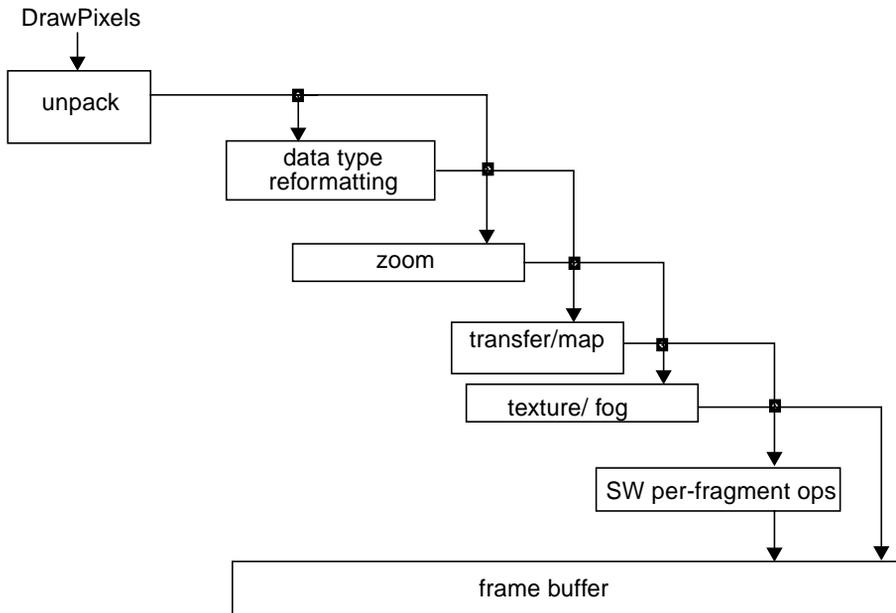


FIGURE 3-3 Sun OpenGL for Solaris Architecture for Drawing Pixels

Conditions That Result in VIS Optimization on UltraSPARC Systems

In general, for `DrawPixels`, `CopyPixels`, and `Bitmap`, the use of texture mapping or nonlinear fog (except in RGBA mode on Creator3D Series 2) will force the use of the generic software rasterizer, resulting in slow performance. In addition, if the hardware does not support the per-fragment operations that the application has enabled, the generic software rasterizer is used. See the OpenGL documentation or the “OpenGL Machine” diagram for a list of per-fragment operations.

For the Creator3D system, if the following conditions are true, pixel operations are optimized. If these conditions are not true, the generic software rasterizer is used.

`glDrawPixels` Command

- Pixel format is `GL_RGBA`, `GL_RGB`, `GL_ABGR_EXT`, `GL_RED`, `GL_GREEN`, `GL_BLUE`, `GL_LUMINANCE`, and `GL_LUMINANCE_ALPHA`.
- Data type is `GL_UNSIGNED_BYTE`. (For `GL_LUMINANCE` the data type can also be `GL_SHORT`).
- For the format of `GL_DEPTH_COMPONENT`, the types `GL_INT`, `GL_UNSIGNED_INT`, and `GL_FLOAT` are optimized for the case with no pixel transfer.

- Texturing is disabled.
- Pixel unpacking is unnecessary.
- For the formats listed in the first line, the pixel transfer operations for scale/bias, pixel map, SGI color table, convolution, post convolution color table, histogram, and minmax may be enabled.
- Pixel Zoom may be done if its zoom factors are other than the default values.
- Pixel transform may be done if its current matrix is other than the identity matrix.

`glReadPixels` *Command*

- Pixel format is GL_RGBA, GL_RGB, GL_ABGR_EXT, GL_RED, GL_GREEN, GL_BLUE, GL_LUMINANCE, and GL_LUMINANCE_ALPHA.
- Data type is GL_UNSIGNED_BYTE.
- For the format of GL_DEPTH_COMPONENT, the types GL_INT, GL_UNSIGNED_INT, and GL_FLOAT are optimized for the case with no pixel transfer.
- Pixel packing is unnecessary.
- For the formats listed in the first line, the pixel transfer operations for scale/bias, pixel map, SGI color table, convolution, SGI post convolution color table, histogram, and minmax may be enabled.

`glCopyPixels` *Command*

- Pixel type is GL_COLOR.
- Texturing is disabled.
- Pixel zooming is in the default state.
- The pixel transfer operations for scale/bias, pixel map, SGI color table, convolution, SGI post convolution color table, histogram, and minmax may be enabled.

`glBitmap` *Command*

- Texturing is not enabled.
- Blending is not enabled.

Conditions That Result in VIS Optimization

For the UltraSPARC the Creator and non-Creator frame buffers, if the following conditions are true, pixel operations are optimized. If these conditions are not true, the generic software rasterizer is used.

`glDrawPixels` *Command*

- For `GL_LUMINANCE` with data types `GL_UNSIGNED_BYTE` and `GL_SHORT`, there are special VIS optimized routines for:
 - drawing directly to the frame buffer (or pbuffer).
 - performing pixel transfer (i.e., scale/bias, pixel map, SGI color table, convolution, SGI post convolution color table, histogram, and minmax) then displaying directly to the frame buffer (or pbuffer).
 - performing the pixel transform extension, then drawing directly to the frame buffer (or pbuffer).
 - performing pixel transfer followed by the pixel transform extension, then finally drawing directly to the frame buffer (or pbuffer).
- Pixel format is `GL_RGBA`, `GL_RGB` or `GL_ABGR_EXT`.
- Data type is `GL_UNSIGNED_BYTE`.
- Texturing is disabled.
- Pixel unpacking is unnecessary.
- If depth test is enabled, and if `glDepthFunc(3gl)` is on, enabling any Z comparison other than `GL_LESS` or `GL_LEQUAL`.

`glReadPixels` *Command*

- For `GL_RED` with the data type `GL_UNSIGNED_BYTE`, there is one special VIS optimized routine for extracting the red channel from an ABGR frame buffer or pbuffer.
- If `glReadPixels` format is `GL_RGBA`, `GL_RGB`, or `GL_ABGR_EXT`, and the pixel type is `GL_UNSIGNED_BYTE`, `glReadPixels` is optimized.
- If `glReadPixels` format is `GL_DEPTH_COMPONENT`, these pixel types are optimized: `GL_INT`, `GL_UNSIGNED_INT`, `GL_FLOAT`.
- Pixel packing is unnecessary.

`glCopyPixels` *Command*

- Pixel type is `GL_COLOR`.
- Texturing is disabled.
- Enabling any Z comparison other than `GL_LESS` or `GL_LEQUAL`.

`glBitmap` *Command*

- Texturing is disabled.
- If depth test is enabled, and if `glDepthFunc` is on, enabling any Z comparison other than `GL_LESS` or `GL_LEQUAL`.

Pixel Transfer Pipeline (ARB) Imaging Extensions and the Pixel Transform

The Pixel Transfer Pipeline consists of a small set of image processing functions that operate on most rectangular imagery with OpenGL. These operations are performed whenever Pixel Transfer operations can occur within OpenGL (that is, `glDrawPixels`, `glReadPixels`, `glCopyPixels`, `glTexImage2D`, `glTexImage3D`, and so on).

This pipeline has been fine tuned for maximum performance on `GL_LUMINANCE` formatted data for the data types `GL_UNSIGNED_BYTE` and `GL_SHORT`. Other formats have been accelerated as well; however, `GL_LUMINANCE` gains the most in performance with this Implementation of the Pipeline.

