XIL Programmer's Guide

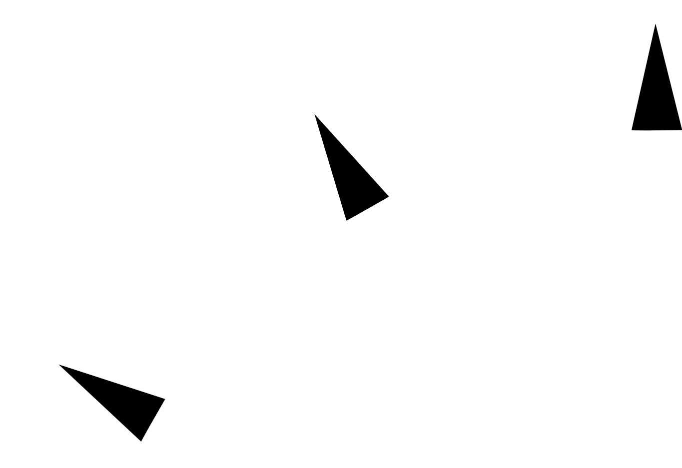

2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.

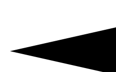

© 1994 Sun Microsystems, Inc.—Printed in the United States of America.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

All rights reserved. This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® and Berkeley 4.3 BSD systems, licensed from UNIX System Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and the University of California, respectively. Third-party font software in this product is protected by copyright and licensed from Sun's Font Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the United States Government is subject to the restrictions set forth in DFARS 252.227-7013 (c)(1)(ii) and FAR 52.227-19.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

TRADEMARKS
Sun, the Sun logo, Sun Microsystems, Sun Microsystems Computer Corporation, SunSoft, the SunSoft logo, Solaris, SunOS, OpenWindows, DeskSet, ONC, ONC+, NFS, and XIL are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and certain other countries. UNIX is a registered trademark of Novell, Inc. in the United States and other countries; X/Open Company, Ltd., is the exclusive licensor of such trademark. OPEN LOOK® is a registered trademark of Novell, Inc. PostScript and Display PostScript are trademarks of Adobe Systems, Inc. Photo CD, the Photo CD logo, PhotoYCC, and Kodak are trademarks or registered trademarks of Eastman Kodak Company. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCstorage, SPARCware, SPARCcenter, SPARCclassic, SPARCcluster, SPARCdesign, SPARC811, SPARCprinter, UltraSPARC, microSPARC, SPARCworks, SPARCompiler, ProWorks, and ProCompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interfaces were 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.

X Window System is a trademark and product of the Massachusetts Institute of Technology.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Please
Recycle

Contents

Preface. xxv

New Features xxxi

1. Introduction to the XIL Library 1

Functions in the XIL Library 1

Image-Processing Functions. 1

Digital-Video Functions 2

Additional XIL Functions 5

The XIL Library: A Foundation Library 5

The XIL Library is Multithread Unsafe 7

Verifying Installation and Setting Environment Variables 7

XIL Packages 8

XIL Directory Structure. 9

Setting Environment Variables. 10

Using the XIL Manual Pages 10

2. Basic XIL Program **11**
Running the Program 12
Including Header Files 12
Initializing the Library 13
Acquiring an Input Image 13
Step 1: Reading the File's Header 14
Step 2: Creating an XIL Image 16
Step 3: Exporting the Image 17
Step 4: Copying Data from a File to Your XIL Image 19
Step 5: Importing the Image 22
Creating an Output Image 22
Creating an X Window 23
Creating a Display Image 24
Processing an Image 24
Making Source and Destination Images Compatible 25
Additional Processing 27
Closing the Library 28
Building an XIL Program 29
Conditionally Compiling Code for Different XIL Versions . . . 31
3. XIL Images **33**
Basic XIL Image Attributes 34
Width, Height, and Number of Bands 34
Data Type 35
Exporting XIL Images 36

Memory Formats for XIL Images 37
XIL_BYTE and XIL_SHORT Images 37
XIL_BIT Images 40
Memory Formats for Images of Different Color Spaces . . . 42
Types of XIL Images 43
Memory Images 43
Device Images 43
Display Images. 44
Additional XIL Image Attributes 46
Origin 46
Region of Interest. 48
Color Space. 52
Parent 53
Image Type 54
Synchronization Flag. 55
Readable and Writable Flags 55
Name 56
4. Handling Input and Output 57
Moving Image Data from a File to an XIL Image. 57
Moving Image Data from an XIL Image to a File. 62
Sending Output to (and Reading Input from) the Display. . . . 64
Possible Complications 66
Reading a Display Image 66
Reading and Writing Devices Other than Displays. 67

Initializing a Device's Attributes 68
Creating a Device Image 71
Destroying a Device Object 73
5. Reading Kodak Photo CD Images 75
The Photo CD Technology 75
The Photo CD Imaging Workstation 75
How Images Are Stored 76
Reading Photo CD Images Using the XIL Library 77
Creating a Device Image 78
Setting Device-Image Attributes 79
Capturing an Image from a Photo CD Disk 82
Converting the Image's Color Space 83
6. Preparing Images for Display 87
Running the Sample Display Program 88
Converting a Single-Band Image to a Multiband Image 88
Passing the Source Image Through a Lookup Table 88
Replicating the Source Image in the Bands of the Destination 90
Converting an XIL_SHORT Image to an XIL_BYTE Image . . . 91
Converting an RGB Image to an Indexed-Color Image and a Colormap 92
Converting a 24-Bit Image to a 1-Bit Image 93
Converting an 8-Bit Image to a 1-Bit Image 94
Displaying a 1-Bit Image on a Monochrome Display 94
Types of Images Displayed 94

7. Presentation Functions **97**
Copying an Image to the Display 97
Copying All Bit Planes 98
Copying Only the Planes Defined in a Plane Mask 99
Rescaling an Image 102
Casting an Image from One Data Type to Another 103
Dithering an Image 105
What Is Dithering? 105
Methods of Dithering 107
When to Use Each Dithering Function 121
Color-Space Conversion 122
Black Generation 125
8. Error Handling **127**
Writing an Error Handler 128
Functions You Can Call in Your Error Handler 129
An Example 135
Installing and Chaining Error Handlers 136
Installing Error Handlers 136
Chaining Error Handlers 139
9. Arithmetic, Relational, and Logical Functions **141**
Arithmetic Functions 141
Relational Functions 143
Logical Functions 144
Operations with Constants 144

Arithmetic and Logical Operations with Bit Images 147
10. Geometric Functions 149
Interpolation Options 150
Nearest Neighbor Interpolation. 150
Bilinear Interpolation 151
Bicubic Interpolation. 151
General Interpolation 151
Translating Images. 160
Scaling and Subsampling Images 161
xil_scale() 162
xil_subsample_adaptive(). 164
xil_subsample_binary_to_gray() 165
Rotating Images 167
Performing General Affine Transforms. 169
Warping Images 172
Transposing Images. 174
11. Miscellaneous Image Processing Functions. 177
Finding the Minimum and Maximum Values in an Image. . . . 178
Producing a Histogram for an Image 179
Creating a Histogram 179
Writing Level Information to the Histogram Structure . . . 180
Reading Data from a Histogram 181
Destroying a Histogram 182
Thresholding an Image 183

Filling an Area in an Image 183
xil_fill() 183
xil_soft_fill() 185
Filtering an Image 187
Creating a Convolution Kernel 188
Filtering an Image 190
Destroying a Convolution Kernel 191
Additional Kernel-Related Functions 192
Detecting Edges in an Image 192
Dilating or Eroding an Image. 194
Creating a Structuring Element 196
Dilating or Eroding an Image. 198
Destroying a Structuring Element 198
Additional Structuring-Element Functions 199
Passing an Image Through a Lookup Table 199
Creating a Lookup Table. 200
Passing an Image Through the Table 204
Destroying a Lookup Table. 205
Additional Lookup-Table Functions 206
Linear Combination. 207
Performing a Linear Combination. 209
How to Use Linear Combinations 210
Blending Images. 211
Painting on an Image. 213

Setting and Getting the Values of Pixels in an Image 215
xil_set_pixel() and xil_get_pixel() 216
xil_set_value() 216
Copying a Pattern to an Image. 217
12. Compressing and Decompressing Sequences of Images. . . . 219
Creating a JPEG Datastream. 222
Building and Running the Example 223
Creating a CIS 223
Compressing Video Frames and Writing Compressed Data to a File. 224
Performing Any Outstanding Compression Operations . . 229
Playing a JPEG Movie 230
Running the Movie Player 231
Memory Mapping the Movie 232
Creating a CIS 233
Putting Compressed Data in a CIS 233
Creating a Display Image 234
Creating an Image to Hold Decompressed Frames. 236
Initializing Parameters to Be Used with the Dither Function 238
Installing an X Colormap 238
Playing the Movie 240
Playing Cell Movies. 241
Installing an X Colormap 242
Creating an Image to Hold Decompressed Frames. 245

Playing the Movie 246

13. Compressed Image Sequences 249

Basic CIS Management 249

Creating and Destroying a CIS. 250

Putting Compressed Data into a CIS. 251

Reading Data from a CIS 254

General CIS Attributes 257

Compressor and Compression Type 258

Input and Output Image Type 259

Random Access Flag 260

Start Frame, Read Frame, Write Frame 260

Maximum Frames and Keep Frames. 262

Error-Recovery Flag 265

Name 266

CIS Error Recovery 267

14. Cell Codec 271

How the Cell Codec Works. 272

Choosing a Colormap 274

Cell Compression Ratios. 275

Image Types 275

Creating a Cell CIS. 275

Cell Codec Attributes 276

Compression Attributes 276

Decompression Attributes 284

Cell Molecules 288
Rules for Calling Decompression Molecules 289
Calling Cell Molecules 290
15. CellB Codec 297
How the Codec Works. 298
Cell Codes. 298
Skip Codes 300
Creating a CellB CIS 301
CellB Decompression Attributes 301
CellB Molecules 303
16. JPEG Baseline Sequential Codec 305
How the JPEG Baseline Sequential Codec Works 306
Discrete Cosine Transform 307
Quantization. 308
Entropy Coding 309
Creating a JPEG Baseline Sequential CIS 310
JPEG Baseline Sequential Codec Attributes 310
Compression Attributes 310
Decompression Attributes 324
JPEG Molecules 326
17. JPEG Lossless Codec 327
How the JPEG Lossless Codec Works 328
Prediction 329
Entropy Coding 330

JPEG Lossless Compressor Attributes. 331
18. H.261 Codec 337
How an H.261 Codec Works. 338
Source Images 338
Basic Encoding Scheme. 340
Bit-Rate Control 345
Provisions for Multipoint Conferencing. 346
Creating an H.261 CIS 346
H.261 Codec Attributes. 347
Compression Attributes 347
Decompression Attributes 354
H.261 Molecules. 358
19. MPEG-1 Codec 359
How an MPEG-1 Codec Works 360
Similarities Between MPEG-1 and H.261 360
Differences Between MPEG-1 and H.261 361
How MPEG-1 Organizes a Video Sequence 364
Creating an MPEG-1 CIS. 365
MPEG-1 Codec Attributes 365
Compression Attributes 366
Decompression Attributes 384
MPEG-1 Molecules 390

20. CCITT Group 3 and Group 4 Codecs 391
How CCITT Group 3 and Group 4 Codecs Work 391
CCITT Group 3 and Group 4 Decompressor Attributes 392
21. Acceleration in XIL Programs 395
What Is Deferred Execution? 395
XIL Molecules 398
Rules for Executing Molecules 398
Video Decompression Molecules 399
CCITT Group 4 Decompression Molecule 407
Image-Filtering Molecule 408
SPARC Molecules That Result in a Display 409
Troubleshooting Molecules 410
Determining Whether Molecules Are Executing 410
Determining Why a Molecule Is Not Executing 411
Side Effects of Executing Molecules 414
XIL Functions That Relate to Deferred Execution 416
A. XIL Molecules 417
Key to the Names Used in the Molecule Definitions 418
Molecule Descriptions 421
Cell Decompression 421
CellB Decompression 421
JPEG Baseline Sequential Decompression 422
H.261 Decompression 422
MPEG-1 Decompression 422

FaxG4 Decompression. 422
Image Filtering. 422
SPARC Other 423
B. XIL Error Messages. 425
C. XIL-XGL Interoperability 459
D. Cell and CellB Bytestream Definitions. 463
Introduction to Cell 463
Encoding Images for Cell 464
Cell Bytestream Description. 466
Key-Frame Header and Key Parameters 467
Cell Code 469
Run Length Code. 470
Escape Codes 470
Summary of Cell Codes 473
CellB Bytestream Description. 474
Cell Code 475
C_b/C_r Quantization Table 476
Y/Y Quantization Table 477
Skip Code 479
New Y/Y Table 479
Default CellB Quantization Tables 480
E. Bibliography 489
Glossary 493
Index 509

Figures

Figure 1-1 Foundation Libraries and Application Programming Interfaces 6

Figure 1-2 Directory Structure of XIL Developer's Kit. 9

Figure 3-1 Digitized Image 33

Figure 3-2 Memory Format for a 3-Band Image Containing 8-Bit Data Elements 38

Figure 3-3 Memory Format for a 3-Band Image Containing 16-Bit Data Elements 38

Figure 3-4 Memory Format for a 3-Band Image Containing 1-Bit Data Elements 41

Figure 3-5 Copying Data from a Temporary Image to the Display (Not Necessary). 44

Figure 3-6 Writing the Output of an Operation to a Display Image. 45

Figure 3-7 Image Origins. 47

Figure 3-8 Regions of Interest. 48

Figure 3-9 Regions of Interest and Origins. 49

Figure 6-1 Lookup Table 89

Figure 7-1 XIL Lookup Operation 106

Figure 7-2 Dithering an Image 107

Figure 7-3 Colorcube for Dithering a True-Color Image to a Pseudocolor Image 112

Figure 7-4 Error-Distribution Kernel 116

Figure 7-5 Using xil_error_diffusion() to Dither an Image 116

Figure 7-6 Error Diffusion 117

Figure 7-7 Dither Mask Replicated over a Source Image 120

Figure 8-1 List of Error Handlers 137

Figure 8-2 Adding to the List of Error Handlers 138

Figure 10-1 Conceptual Model of a General Interpolation. 153

Figure 10-2 Determining the Kernel to Use for a General Interpolation . . 156

Figure 10-3 Zooming the Upper-Left Corner of an Image 163

Figure 10-4 Zooming the Center of an Image 163

Figure 10-5 Subsampling Bit Images 166

Figure 10-6 Rotating an Image Around Its Default Origin 168

Figure 10-7 Rotating an Image Around Its Center 169

Figure 10-8 Shearing an Image Along Its *x* Axis 172

Figure 10-9 Flipping and Rotating Images Using xil_transpose() . . . 176

Figure 11-1 Boundary Fill 184

Figure 11-2 Convolution Operation. 188

Figure 11-3 High-Pass Filters 189

Figure 11-4 Low-Pass Filters 190

Figure 11-5 Filters Used by the XIL_EDGE_DETECT_SOBEL Algorithm . 193

Figure 11-6 Dilating and Eroding Images. 194

Figure 11-7 Dilating an Image 195

Figure 11-8 Eroding an Image 196

Figure 11-9 Single Lookup Table 201

Figure 11-10 Interband Linear Combination 208

Figure 11-11 Linear Combination Matrix 209

Figure 11-12 RGB-to-CMY Conversion Using `xil_band_combine()` . . . 210

Figure 11-13 RGB-to-Y Conversion Using `xil_band_combine()` 210

Figure 11-14 Calculating the Normalized Sum of an Image 211

Figure 11-15 Blending Images 213

Figure 11-16 Painting on an Image 214

Figure 11-17 Replicating a Source Image 218

Figure 12-1 Compressing and Decompressing XIL Images 219

Figure 12-2 Decompressing and Dithering a Frame of Video 237

Figure 14-1 Cell Compression 273

Figure 15-1 Cell Code. 298

Figure 15-2 Vectors in Chrominance Table 299

Figure 15-3 Vectors in Luminance Table. 299

Figure 16-1 JPEG Baseline Sequential Compressor 306

Figure 16-2 Output of the Discrete Cosine Transform 307

Figure 16-3 Quantization in the JPEG Encoder 308

Figure 16-4 Zigzag Sequencing in JPEG Encoder 309

Figure 17-1 JPEG Lossless Compressor. 328

Figure 17-2 Predicting Values in the JPEG Lossless Compressor 329

Figure 18-1 Macroblock 339

Figure 18-2 Flow Diagram for H.261 Encoding 341

Figure 18-3 Motion Compensation in H.261 343

Figure 18-4 Encoding of YC_bC_r or Difference Values in H.261 344

Figure 19-1 An MPEG-1 Bitstream Containing I and P Pictures 361

Figure 19-2 MPEG-1 Display Order Versus Decoding Order 362
Figure 19-3 Bidirectional Prediction in MPEG-1 363
Figure 19-4 Zigzag Ordering of Quantization Table Values 373
Figure 19-5 Sample Group of Pictures. 377
Figure 21-1 Stored Atomic Operations 396
Figure 21-2 Replacing Atomic Functions with a Molecule 397
Figure D-1 Cell. 464
Figure D-2 Encoding a Cell 466
Figure D-3 Default CellB Chrominance Quantization Table 476
Figure D-4 Default CellB Luminance Quantization Table 478