This pipeline has been accelerated using the Visual Instruction Set, which is only available on those systems with the UltraSPARC processor. The Pixel Transfer Pipeline with VIS acceleration is not supported on Non-UltraSPARC processors; however, the original Pixel Transfer Functionality is still there, minus the new imaging extensions.

Implementation

FIGURE 3-4 shows the functions and the order of execution (from top to bottom) of these functions in the Pixel Transfer Pipeline.

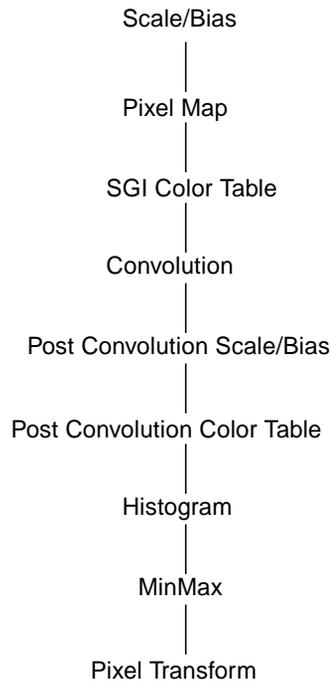


FIGURE 3-4 Pixel Transfer Pipeline Functions and Order of Execution

All functions in the pipeline have been accelerated using VIS whenever possible. The imaging extensions within this pipeline are convolution, post convolution scale/bias, post convolution color table, histogram, minmax, and pixel transform. The last one, pixel transform, is not really part of the pixel transfer pipeline, but is instead considered part of the pixel rasterizer. Also, pixel transform is only executed in the `glDrawPixels` interface.

Another optimization that is worth noting here is that direct output to the display, via the `glDrawPixels` interface, or into a pbuffer has been optimized for `GL_LUMINANCE` format with `GL_UNSIGNED_BYTE` and `GL_SHORT` data types. For `GL_UNSIGNED_BYTE`, while the frame buffer is in TrueColor mode (rgb mode), the luminance pixels are expanded to XBGR format and then written directly to the frame buffer memory using VIS for optimal throughput. For `GL_LUMINANCE`, `GL_SHORT` data, the conversion of `GL_SHORT` data to `GL_UNSIGNED_BYTE` and then expansion to XBGR for direct display has been optimized for maximum throughput using VIS.

When the input format is `GL_LUMINANCE` and the input data type is `GL_SHORT` the Pixel Transfer Pipeline has been made so that it will process the data from the beginning to end of the pipe as `GL_SHORT` data. This maintains the accuracy and integrity of the data from one stage of the pipeline to the next. Only just before rendering into the frame buffer or pbuffer does the data get scaled down and clamped to `[0, 255]`.

In this pipeline none or all of these processing blocks can be enabled. Any time the Pixel Transfer Pipeline is used, there is only one pass through the pipe, and the order of execution does not change from that represented in FIGURE 3-4.

How To Use the Pixel Transfer Pipeline and Pixel Transform

For the most part, OpenGL operates on RGBA colors. Therefore, to be specification compliant in OpenGL, if a user of OpenGL wants to do pixel transfer operations on `GL_LUMINANCE` data, that data should first be expanded to `GL_RGBA` format, (or `GL_ABGR_EXT` format) before doing any processing. However, depending on the OpenGL pixel transfer state parameters, it may not be necessary to expand the image data before processing in the pixel transfer pipeline. That is, if we expand the data from `GL_LUMINANCE` to `GL_RGBA` first, process the image as four-banded data in the Pixel Transfer Pipeline, and then display, or if we process the `GL_LUMINANCE` data as a single banded image in the Pixel Transfer Pipeline, then expand the data at the end of the pipeline, then display the data; if the result would be the same using either of the two paths, it makes sense to use the faster path, which, in this case, would be the latter path.

This takes about 1/4th the time, (or less) to do the correct desired operation. The Pixel Transfer Pipeline evaluates the various states of the pixel transfer functions and determines if it needs to do format expansion, before, during, or after processing, but expansion always occurs, if needed, just before rendering to the frame buffer or pbuffer.

The only case were format expansion can occur inside the Pixel Transfer Pipeline is within the “pixel map” block. If you want optimal throughput for `GL_LUMINANCE` data, do not use pixel map, instead use SGI color table if you need to use a color table at this stage in the pipeline.

The following sections explain each stage of the Pixel Transfer Pipeline. The example code provided shows you how to set the state parameters for the given stage so that `GL_LUMINANCE` data is not expanded until the very end of the pipeline, just before rendering to the frame buffer's window or the pbuffer.

Scale/Bias

This operation multiplies all pixels by a given scale value then adds a bias value. Scale and Bias values can be set differently for each color component of a pixel. These values are set as follows:

```
glPixelTransferf (GL_RED_SCALE,    red_scale_value);
glPixelTransferf (GL_GREEN_SCALE,  green_scale_value);
glPixelTransferf (GL_BLUE_SCALE,   blue_scale_value);
glPixelTransferf (GL_ALPHA_SCALE,  alpha_scale_value);
glPixelTransferf (GL_RED_BIAS,     red_bias_value);
glPixelTransferf (GL_GREEN_BIAS,   green_bias_value);
glPixelTransferf (GL_BLUE_BIAS,    blue_bias_value);
glPixelTransferf (GL_ALPHA_BIAS,   alpha_bias_value);
```

If any of these deviate from their default values, (1.0 for scale and 0.0 for bias), the Scale/Bias block in the Pixel Transfer Pipeline is enabled. If any of the red, green, blue, or alpha components differ from each other for either scale or bias, and if the input format can be expanded to `GL_RGBA` or `GL_ABGR_EXT` format, the expansion will occur before processing starts in the pixel transfer pipeline. If the red, green, blue and alpha scale values are all the same or alpha scale is 1.0, and the red, green, blue and alpha bias values are the same or the alpha bias is 0.0, but the red, green, and blue components are different from their default values, expansion does not need to occur. Hence, if you do a `glDrawPixels` operation and pass in `GL_LUMINANCE` data, the red component will be used to do the scale and bias, and the output will be a `GL_LUMINANCE` format image. Hence, the following OpenGL calls will set up Scale/Bias to process `GL_LUMINANCE` without format expansion:

```
glPixelTransferf (GL_RED_SCALE,    scale_value);
glPixelTransferf (GL_GREEN_SCALE,  scale_value);
glPixelTransferf (GL_BLUE_SCALE,   scale_value);
glPixelTransferf (GL_ALPHA_SCALE,  scale_value);
glPixelTransferf (GL_RED_BIAS,     bias_value);
glPixelTransferf (GL_GREEN_BIAS,   bias_value);
glPixelTransferf (GL_BLUE_BIAS,    bias_value);
glPixelTransferf (GL_ALPHA_BIAS,   bias_value);
```

or

```
glPixelTransferf (GL_RED_SCALE,    scale_value);
glPixelTransferf (GL_GREEN_SCALE,  scale_value);
glPixelTransferf (GL_BLUE_SCALE,   scale_value);
glPixelTransferf (GL_ALPHA_SCALE,  1.0);
glPixelTransferf (GL_RED_BIAS,     bias_value);
```

```
glPixelTransferf (GL_GREEN_BIAS, bias_value);
glPixelTransferf (GL_BLUE_BIAS, bias_value);
glPixelTransferf (GL_ALPHA_BIAS, 0.0);
```

To disable scale/bias, just reset the scale/bias values back to their default values as shown below:

```
glPixelTransferf (GL_RED_SCALE, 1.0);
glPixelTransferf (GL_GREEN_SCALE, 1.0);
glPixelTransferf (GL_BLUE_SCALE, 1.0);
glPixelTransferf (GL_ALPHA_SCALE, 1.0);
glPixelTransferf (GL_RED_BIAS, 0.0);
glPixelTransferf (GL_GREEN_BIAS, 0.0);
glPixelTransferf (GL_BLUE_BIAS, 0.0);
glPixelTransferf (GL_ALPHA_BIAS, 0.0);
```

Pixel Map

When in true color mode (RGB mode), if the input image data format is not `GL_RGBA` or `GL_ABGR_EXT`, expansion is always forced if pixel map is enabled using `glPixelTransfer (GL_MAP_COLOR, GL_TRUE)`. If the input image format is `GL_COLOR_INDEX` and the current display mode is RGB, Pixel Map is called automatically whether it was enabled or not to do the conversion from color index to RGBA. In terms of performance for `GL_LUMINANCE`, this case is not optimal and you should use SGI color table instead.

To learn how to use Pixel Map consult the “*OpenGL Reference Manual*,” by the OpenGL Architecture Review Board, known as the blue book. Read the sections on `glPixelTransfer`, and `glPixelMap`.

Color Table

This extension is very useful for accelerating color lookup for `GL_LUMINANCE` data. Other formats are accelerated as well; however, `GL_LUMINANCE` benefits the most. The following code fragment shows how to correctly setup color table to perform a color lookup for `GL_LUMINANCE` data:

```
int unpack_row_length;
int unpack_skip_pixels;
int unpack_skip_rows;
int unpack_alignment;
int lut_size;
```

```

void *lut;
/* Turns on color table. */
glEnable (GL_COLOR_TABLE);
/* The current pixel storage modes also affect color table */
/* definition at the time the color table is created. We */
/* need to grab the current values, set the row length, */
/* skip pixels and skip rows to the defaults and */
/* set unpack alignment to 1. When finished defining the */
/* color table, restore the original values. */
glGetIntegerv (GL_UNPACK_ROW_LENGTH, (long *) &unpack_row_length);
glGetIntegerv (GL_UNPACK_SKIP_PIXELS, (long *) &unpack_skip_pixels);
glGetIntegerv (GL_UNPACK_SKIP_ROWS, (long *) &unpack_skip_rows);
glGetIntegerv (GL_UNPACK_ALIGNMENT, (long *) &unpack_alignment);
glPixelStorei (GL_UNPACK_ROW_LENGTH, 0);
glPixelStorei (GL_UNPACK_SKIP_PIXELS, 0);
glPixelStorei (GL_UNPACK_SKIP_ROWS, 0);
glPixelStorei (GL_UNPACK_ALIGNMENT, 1);
/* Define the color table for GL_LUMINANCE. */
/* If data type is GL_UNSIGNED_BYTE create a lookup table with */
/* 256 entries. Each entry is of type GL_UNSIGNED_BYTE. */
/* Range of values for any entry is [0, 255]. */
/* For a GL_SHORT lookup table, generate a table of 65536 entries */
/* ranging from -32768 to 32767. */
if (data_type == GL_UNSIGNED_BYTE) {
    lut_size = 256;
    lut = generate_unsigned_byte_lut();
}
else if (data_type == GL_SHORT) {
    lut_size = 65536;
    lut = generate_short_lut();
}
glColorTable (GL_COLOR_TABLE,
             GL_LUMINANCE, /* Need to specify internal format.
*/
             lut_size,
             GL_LUMINANCE, /* Format of lut passed in. */

```

```

        data_type,          /* Data type of lut passed in.  */
        lut);              /* Actual pointer to lut arrayl. */
/* Restore original Pixel Storage values in case something else */
/* needed these values.                                          */
glPixelStorei (GL_UNPACK_ROW_LENGTH,  unpack_row_length);
glPixelStorei (GL_UNPACK_SKIP_PIXELS, unpack_skip_pixels);
glPixelStorei (GL_UNPACK_SKIP_ROWS,  unpack_skip_rows);
glPixelStorei (GL_UNPACK_ALIGNMENT,  unpack_alignment);

```

Convolution, Post Convolution Scale/Bias and Post Convolution Color Table

Convolution comes in three flavors: 1D convolution (applies to 1D textures only), 2D general convolution, and 2D separable convolution. Special effort has been made to maximize throughput for 2D general and separable convolutions for `GL_LUMINANCE` format for `GL_UNSIGNED_BYTE` and `GL_SHORT` data types via the `glDrawPixels` interface.

Convolution allows you to set scale and bias values that are applied to the convolution filter kernel before it is used for convolving the image. This is different from post convolution scale/bias (below) in that the bias is applied to the filter itself before processing, where as with post convolution scale/bias, the bias is added to the final convolution result before clamping for the given data type (`GL_UNSIGNED_BYTE` or `GL_SHORT`).

Convolution and post convolution scale/bias have been combined into one operation. The kernel values for convolution are multiplied by the scale value of the post convolution scale/bias, then after each pixel is convolved the bias is added. Since this is all done in VIS, there is no loss in performance when compared with an ordinary convolve implemented in VIS.

The Sun OpenGL for Solaris implementation of convolution only supports 1×3 , 1×5 , and 1×7 convolves for 1D convolves, and 3×3 , 5×5 , and 7×7 for 2D convolves. Also, the source image must be three times larger than the size of the convolve kernel to be used.

The Sun OpenGL for Solaris convolution also supports the following border modes: `GL_REDUCE`, `GL_IGNORE_BORDER`, `GL_CONSTANT_BORDER`, `GL_WRAP_BORDER`, `GL_REPLICATE_BORDER`.

SGI post convolution color table is set up exactly the same way as color table. The only difference being the target value when defining the table.

The code fragment below shows how to setup 2D convolution for both the general and separable cases for a 3×3 convolve on GL_LUMINANCE format image data. The setup is the same for either GL_UNSIGNED_BYTE or GL_SHORT data. It also prepares for using the GL_CONSTANT_BORDER mode, uses the GL_CONVOLUTION_FILTER_SCALE and the GL_CONVOLUTION_FILTER_BIAS, sets up for post convolution scale/bias, then finally sets up the SGI post convolution color table.

```

int unpack_row_length;
int unpack_skip_pixels;
int unpack_skip_rows;
int unpack_alignment;
int lut_size;
void *lut;
float kernel3x3[9] = { 0.111111111, 0.111111111, 0.111111111,
                      0.111111111, 0.111111111, 0.111111111,
                      0.111111111, 0.111111111, 0.111111111};
float sepkernel3[3] = { 0.333333333, 0.333333333, 0.333333333};
float const_color[4] = { 0.5, 0.5, 0.5, 0.5 };
float kernel_scales[4] = { 0.8, 0.8, 0.8, 0.8 };
float kernel_biases[4] = { 0.2, 0.2, 0.2, 0.2 };
float post_conv_scales[4] = { 0.75, 0.75, 0.75, 0.75 };
float post_conv_biases[4] = { 0.25, 0.25, 0.25, 0.25 };
/* The current pixel storage modes affect convolve kernel */
/* destination at the time the kernels are created. */
/* We need to grab the current values, set the row length, */
/* skip pixels and skip rows to the defaults and set unpack */
/* alignment to 1. */
/* When finished defining the color table, restore the */
/* original values. */
glGetIntegerv (GL_UNPACK_ROW_LENGTH, (long *) &unpack_row_length);
glGetIntegerv (GL_UNPACK_SKIP_PIXELS, (long *) &unpack_skip_pixels);
glGetIntegerv (GL_UNPACK_SKIP_ROWS, (long *) &unpack_skip_rows);
glGetIntegerv (GL_UNPACK_ALIGNMENT, (long *) &unpack_alignment);
glPixelStorei (GL_UNPACK_ROW_LENGTH, 0);
glPixelStorei (GL_UNPACK_SKIP_PIXELS, 0);
glPixelStorei (GL_UNPACK_SKIP_ROWS, 0);
glPixelStorei (GL_UNPACK_ALIGNMENT, 1);