Tables

Table 0-1 Typographic Conventions xxix

Table 1-1 Finding Information About Image-Processing Functions. . . . 2

Table 1-2 Finding Information About Compression Techniques 4

Table 1-3 Finding Information About Additional XIL Functions 5

Table 1-4 Example Programs Provided with the XIL Developer’s Release 8

Table 2-1 Libraries Used in Linking XIL Programs 30

Table 2-2 Examples of Major and Minor Version Numbers. 31

Table 3-1 Functions for Reading an Image’s Width, Height, or Number of Bands 35

Table 3-2 Functions for Importing and Exporting Images 37

Table 3-3 Functions for Reading or Setting an Image’s Origin 47

Table 3-4 Functions Used to Build an ROI 50

Table 3-5 ROI Naming Functions. 51

Table 3-6 Parameters to `xil_create_child()` 53

Table 3-7 Image Type Utility Functions 55

Table 3-8 Image Naming Functions. 56

Table 4-1 Cases Handled by the Function `load_file()` 61

Table 5-1	Resolutions of Photo CD Images	77
Table 5-2	Photo CD Image Attributes	79
Table 5-3	Converting PhotoYCC Data to Another Color Space	84
Table 6-1	Source Files for `display`	87
Table 6-2	Cases Handled by the `display` Program	95
Table 7-1	Matching Source and Display Images	98
Table 7-2	Plane Masks for an Overlay	100
Table 7-3	Functions for Managing Colorcubes	114
Table 7-4	Utility Functions for Dither Masks	121
Table 7-5	Review of Dithering Operations	121
Table 7-6	Strings Used to Specify Color Spaces	123
Table 8-1	XIL Error Categories	130
Table 9-1	Valid Values for Each XIL Data Type	142
Table 9-2	Arithmetic Operations Using a Source Image and a Constant	146
Table 9-3	Logical Operations Using a Source Image and a Constant	147
Table 10-1	Types of Interpolation	150
Table 10-2	Constants in the Enumeration `XilFlipType`	175
Table 11-1	Parameters to `xil_histogram_create()`	180
Table 11-2	Additional Histogram Functions	182
Table 11-3	Parameters to `xil_soft_fill()`	186
Table 11-4	Handling Edges in a Convolution Operation	191
Table 11-5	Utility Functions for Convolution Kernels	192
Table 11-6	Utility Functions for Structuring Elements	199
Table 11-7	Parameters to `xil_lookup_create()`	200
Table 11-8	Additional Functions for Lookup Tables	206

Table 11-9 Parameters to xil_paint() 215
Table 12-1 Command-Line Options for xilcis_example 232
Table 12-2 Parameters to xil_cis_put_bits_ptr() 234
Table 13-1 Types of Images Supported by XIL Compressors 252
Table 13-2 Compressors and Compressor Types. 259
Table 13-3 CIS Naming Functions 266
Table 16-1 Image Bands and Huffman Tables 311
Table 16-2 Default Huffman Tables 318
Table 17-1 JPEG Lossless Prediction Methods 335
Table 18-1 Sizes of CIF- and QCIF-Format Images 338
Table 19-1 Two Command Sequences: CIS Pattern = IPBB 368
Table 19-2 Releasing a Frame: CIS Pattern = All I Frames 371
Table 19-3 Releasing a Frame: CIS Pattern = IPB. 372
Table 19-4 Releasing a Frame: CIS Pattern = IPBBPBB 372
Table 19-5 Characters Representing Picture Types 377
Table 21-1 Exceptions to the General Decompression-Molecule Rules . . 405
Table 21-2 Type of Flip Designated on the Call to xil_transpose() . 407
Table 21-3 Functions That Affect Deferred Execution 416
Table A-1 Key to Names Used in Molecule Definitions 418
Table B-1 XIL Library Error Messages 426
Table C-1 XIL-XGL Interoperability 460
Table D-1 Cell Bytestream Codes 473
Table D-2 Default Y/Y Table 480
Table D-3 Default C_b/C_r Table. 484

Preface

The XIL™ Imaging Library provides a set of key functions from the fields of image processing and digital video. The purpose of this book is to explain how to use these functions in developing application programming interfaces (APIs) and end-user applications.

Who Should Use This Book

The XIL library was designed to meet the needs of developers creating APIs and applications for a number of markets. These markets include:

- Markets that require digital video technology
- Commercial document imaging
- Technical document imaging
- Desktop publishing
- Color prepress
- Graphics arts
- Technical imaging

How This Book Is Organized

This bulk of this book is divided into three parts. Chapters 1 through 8 explain how to create the framework for an XIL program. They discuss such things as input, output, and error handling. Chapters 9 through 11 discuss the XIL library's image-processing functions, functions that would be used inside the framework. Chapters 12 through 20 explain how the XIL library enables you to compress and decompress sequences of digital images.

A chapter-by-chapter description of the book follows:

Chapter 1, "Introduction to the XIL Library," provides an overview of the functions in the XIL library and discusses how the library relates to other Sun, and to third-party, libraries. It also includes information about verifying the installation of the XIL software developer's kit and setting environment variables.

Chapter 2, "Basic XIL Program," introduces programming with the XIL library by looking at a simple XIL program that reads an 8-bit grayscale image from a file and displays it in an X window.

Chapter 3, "XIL Images," discusses the different types of XIL images, what attributes XIL images have, and how images are stored in memory.

Chapter 4, "Handling Input and Output," takes a systematic look at how you handle the reading and writing of images in an XIL program, including input/output (I/O) with files, displays, and other devices.

Chapter 5, "Reading Kodak Photo CD Images," discusses the library's device handler for reading and decoding the Eastman Kodak Company's Photo CD™ format.

Chapter 6, "Preparing Images for Display," discusses some of the issues that arise when you need to display different types of XIL images on different types of displays.

Chapter 7, "Presentation Functions," discusses a group of functions that are useful in preparing images for display. The topics covered in the chapter include dithering and color-space conversion.

Chapter 8, "Error Handling," discusses the XIL library's default error handler and explains how to write and install a custom error handler.

Chapter 9, “Arithmetic, Relational, and Logical Functions,” covers the XIL library’s arithmetic and logical functions. These functions enable you to add two images, take the logical AND of two images, multiply an image by a constant, and so on.

Chapter 10, “Geometric Functions,” discusses the XIL library’s geometric functions. These functions enable you to do such things as scale images, rotate images, and transpose images.

Chapter 11, “Miscellaneous Image Processing Functions,” presents the remaining image-processing functions in the XIL library. Among other things, these functions give you the ability to filter images, pass images through lookup tables, and dilate or erode images.

Chapter 12, “Compressing and Decompressing Sequences of Images,” introduces the subject of compressing and decompressing image sequences by presenting an example movie maker and an example movie player.

Chapter 13, “Compressed Image Sequences,” discusses the data structure (called a CIS) in which you store compressed image data. The chapter goes over basic CIS operations, such as creating a CIS and writing data to it, and also discusses what attributes CISs have and how you set and read the values of those attributes.

Chapter 14, “Cell Codec,” deals with the Cell compressor/decompressor. (The Cell bytestream definition was developed at Sun.) The chapter briefly explains how the compressor works, discusses attributes that are specific to a Cell codec, and explains how to call optimized routines to play back Cell-encoded movies.

Chapter 15, “CellB Codec,” discusses the library’s CellB codec, which was derived from the Cell codec for use in videoconferencing applications. The chapter explains how the codec works and what special attributes a CellB CIS has.

Chapter 16, “JPEG Baseline Sequential Codec,” discusses how the library’s JPEG baseline sequential codec works and what special attributes a JPEG CIS has

Chapter 17, “JPEG Lossless Codec,” explains how the JPEG lossless codec performs its job and presents information about attributes that are specific to a CIS associated with a JPEG lossless codec.

Chapter 18, “H.261 Codec,” covers the XIL interface to the codec specified by the CCITT in Recommendation H.261. The chapter discusses how such a codec works and lists the CIS attributes that apply specifically to a CIS associated with an H.261 compressor or decompressor.

Chapter 19, “MPEG-1 Codec,” discusses the XIL interface to the MPEG-1 codec specified by the Moving Pictures Expert Group. The chapter discusses how an MPEG-1 codec works and lists the CIS attributes that apply specifically to a CIS associated with an MPEG-1 compressor or decompressor.

Chapter 20, “CCITT Group 3 and Group 4 Codecs,” discusses how these document image compressors do their jobs and lists the CIS attributes that apply only to this class of codec.

Chapter 21, “Acceleration in XIL Programs,” explains how the XIL runtime system defers the execution of functions called in XIL programs as long as possible so that it can replace certain sequences of functions with optimized routines.

Appendix A, “XIL Molecules,” lists the *molecules* provided with the current release of the library. These molecules are optimized routines that the library can execute in lieu of executing a predefined sequence of functions from the API.

Appendix B, “XIL Error Messages,” provides a list of XIL error messages. For each message, the appendix specifies an error ID and a list of functions that can generate the error.

Appendix C, “XIL-XGL Interoperability,” explains how a single program can use both XIL and XGL™ functions to process an image.

Appendix D, “Cell and CellB Bytestream Definitions,” presents the information about Cell and CellB bytestreams that you would need to implement a Cell or CellB compressor or decompressor.

Appendix E, “Bibliography,” lists some books and articles to consult for further information on such subjects as image-processing operations, JPEG compression, dithering, and color models.

The book also contains a glossary of terms from the fields of image processing and digital video.

Related Books

The primary companion to this book is the *XIL Reference Manual*. The reference manual contains man pages for all the functions in the XIL library.

Because programming with the XIL library can be closely tied to programming with the X library, you may also find it useful to consult the *Xlib Programming Manual* and the *Xlib Reference Manual*.

What Typographic Changes Mean

The following table describes the type faces used in this book.

Table 0-1 Typographic Conventions

Typeface	Meaning	Example
`AaBbCc123`	The names of commands, functions, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `system% You have mail.`
`AaBbCc123`	What you type, contrasted with on-screen computer output	`system%` **`su`** `password:`
AaBbCc123	Command-line placeholder: replace with a real name or value	To delete a file, type `rm` *filename*.
AaBbCc123	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in *User's Guide*. These are called *class* options. You *must* be root to do this.

New Features

This document list the new features in XIL 1.2.

New Imaging Operations

XIL 1.2 contains new imaging operations that let you: store the absolute values of a source image's pixels; detect a source image's edges; and find the larger or lesser of pixel values in two source images. The new imaging operations are:

- `xil_absolute()`
- `xil_edge_detection()`
- `xil_max()`
- `xil_min()`

Using a Plane Mask for Copying Images

The `xil_copy()` function copies a source image to a destination image, copying all source-image *planes* (bits) to the destination. XIL 1.2 contains a new `xil_copy_with_planemask()`function, which lets you determine which source-image planes are copied to the destination. Copying with plane-mask control is typically used for overlaying images. It is also useful for double buffering on hardware that doesn't have separate memory buffers; double buffering is useful for animation because you can render an image in a hidden memory buffer while you display another image in a second buffer.

Multiband Lookup Tables

XIL 1.2 supports multiband images as input to `xil_lookup()`; earlier releases supported only single-band input images. With multiband input images, the `XilLookup` object contains a separate data array for each band in the input image; thus, it can be used to convert a multiband image of any data type to a multiband image of any data type.

The new multiband lookup functions are:

- `xil_lookup_create_combined()`
- `xil_lookup_get_input_nbands()`
- `xil_lookup_get_band_lookup()`

Interpolation Tables

When geometric transformations result in non-integer pixel locations, the XIL library provides nearest-neighbor, bilinear, and bicubic options for interpolating pixel values in a destination image; for these options, the interpolation filter is provided by the library. XIL 1.2 introduces a *general* interpolation option, for which your application defines the filter.

To support general interpolation, the XIL library provides a new `XilInterpolationTable` object, which is an array of 1xn kernels that represents the interpolation filter in either the horizontal or vertical direction. Interpolation tables let an application determine how many pixels are used to calculate interpolated pixel values, weight the pixels used in the calculations, and divide the space between adjacent pixels into multiple sub-locations so a destination pixel's value can be interpolated differently, depending on which sub-location it falls in. The interpolation tables are set on an `XilSystemState` object, and they affect all general interpolation operations using images created from this `XilSystemState`.

The following new functions support general interpolations:

- `xil_interpolation_table_create()`
- `xil_interpolation_table_get_data()`
- `xil_interpolation_table_get_kernel_size()`
- `xil_interpolation_table_get_subsamples()`
- `xil_interpolation_table_destroy()`
- `xil_state_get_interpolation_tables()`
- `xil_state_set_interpolation_tables()`

Warping an Image

The XIL library now contains three general functions that let you customize a geometric transformation by creating *warp* tables that displace pixels horizontally, vertically, or in both directions. The warp tables are XIL images whose pixel values define the backward mapping from a pixel in the destination to a pixel in the source.

Warp tables are typically used to stretch an image according to predefined rules and are most useful for performing nonlinear transformations. For example, they might be used to correct distortions that were imposed on an image by the equipment used to capture it. Or they might be used for cartographic projection of an image.

The new warping functions are:

- `xil_tablewarp()`
- `xil_tablewarp_horizontal()`
- `xil_tablewarp_vertical()`

XilDevice Objects

Typically when working with device images, an application creates the device image, then sets its attributes. Some device types, however, require that the device attributes be defined when the device image is created. For these device types, XIL 1.2 has a new `XilDevice` object that you can create and associate with the device type, then use to store device attributes. After storing all required attributes on the device object, you pass the object as an argument on the `xil_create_from_device()`function that creates the device image, thereby creating the device image and simultaneously setting its required attributes.

The device objects have the following new man pages:

- `xil_device_create()`
- `xil_device_destroy()`
- `xil_device_set_value()`

In addition, `xil_create_from_device()` has been modified to accept the device object as an argument on its last parameter.

Note – Device objects can be used only to initialize attributes before a device image is created. They cannot be used to modify the attributes of an existing device image. Devices that don't require initialized attributes may not recognize the device object; for these device types, you still pass `NULL` as the final argument on `xil_create_from_device()`.

Reading Kodak Photo CD Images

The XIL library includes a device handler that reads and decodes images stored in Eastman Kodak Company's Photo CD™ format. This feature was previously available as a patch to XIL 1.1.

Conditionally Compiling Code for Different XIL Versions

Beginning with release 1.2, the XIL library lets you conditionally compile code so you can take advantage of new interfaces in a current release while still supporting earlier releases of the library. To permit compile-time decisions, the library uses preprocessor directives to define symbols that identify the current XIL major and minor version numbers; these symbols—`XIL_API_MAJOR_VERSION` and `XIL_API_MINOR_VERSION`—are available to your application when you include the `xil.h` header file.