```

```

/* Now, setup convolution with constant color border mode. */
if (convolve_type == GL_CONVOLUTION_2D) {
    glEnable (GL_CONVOLUTION_2D);
    glConvolutionFilter2D (GL_CONVOLUTION_2D,
        GL_LUMINANCE, /* Internal format. */
        3, 3, /* Kernal dimensions. */
        GL_LUMINANCE, /* Input kernel data format
*/
        GL_FLOAT, /* Data type for kernel */
        /* entries. */
        (void *) kernel3x3); /* Pointer to kernel
.*/

    glConvolutionParameteri(GL_CONVOLUTION_2D,
        GL_CONVOLUTION_BORDER_MODE,
        GL_CONSTANT_BORDER);

    glConvolutionParameterfv(GL_CONVOLUTION_2D,
        GL_CONVOLUTION_BORDER_COLOR,
        const_color);

    glConvolutionParameterfv(GL_CONVOLUTION_2D,
        GL_CONVOLUTION_FILTER_SCALE,
        kernel_scales);

    glConvolutionParameterfv(GL_CONVOLUTION_2D,
        GL_CONVOLUTION_FILTER_BIAS,
        kernel_biases);
}
else if (convolve_type == GL_SEPARABLE_2D) {
    glEnable (GL_SEPARABLE_2D);
    glSeparableFilter2D (GL_SEPARABLE_2D,
        GL_LUMINANCE,
        3, 3,
        GL_LUMINANCE,
        GL_FLOAT,
        sepkernel3, /* Horizontal Kernal Values. */
        sepkernel3); /* Vertical Kernal Values. */
    glConvolutionParameteri(GL_SEPARABLE_2D,
        GL_CONVOLUTION_BORDER_MODE,
        GL_CONSTANT_BORDER);
    glConvolutionParameterfv(GL_SEPARABLE_2D,
        GL_CONVOLUTION_BORDER_COLOR,

```

```

        const_color);
    glConvolutionParameterfv(GL_SEPARABLE_2D,
        GL_CONVOLUTION_FILTER_SCALE,
        kernel_scales);
    glConvolutionParameterfv(GL_SEPARABLE_2D,
        GL_CONVOLUTION_FILTER_BIAS,
        kernel_biases);
}
glPixelTransferf(GL_POST_CONVOLUTION_RED_SCALE,
    post_conv_scales[0]);
glPixelTransferf(GL_POST_CONVOLUTION_GREEN_SCALE,
    post_conv_scales[1]);
glPixelTransferf(GL_POST_CONVOLUTION_BLUE_SCALE,
    post_conv_scales[2]);
glPixelTransferf(GL_POST_CONVOLUTION_ALPHA_SCALE,
    post_conv_scales[3]);
glPixelTransferf(GL_POST_CONVOLUTION_RED_BIAS,
    post_conv_biases[0]);
glPixelTransferf(GL_POST_CONVOLUTION_GREEN_BIAS,
    post_conv_biases[1]);
glPixelTransferf(GL_POST_CONVOLUTION_BLUE_BIAS,
    post_conv_biases[2]);
glPixelTransferf(GL_POST_CONVOLUTION_ALPHA_BIAS,
    post_conv_biases[3]);
/* Turns on post convolution color table. */
glEnable (GL_POST_CONVOLUTION_COLOR_TABLE);
/* Define the color table for GL_LUMINANCE. */
/* If data type is GL_UNSIGNED_BYTE create a lookup table with */
/* 256 entries. Each entry is of type GL_UNSIGNED_BYTE. */
/* Range of values for any entry is [0, 255]. */
/* For a GL_SHORT lookup table, generate a table of 65536 entries */
/* ranging from -32768 to 32767.*/

if (data_type == GL_UNSIGNED_BYTE) {
    lut_size = 256;
    lut = generate_unsigned_byte_lut();
}

```

```

}
else if (data_type == GL_SHORT) {
    lut_size = 65536;
    lut = generate_short_lut();
}
glColorTable (GL_POST_CONVOLUTION_COLOR_TABLE,
              GL_LUMINANCE,      /* Need to specify internal
format. */
              lut_size,
              GL_LUMINANCE,     /* Format of lut passed in. */
              data_type,       /* Data type of lut passed
in. */
              lut);           /* Actual pointer to lut
arrayl. */
/* Restore original Pixel Storage values in case something else */
/* needed these values. */
glPixelStorei (GL_UNPACK_ROW_LENGTH,  unpack_row_length);
glPixelStorei (GL_UNPACK_SKIP_PIXELS,  unpack_skip_pixels);
glPixelStorei (GL_UNPACK_SKIP_ROWS,    unpack_skip_rows);
glPixelStorei (GL_UNPACK_ALIGNMENT,    unpack_alignment);

```

Histogram and Minmax

The Histogram and Minmax operations come at the end of the Pixel Transfer Pipeline. When used, both can have their own “sink” values. If sink is enabled (`GL_TRUE`), processing of image data stops here, and does not continue down the pipeline and no output is generated. If the histogram's sink value is true, minmax is not executed. (See the man pages for more information about the sink behavior of these operations).

The code below gives an example of getting a histogram for `GL_LUMINANCE` and data for both `GL_UNSIGNED_BYTE` and `GL_SHORT`. Notice below that the requested width of the histogram definition for `GL_SHORT` has been specified to be 32768 instead of 65536. The reason is that, for `GL_SHORT` data, the data is effectively clamped in the range [0, 32767]. That is, if any of the `GL_SHORT` values are negative, they will contribute to the very first histogram bin counter value for 0. Specifying a larger width is pointless since only every other histogram bin would have a value in it. Histogram widths, in general, may be any value which is a power of 2 in the range [0, 65536]. However, for those cases where you want to actually display the computed histogram, you can specify a smaller width for `GL_SHORT` data type, say 256, 512, or 1024. This saves you the time because you do not have to do the code. By

requesting a smaller histogram width, histogram bins are added together. For example, for `GL_SHORT`, if you requested a width of 256, each returned bin value in the histogram image would have 128 bins added together. Hence, all values in the range [0, 127] would be in bin 0. All values in the range [128, 255] would be in bin 1, and so on.

Minmax uses the histogram to compute its values. It gets the minmax values using the histogram for the full width of the positive values for `GL_UNSIGNED_BYTE` and `GL_SHORT`. Therefore, if the histogram is taken of `GL_UNSIGNED_BYTE`, the possible range of minmax values is [0, 255]. For `GL_SHORT`, the possible range of minmax values is [0, 32767].

```
int                minmax[2];
int                histogram[32768];
unsigned char      *uc_buff;
short             *s_buff;

glEnable(GL_HISTOGRAM);
glEnable(GL_MINMAX);

/* Allocate enough space for 64 x 64 GL_LUMINANCE images. */
uc_buff = (unsigned char *) malloc (4096*sizeof(unsigned char));
s_buff  = (short *)          malloc (4096*sizeof(short));

/* First, do it for GL_UNSIGNED_BYTE with GL_LUMINANCE format. */
glHistogram(GL_HISTOGRAM, 256, GL_LUMINANCE, GL_FALSE);
glMinmax(GL_MINMAX, GL_LUMINANCE, GL_FALSE);
glDrawPixels(64, 64, GL_LUMINANCE, GL_UNSIGNED_BYTE, uc_buff);
/* Since the call to glHistogram defined a width of 256, */
/* 256 entries of the histogram array will be filled in. */
/* The remaining entries in the array are untouched. */
glGetHistogram(GL_HISTOGRAM, GL_TRUE, GL_LUMINANCE, GL_INT,
               histogram);

glGetMinmax(GL_MINMAX, GL_TRUE, GL_LUMINANCE, GL_INT,
            minmax);

/* Do something with the histogram and minmax. */
/* Now, do GL_SHORT data. */
glHistogram(GL_HISTOGRAM, 32768, GL_LUMINANCE, GL_FALSE);
glMinmax(GL_MINMAX, GL_LUMINANCE, GL_FALSE);
glDrawPixels(64, 64, GL_LUMINANCE, GL_SHORT, s_buff);
/* Since the call to glHistogram defined a width of 32768, */
/* 32768 entries of the histogram array will be filled in. */
```

```
glGetHistogram(GL_HISTOGRAM, GL_TRUE, GL_LUMINANCE, GL_INT,
              histogram);
glGetMinmax(GL_MINMAX, GL_TRUE, GL_LUMINANCE, GL_INT,
           minmax);
```

Pixel Transform

Pixel Transform, while shown at the end of the Pixel Transfer Pipeline, is not part of it. Pixel Transform is in the Pixel Rasterizer, and it only works through the `glDrawPixels` interface.

Pixel Transform has been especially optimized for applying affine transformation warping to an input image on its way to the frame buffer or pbuffer. It has been specially tuned for handling `GL_LUMINANCE` format and the `GL_UNSIGNED_BYTE` and `GL_SHORT` data types. For `GL_SHORT`, the data is scaled and clamped to `[0, 255]` and then warped into the frame buffer or pbuffer. On the way to the frame buffer, the data is also expanded from `GL_LUMINANCE` data to XBGR format, which is the native format of the frame buffer while in `rgb` mode.

Pixel Transform has its own matrix mode with its own matrix stack 32 deep.

```
glMatrixMode(GL_PIXEL_TRANSFORM_2D_EXT);
```

Pixel Transform is always enabled; however, if its current matrix is the identity matrix, the pixel transform is not performed. Only when the current matrix is not the identity matrix will pixel transform be performed.

You can use all of the existing API calls available for matrix operations in OpenGL. These will operate on the current matrix of the `GL_PIXEL_TRANSFORM_2D_EXT` matrix mode (that is, `glLoadMatrix`, `glTranslate`, `glRotate`, `glScale`, `glLoadIdentity`, `glPushMatrix`, `glPopMatrix`, `glMultMatrix`, and so on). When using these matrix operators on the current matrix, after the operation is performed, only the affine components are kept. Entries in the matrix which apply to the `z` and `w` components are left like they were initialized with the identity matrix.

The pixel transform extension operates as if the current raster position is the origin of the coordinate system. To simplify, set the current raster position to be located in the lower left corner of the display window, then figure out your operations. If you want to translate the image, you can use `glTranslate`, or move the current raster position. The difference is that `glTranslate` will be integrated into the total transformation for pixel transform, while moving the raster position will translate the image regardless of the current matrix contents of the pixel transform matrix.

`glPixelZoom` also affects the pixel transform current matrix; however, only if the current matrix mode is set to `GL_PIXEL_TRANSFORM_2D_EXT`. Also, if `glPixelZoom` is called, it replaces the contexts of the current matrix as shown below:

$$\begin{bmatrix} x_zoom & 0 & 0 & 0 \\ 0 & y_zoom & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

If the current matrix mode is not `GL_PIXEL_TRANSFORM_2D_EXT`, the current matrix of `GL_PIXEL_TRANSFORM_2D_EXT` is not replaced. However, pixel zoom will still be set.

If the current matrix of `GL_PIXEL_TRANSFORM_2D_EXT` has been set to something different than identity, and `glPixelZoom` has been set, the pixel transform will override the `glPixelZoom` operation.

If you want to do any image warping, use the pixel transform extension. Do not use the `glPixelZoom` interface. Instead, use `glScale` to set up a zoom matrix. If you are using multiple matrix operations on the pixel transform's current matrix, do not use `glPixelZoom` in the middle or end of the list of operations since it will reset the matrix (shown above) and remove the affect of any previous operations. Instead, use `glScale`.

Pixel Transform supports four types of resampling for minification and three types for magnification. `GL_NEAREST`, `GL_LINEAR`, and `GL_CUBIC_EXT` are shared by minification and magnification. `GL_AVERAGE_EXT` is only supported for minification.

The code fragment below demonstrates how to prepare a pixel transform matrix to do an arbitrary rotation of “angle” degrees about the center of the input image in the center of the frame buffer display window. It assumes the image is `GL_LUMINANCE` data and `GL_UNSIGNED_BYTE`. It also sets up the resampling method to be `GL_LINEAR` for minification and `GL_CUBIC_EXT` for magnification and sets the `GL_CUBIC_WEIGHT_EXT` to have the value `-0.5`.

```
double rotation_angle;
int    window_width, window_height;
int    image_width, image_height;
unsigned char *image_data;
/* Grab needed values for placing image in center. */
window_width  = get_window_width();
window_height = get_window_height();
image_width   = get_image_width();
image_height  = get_image_height();
image_data    = get_image_data();
```

```

rotation_angle = get_rotation_angle_between_0_and_360_degrees();
/* Prepare current pixel transform matrix. */
glMatrixMode(GL_PIXEL_TRANSFORM_2D_EXT);
glLoadIdentity();
glTranslated(window_width/2.0, window_height/2.0, 0.0);
glRotated(rotation_angle, 0.0, 0.0, 1.0);
glTranslated (-image_width/2.0, -image_height/2.0, 0.0);
/* Set up resampling methods. */
glPixelTransformParameteriEXT(GL_PIXEL_TRANSFORM_2D_EXT,
                               GL_PIXEL_MIN_FILTER_EXT,
                               GL_LINEAR);
glPixelTransformParameteriEXT(GL_PIXEL_TRANSFORM_2D_EXT,
                               GL_PIXEL_MAG_FILTER_EXT,
                               GL_CUBIC_EXT);
glPixelTransformParameterfEXT(GL_PIXEL_TRANSFORM_2D_EXT,
                              GL_PIXEL_CUBIC_WEIGHT_EXT,
                              -0.5);
/* Finally, render the image to the screen. */
glDrawPixels (image_width, image_height, GL_LUMINANCE,
             GL_UNSIGNED_BYTE,
             image_data);

```

Software Performance

Software performance is affected by attributes that force the use of the generic software rasterizer:

1. Texturing Attributes

- a. Only triangles are optimized. Texturing of points and lines is handled by the generic software.
- b. Texture environment mode `glTexEnv(3gl) GL_TEXTURE_ENV_MODE` is `GL_BLEND`.