Shared Memory

XIL 1.2 supports the X Shared Memory Extension for displaying images with a new I/O device. Given an X window by `xil_create_from_window()`, the XIL library searches for an XIL I/O device for the window to associate with the image.

Previously, the XIL library used DGA to locate a device-specific I/O device, and if no device-specific I/O device could be loaded, it used an Xlib I/O device. For XIL 1.2, the XIL library tries to use a device-specific I/O device, and if no device-specific I/O device can be loaded, it attempts to load an XShm I/O device. If both the device-specific and the XShm I/O device fail, the Xlib I/O device is loaded. These changes were implemented to assure you of the following:

- The XIL library always chooses the fastest available way to display an image
- The display mechanism is network independent

Performance Enhancements

In addition to supporting the X Shared Memory Extension for speeding up image displays, many XIL functions have been optimized to improve the overall performance of XIL applications.

Installing XIL Packages

As discussed in the *Software Developer Kit Installation Guide*, the best way to install XIL packages is to use the `pkgadd/pkrm` commands. However, if you prefer, you can still use `swmtool`. You can no longer use the `swm` command to install XIL packages.

Introduction to the XIL Library 1

This introduction discusses three subjects:

- What functionality the XIL library provides
- How the library is meant to work with other Sun, and third-party libraries
- How to verify that the XIL software developer's kit (SDK) has been installed correctly and to set some environment variables that will make using the library easier

Functions in the XIL Library

The functions in the XIL library, which you call as C subroutines, fall into two main categories: image processing and image compression. The next two sections introduce these groups of functions, and a third section summarizes the remaining library functions.

Image-Processing Functions

The image-processing functions in the XIL library can be broadly grouped under the following headings: arithmetic, relational, and logical functions, geometric functions, and some miscellaneous functions. The arithmetic, relational, and logical functions include functions that enable you to add two images, take the maximum pixelwise values of two images, take the logical

AND of two images, multiply an image by a constant, and so on. The geometric functions include, among others, routines to scale, rotate, and transpose images. The miscellaneous functions enable you to:

- Find the minimum and maximum values in an image
- Produce a histogram for an image
- Threshold an image
- Fill an area in an image
- Filter an image
- Detect image edges
- Dilate or erode an image
- Pass an image through a lookup table
- Perform an interband linear combination
- Blend images
- Paint on an image
- Set and get the values of pixels in an image
- Copy a pattern to an image

Table 1-1 indicates where in this book to turn to find more information about the different classes of image-processing functions.

Table 1-1 Finding Information About Image-Processing Functions

Type of Image-Processing Functions	Discussed in This Chapter
Arithmetic, relational, and logical functions	Chapter 9, "Arithmetic, Relational, and Logical Functions"
Geometric functions	Chapter 10, "Geometric Functions"
Other functions	Chapter 11, "Miscellaneous Image Processing Functions"

Digital-Video Functions

The digital-video functions in the library enable you to compress and decompress images and sequences of images. Several compression formats are supported:

- Cell
- CellB
- JPEG baseline sequential
- JPEG lossless
- H.261

- MPEG-1
- CCITT Group 3 and Group 4

The Cell image compression technology, which was developed by Sun, has been optimized for the rapid decompression and display of images on simple hardware. Therefore, the Cell codec is able to achieve reasonable display quality on indexed-color frame buffers. The initial focus of the Cell technology is on Sun-to-Sun communications, where the benefits of fast decoding outweigh the benefits of standards. Possible areas of application include media distributions on CD-ROM and multimedia mail.

The CellB codec, which derives from its Cell counterpart, is intended for use primarily in videoconferencing applications. It features greater balance between the time spent compressing and decompressing images than the Cell codec. The CellB codec's strengths include

- Software compression at interactive rates
- Very fast decoding and display, especially on indexed-color frame buffers
- Low rates of CPU use
- Good quality output

The JPEG compression standards were developed by the Joint Photographic Experts Group to support the compression of still images, both grayscale and color. Although not specifically designed for the compression of sequences of images, or movies, JPEG compressors are also used frequently for that purpose. The JPEG baseline sequential compressor is a *lossy* compressor, which means that it compresses an image in such a way that when the compressed data is decompressed, the decompressed image and the original image may not match exactly. On the other hand, with a *lossless* JPEG compressor, the decompressed image does match the original image pixel for pixel.

The H.261 compression-decompression scheme was developed by the International Telegraph and Telephone Consultative Committee (CCITT). An H.261 video codec is intended to be used to compress and decompress video data sent over Integrated Services Digital Network (ISDN) lines. Thus, the codec is suitable for use in video telephony and videoconferencing applications. The current release of the XIL library includes an H.261 *decompressor*. A compressor can be obtained from a third party.

The MPEG-1 video compression standard was developed by the Moving Picture Experts Group. The group's goal was to compress full-motion video and the associated audio at the rate of about 1.5 Mbits/s. This is approximately the rate at which data can be read from a CD-ROM, so MPEG-1 compressed

video is a good choice for use in interactive multimedia applications. The current release of the XIL library includes an MPEG-1 *decompressor*. A compressor can be obtained from a third party.

The CCITT Group 3 and Group 4 compression standards were developed by the International Telegraph and Telephone Consultative Committee to enable facsimile machines to compress and decompress digitized documents. Now, Group 3 and Group 4 compressors and decompressors are also used for general document storage and retrieval.

Table 1-2 tells you where in this book to look for further information about compressing and decompressing images.

Table 1-2 Finding Information About Compression Techniques

Compression Type	Discussed in This Chapter
Cell	Chapter 14, “Cell Codec”
CellB	Chapter 15, “CellB Codec”
JPEG baseline sequential	Chapter 16, “JPEG Baseline Sequential Codec”
JPEG lossless	Chapter 17, “JPEG Lossless Codec”
H.261	Chapter 18, “H.261 Codec”
MPEG-1	Chapter 19, “MPEG-1 Codec”
CCITT Group 3 and Group 4	Chapter 20, “CCITT Group 3 and Group 4 Codecs”

Additional XIL Functions

The image-processing and digital-video functions mentioned above are the heart of the XIL library. However, the library also contains functions that perform the tasks listed in Table 1-3. Table 1-3 also indicates where in this book you can look for further information about the functions that perform these tasks.

Table 1-3 Finding Information About Additional XIL Functions

Additional Tasks You Can Perform	Discussed in This Section or Chapter
Copy an image to the display	"Copying an Image to the Display" on page 97
Rescale an image	"Rescaling an Image" on page 102
Cast an image from one data type to another	"Casting an Image from One Data Type to Another" on page 103
Dither an image	"Dithering an Image" on page 105
Convert an image from one color space to another	"Color-Space Conversion" on page 122
Perform undercolor removal	"Black Generation" on page 125
Handle XIL errors	Chapter 8, "Error Handling"

As you can see, the XIL library contains a broad range of functions. The next section explains why the XIL library was designed this way and how it relates to other Sun, and to third-party, libraries.

The XIL Library: A Foundation Library

The XIL library is what Sun calls a foundation library. Such a library is an application programming interface (API), but has special characteristics that distinguish it from other APIs. For example, a foundation library must deal with hardware dependencies, while most APIs are hardware independent. In addition, a foundation library is geared toward a broad area of application, such as imaging and video (the XIL library) or graphics (the XGL library), while most APIs have a narrower scope. For example, APIs based on the XIL library might address such areas as document imaging, facsimile applications, and digital video (see Figure 1-1).

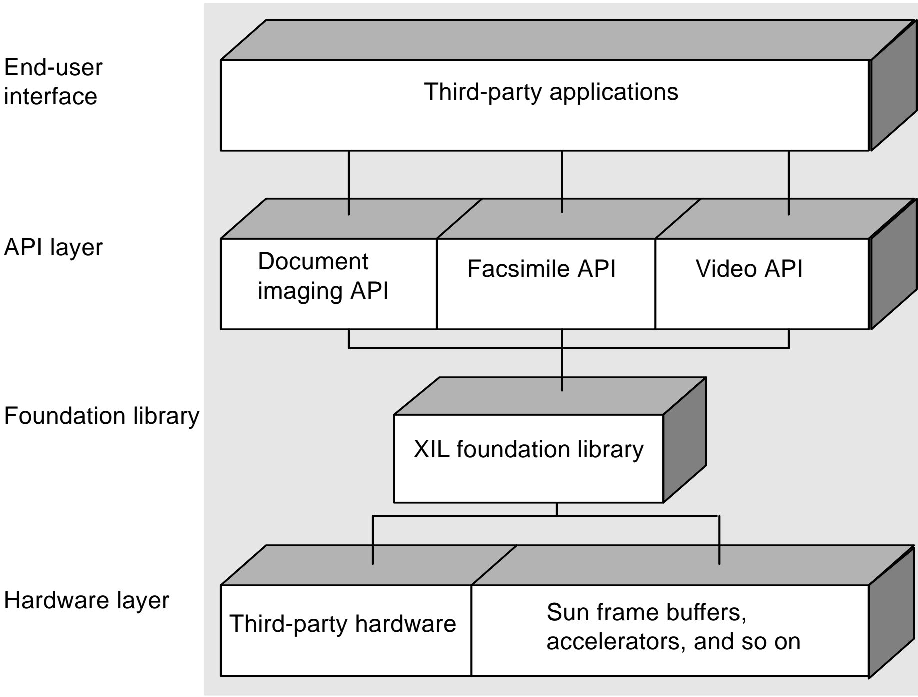

Figure 1-1 Foundation Libraries and Application Programming Interfaces

Why is this extra layer of software desirable? One reason is that a number of APIs—developed either by Sun or third parties—may require the same functionality. For instance, an API directed at the commercial document imaging market and another directed at the facsimile market both must be able to compress and decompress image data according to CCITT standards. If the code that handles this compression is part of a foundation library, the two APIs don't have to create two functions that do the same thing.

Second, a foundation library provides a single interface to the hardware with which the library interacts. For example, the XIL library provides an interface to video cards, scanners, imaging accelerators, frame buffers, and printers. Any applications or APIs written on top of the XIL library access these devices using XIL routines.

Third, a foundation library allows multiple APIs to share objects in a simple way. For instance, because XIL defines memory formats for different types of images, all APIs whose imaging functionality is based on XIL routines can operate on these images without having to convert the format of any image data.

Because an API may also need to provide capabilities that the XIL library does not support, it is also possible to *export* data from the XIL domain into *application space*. There, an application or API can do whatever processing of the data it needs to do. When this processing is complete, the data can be *imported* so that the program can process the data using only XIL functions.

The XIL Library is Multithread Unsafe

In this release, the XIL library doesn't support multithread programs and is therefore multithread-unsafe. Because this is true for all XIL functions, it is not stated in the man pages.

In the next release, the XIL library will support multithread programs. If you want to take advantage of that thread-enabled environment, you must link with the `thread` library when you compile your current XIL programs. Linking to the `thread` library when compiling a current XIL application won't provide any thread-enabled capabilities now, but it will ensure full binary compatibility with the thread-enabled environment later. Applications that don't link `-lthread` will run fine, but they won't take advantage of any thread support included in the release. Linking to the `thread` library now will not degrade your application's performance.

Note – "Building an XIL Program" on page 29 shows the Makefile delivered with one of the XIL example programs; the Makefile shows one way you can link with the `thread` library. The section also discusses the considerations you must make when linking to the library, depending on which compiler and debugger you use.

Verifying Installation and Setting Environment Variables

This section discusses

- The XIL packages that should have been installed
- What the XIL directory structure should look like

- The XIL environment variables you should set
- How to use the XIL man pages

XIL Packages

The default installation directory for the XIL SDK software package (SUNWxilh) is /opt/SUNWits/Graphics-sw/xil; the default installation directory for the XIL AnswerBook package (SUNWAxi) is /opt/SUNWAxi.In the SUNWxilh package, you will find the following subdirectories included at the top level: examples, include, and man. The content of these directories is described in the following sections.

examples/

The examples directory contains several examples that are discussed later in this book.

Table 1-4 Example Programs Provided with the XIL Developer's Release

Example Program	**What It Does**
example1	Reads a grayscale image from a file and displays the image in an Xlib™ window
display	Demonstrates how to display images of different data types and with different numbers of bands in an Xlib window
encode	Demonstrates how to create movies using XIL functions
xilcis_example	Demonstrates how to play back movies using XIL functions

include/

This directory has one subdirectory, xil, containing the include (*.h) files required by the XIL library.

`man/`

This directory has one subdirectory, `man3`, containing the on-line manual pages (accessible via the `man(1)` command) for the XIL library. These manual pages are also part of the XIL AnswerBook on-line documentation.

AnswerBook Package

The `SUNWAxi` package contains the XIL AnswerBook on-line documentation, including:

- *XIL Programmer's Guide*
- *XIL Reference Manual*

XIL Directory Structure

Figure 1-2 shows the directory structure of the XIL developer's kit. Directories shown in bold print are part of the SDK packages. The `lib` directory, containing the XIL Runtime Environment, is released on the Solaris™ CD-ROM disk.

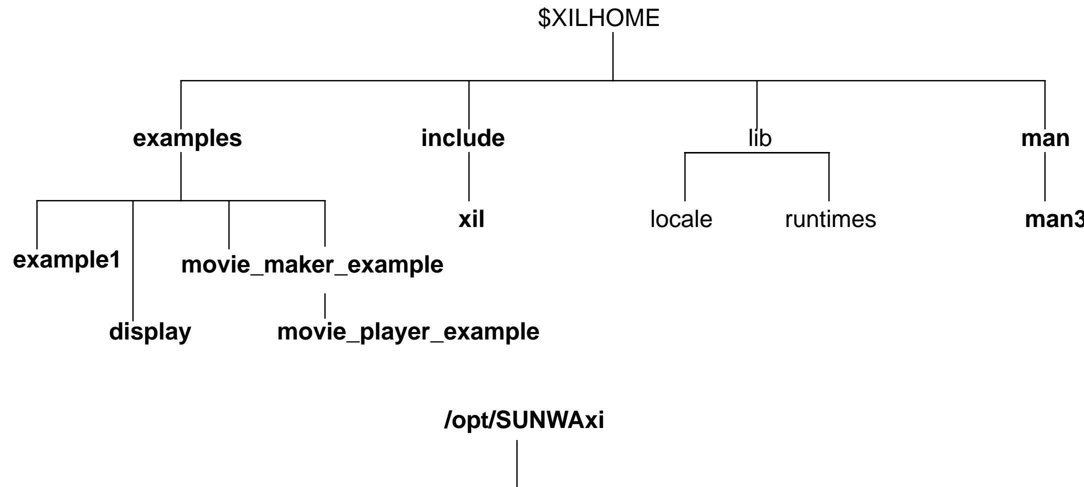

Figure 1-2 Directory Structure of XIL Developer's Kit.

Setting Environment Variables

You need to set the XILHOME and LD_LIBRARY_PATH environment variables as follows. (The example instructions assume that the XIL library is installed in the default directory.)

1. **Set the XILHOME environment variable to the directory containing the XIL libraries:**

```
% setenv XILHOME /opt/SUNWits/Graphics-sw/xil
```

2. **Set LD_LIBRARY_PATH to include $XILHOME/lib:**

```
% setenv LD_LIBRARY_PATH $XILHOME/lib:$LD_LIBRARY_PATH
```

Using the XIL Manual Pages

To access the XIL reference manual pages, you must modify the MANPATH environment variable to include the directory where the man pages are located. The optional catman(1M) command creates preformatted versions of the man pages and recreates the windex database.

1. **Set the MANPATH environment variable.**

   Assuming you have already set the XILHOME environment variable, set MANPATH as follows:

```
% setenv MANPATH $MANPATH:$XILHOME/man
```

2. **(Optional) Format the man pages.**

```
% su root
Password: <enter your root password>
# /usr/bin/catman $XILHOME/man 3
```

Basic XIL Program 2

To introduce programming with the XIL library, this chapter takes a look at a simple XIL program called `example1`. This program reads an 8-bit grayscale image from a file and displays it in an X window. The program also enables you to take the one's complement of the image displayed in the window, which produces a negative of the image.