2. Fragment Attributes

- a. Stencil operations
- b. Logic operations
- c. Any blending operation
- d. Linear or nonlinear fog
- e. Enabling any Z comparison other than `GL_LESS` or `GL_LEQUAL`

X Visuals for Sun OpenGL for Solaris

Programming With X Visuals for Sun OpenGL for Solaris Software

OpenGL rendering is supported on a subset of the visuals exported by the Solaris X window server on the Creator and Creator3D workstations. Because GLX overloads the core X visual classes with a set of attributes that indicate frame buffer capabilities, such as double buffer mode or stereo capabilities, the number of visuals supported by an OpenGL-capable X server is potentially large. For example, for the 24-bit TrueColor visual, the Solaris X window server on the Sun XVR-500 and Sun XVR-1000 graphics accelerator workstations exports the following types of GLX visuals: double buffer, single buffer, monoscopic, stereoscopic, and multisample.

This approach of exporting multiple GLX visuals for each X protocol core visual is colloquially referred to as the *GLX expansion* (or *visual explosion*). For each different type of GLX visual that is exported, there is a corresponding X protocol core visual. Thus, there are multiple GLX visuals whose core X visual attributes are all identical.

Various OpenGL-capable visuals are supported in various releases of the Solaris operating environment. These are the visuals that an OpenGL program can use. This information applies to Sun Creator, Creator3D, Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators.

The advantage to the overloading of X visuals is that the X server can be specific about the frame buffer configurations that the graphics hardware provides. This approach also enables the OpenGL implementation to better manage resources. Instead of allocating the maximal amount of resources for each window, the OpenGL

implementation only needs to allocate the resources necessary for the GLX visual the application has selected. Thus, the application has more direct control over resource allocation.

Using the `glXGetConfig(3gl)` and `glXChooseVisual(3gl)` routines, applications can get information on the supported visuals and choose the appropriate visual. For helpful information on GLX programming, refer to *OpenGL Programming for X Windows Systems* by Mark Kilgard and *OpenGL Programming Guide*.

When the frame buffer video mode is monoscopic, only `GL_MONO` versions of these visuals are supported. In a stereoscopic video mode, both `GL_MONO` and `GL_STEREO` versions of these visuals are supported.

Colormap Flashing for OpenGL Indexed Applications

With the visuals exploded, there is greater potential for colormap flashing to occur for OpenGL indexed applications. This is because applications are forced to create private colormaps in order to create windows on the GLX visual they choose. In the post Solaris 2.5.1 releases, the colormap flashing problem is eased by the colormap equivalence feature. This feature allows OpenGL color indexed applications to be written in a way that creates less flashing.

Colormap equivalence allows a program to assign a colormap of one visual to a window that was created with a different visual, as long as the two visuals are colormap equivalent. This means, in general, that they share the same plane group and have the same number of colormap entries. The standard X11 protocol does not let programs mix visuals of colormaps and windows in this way. For more information on colormap equivalence, see the `XSolarisCheckColormapEquivalence(3)` man page.

Colormap equivalence is useful for OpenGL programs because the GLX visual expansion creates up to four different variants of each base `GL_CAPABLE` visual. So, for example, instead of one 8-bit default PseudoColor colormap, there may be a double-buffered variant, a stereo variant, and so on. Without colormap equivalence, an application cannot assign the default colormap to windows of these variant visuals, and this will result in more colormap flashing. With colormap equivalence, windows of all variants can share a colormap that was created using the base visual, and less colormap flashing will occur.

GL Rendering Model and X Visual Class

OpenGL RGBA rendering is supported on the 24-bit TrueColor and DirectColor visuals. OpenGL indexed rendering is supported on the 8-bit PseudoColor visuals and on the indexed or 224-color overlay visuals.

Note – For the Sun PGX64 graphics frame buffer, OpenGL indexed rendering is supported when the frame buffer is configured with depth 8 mode.

Stereo

Note – This section is specific to Sun Creator, Creator3D, Elite3D, Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 graphics accelerators.

To run a stereo application in stereo mode, the frame buffer must be configured for stereo operation.

▼ To Set Up the Frame Buffer for Stereo Operation (Sun Expert3D, Expert3D-Lite, Sun XVR-500 and Sun XVR-1000 Graphics Accelerators)

1. Exit the window system.
2. Type this command:

```
# /usr/sbin/fbconfig -res stereo
```

Note – This command must be run under superuser permissions or sys admin permissions.

3. Restart the window system.

Application can now use the stereo hardware buffers.

▼ To Set Up the Frame Buffer for Stereo Operation (Creator and Creator3D)

1. Exit the window system.

2. Type this command:

```
# /usr/sbin/ffbconfig -res stereo
```

Note – This command must be run under superuser permissions or sys admin permissions.

3. Restart the window system.

Application can now use the stereo hardware buffers.

▼ To Set Up the Frame Buffer for Stereo Operation (Elite3D)

1. Exit the window system.

2. Type this command:

```
# /usr/sbin/afbconfig -res stereo
```

3. Restart the window system.

Application can now use the stereo hardware buffers.

Rendering to DirectColor Visuals

The OpenGL API has no support for color mapping. The only way to get a DirectColor visual is to implement visual selection in the application using `XGetVisualInfo(3gl)` and `glXGetConfig`. If you request a visual with `glXChooseVisual`, you will get a 24-bit TrueColor visual for RGBA rendering and an 8-bit PseudoColor visual for index rendering.

When rendering to DirectColor visuals, the GL system calculates pixel values in the same way as it does for TrueColor visuals. The application is responsible for loading the window colormap with cells that make sense to the application.

Overlays

The overlay visual GLX level is greater than zero (`GLX_LEVEL > 0`). Visuals with a GLX level less than or equal to zero are underlay visuals.

Server Overlay Visual (SOV) Convention

Server Overlay Visual (SOV) is an API for rendering transparent pixels in an overlay window. A transparent pixel is a special pixel code that allows the contents of underlay windows underneath to show through. SOV derives its name from the X property that informs the user of the special transparent pixel value: `SERVER_OVERLAY_VISUALS`. This value can be used as the input value to `glIndex*` calls so that the transparent pixel can be rendered into the overlay.

The SOV API, while not an X11 standard, is a convention that is supported by many X11 vendors. It is described at length in the book *OpenGL Programming for the X Window System* by Mark J. Kilgard. This section describes Sun-specific aspects of the SOV implementation.

Note – In this section, the term underlay is used as a synonym for the normal planes referred to in *OpenGL Programming for the X Window System*.

The `SERVER_OVERLAY_VISUALS` property describes visuals with transparent pixels (`TransparentType = TransparentPixel`), and also completely opaque visuals (`TransparentType = None`). If you need an overlay visual with a transparent pixel, make sure that you check the `TransparentType` field of the entries in this property. The remainder of this section will discuss only the `TransparentPixel` SOV visuals.

Compatibility of SOV with other Overlay Models

Programs that use SOV visuals may coexist on the same screen with programs that use OVL, the Sun-specific overlay extension. But the two may not be used simultaneously with the same window.

Some OpenGL 1.0 programs are written to use the SOV transparent pixel if the SOV property is present, and to use XOR rendering in the default underlay visual if the SOV property is not present. These programs may not behave properly when the SOV property is present. When the SOV property is not present and the underlay is being used, a program may simply attach the default colormap to the default visual underlay window. In the presence of the SOV visual, the program will switch to using the SOV overlay visual but may continue to use the default colormap. Since the SOV overlay visual is usually not the same as the default visual, this will result in an X11 `BadMatch` error when the program attempts to attach the colormap to the overlay window. Care should be taken to write programs that always attach colormaps of the proper visual to overlay windows. In this case, the program should have created a colormap using the SOV visual instead of trying to use the default colormap.

Programs that use SOV can also coexist with programs using the Sun visual overlay capability `glXGetTransparentIndexSUN`. However, `glXGetTransparentIndexSUN` is deprecated. It is provided only for compatibility for existing programs that use it. Newly written transparent overlay programs should use `SERVER_OVERLAY_VISUALS` instead.

For information on using the Sun visual overlay capability, see the `glXGetTransparentIndexSUN` man page.

Gamma Correction

On some Sun frame buffers, two 24-bit TrueColor visuals are exported. One is gamma corrected; the other is not. To support imaging and Xlib applications, the nonlinear (not gamma-corrected) visuals are listed before linear visuals. However, to provide linear visuals for graphics applications running under the Sun OpenGL for Solaris software, the `glXChooseVisual()` call was modified to return a linear visual.

If you want to use a nonlinear TrueColor visual, you need to get the visual list from Xlib. Use the Solaris API `XSolarisGetVisualGamma(3)` to query the linearity of the visual. To determine whether a visual supports OpenGL, call `glXGetConfig` with *attrib* set to `GLX_USE_GL`.

Tips and Techniques

This chapter presents miscellaneous topics that you may find useful as you port your application with Sun OpenGL for Solaris.

Avoiding Overlay Colormap Flashing

Colormap flashing may occur when your application uses overlay windows. This problem stems from several characteristics of the graphics frame buffers: the overlay visual is not the default visual, the Creator3D and Elite3D are a single hardware colormap device, and X11 allocates colormap cells from pixel 0 upward. When the application renders to the overlay window, it must use a non-default visual, and a non-default colormap is loaded. In this case, colormap flashing between the default and non-default colormaps can occur.

The best solution to this problem is to allocate the overlay colors at the high end of the overlay colormap. In other words, if you have n colors to allocate, allocate them in the positions $colormap_size - n - 1$ to $colormap_size - 1$. This avoids the colors in the default colormap, which are allocated upward starting at 0. To allocate n colors at the top of the overlay colormap, first allocate $colormap_size - n$ read/write placeholder cells using `XAllocColorCells`. Then allocate the n overlay colors using `XAllocColor`. Finally, free the placeholder cells. This solution is portable; it works on both single- and multiple-hardware colormap devices.

Changing the Limitation on the Number of Simultaneous GLX Windows

There is a limitation on the number of GLX windows that an application can use simultaneously. Each GLX window that has an attached GLX context uses a file descriptor for DGA (Direct Graphics Access) information. You can find the current number of open file descriptors using the `limit(1)` command:

```
% limit descriptors
descriptors 64
```

The system response tells you that you have up to 64 direct GLX contexts, assuming that you have no other processes concurrently using file descriptors.

You can increase the per-process maximum number of open file descriptors using the `limit` command as follows:

```
% limit descriptors 128
```

This command changes the number of file descriptors available for DGA and other uses to 128. Use the `sysdef(1M)` command to determine the maximum number of file descriptors for your system.

Hardware Window ID Allocation Failure Message

When a program calls `glXMakeCurrent(3gl)` to make a window the current OpenGL drawable, the system will attempt to allocate a unique hardware window ID (WID) for the window. This allows double buffering and hardware WID clipping to be used. Because hardware WIDs are a scarce resource and can be used for other purposes, there might not be any WIDs available when `glXMakeCurrent` is called. If this should happen, the following message is displayed:

```
Warning: unable to allocate hardware window ID
```

In this situation, double buffering will not be provided for the window, and the window will be treated as a single-buffered window.

Getting Peak Frame Rate

The frame rate that `ogl_install_check` prints out is synchronized to monitor frequency. It measures the time it takes to render the frame, wait for `vblank`, then swap the buffers. Since accelerators can render the `ogl_install_check` image very quickly, even on a Creator3D UltraSPARC 167 Mhz machine, the bottleneck is waiting for the monitor *vblank*. So, under normal circumstances, `ogl_install_check` is never going to be able to get a frame rate faster than the monitor frequency.

However, there is an environment variable called `OGL_NO_VBLANK` that you can set to see the peak, unsynchronized frame rate. When set, this environment variable swaps buffers immediately, without waiting for *vblank*.

Identifying the Release Version

You can identify the Release Version Number of the Sun OpenGL Library by:

1. Using the `what(1)` command:

```
% what /usr/openwin/lib/libGL.so.1
```

2. Programmatically, by calling `glGetString (GL_VERSION)`
(see the `glGetString` man page for more details)
3. Running the Sun OpenGL for Solaris `install_check` demo program:

```
% /usr/openwin/demo/GL/ogl_install_check
```

Determining Visuals Supported by a Specific Frame Buffer

To determine what visuals a specific frame buffer supports, use

```
/usr/openwin/demo/GL/xglnfo
```

Developing Applications for 64-bit

You should use Sun's DevPro 5.x tools (compiler, workshop, and so on), or subsequent compatible release, to develop your applications on 64-bit. To develop your applications on 64-bit, run `lint` on your C source files using the `-errchk=longptr64` option.

Before you compile your source files, be sure to resolve all 64-bit warnings. Then compile your source files using the `-xarch=v9` option.

You can determine if your application is using 64-bit by running the `file` command.

For example, if your application is 64-bit, you will see the following result when you run the `file` command:

```
ELF 64-bit MSB dynamic lib SPARCv9
```

Common 64-bit Application Development Errors

This section identifies known 64-bit application development errors, describes ways to detect them, and suggests solutions.

Subtracting With Unsigned Int

A possible problem exists when subtracting with unsigned int. You may need to change unit to int in the following example:

```
long_type = long_type - uint_type
```

This is especially true when doing pointer arithmetic.

Functions Without Prototypes

Functions without prototypes return integer values. Consequently, you must include `#include <malloc.h>` or the correct prototypes, otherwise returned pointer values truncate to integers. To show the warnings where function prototypes and usage do not match correctly, compile your applications using the `-fd` option.

Pointer Alignment

If you have a pointer alignment problem, it probably means internal data structures are not 64-bit aligned. For example, if you have this problem, you will see lines similar to the following ones when you run the debugger:

```
signal SEGV (no mapping at the fault address in my_func at 0xfe43afac
0xfe43afac: my_func+0x0004:      ld    [%i2], %o0
```

The following procedure suggests a way to debug this problem.

- Compile the file containing `my_func` using the `-s` option to generate assembly output. Be sure to maintain the same optimization levels.
- In the `myfile.s` assembly output file, look for:

```
! SUBROUTINE my_func
```

or possibly:

```
.global my_func
```

- In `my_func()`, look for the offset (in this case the offset is `0x0004`). In this example, we have the following lines:

```
/* 000000 457 */ add %g2,%o0,%g2
/* 0x0004 */ ld [%i2,%o0
/* 0x0008 */ add %i0,2520,%o0
```

A problem exists with line 457 of the source code.

Alignment errors typically appear with non-debug compilations. Also, in some cases you can debug the alignment problem from within the C source code by following `cc -g` with `-xo0`.