The source files for this program—`example1.c`, `fileio.c`, and `window.c`—can be found in the directory `$XILHOME/examples/example1`. This directory also contains an image file that the program should use as input and a makefile you can use to build the program.

After briefly discussing how to build and run `example1`, this chapter walks you through the program's code, pointing out some of the features of the program that will be part of most or all XIL programs. The topics covered are:

- Including header files
- Initializing the library
- Acquiring an input image
- Creating an output image
- Processing the image
- Closing the library

The chapter also talks briefly about building XIL programs.

Running the Program

If you want to run example1 before reading the discussion of how it is written, follow these steps:

1. **Change your working directory to** $XILHOME/examples/example1**. (You should have set the environment variable** $XILHOME **when you installed the XIL software.)**
2. **Build the program by typing** make**.**
3. **Execute the program from the directory in which you built it by typing**

   ```
   % example1 brainscan.header
   ```

   At this point, you will see an X-ray of a brain displayed in an X window.
4. **To take the one's complement of the image in the window, move your pointer into the program's window and press any key on your keyboard. A second keypress will restore the image to its original state (by causing a second bitwise negation to be performed).**
5. **Exit the program by moving your pointer into the program's window and clicking any mouse button.**

The following sections provide an explanation of the XIL-related code in example1's source files.

Including Header Files

The first XIL-related item in example1.c is a preprocessor directive that includes the header file xil.h. All XIL programs must include this header file and, therefore, must include the line

```
#include <xil/xil.h>
```

This header file defines a number of symbolic constants and enumeration constants and contains function prototypes for all of the functions in the XIL library.

Initializing the Library

If you skip down to the main routine in example1.c, you'll note that after reading command-line arguments, the example calls the routine xil_open().

```
XilSystemState state;

state = xil_open();
if (state == NULL) {
    /* XIL's error handler will print an error msg on stderr */
    exit(1);
}
```

All XIL programs must call this function to initialize the library before using other functions from the XIL library.

As part of the initialization process, xil_open() returns a data structure of type XilSystemState called state. The library uses this data structure to keep track of information such as what XIL images (and other XIL data structures) have been created during the current session and what the program's error-handling routine is. You'll notice later in the program that state must be passed to any XIL function that creates an image.

Note – Since someone might want to use XIL routines and routines from an API built on top of the XIL library in the same program, it is possible to have multiple XilSystemState structures.

Acquiring an Input Image

All XIL applications process images in some way. As you will see later, these images can come from a variety of sources, and how you bring an image into an application depends on its source. The program example1, however, only considers the case in which a program gets its input, or *source*, image from a disk file. To read such a file, you must perform the series of steps outlined below:

1. Read the image file's header (or take whatever steps are appropriate) to determine the image's width, height, number of bands, and data type. You will need this information when you perform the next step.

2. Create an XIL image of the correct dimensions and data type. This image will serve as an empty container in which you can place the pixel values from your file.

3. *Export* the XIL image to obtain information about how the image's pixel values will be stored in memory.

   At all times, an image is either in *XIL space* or *application space*. While it is in XIL space (or *imported*), you can operate on the image only by using XIL functions; you do not know the address of the image or its pixel values. On the other hand, when the image is exported, you can get a pointer to the storage allocated to hold the image's pixel values and write data in that memory (or modify existing data). You can also change the pointer to the storage.

4. Read the pixel values from your file, and write them to the storage you received a pointer to in the preceding step. *It may be necessary to convert the format of this data as you copy it.* The XIL library supports images of several data types, and for each type, it specifies a memory format. These memory formats are discussed in the section "Exporting XIL Images" on page 36.

5. Import the XIL image so that all subsequent operations on the image must be performed using XIL functions.

The sections below discuss each of these steps in more detail and show how the steps are performed in `example1`. The code fragments shown are taken from the source file `fileio.c`.

Step 1: Reading the File's Header

As mentioned above, you need to read the header of your image file to determine the following image attributes:

- Width of the image in pixels
- Height of the image in pixels
- Number of bands, or channels, in the image
- Data type of the image

You need this information to create an XIL image in which to store the file's pixel values.

How you get this information from your file's header depends on the format of that header. The function `load_file()` shows you how this task was handled in one case. The header for the image file supplied for use with the example is shown below.

```
512 512 8 1
brainscan.data
```

The code the example uses to get the necessary image attributes from the header looks like this.

```
XilImage
load_file (XilSystemState state, char *filename)
{
   unsigned int width, height, nbands, bits;
   XilDataType datatype;
   FILE *in_file;
   char datafile[256];

   /* Open the header file */
   if ((in_file = fopen(filename, "r")) == NULL) {
      fprintf(stderr,"Could not open %s\n", filename);
      exit(1);
   }

   /* Get header information */
   fscanf(in_file, "%d%d%d%d", &width, &height, &bits, &nbands);
   fscanf(in_file, "%s", datafile);
   fclose(in_file);

   if ((bits == 8) && (nbands == 1)) {
      datatype = XIL_BYTE;
   }
   else {
     fprintf(stderr,"Program requires a 1 band, 8 bit image\n");
      exit(1);
   }
```

The routine `load_file()` is passed the name of the image file, opens the file, and reads the four integers in the first line of the header. These integers represent the width of the image, its height, the number of bits used to represent one band of one pixel, and the number of bands in the image.

You can pass the width, height, and number-of-bands arguments directly when you call the routine that creates your XIL image (in step 2). However, before making that call, you must use the number-of-bits value to set a variable of the enumeration type `XilDataType`. You set this variable to one of the following constants: `XIL_BIT`, `XIL_BYTE`, or `XIL_SHORT`. These enumeration constants correspond to the three types of data the XIL library supports: 1-bit data, 8-bit unsigned data, and 16-bit signed data. The variable is set in the line

```
datatype = XIL_BYTE;
```

Since this program expects a single-band, 8-bit image as input, an error occurs if `bits` does not equal 8.

Note – The other item read from the file is the name of the file that contains the actual pixel values. This file is opened and read in step 4.

Step 2: Creating an XIL Image

Once you have read the appropriate values from your image file's header, you can create your XIL image by calling `xil_create()`.

```
XilImage image;

/* Create the image to read the data into */
image = xil_create(state, width, height, nbands, datatype);
if (image == NULL) {
   /* XIL's error handler will print an error msg to stderr */
   exit(1);
}
```

With the exception of `state`, the parameters being passed to `xil_create()` are the variables set in the preceding step. The `state` parameter is the data structure returned by the call to `xil_open()` earlier in the program.

The return value of `xil_create()` is a handle to an XIL image (a data structure of type `XilImage`). This is the image you export in the next step.

Step 3: Exporting the Image

You now need to export the image you just created and to get some information about how the image's pixel values will be stored in memory. The example uses the following code to accomplish these tasks:

```
XilImage image;
XilMemoryStorage storage;
Xil_boolean status;

/* Get the pointer to the image data */
if (xil_export(image) == XIL_FAILURE) {
    /* XIL's error handler will print an error msg to stderr */
    exit(1);
}
status = xil_get_memory_storage(image, &storage);
if (status == FALSE) {
   /* XIL's error handler will print an error msg to stderr */
   exit(1);
}
```

The call to `xil_export()`, whose only parameter is a handle to the image, marks the image as exported. By exporting the image, the application gains access to information about the image that enables it to modify the image directly. For example, the application can get a pointer to the location at which the image's pixel values will be stored. For further information about exported images, see the section "Exporting XIL Images" on page 36.

The example program exports the image so that it can find out the address at which the image's pixel values will be stored, the number of bytes between pixels, and the number of bytes between scanlines. To get this information, the example calls the function `xil_get_memory_storage()`.

This function takes as its parameters a handle to your XIL image and the address of a union of structures named `storage`. This union is of the defined type `XilMemoryStorage`.

```
typedef struct {
    Xil_unsigned8 *data;
    unsigned int scanline_stride;
    unsigned long band_stride;
    unsigned char offset;
} XilMemoryStorageBit;

typedef struct {
    Xil_unsigned8 *data;
    unsigned long scanline_stride;
    unsigned int pixel_stride;
} XilMemoryStorageByte;

typedef struct {
    Xil_signed16 *data;
    unsigned long scanline_stride;
    unsigned int pixel_stride;
} XilMemoryStorageShort;

typedef union {
    XilMemoryStorageBit bit;
    XilMemoryStorageByte byte;
    XilMemoryStorageShort shrt;
} XilMemoryStorage;
```

Note – `Xil_unsigned8` is a data-type name created in `xil.h`. You should use this name (instead of `unsigned char`) whenever you need to get the size of a single pixel value in an 8-bit XIL image. There is a corresponding type name, `Xil_signed16`, for 16-bit data.

Thus, the example can refer to the address of the pixel values for the exported image by writing `storage.byte.data`. Similarly, it can specify the number of bytes between pixels and the number of bytes between scanlines by writing `storage.byte.pixel_stride` and `storage.byte.scanline_stride` respectively.

Note – Any time you plan to use an image's data, you must export the image by calling `xil_export()`, then get pointers to its data with `xil_get_memory_storage()`. These pointers to the data are valid only until the image is imported. For more information, see "Exporting XIL Images" on page 36.

Step 4: Copying Data from a File to Your XIL Image

The example now knows the location of the pixel values it wants to read (they're stored in `datafile`) and the address (`storage.byte.data`) at which to write those values. The principal remaining question is whether the format of the pixel data in the file matches the memory format for the XIL image. Since this example reads a single-band (grayscale) image, the formats will match. The XIL image containing 8-bit data requires that data be stored in a

pixel-sequential format, and the values in the file are stored in a pixel-sequential format. Given this situation, the example can copy the file data to the proper location in memory using the code shown below:

```
unsigned int width, height, nbands;
XilMemoryStorage storage;
FILE *in_file;
int size;

/* Open the data file */
if ((in_file = fopen(datafile, "r")) == NULL) {
   fprintf(stderr, "Could not open %s\n", datafile);
   exit(1);
}

/* Load the image data */
if ((storage.byte.pixel_stride == nbands) &&
      (storage.byte.scanline_stride == nbands * width)) {
   size = nbands * width * height;
   if (fread((char *)storage.byte.data, sizeof(Xil_unsigned8),
         size, in_file) != size) {
      fprintf(stderr, "Error reading data in %s\n", datafile);
      exit(1);
   }
} else {
   int h, w, n;
   Xil_unsigned8 *scanline = storage.byte.data;
   for ( h = 0; h < height; h++) {
      Xil_unsigned8 *row = scanline;
      for (w = 0; w < width; w++) {
         Xil_unsigned8 *pixel = row;
         for (n = 0; n < nbands; n++) {
            *pixel++ = fgetc(in_file);
         }
         row += storage.byte.pixel_stride;
      }
      scanline += storage.byte.scanline_stride;
   }
}
```

Note that if the pixel stride equals the number of bands in the image (that is, there are no unused bytes between pixels in the XIL image), the data in the file `datafile` can be copied to the image in a block. Otherwise, the example must

use the values of `storage.byte.pixel_stride` and `storage.byte.scanline_stride` to determine where to write each pixel value it reads from the file.

What if the example were reading an RGB image from a file and the 8-bit RGB values in the file were stored in band-sequential format? In that case, the example would need not only to copy the data from the file to an XIL image, but to convert the format of that data as well. See the example code below.

Note – This code assumes the file contains blue values, followed by green values, followed by red values. Because this image is being read from a file, this is a valid assumption. However, you can never make this assumption about an `XilImage` in XIL memory space; in memory, an `XilImage`'s format is arbitrary and must be accessed only through `xil_get_memory_storage()`.

```
unsigned int width, height, nbands;
XilMemoryStorage storage;
FILE *infile;
int i, h, w;

/* Load the image data */
for (i = 0; i < nbands; i++) {
    Xil_unsigned8 *scanline = storage.byte.data + i;
    for (h = 0; h < height; h++) {
        Xil_unsigned8 *pixel = scanline;
        for (w = 0; w < width; w++) {
            *pixel = getc(in_file);
            pixel += storage.byte.pixel_stride;
        }
        scanline += storage.byte.scanline_stride;
    }
}
```

Step 5: Importing the Image

Once you have read image data from a file into an XIL image, you should import the image so that all further operations on it must be performed using XIL functions. You import the image with a call to `xil_import()`:

```
XilImage image;

xil_import(image, 1);
```

Note – Importing an image is the opposite of exporting it. When you import an image, any pointer to the image's pixel values you had while the image was exported becomes invalid, so you can no longer modify the image directly. You must make all modifications using XIL functions. To access the data again, you must export the image again, and make another call to `xil_get_memory_storage()` to get new pointers to the data. For more information, see "Exporting XIL Images" on page 36.

The argument `image` is a handle to the image being imported, and the `1` is a change flag. This flag indicates that the image was modified while it was exported; in this case, it was filled with pixel values.

Creating an Output Image

Because `example1` displays its output, its output image must be a special type of image called a *display image*. This type of image is written to a screen's frame buffer and is displayed in an X window.

To create this display image, the example program:

1. Creates an X window the size of the image it wants to display.
2. Uses the function `xil_create_from_window()` to create a display image from the X window discussed in step 1. This routine turns the X window into an XIL image that can be used as a destination for an XIL operation.

The sections below discuss these steps in more detail.

Creating an X Window

Most of the Xlib code in this example appears in the function `open_window()`, which is defined in the source file `window.c`. This function:

- Establishes a connection with the X server
- Finds the best X visual to use when displaying a grayscale image on a particular display
- Creates the X colormap that will be used when the image is displayed
- Creates the X window in which the image will be displayed

To create the X window, the function calls `XCreateWindow()`.

```
Display *xdisplay;
Window xwindow;
XVisualInfo *vinfo;

xwindow = XCreateWindow(xdisplay, DefaultRootWindow(xdisplay),
    20, 20, width, height, 0, vinfo->depth, InputOutput,
    vinfo->visual, CWColormap, &setwinattr);
if (xwindow == NULL) {
   fprintf(stderr, "Unable to create window\n");
   return (0);
}
```

The return value of `XCreateWindow()` is an X window ID, which is used in the creation of the XIL display image.

Creating a Display Image

Before you can use an XIL function to display an image, you must create a display image. As mentioned earlier, to do this, you call the XIL function `xil_create_from_window()`. This routine turns an X window into a special type of XIL image so that the window can be used as the source or (as in this case) the destination image for an operation.

```
Display *xdisplay;
Window xwindow;
XilSystemState state;
XilImage display;

display = xil_create_from_window (state, xdisplay, xwindow);
```

This call returns a handle to the newly created display image. The parameters you pass to the routine are a handle to the system-state data structure returned by `xil_open()`, a pointer to a structure describing your display, and the ID of your X window.

At this point, the program is ready to do some image processing of the source image it read earlier.

Processing an Image

Once `example1` has obtained handles to a source and a destination image, it is ready to use XIL functions to process the source image and eventually write the processed image to the display. The functions used in this part of the program can be thought of as serving two purposes.

One purpose is to alter the source image so that it can be written to the destination image. For example, in the sample program, the source image contains 8-bit data, but the display image may just be 1 bit deep (if the screen on which the image is to be shown is monochrome). In this case, the source image must be converted to 1-bit data before it can be copied to the display image. This kind of conversion may require that you create an intermediate image.

The second purpose is to change the source image to alter its appearance in some way. For instance, you might darken the image or enhance the edges in the image. The `example1` program enables you to take the one's complement of its source image.

Making Source and Destination Images Compatible

The program `example1` handles only one type of source image: an 8-bit grayscale image. However, it handles several types of display images. It may write its output to:

- A 3-band `XIL_BYTE` display image whose X colormap is read only.
- A 1-band `XIL_BYTE` display image. There are actually two cases here because the image's X colormap (grayscale) may be read-write or read only.
- A 1-band `XIL_BIT` display image whose X colormap is read only.

In only one of the cases mentioned above can the source image be copied to the display image as is. The following sections briefly describe how each case must be handled.