Index

NUMERICS

1D convolution, 47

2D

convolution, 48

general convolution, 47

separable convolution, 47

3D

optimized cases, 35

texture mapping, 26

64-bit

application development errors, 68

developing applications for, 68

A

affine transformation warping, 53

alignment errors, 70

alpha test, 27, 37

application development errors, 64-bit, 68

application tuning, 13

architecture, 7 to 12

graphics hardware, 21

arrays, vertex, 17

attributes

Creator3D performance, 28

Elite3D graphics, 25

fragment, 31, 37

software rasterizer, 36

texture load time, 34

texturing, 31

B

batching primitives, 20

blending, 27

fragment, 31

stencil, 31

C

clipping, 9

color mapping, 61

color table, SGI, 33

colormap equivalence, 58

colormap flashing, 58

avoiding, 65

compatibility issues, 3

compression, display list, 16

consistent data, 18, 27

Constant Data Extension, 35

convolution, 39, 42, 47

1D, 47

2D, 48

2D general, 47

2D separable, 47

convolve kernel, 47

coordinate transformations, 9

Creator, 13

Creator graphics, 5, 9

performance, attributes affecting, 36

Creator3D, 13

Creator3D graphics, 5, 9, 21

D

- data types
 - consistent, 18
 - optimized, 21
- default colormap, 58
- depth test `glEnable`, 37
- device coordinate (DC) devices, 10, 21
- Direct Graphics Access (DGA), 66
- DirectColor visuals, rendering to, 61
- display list compression, 16
- display list mode, 17
- documentation, where to find, 5
- double buffering, 66

E

- Elite3D, 5, 13
- Elite3D graphics, 5, 9
 - MIP mapping, 34
 - performance, 25 to 27
 - pixel operations, 27
- env mode, 33
- environment mapping, 26
- environment variables
 - Creator 3D graphics performance, 30
 - performance, 25
- Expert3D, 5, 9, 10, 13, 21
- Expert3D graphics, 9
 - performance, 23 to 25
- Expert3D, *See* Sun Expert3D
- Expert3D-Lite, 5, 9, 10, 13
- Expert3D-Lite, *See* Sun Expert3D-Lite
- extensions
 - determining, 3

F

- fog
 - in indexed color mode, 31
 - on Creator3D, 37
- fragment attributes
 - Creator graphics, 37
 - Creator3D graphics, 31
 - Elite3D graphics, 27

GX, 56

- fragment blending, 31
- fragment processing, 21
- fragments, 8
- frame rate, peak, 67
- functions without prototypes, 69

G

- gamma correction, 63
- GL rendering model, 59
- `glBitmap` command, 39, 41
- `glCopyPixels` command, 39, 41
- `glDrawPixels` command, 38, 40
- `glReadPixels` command, 39, 40
- `glTexEnv` texture base, 36
- GLX
 - contexts, 66
 - visuals, 57, 58
 - windows, limitation on number of, 66
- graphics hardware architecture, 21

H

- hardware
 - acceleration, 13
 - rasterizer, 12
 - rasterizer, Creator3D graphics, 28
 - window ID (WID), 66
- histogram bin, 51
- histogram operations, 51
- homogeneous coordinates, 25

I

- image warping, 54
- immediate mode, 17
 - texture mapping, 26
- indexed
 - applications, colormap flashing, 58
 - color antialiasing, 26
 - color exponential fog, 27
 - color linear fog, 25

- color mode, 30, 37
- install_check demo program, 67
- issues, compatability, 3

L

- line stipple scale factor, 26
- linear visuals, 63
- logical operations, 37
- low batching, 20

M

- magfilter, 33
- memory mappable device, 10
- memory usage, texture, 34
- minfilter, 33
- minification filter, 32
- minmax operations, 51
- model coordinate (MC) devices, 9, 21
- monoscopic mode, 58
- Motif Drawing Widget libraries, 3
- MultiDrawArrays, 17
- multiple rendering threads, 4
- multisampling, 15
- multi-screen Xinerama, 14
- multithread, 4

N

- non-antialiased points, 30
- nonlinear fog, 38
- nonlinear TrueColor visual, 63
- normal planes, 61

O

- OGL_NO_VBLANK environment variable, 67
- OpenGL
 - architecture, 8
 - indexed rendering, 59
 - product functionality, 1

- references, 5
- release version number, 67
- RGBA rendering, 59
- software architecture, 9
- optimized data types, 21
- overlay window, 61, 65
- overlays, 61

P

- packed representation, 35
- parallelism, 4
- peak frame rate, 67
- performance, 13 to 56
 - Creator and Creator3D graphics, 28 to 41
 - Elite3D graphics, 25 to 27
 - Expert3D graphics, 23 to 25
 - GX, 56
- pixel map, 45
- pixel operations, 38
 - Creator3D, 37
 - Elite3D graphics, 27
- pixel store, 35
- pixel transfer pipeline, 43
 - imaging extensions, 41
- pixel transform, 42, 53
 - extension, 53
 - matrix, 54
- platforms supported, 5
- pointer alignment problem, 69
- pointer arithmetic, 69
- polygon anti-aliasing, 30
- polygon offset, 36
- post convolution
 - color table, 42, 47
 - scale and bias, 47
- primitive types, 29
- projection type, 33

R

- rasterization and fragment processing, 12
- rasterization stage, 8
- rasterizer

- hardware, 12
- software, 12
- references, OpenGL, 5
- release version number, OpenGL, 67
- rendering threads, 4

S

- scale and bias values, 44
- server overlay visual (SOV), *See* SOV visuals
- `SERVER_OVERLAY_VISUALS` property, 62
- SGI color table, 39, 43, 45
- SGI post convolution color table, 47, 48
- SGI Texture Color Table, 37
- software optimization, 13
- software performance, 56
- software rasterizer, 12
 - Creator and Creator3D graphics, 28
- SOV visuals, 61
 - compatibility with other overlay models, 62
 - overlay visual, 62
 - property, 62
- stencil test, 31
- stereo, 59
- stereo operation, 59
- stippled lines, 26, 30
- subtracting with unsigned int, problem with, 69
- Sun Expert3D, 21
- Sun Expert3D-Lite, 5
- Sun visual overlay, 62
- Sun XVR-1000 graphics accelerator, 5, 9, 10, 13, 21, 22
- Sun XVR-500 graphics accelerator, 5, 9, 10, 13, 21, 23, 24
- `SUN_OGL_COMPRESSED_DISPLAY_LISTS`, 16
- supported platforms, 5
- surface antialiasing, 26

T

- texture
 - color lookup table, 32, 33
 - coordinate

- classification, 33
- generation, 9
- environment mode, 36
- load time attributes, 34
- mapping, 36, 38
- memory, 8
- memory usage, 34
- texturing
 - attributes
 - Creator3D graphics, 31, 32, 36
 - Elite3D graphics, 27
 - GX, 56
 - rasterization, 9
 - speed, 31
- tips and techniques, 65 to 70
- transparent pixels, 61
- TrueColor visuals, 63

V

- vertex array commands, 17
- vertex array mode, 17
- vertex arrays, 17
- vertex commands, 3
- vertex processing, 12, 21, 25
 - architecture, 12
 - Creator3D graphics, 28
 - optimization, 18
 - overhead, 36
 - tips, 17
- view frustum culling, 15
- VIS
 - instruction set, 9, 12, 37
 - optimization, conditions that result in, 38, 40
 - rasterization, 36
- visual overlay, 62
- visuals
 - OpenGL capable, 57
 - supported by a specific frame buffer, 68

W

- WID clipping, hardware, 66
- window ID allocation failure, 66
- wrap mode, 33

X

X visuals, programming with, 57 to ??

Xinerama, 14

Z

Z buffering, 8, 21