Three-Band, Eight-Bit Destination—Read-Only Colormap

In this case, the source image contains one band, and the display image contains three. To solve this problem, `example1` creates a temporary 3-band `XIL_BYTE` image, called `retained_image`, and copies the source image to each band of the temporary image.

```
if (vinfo.class == TrueColor) {
    /* Copy the grayscale image into each of the bands */
    retained_image = xil_create(state, width, height, 3,
        XIL_BYTE);
    band0 = xil_create_child(retained_image, 0, 0, width, height,
        0, 1);
    band1 = xil_create_child(retained_image, 0, 0, width, height,
        1, 1);
    band2 = xil_create_child(retained_image, 0, 0, width, height,
        2, 1);
    xil_copy(image, band0);
    xil_copy(image, band1);
    xil_copy(image, band2);
}
```

Note that the example actually creates three band children of the temporary image and copies the source image to those children. Once this has been done, `retained_image` can be copied to the display image.

One-Band, Eight-Bit Destination—Read-Write Colormap

In this case, the program could write the source image directly to the display image. However, doing that would require that the program create and use a colormap with 256 entries. If such a colormap were installed, it might cause other X clients, including the window manager, to display their output in false colors.

To avoid this problem, the example creates an X colormap of 256 entries, but writes its grayscale ramp in the last 216 entries. The program then uses a call to `xil_rescale()` to ensure that all the values in the image to be displayed fall between 40 and 255.

One-Band, Eight-Bit Destination—Read-Only Colormap

In this case, the example can write the source image directly to the display image because the program's X colormap is a read-only grayscale ramp.

One-Band, One-Bit Destination—Read-Only Colormap

In this case, the example creates a single-bit XIL image (using `xil_create()`) and then uses the function `xil_error_diffusion()` to dither the 8-bit image to a 1-bit image.

Additional Processing

Besides processing the source image so that it can be written to the display, `example1` enables you to take the one's complement of the brain scan by pressing any key while your pointer is in the program's window. The code used to perform this operation is shown below:

```
Display *xdisplay;
XilImage retained_image, display;
unsigned int width, height;
XEvent event;

/* Refresh window as necessary till program terminates */
while (1) {
    XNextEvent(xdisplay, &event);
    if (event.xany.type == KeyPress) {
        xil_not(retained_image, retained_image);
        if (vinfo.class == GrayScale)
            xil_add_const(image, offset, image);
        xil_copy(retained_image, display);
    }
    else if (event.xany.type == Expose)
        xil_copy(retained_image, display);
```

If the example gets a keypress event, it calls `xil_not()` to take the one's complement of the source image. This single statement works for all cases except the case where the program is using a `GrayScale` X visual. In that case, the values in `retained_image` before the call to `xil_not()` fall in the range 40 to 255. After the call to `xil_not()`, they fall in the range 0 to 215. Therefore, before the one's complement is displayed, the example adds 40 to each value in the image using the function `xil_add_const()`.

You may want to experiment with removing from the example the code used to take the one's complement of the brain scan and replacing it with code that performs another image-processing operation. You might try any of the arithmetic and logical operations that work with one source image and a constant. These operations are discussed in Chapter 9, "Arithmetic, Relational, and Logical Functions." Or you might try replacing the logical negation with a geometric operation. The XIL geometric operations are the subject of Chapter 10, "Geometric Functions."

Closing the Library

Before exiting, all XIL programs should deallocate the memory associated with any XIL data structures that haven't already been destroyed. Because the system-state data structure keeps a list of all the data structures XIL has created (and not destroyed), you can destroy all of these data structures by calling `xil_close()`, which takes the system state as its only parameter.

```
XilSystemState state;

xil_close(state);
```

Note – If your program creates a display image and you do not destroy it with a call to `xil_destroy()`, you must close the XIL library (`xil_close()`) before you break your connection with the X server and display (`XCloseDisplay()`).

Building an XIL Program

The directory `$XILHOME/examples/example1`, which contains the source files for `example1`, also contains the Makefile you use to build the program. The Makefile's contents are shown below.

```
# You may get better performance with these 2.0.1 flags
# OPTFLAGS = -xcg89 -Wa,-cg92
CFLAGS = $(OPTFLAGS) -I$(XILHOME)/include -I$(OPENWINHOME)/include

LIBS = -L$(XILHOME)/lib -L$(OPENWINHOME)/lib -R/opt/SUNWits/Graphics-
sw/xil/lib:usr/openwin/libs

install := LIBS += -lxil -lX11 -ldl -ldga -lthread
debug_install := LIBS += -lxil -lX11 -ldl -ldga
debug_install := CFLAGS += -g

install: example1 brainscan.header brainscan.data
debug_install: example1 brainscan.header brainscan.data

debug: debug_install

example1: example1.o window.o fileio.o
        $(CC) -o example1 example1.o window.o fileio.o $(LIBS)

clean:
        rm -f a.out core *.o example1
...
```

The main thing to note here is the list of libraries being used to link the program. When you link an XIL program, you must use these libraries (with the possible exception of the X11 library) *in this order* in your makefile or on your command line.

Table 2-1 below provides a brief explanation of what each library is used for.

Table 2-1 Libraries Used in Linking XIL Programs

Library Name	Library Referred To
`-lxil`	The XIL library.
`-lX11`	The X library. If your program doesn't display any output in an X window, you can omit this library.
`-ldl`	The dynamic-linking library.
`-ldga`	The Direct Graphics Access library.
`-lthread`	The multithread library. Linking to this library now ensures full binary compatibility with the next XIL release, which will be thread-enabled. Applications that don't link with `-lthread` will run fine, but they won't take advantage of any thread support included in the next XIL release (see "The XIL Library is Multithread Unsafe" on page 7).

SPARC – The SPARCompiler™ C 3.0 or later compiler and SPARCworks™ C 3.0 or later debugger are the recommended tools for building XIL applications; each can be used to build and debug XIL applications that have been linked with `-lthread`. However, linking with `-lthread` doesn't work with the SPARCworks C 2.0.1 debugger (it does work with the SPARCompiler C 2.0.1 compiler); see the **Note** below.

---

x86 – The ProCompiler™ C 2.0.1 or later compiler and the ProWorks™ C 2.0.1 or later debugger are the recommended tools for building XIL applications. The ProCompiler C 2.0.1 compiler can compile XIL applications that link with `-lthread`, but linking with `-lthread` doesn't work with the ProWorks C 2.0.1 debugger; see the **Note** below.

Note – If your debugger doesn't work with `-lthread`, you may have to link without `-lthread` to debug your application (as shown on page 29), and then link with `-lthread` when the application is complete. Be aware, however, that compiling without `-lthread` generates a different binary file than compiling with `-lthread`. Thus, a feature that works fine in the binary file that isn't linked with `-lthread` may not work in the binary file that is linked with `-lthread`; or, the feature may not work in the binary file that is linked whereas it does work in the binary file that isn't linked.

Conditionally Compiling Code for Different XIL Versions

Beginning with release 1.2, the XIL library lets you conditionally compile code so you can take advantage of new interfaces in a current release while still supporting earlier releases of the library. To permit compile-time decisions, the library uses preprocessor directives to define symbols that identify the current XIL major and minor version numbers; these symbols—`XIL_API_MAJOR_VERSION` and `XIL_API_MINOR_VERSION`—are available to your application when you include the `xil.h` header file (see "Including Header Files" on page 12).

By convention, the major version number is the numeral preceding the decimal point in the XIL version number, and the minor version number is the numeral following the decimal point. Table 2-2 shows what the major and minor version numbers would be for various XIL releases.

Table 2-2 Examples of Major and Minor Version Numbers

XIL Version Number	Major Version Number	Minor Version Number
1.2	1	2
1.3	1	3
2.0	2	0
2.1	2	1

By using these symbols, you can maintain a single source-code file that can be compiled against different versions of the XIL library. For example, version 1.2 of the XIL library introduced `XilDevice` objects, which can be used to initialize device attributes before creating the corresponding device image; earlier versions of the library can't use this object. The following code fragment

shows how you might use the XIL_API_MAJOR_VERSION and XIL_API_MINOR_VERSION symbols to write code that uses an XilDevice object when you compile your application with the 1.2 version of the library, but that omits the object when you compile with earlier versions:

```
#if (XIL_API_MAJOR_VERSION == 1 && XIL_API_MINOR_VERSION >= 2) ||
    (XIL_API_MAJOR_VERSION > 1)
{
  /*create a device object */
  XilDevice device;
  device = xil_device_create(state, "SUNWrtvc");
  /* initialize the DEVICE_NAME attribute */
  xil_device_set_value(device, "DEVICE_NAME", (void *)devname);
  /* create a device image with the initialized attribute */
 rtvc_image = xil_create_from_device(state, "SUNWrtvc", device);
  xil_device_destroy(device);
}
#else
  /* create a device image without initialized attributes */
  rtvc_image = xil_create_from_device(state, "SUNWrtvc",
                                        (void *) devname);
#endif
```

In this example, a device image with an initialized attribute is created only for XIL versions 1.2 and later; otherwise, the device image is created without initialized attributes.

See "New Features" on page xxxi for a discussion of the new features introduced in the current release of the library.

XIL Images 3

Because all XIL programs work with XIL images, you should not only know how to create an XIL image, but understand exactly what an XIL image is and what its chief characteristics are. When you first think about a computer image, you probably think of a picture or scene that has been digitized and is represented as a two-dimensional array of numbers or vectors. For example, the upper-left corner of an 8-bit, single-band image might look like this:

255	255	255	255	255	255	255	255	255	255	255	255
255	254	254	254	254	254	254	254	254	254	254	254
255	254	253	253	253	253	253	253	253	253	253	253
255	254	253	252	252	252	252	252	252	252	252	252
255	254	253	252	251	251	251	251	251	251	251	251
255	254	253	252	251	250	250	250	250	250	250	250
255	254	253	252	251	250	249	249	249	249	249	249
255	254	253	252	251	250	249	248	248	248	248	248

Figure 3-1 Digitized Image

One component of an XIL image *is* a set of pixel values like those shown above. However, an XIL image is more than that. It is a data structure containing many members, one of which is a pointer to a set of pixel values.

The upcoming sections explain some of the key things to understand about XIL images, such as:

- The most basic attributes of an XIL image
- Exporting XIL images so that an application can process pixel values using non-XIL functions
- The memory formats used for XIL images
- Special types of images, such as display images

You should understand the concepts presented in these sections before you proceed too far with your XIL programming.

The last section in this chapter completes the discussion of XIL images by talking about the image attributes that haven't already been discussed.

Basic XIL Image Attributes

When you use the function `xil_create()` to create an XIL image, you must specify four attributes of the image being created: its width, height, number of bands, and data type.

Width, Height, and Number of Bands

These three attributes define the dimensions of an image. An image's width and height are given in pixels. Its number of bands is the number of values needed to describe a single pixel; for example, in a grayscale image, one value describes the gray level of each pixel, so the image is said to have one band. In an RGB image, three values (red, green, and blue) are required to describe a single pixel, so such an image has three bands. Each band in a multiband image must have the same width and height.

Each of these image attributes is stored as an unsigned 32-bit integer. In this release of the XIL library, the possible values for each attribute are 0 to 65,535.

One other note. These three attributes are set when you create an XIL image and cannot be changed afterwards. The functions you use to read the values of these attributes are shown in Table 3-1.

Table 3-1 Functions for Reading an Image's Width, Height, or Number of Bands

Function Name	What the Function Does
xil_get_width	Gets the width of an image in pixels
xil_get_height	Gets the height of an image in pixels
xil_get_size	Gets the width and height of an image in pixels
xil_get_nbands	Gets the number of bands in an image

Data Type

This image attribute specifies the type of data used to represent the value of one band of one pixel. The XIL library supports three data types: 1-bit data, 8-bit unsigned data, and 16-bit signed data. You establish an image's data type when you create the image by passing to the `xil_create()` function one of the following enumeration constants: `XIL_BIT`, `XIL_BYTE`, or `XIL_SHORT`. These enumerators are of type `XilDataType`.

Note – IEEE single-precision floating-point images are defined and may be created, destroyed, imported, and exported, but image-processing operations on them are not implemented. To create a floating-point image, you specify a data type of `XIL_FLOAT`.

The library enables you to convert an image from one supported data type to another by providing the function `xil_cast()`.

To read an image's data type, you use the function `xil_get_datatype()`. The XIL library also includes the function `xil_get_info()`, which reads an image's width, height, number of bands, and data type.

3

Exporting XIL Images

When you first create an XIL image using `xil_create()`, you receive a handle to the image, and the image is in a state called *imported*. While the image is in this state, you can process it only with XIL functions because you don't know, and can't find out, the address at which the image's pixel values are stored. Nor do you know the exact layout of those values in memory.

However, you do have the option of *exporting* the image, using the function `xil_export()`. Once an image is exported, you can obtain a complete description of how its pixel values are stored in memory using the function `xil_get_memory_storage()`. This information enables you to process the image in ways that the XIL library does not support. You can also continue to process the image using XIL functions while it is exported; however, you may lose some acceleration.

While an image is exported, you can also *set* its memory storage using the function `xil_set_memory_storage()`. This involves setting a pointer to the image's pixel values and providing other information such as the distance between pixels and the distances between scanlines. (The exact values you set depend on the data type of the image. See the section "Memory Formats for XIL Images" on page 37 for more information on this subject.) Defining a image's memory storage in this way would be appropriate, for example, if you had exported an empty image and wanted to fill it with pixel values by setting a pointer to some pixel values residing in system memory.

After you've finished the work you needed to perform with the image exported, you should import the image using the function `xil_import()`. As mentioned above, while imported, the image must be processed using XIL functions.

Caution – The description of memory storage that you read using `xil_get_memory_storage()` or set using `xil_set_memory_storage()` is valid only as long as your image remains exported. Once you import the image, both the address at which the image's pixel values are located and their layout in memory may change. Trying to access pixel values using an invalid pointer or invalid information about their layout can cause serious problems in your application.

The XIL functions related to importing and exporting XIL images are shown in Table 3-2.

Table 3-2 Functions for Importing and Exporting Images

Function Name	What the Function Does
xil_import	Imports an image so that it must be processed using XIL functions
xil_export	Exports an image so that you can get or set the image's pointer to its pixel values
xil_get_exported	Determines whether an image is currently exported
xil_get_memory_storage	Reads information about how an image is stored in memory (if the image is exported)
xil_set_memory_storage	Defines how an image's pixel values will be stored in memory (if the image is exported)

Memory Formats for XIL Images

The memory format for an XIL image depends primarily on its data type. Images containing 8-bit or 16-bit data elements are stored in a pixel-sequential format, and images containing 1-bit data elements are stored in a band-sequential format.

XIL_BYTE and XIL_SHORT Images

Figure 3-2 illustrates how a 3-band image with 8-bit data elements might be stored in memory.

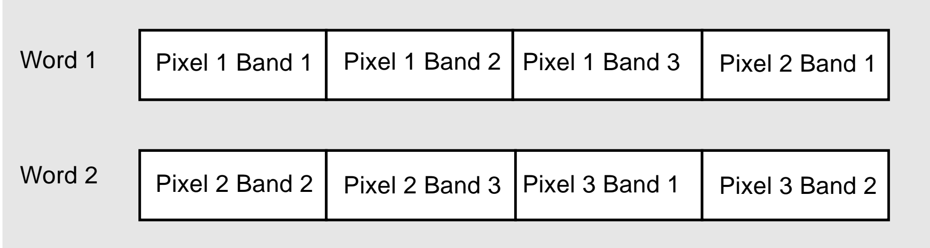

Figure 3-2 Memory Format for a 3-Band Image Containing 8-Bit Data Elements

Similarly, Figure 3-3 shows how a 3-band image with 16-bit data elements might be stored.

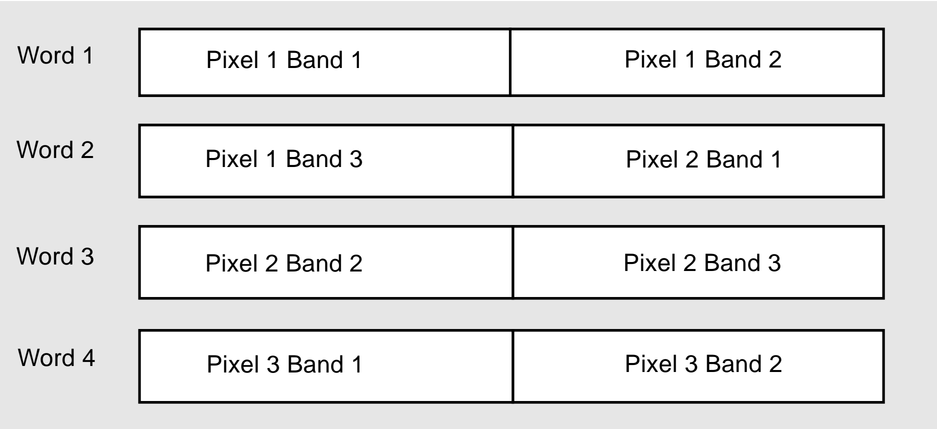

Figure 3-3 Memory Format for a 3-Band Image Containing 16-Bit Data Elements

It is important to realize, however, that there may be unused bytes between pixels and between scanlines.

As long as you're processing an image using XIL functions, you don't need to know the details about how the image is stored in memory. However, there may be occasions when you want to process an image directly. To do this, you

must export the image and then call `xil_get_memory_storage()` to obtain the details about how the pixels values are stored. This function returns a union of structures of the type defined below.

```
typedef struct {
    Xil_unsigned8 *data;
    unsigned int scanline_stride;
    unsigned long band_stride;
    unsigned char offset;
} XilMemoryStorageBit;

typedef struct {
    Xil_unsigned8 *data;
    unsigned long scanline_stride;
    unsigned int pixel_stride;
} XilMemoryStorageByte;

typedef struct {
    Xil_signed16 *data;
    unsigned long scanline_stride;
    unsigned int pixel_stride;
} XilMemoryStorageShort;

typedef union {
    XilMemoryStorageBit bit;
    XilMemoryStorageByte byte;
    XilMemoryStorageShort shrt;
} XilMemoryStorage;
```

The structures that contain information about `XIL_BYTE` and `XIL_SHORT` images both contain an element called `data`. This is a pointer to the beginning of the pixel data. The structures also contain elements called `pixel_stride` and `scanline_stride`. For an `XIL_BYTE` image, these elements indicate the distance in bytes between pixels and between scanlines, and for an `XIL_SHORT` image, they indicate the distance in 16-bit units between pixels and between scanlines.

This information about an image's memory storage remains valid only as long as your image is exported. Once you reimport the image, the location of your image's pixel values in memory and the layout of those values may change (except that the pixels will always remain in a pixel-sequential order).

XIL_BIT Images

Unlike images containing 8- and 16-bit data elements, multiband images containing 1-bit data elements are stored in a band-sequential format. In addition, scanlines are padded to the nearest byte boundary if necessary. For instance, Figure 3-4 shows how a 3-band, 3-by-3 image containing 1-bit data elements might look in memory.

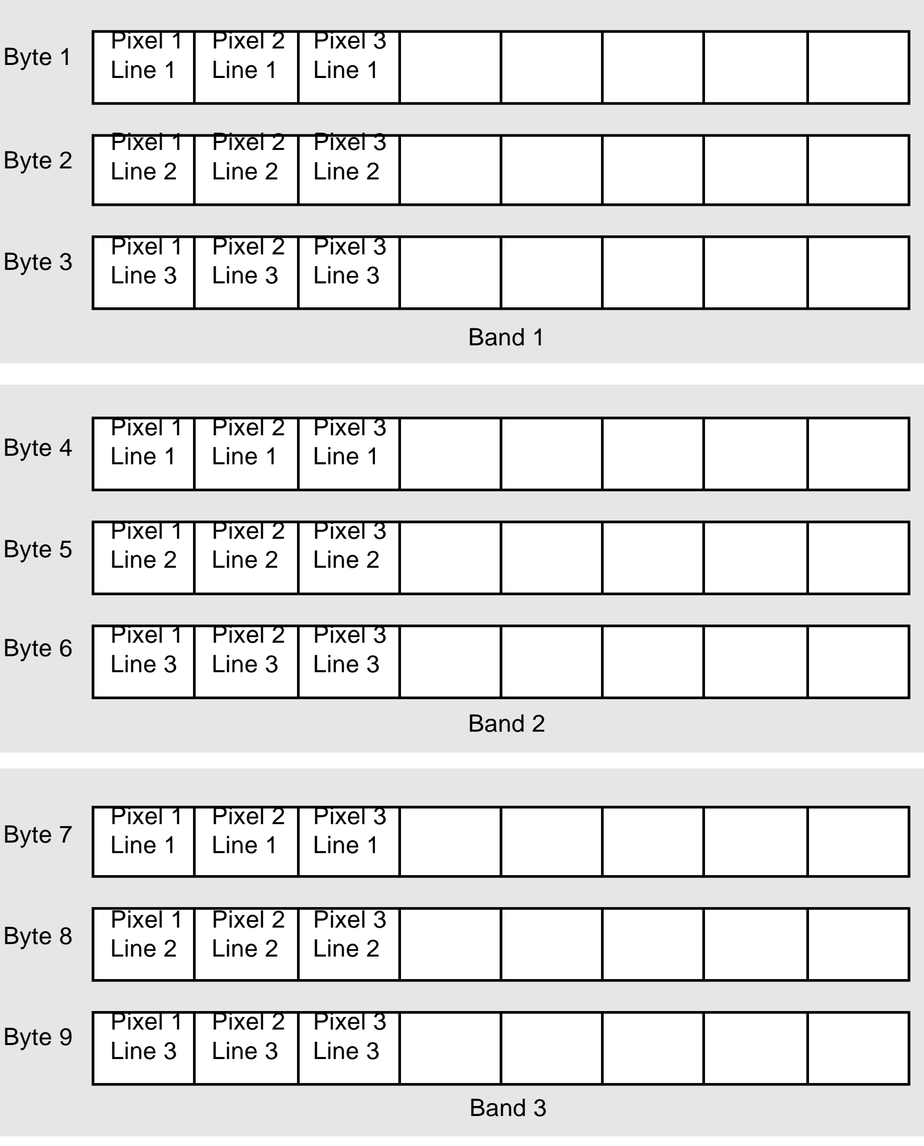

Figure 3-4 Memory Format for a 3-Band Image Containing 1-Bit Data Elements

To obtain a complete description of how an `XIL_BIT` image is stored in memory, however, you must export the image and call `xil_get_memory_storage()`. The function will give you access to the structure members shown below.

```
typedef struct {
    Xil_unsigned8 *data;
    unsigned int scanline_stride;
    unsigned long band_stride;
    unsigned char offset;
} XilMemoryStorageBit;
.
.
.
typedef union {
    XilMemoryStorageBit bit;
    XilMemoryStorageByte byte;
    XilMemoryStorageShort shrt;
} XilMemoryStorage;
```

The member `data` is a pointer to the first byte of pixel data. The member `scanline_stride` indicates the number of bytes from the first byte of one scanline to the first byte of the next, and `band_stride` indicates the number of bytes from the beginning of one band to the beginning of the next. The member `offset` specifies the number of ***bits*** from the beginning of a scanline to the first bit that contains a pixel value.

Remember that the values of all these structure members remain valid only as long as the image is exported.

Memory Formats for Images of Different Color Spaces

Besides an image's data type, its color space can play a part in how the image is stored. Consequently, you need to be aware of the ordering of bands for multiband images of different color spaces, such as RGB, YC_bC_r, CMY, and CMYK. For an RGB image, blue values are stored in the first band, green values in the second, and red values in the third. For YC_bC_r images, luminance (Y) values are stored in the first band, C_b values in the second, and C_r values in the third. For CMY and CMYK images, cyan values are stored in the first band, magenta values in the second, and yellow values in the third. For CMYK images, black values are stored in a fourth band.

Types of XIL Images

There are two basic types of XIL images:

- Memory images
- Device images

As you will see in the upcoming section “Device Images,” device images can be further subdivided into images that represent displays and images that represent other devices. Images that represent displays are referred to as *display* images. The main difference between these two subtypes of device image is the way in which they are created.

Memory Images

The memory image is the most common type of XIL image. When you create this type of image, using the function `xil_create()`, system memory is allocated to store both the image structure that describes the image and the image’s pixel values. All operations that do not read or write a device use memory images exclusively.

Device Images

XIL device images represent devices such as displays, cameras, frame grabbers, scanners, and printers. For each device image, the structure describing the image is stored in system memory, but the image’s pixel values are read from or written to the device. To create a device image, you generally pass the name of a device to the function `xil_create_from_device()`. The exception to this rule is that to create a display image, you call the function `xil_create_from_window()`.

You can use a device image as either a source image or a destination image in an XIL operation. When you use a device image as a source image, data is read from the device; similarly, when you use a device image as a destination image, data is written to the device.

The handler for each device is responsible for moving XIL images to or from the device. The current release of the XIL Imaging Library includes a group of these handlers that support standard devices.

Display Images

Display images are a special subclass of device images. They are very valuable because they eliminate the need for the scenario shown in Figure 3-5.

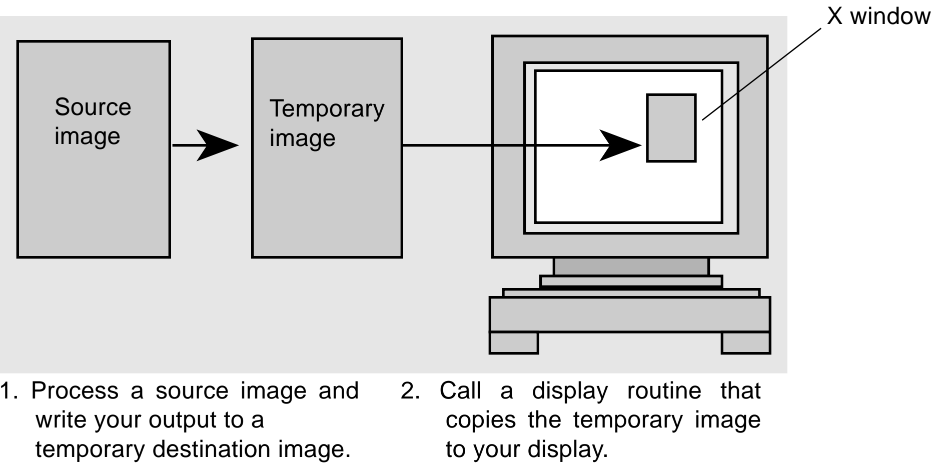

Figure 3-5 Copying Data from a Temporary Image to the Display (Not Necessary)

Display images enable the XIL library to process a source image and write the results of the processing directly to an X window by using a display image as the destination image for the operation. In other words, functions can write directly to frame-buffer memory, so the need for the temporary image is eliminated (see Figure 3-6).

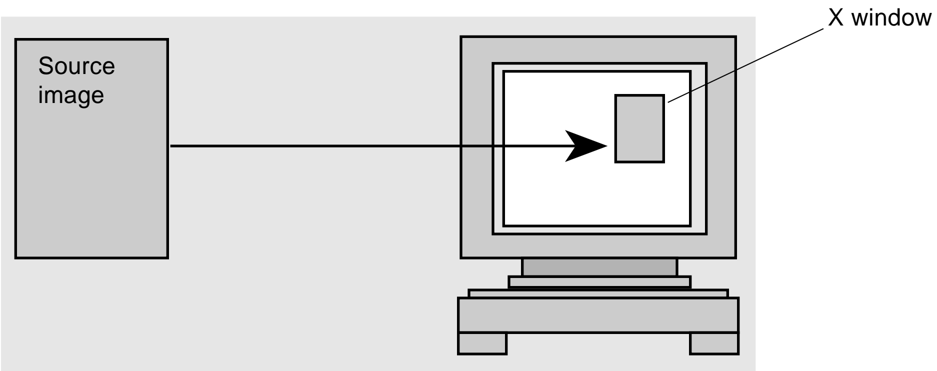

1. Process a source image and write your output directly to a display image.

Figure 3-6 Writing the Output of an Operation to a Display Image

To create a display image, you first create an X window. Then, you call the function `xil_create_from_window()`, passing information to it about which display and which window to associate with the display image.

Note – A display image cannot be created from an X window associated with an 8-bit `StaticColor`, `TrueColor`, or `DirectColor` visual.

A display image can also serve as a source image in most cases (if the display is readable). Any pixels that are obscured by another window when a read takes place are set to 0.

There are some restrictions on the types of images that can be display images. The only possibilities are:

- 1-band images containing 1-bit data elements
- 1-band images containing 8-bit data elements
- 3-band images containing 8-bit data elements

Additional XIL Image Attributes

The remainder of this chapter discusses some additional image attributes:

- Origin
- Region of interest
- Color space
- Parent
- Image type
- Synchronization flag
- Readable and writable flags
- Name

Note – It is possible for applications to create additional image attributes. To create and set an attribute, you use the function `xil_set_attribute()` and to retrieve the value of an application-defined attribute, you use the function `xil_get_attribute()`.

Origin

An image's *origin* is a pair of floating-point numbers that define a point in the image's coordinate plane. By default, an image's origin is in its upper-left corner (0.0, 0.0), but you can change the origin using the function `xil_set_origin()`. When an operation is performed, the origins of the source image or images and the destination image are aligned. (The floating-point origin values are rounded to integers for this purpose.) Then, for all nongeometric operations, the intersection of the source and destination images determines which pixels in the destination image will be modified. See Figure 3-7.

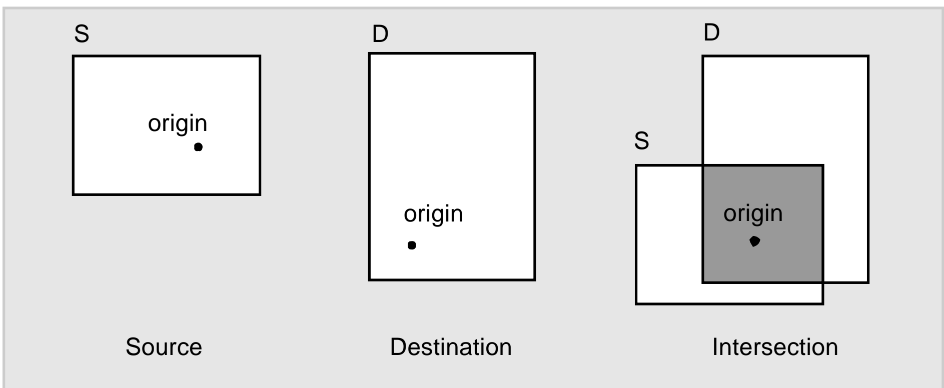

Figure 3-7 Image Origins

Only the shaded area shown in Figure 3-7 is modified in the destination image, and only the shaded area in the source is used by the operator.

Table 3-3 lists the XIL functions you use to read or set an image's origin.

Table 3-3 Functions for Reading or Setting an Image's Origin

Function Name	**What the Function Does**
`xil_get_origin`	Gets both coordinates of an image's origin
`xil_get_origin_x`	Gets the *x* coordinate of an image's origin
`xil_get_origin_y`	Gets the *y* coordinate of an image's origin
`xil_set_origin`	Sets an image's origin

Region of Interest

An XIL region of interest (ROI) is a data structure of type `XilRoi` that describes a single-bit mask for an image. If the region of interest is an attribute of a destination image, it determines which pixels in the destination may be written. Where there are 0's in the ROI, the destination image may not be modified, and where there are 1's, it may be modified. If the ROI is an attribute of a source image, it determines which pixels may be used as input to an operation.

Since an ROI is an attribute of an image, as opposed to a parameter for an operation, two or three ROIs may be involved in a single operation; for example, the operation's source image may have one ROI, and the operation's destination image another. In this case, the ROI used for the operation is the intersection of the ROIs associated with the source and destination images, as shown in Figure 3-8.

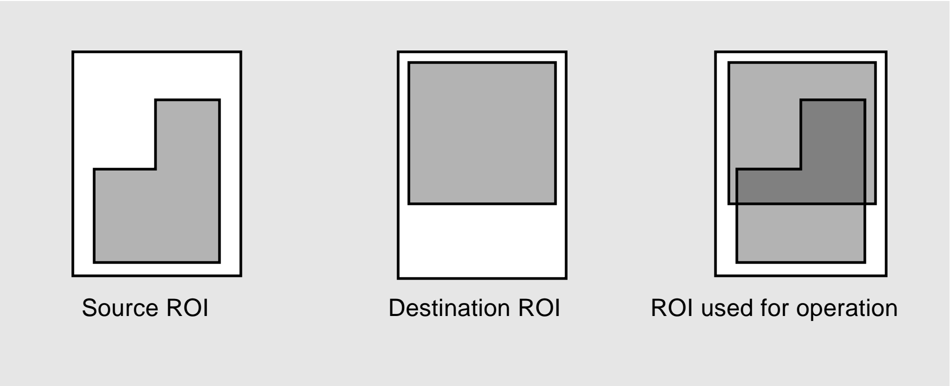

Figure 3-8 Regions of Interest

The dark gray area in the image on the right represents the ROI used for the operation.

If the origin of the source or destination has been set to something other than 0.0, the origins are first aligned; then, the ROI is determined. See Figure 3-9.

Figure 3-9 Regions of Interest and Origins

The dark gray area represents the ROI used for the operation.

Creating and Destroying an ROI

There are several ways to create an ROI.

- You can create a new ROI data structure using the function `xil_roi_create()`, which returns a handle to the ROI. When first created, the ROI is empty. For information on how to set values in the ROI, see the next section "Building an ROI."
- You can create a copy of an existing ROI using the function `xil_roi_create_copy()`.
- You can get a copy of the ROI associated with a image using the function `xil_get_roi()`.
- You can create an ROI by taking the union or intersection of two existing ROIs. The functions you use for these operations are `xil_roi_unite()` and `xil_roi_intersect()`.

When you are finished with an ROI, you destroy it using the function `xil_roi_destroy()`.

Building an ROI

After creating an empty ROI, you can build a region of interest using the functions shown in Table 3-4.

Table 3-4 Functions Used to Build an ROI

Function Name	What the Function Does
`xil_roi_add_rect`	Adds a rectangle to the ROI. You specify the width and height of the rectangle and its coordinates.
`xil_roi_subtract_rect`	Subtracts a rectangle from an ROI.
`xil_roi_add_image`	Adds an `XIL_BIT` image to the ROI. Bits that are set in the image are added to the ROI.
`xil_roi_add_region`	Adds an X region to an ROI.

Note – It's also possible to convert an ROI to a single-bit XIL image or an X region. To perform these jobs, you use the functions `xil_roi_get_as_image()` and `xil_roi_get_as_region()`, respectively.

Setting an Image's ROI Attribute

Once you have created an ROI and established its contents, you make it an attribute of an image using the function `xil_set_roi()`. You can get a copy of the ROI associated with a particular image by using the function `xil_get_roi()`.

Translating an ROI

The library includes a function `xil_roi_translate()` that moves all the pixels in an ROI that have been set to 1 left or right and/or up or down. The prototype for this function is shown below.

```
XilRoi xil_roi_translate(XilRoi roi, int xoffset, int yoffset);
```

The parameter `roi` is the ROI whose pixels you want to translate. The parameters `xoffset` and `yoffset` are integers that represent the number of pixels the ROI should be moved horizontally and vertically. If `xoffset` is positive, the ROI is moved to the right, and if it is negative, the ROI is moved to the left. If `yoffset` is positive, the ROI is moved down, and if it is negative, the ROI is moved up.

The function returns a handle to the translated ROI.

Naming an ROI

The XIL library enables you to specify a string that will serve as the name for an ROI. This naming is useful because it enables you to later get a handle to an ROI by using its name. The functions that allow you to perform these tasks are listed in Table 3-5.

Table 3-5 ROI Naming Functions

Function Name	What the Function Does
`xil_roi_set_name`	Sets the name of an ROI
`xil_roi_get_name`	Returns a copy of an ROI's name
`xil_roi_get_by_name`	Returns a handle to the ROI that has the name you specify

Note – When you call `xil_roi_get_by_name()`, you are returned a pointer to the named object, not a copy of it. Therefore, you should not destroy an ROI obtained in this way.

Color Space

A newly created image does not have a color space associated with it, but you can assign one to it using the function `xil_set_colorspace()`. The possible color spaces are:

- CCIR Rec. 709 RGB
- A linear version of CCIR Rec. 709 RGB
- CCIR Rec. 709 YC_bC_r
- A luminance-only space derived from CCIR Rec. 709 YC_bC_r
- A linear version of the luminance-only space mentioned above
- CCIR Rec. 601 YC_bC_r
- A luminance-only space derived from CCIR Rec. 601 YC_bC_r
- A YC_bC_r color space defined by Kodak for PhotoCD
- A linear CMY
- A linear CMYK

Images are usually assigned color spaces so that an image can be converted from one color space to another (using the `xil_color_convert()` function). For example, if you had an RGB image and wanted to convert it to the YC_bC_r color space, you would follow this procedure:

1. Set the color-space attribute of your RGB image to one of the RGB color spaces listed above.

2. Create an image whose width, height, number of bands, and data type match those of the RGB image.

3. Set the new image's color-space attribute to one of the YC_bC_r color spaces listed above.

4. Use the function `xil_color_convert()` to convert the RGB data from the first image to its YC_bC_r counterpart and write the converted data to the newly created image.

For more information about color-space conversion, see the section "Color-Space Conversion" on page 122.

Parent

The XIL library enables you to create a child image (subimage) from an image (the child's parent). Such a child image has an attribute that identifies its parent image. To retrieve this attribute, a handle to the parent image, use the function `xil_get_parent()`.

To create the child image, you use the function `xil_create_child()`, whose prototype is shown below.

```
XilImage xil_create_child(XilImage src, unsigned int xstart,
    unsigned int ystart, unsigned int width, unsigned int height,
     unsigned int startband, unsigned int numbands);
```

This function returns a handle to the child image. The parameters to the function are defined in Table 3-6.

Table 3-6 Parameters to `xil_create_child()`

Parameter	How It Is Used
`src`	A handle to the parent image
`xstart`	The horizontal offset in pixels from the upper-left corner of the source image to the upper-left corner of the child
`ystart`	The vertical offset in pixels from the upper-left corner of the source image to the upper-left corner of the child
`width`	The width of the child image in pixels
`height`	The height of the child image in pixels
`startband`	The offset in bands from the first band in the parent to the first band in the child
`numbands`	The number of bands in the child image. These bands must match contiguous bands in the parent.

Note – A child image does not get a copy of part of the parent's pixel data, but a pointer to that data.

If overlapping but not coincident sibling images (children of the same parent) are specified as the source and destination for an operation, the operation is performed. However, the library generates a warning message and the results of the operation are undefined.

Note – An exception to this behavior is `xil_copy()`, which detects the overlap and correctly generates the destination image.

In general, you read the attributes of a child image using the same functions you use to read the attributes of a parent. However, the library does contain a routine `xil_get_child_offsets()` that is designed for use only with child images. This routine returns the values of `xstart`, `ystart`, and `startband` that were used in creating the child image.

```
void xil_get_child_offsets(XilImage child,
    unsigned int *offset_x, unsigned int *offset_y,
    unsigned int *offset_band);
```

Image Type

For every XIL image, there is a subset of attributes that constitutes its image type. These attributes are its width, height, number of bands, data type, and color space. If you call the function `xil_get_imagetype()` and pass it a handle to an image, the routine will return a data structure of type `XilImageType` that contains only these data elements. You can use this structure to create an image of the same type as the one whose image type you just ascertained. To create an image in this way, you use the function `xil_create_from_type()`.

The XIL library also contains a number of utility functions that affect image types. These are listed in Table 3-7.

Table 3-7 Image Type Utility Functions

Function Name	What the Function Does
xil_imagetype_get_width	Gets the width of an image type in pixels
xil_imagetype_get_height	Gets the height of an image type in pixels
xil_imagetype_get_size	Gets the width and height of an image type in pixels
xil_imagetype_get_nbands	Gets the number of bands in an image type
xil_imagetype_get_datatype	Gets the data type of an image type
xil_imagetype_get_info	Gets the width, height, number of bands, and data type of an image type
xil_imagetype_set_name	Sets the name of an image type
xil_imagetype_get_name	Returns a copy of an image type's name
xil_imagetype_get_by_name	Returns a handle to an image type that has the name you specify

Synchronization Flag

XIL operations that affect images may be deferred similar to the way in which Xlib client requests are buffered before being sent to the X server. However, if an image's synchronization flag is set, operations on that image will not be deferred. You set this flag using the function `xil_set_synchronize()`. You can determine whether the flag is currently set by calling the function `xil_get_synchronize()`.

Readable and Writable Flags

Memory images can always be read and written to; however, certain device images may be read-only or write-only. To determine whether a device image is readable, you use the function `xil_get_readable()`.

```
Xil_boolean xil_get_readable(XilImage device_image);
```

This function returns TRUE if the device image is readable.

To determine whether a device image is writable, you use the function xil_get_writable).

```
Xil_boolean xil_get_writable(XilImage device_image);
```

This function returns TRUE if the device image is writable.

Name

The library enables you to assign a name (char *) to an image. This type of naming is useful because it enables you to get a handle to an image later in your program by specifying the name of the image. The functions that allow for the naming of images are shown in Table 3-8.

Table 3-8 Image Naming Functions

Function Name	What the Function Does
xil_set_name	Sets the name of an image
xil_get_name	Returns a copy of an image's name
xil_get_by_name	Returns a handle to the image that has the name you specify

Handling Input and Output 4

This chapter takes a systematic look at how you handle the reading and writing of images in an XIL program. Basically, input can come from, and output can go to, three places:

- A file
- A display
- A device such as a scanner, frame grabber, video card, or printer

The following sections explain the different possibilities.

Note – This chapter deals exclusively with reading and writing single images. As we'll discuss in Chapter 12, "Compressing and Decompressing Sequences of Images," XIL programs can also read and write data streams that represent compressed video or multipage documents. These data streams are stored in an XIL data structure called a compressed image sequence. Moving data into and out of these structures is discussed in Chapter 12.

Moving Image Data from a File to an XIL Image

The XIL library does not include any single functions that you use to read image data from a particular type of image file into an XIL image. Instead, the library provides routines that enable you to use the following general procedure to load images from files.

1. **Read the image file's header (or take whatever action is appropriate) to determine the dimensions and data type of the image.**
   Specifically, you need to know the image's width in pixels, its height in pixels, the number of bands in the image, and the type of data used to represent the value of one band of one pixel. You need this information to create an XIL image in which you can store the image's pixel values.

2. **Create an XIL image in which to store the pixel values you will read from the image file.**
   To create your XIL image, you use the function `xil_create()`, whose function prototype is shown below.

```
XilImage xil_create(XilSystemState state, unsigned int width,
    unsigned int height, unsigned int nbands,
    XilDataType datatype);
```

   The parameter `state` is the system-state data structure that was returned when you called `xil_open()` to initialize the library. The parameters `width`, `height`, and `nbands` are integers (`unsigned int`) representing the image's width, height, and number of bands. The final parameter `datatype` is an enumeration constant of type `XilDataType`, whose definition is shown below.

```
enum XilDataType { XIL_BIT, XIL_BYTE, XIL_SHORT };
```

   The constant `XIL_BIT` indicates that the image contains 1-bit data; the constant `XIL_BYTE` indicates that the image contains unsigned 8-bit data; and the constant `XIL_SHORT` indicates that the image contains signed 16-bit data. These are the only data types supported by the XIL library.

   When the image has been created, the library returns a handle to the XIL image.

3. **Export the XIL image you just created.**
   You export the image using the function `xil_export()`, whose function prototype is shown below.

```
int xil_export(XilImage image);
```

   You pass this routine the handle to the image you created in the last step. The return value—`XIL_SUCCESS` or `XIL_FAILURE`—indicates whether the export operation was successful.

   The reason for exporting the image is to gain access to information about how the image is stored in memory. This information is not available to the application while the image is imported.

4. **Get information about how your XIL image's pixel values are stored in memory.**
   To obtain this information, you call the function `xil_get_memory_storage()`, whose function prototype is shown below.

```
Xil_boolean xil_get_memory_storage(XilImage image,
    XilMemoryStorage *storage);
```

The function's return value indicates whether the routine succeeded (`TRUE`) or failed (`FALSE`). The parameter `image` is a handle to an XIL image. The argument `&storage` is a pointer to a union of structures defined like this:

```
typedef struct {
    Xil_unsigned8 *data;  /* pointer to first byte of image */
    unsigned int scanline_stride;  /* number of bytes between scanlines */
    unsigned long band_stride;  /* number of bytes between bands */
    unsigned char offset;  /* number of bits to the first pixel */
} XilMemoryStorageBit;

typedef struct {
    Xil_unsigned8 *data;  /* pointer to first byte of image */
    unsigned long scanline_stride;  /* number of bytes between scanlines */
    unsigned int pixel_stride; /* number of bytes between pixels */
} XilMemoryStorageByte;

typedef struct {
    Xil_signed16 *data;  /* pointer to first 16-bit word of image */
   unsigned long scanline_stride;  /* no of 16-bit words between scanlines */
    unsigned int pixel_stride;  /* number of 16-bit words between pixels */
} XilMemoryStorageShort;

typedef union {
    XilMemoryStorageBit   bit;
    XilMemoryStorageByte  byte;
    XilMemoryStorageShort shrt;
} XilMemoryStorage;
```

The structure you should look at for information depends on the data type of your image. If it is an `XIL_BIT` image, look at the structure `bit`; if it is an `XIL_BYTE` image, look at the structure `byte`; and if it is an `XIL_SHORT` image, look at the structure `shrt`. The elements in the appropriate structure will provide you with the information you need to write your image to memory in the proper location and in the proper format.

5. **Read the pixel values stored in your image file and write them to memory using the information obtained in the preceding step.**
   How you perform this step depends on several factors:

   - The type of data in your image: 1-bit, 8-bit, or 16-bit.
   - The organization of the data elements in your image file. For multiband images, the elements may be arranged in a pixel-sequential or band-sequential format.
   - The organization of the memory allocated to hold your XIL image. The information you obtained in step 4 describes this organization.

   One of the examples supplied with the XIL library contains a function called `load_file()` that illustrates how to handle this step for the cases listed in Table 4-1.

Table 4-1 Cases Handled by the Function `load_file()`

Data Type	File Format	Memory Layout for XIL Image
`XIL_BIT`	Band sequential	Any
`XIL_BYTE`	Pixel sequential	Any
`XIL_SHORT`	Pixel sequential	Any

You can find this routine in the directory `$XILHOME/examples/display`.

6. **Import your XIL image.**
   You import that image using a call to `xil_import()`, whose function prototype is shown below.

```
void xil_import(XilImage image, Xil_boolean change_flag);
```

   The parameter `image` is a handle to an XIL image, and `change_flag` is a flag that indicates whether the image was modified while it was exported. If the image was modified, the flag should be set to 1; otherwise, it should be set to 0.

Caution – After you import an image you give up control of it, and the pointers you received for it from a previous `xil_export()` are invalid. If you need to use an image's data after the image has been imported, you must call `xil_export()` again to export the image, and then call `xil_get_memory_storage()` again to get new pointers to the image data. After accessing the data, be sure to get the pixel stride and scanline stride when accessing pixel values.

For example, although an RGB image will always contain blue values, followed by green values, followed by red values, the data layout in XIL memory is arbitrary; there may be any number of pixels between the next set of blue, green, and red values. Access the pixel address of the first blue value, or the pixel stride to the next blue value, or the scanline stride to the next scanline, through the values returned by `xil_get_memory_storage()`. For more information, see "Exporting XIL Images" on page 36 and "Memory Formats for XIL Images" on page 37.

Moving Image Data from an XIL Image to a File

Just as the XIL library doesn't contain single functions that enable you to read particular types of image files, it does not include single functions that enable you to write images to files in particular formats. Instead, the XIL library provides a general mechanism for writing image data to files. The procedure to follow is discussed below.

1. **Determine the dimensions and data type of your XIL image.**
   You may already know the following attributes of your XIL image: width, height, number of bands, and data type. But if you don't, you should use the function `xil_get_info()` to obtain that information. You will need to know these values to determine how much image data needs to be written to the file.

   The function prototype for `xil_get_info()` is shown below.

```
void xil_get_info(XilImage image, unsigned int *width,
    unsigned int *height, unsigned int *nbands,
    XilDataType *datatype);
```

   The parameter `image` is a handle to an XIL image. The remaining parameters are the addresses of the variables in which you want the function to return the information you asked for.

   You may want to read other attributes of your XIL image depending on what type of information about the image you must write to your image file. See the section "Additional XIL Image Attributes" on page 46 for a list of image attributes you may be interested in.

2. **Export the image.**
   You export the image using the function `xil_export()`. Again, the reason for exporting the image is to gain access to information about how the image is stored in memory. For example, for an `XIL_BYTE` image, you need to know the address at which the image data begins, the number of bytes between pixels, and the number of bytes between scanlines.

3. **Get information about how your image's pixel values are stored in memory.**
   You get this information by calling the function `xil_get_memory_storage()`. If you haven't read the section "Moving Image Data from a File to an XIL Image" on page 57, see that section for an explanation of what information this function provides.

4. **Write the appropriate header information and pixel values to your file.**
   At this point, you should have all the information that you plan to write to your file's header, a pointer to your XIL image's pixel values, and information about the layout of those values in memory. So your only remaining XIL-related concern should be "Do I want to write the pixel

values to my file in the format in which they are stored in memory?" To answer this question, you need to know how pixel values for XIL images are stored in memory (see the section "Memory Formats for XIL Images" on page 37) and what kind of organization of data your file format demands.

Sending Output to (and Reading Input from) the Display

If you want to display an image in an XIL program, you must display that image in an X window. To perform this task, you first create an X window of the correct size. You then call the function `xil_create_from_window()` to create an XIL display image. This function turns an X window into a legitimate destination image for an XIL operation. You can then call an XIL function such as `xil_rotate()` and name the display image as the function's destination image. If you do, a rotated image will be written to your display's frame buffer.

The steps listed below discuss this procedure in a bit more detail:

1. **Use the Xlib function** `XCreateSimpleWindow()` **or** `XCreateWindow()` **to create your X window.**
   Before you create this window, you need to know—at a minimum—the width and height of the image you want to display. If you don't have this information, call the function `xil_get_info()`. It gives you information about an image's width, height, number of bands, and data type. You can then call one of the Xlib functions mentioned above to create the window.

   For information about using `XCreateSimpleWindow()` and `XCreateWindow()`, see the *Xlib Programming Manual* and the *Xlib Reference Manual*.

Note – A display image cannot be created from an X window associated with an 8-bit `StaticColor`, `TrueColor`, or `DirectColor` visual.

---

2. **Create an XIL display image.**
   As mentioned above, you create this display image by calling the function `xil_create_from_window()`, whose function prototype is shown below:

```
XilImage xil_create_from_window (XilSystemState state,
    Display *display, Window window);
```

   The parameter `state` is the system-state data structure that was returned when you called `xil_open()` to initialize the XIL library. The parameters `display` and `window` are of data types defined in Xlib. The `display` parameter is a pointer to a structure that is returned when you initially connect to the X server. It contains information about the server and the screens the server controls. The parameter `window` is the ID of the X window you created in the previous step. The return value of `xil_create_from_window()` is a handle to your newly created display image.

   Once you have created this display image, you can use it as a destination image for XIL operations just as you would an XIL memory image. When you name the display image as your destination image, the operation's output is written directly to your display's frame buffer and displayed in the X window you created in step 1.

3. **Call an XIL function that processes a source image and writes its output to the display image.**
   The simplest thing to try out here is to use `xil_copy()` to copy a source XIL image to the display. This routine takes two arguments: a handle to the source image and a handle to the display image.

Note – To see a simple example of a program that displays its output in an X window, look in the directory `$XILHOME/examples/example1`. The relevant code can be found in the source files `example1.c` and `window.c`.

To resize a window that contains an `XilImage`, destroy the `XilImage` attached to the window, resize the window, wait for a `ConfigureNotify` event to ensure the `XResizeWindow()` is complete, then call `xil_create_from_window()` to recreate the image in the new window size.

You cannot use an X window's `backing_store` attribute to maintain an image in the window when the window is obscured or unmapped (see the *Xlib Programming Manual*). Thus, your code should always check for an `Expose` event and take the appropriate measures for displaying the image again when the window is exposed.

Possible Complications

One question we ignored in the preceding discussion is "What if the depths of your source image and your X window are different?" What if you have a source image that has three bands and contains 8-bit data (24 bits per pixel) and an X window that is 8 bits deep? The answer to this particular question might be that you call an XIL function to dither the source image from 24 bits to 8 bits before writing it to the window. But other types of mismatches can occur. Chapter 6, "Preparing Images for Display," discusses an example program that handles many of these mismatches.

Another concern is "What happens if the currently installed X colormap is not appropriate for the image you want to display?" The answer this time is that you must solve this problem using Xlib calls. Specifically, you need to create and install an X colormap suitable for displaying your image. For complete information about how to create and install an X colormap, see the *Xlib Programming Manual* and the *Xlib Reference Manual*. For examples of functions that handle this task, look in the directory `$XILHOME/examples`, which contains a number of sample programs. For example, you might be interested in looking at the following routines:

- `window.c`, which is located in the subdirectory `example1`
- `window.c` (a modified version of the routine mentioned above), which is located in the subdirectory `display`
- `xilcis_color.c`, which is located in the subdirectory `movie_player_example`.

Reading a Display Image

Although display images are generally used for output, you can also read a display image. Assuming that you have created such an image, you can read the data in the X window associated with it into an XIL image by taking these steps:

1. **If you don't already have an XIL image that you can read the image data in the frame buffer into, you must create one.**
   You create this image using the function `xil_create()`. The image you create must have the proper width, height, number of bands, and data type.

2. **Perform an XIL operation using the display image as your source image.**
   For example, you could use the function `xil_copy()` to copy an image from your screen to an XIL memory image. Any pixels in the on-screen image that are obscured by another window when the read occurs are cleared to 0.

Reading and Writing Devices Other than Displays

The XIL library enables you to get input from, and write output to, not only displays, but other devices such as scanners, frame grabbers, video cards, and printers. The interface to these devices is very simple. You call a single XIL function to link a device to an XIL device image, which for most practical purposes is treated like an XIL memory image. The only special thing about the device image is that when you perform an XIL operation and your source image is a device image, the operation reads its input from the associated device. Likewise, if an operation's destination image is a device image, output is written to the associated device.

Each device type has attributes associated with it. For example, the Photo CD reader that comes with the XIL library (see Chapter 5, "Reading Kodak Photo CD Images") has a FILEPATH attribute that indicates the Photo CD image you want to read, and a RESOLUTION attribute that specifies a display resolution.

Depending on the device, you either need to create the device image, then set its attributes, or initialize the device attributes *before* creating the device image. The method to use depends on the device. Typically, devices don't require initialized attributes, so you create the device image, then set its attributes. However, some devices do require initialized attributes; this generally occurs when:

- The device needs one or more attributes defined at the time the device is created.
- Multiple attributes are interdependent and need to be set simultaneously for the device.

- Setting the attributes requires you to allocate a substantial block of memory before creating the image.

The group that writes the device handler must indicate whether the device requires initialized attributes.

Initializing a Device's Attributes

As mentioned in the previous section, some device types require you to initialize one or more device attributes before creating an image for the device. To initialize a device's attributes in the XIL library, you create an XIL data structure of type `XilDevice`. Once this device object exists, you use it to store the device-initialization values. You then pass the device object as an argument on the `xil_create_from_device()` function that creates the device image, as discussed in "Creating a Device Image" on page 71; this applies to the device image all the attributes stored by the device object.

Note – Devices that don't require attribute initialization typically don't recognize or support device objects. For these devices, you can't use a device object to initialize device attributes, and you pass `NULL` for the `deviceObj` parameter on the `xil_create_from_device()` function.

After you're done using the device object, you need to destroy it to release the memory allocated to it. You can destroy it immediately after using it, or you can keep the object around for creating other devices of the same type, then destroy the object when you are finally done with it.

The following steps discuss the device object in more detail. These steps can be used only for devices that require or can recognize initialized device attributes.

1. **Create the device object.**
   To create a device object, call the function `xil_device_create()`, whose function prototype is shown below.

```
XilDevice xil_device_create(XilSystemState state,
    char *devicename);
```

The parameter `state` is the system-state data structure that was returned by `xil_open()` when you initialized the XIL library. The `devicename` parameter is the name of the device to be associated with the device image. This name must be provided by the group that writes the device handler that enables the device in the XIL library.

For example, say you plan to write an image to a video card. You might associate a device object with it as shown in the following code fragment:

```
XilSystemState state = xil_open();
XilDevice      deviceObj;

deviceObj = xil_device_create(state, "vidCard");
```

In this example, the device's name is `vidCard`. This name is provided by the group the writes the device handler. The object `deviceObj` is now associated only with a `vidCard` device type; the object cannot subsequently be associated with a different device type.

The device object's only use is to initialize device attributes when you call the `xil_create_from_device()` function that creates the device image. It cannot be used to adjust a device image's attributes after the device image is created; `xil_set_device_attribute()` does that, as discussed on page 72. However, after using the device object to create one device image, you can use the same object to store different initialization attributes, then use the modified device object when you create another device image of the same type.

2. **Set the device-initialization values.**
   To store device-initialization values in a device object, call the function `xil_device_set_value()`, whose function prototype is shown below.

```
void xil_device_set_value(XilDevice deviceObj,
    char *attribute, void *value);
```

The parameter `deviceObj` is the device object associated with the device type. The parameter `attribute` is the name of the attribute you want to set, and `value` is the attribute's value. As with device names, the attribute names and their possible values are defined by the group that writes the device handler.

Only attributes the device understands should be set on the device object. Setting attributes the device doesn't recognize generates an error. Attributes and their associated values may reference data in the application's data space; therefore, any data associated with an `XilDevice` object must remain valid while it is referenced by the device object.

To set multiple attributes, make a separate `xil_device_set_value()` function call for each; although the attributes are set individually on the device object, they are applied simultaneously to the device image when that device image is actually created (see "Creating a Device Image" on page 71). You can set as many attributes as you need to derive all required initialization attributes for the device.

For example, if you created `deviceObj` for a video card as shown in Step 1 on page 69, you might initialize its volume and speed attributes as shown in the following code fragment:

```
int vol_control = 4, speed_control = 5;

xil_device_set_value(deviceObj,
                        "VOLUME", (void*) vol_control);
xil_device_set_value(deviceObj,
                         "SPEED", (void*) speed_control);
```

3. **Create the device image.**
   Do this only after setting all needed device attributes. The next section tells you how to create the image.

4. **Destroy the device object.**
   You can do this immediately after creating the device image. However, if you plan to create other devices of the same type, you may want to keep the device object around to initialize those other devices. If desired, you can set different initialization values on the object. When you're done using it, destroy the object; "Destroying a Device Object" on page 73 tells you how.

Creating a Device Image

As explained in "Initializing a Device's Attributes" on page 68, some device types require that one or more device attributes be initialized before the device image is created; for those devices, initialize the required attributes before following the steps in this section. For devices that don't support initialized attributes, the first step for using the device is to create a device image for it. Once the device image is created, you can set its required attributes, then use the image in an XIL operation.

The following steps show you how to create a device image and set its attributes.

1. **Create a device image.**
   To create a device image, call the function `xil_create_from_device()`, whose function prototype is shown below.

```
XilImage xil_create_from_device(XilSystemState state,
    char *devicename, XilDevice deviceObj);
```

   The parameter `state` is the system-state data structure that was returned to you when you called `xil_open()` to initialize the XIL library. The `devicename` parameter is the name of the device to be associated with the device image; the name of the device must be provided by the group that writes the device handler. If you initialized attributes for this device type, the device name should be the same name specified on the `xil_device_create()` function shown in Step 1 on page 69. The final parameter, `deviceObj`, is the device object you created to initialize the device's attributes. If you didn't create a device object for this device type, pass NULL for this parameter.

Note – Devices that don't require attribute initialization typically don't recognize or support device objects. For these devices, you can't use a device object to set attributes and you must pass NULL for the `deviceObj` argument.

---

2. **Set any required attributes for the device you're reading or writing.**
   For instance, say you want to acquire an image from a frame grabber, and the frame grabber has a brightness control. Before you grab the image, you might want to adjust that control. The XIL function that enables you to make this adjustment is called `xil_set_device_attribute()`. Its prototype is shown below.

```
int xil_set_device_attribute(XilImage image, char *attribute,
    void *value);
```

   The parameters to this function are a handle to your device image, the name of an attribute, and a value for the attribute. For example, you might set the frame grabber's brightness using the function call

```
xil_set_device_attribute(device_image, "BRIGHTNESS", (void *).5);
```

   Like device names, the names of device attributes and their possible values are defined by the group that writes the handler for the device.

   Normally, `xil_set_device_attribute()` returns the value `XIL_SUCCESS`. If the function is unable to set the attribute, it returns `XIL_FAILURE`.

   The XIL library also provides the `xil_get_device_attribute()` function, which reads the value of a device attribute. Its function prototype is shown below.

```
int xil_get_device_attribute(XilImage image,
    char *attribute, void **value);
```

3. **Perform an XIL operation using your device image as either the source or destination for the operation.**
   If the operation's source image is a device image, the operation will read an image from the associated device. For example, the operation might read a document on a scanner or get a frame of video from a video card. If the

operation’s destination image is a device image, the operation will write an image to the associated device. For instance, the operation might print a color image on a color printer.

Destroying a Device Object

When you create a device object, you associate it with a particular device type. You can then use the device object to initialize as many device images as needed of the same device type. When you are done using the device object, you must destroy it to release the memory allocated to it.

To destroy a device object, call the `xil_device_destroy()` function, whose prototype is shown below.

```
void xil_device_destroy(XilDevice deviceObj);
```

The only parameter on `xil_device_destroy()` is the handle to the device object.

Reading Kodak Photo CD Images 5

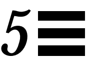

The XIL library includes a device handler that reads and decodes images stored in Eastman Kodak Company's Photo CD™ format. This chapter explains how you use the XIL library to read such images. Before delving into that subject, however, the chapter briefly discusses Photo CD technology in general: how images are scanned in, the PhotoYCC™ color space in which image data is stored, and how the data is stored on the compact disk.

The Photo CD Technology

Today, you can take your 35 mm film or slides to a licensed photofinisher and have your pictures stored on a Photo CD disk. Each disk can accommodate over 100 photographs. You can then view your pictures on a television set using either a Photo CD player or a Philips CD-I system. Or you can view them on your computer using a CD-ROM player, if you have the necessary software.

The Photo CD Imaging Workstation

When you take your film or slides to the photofinisher, the photofinisher puts your pictures on a Photo CD disk using a Photo CD Imaging Workstation. This workstation includes a scanner, a workstation that serves as a "Data Manager," and a CD writer. The scanner scans your 35 mm film or slides and produces digital images.

The data in these images represents RGB values. The Data Manager then

- Converts the RGB image data to the PhotoYCC color space
- Subsamples the chrominance values in each PhotoYCC image so that there is one pair of chrominance values for each two-by-two block of luminance values
- Produces multiple versions of each image with varying resolutions
- Uses the CD writer to write the multiple versions of each image to a compact disk

How Images Are Stored

For XIL programming, the two main things to bear in mind about Photo CD images are

- The images are stored in the PhotoYCC color space
- Each image is stored at multiple resolutions

The PhotoYCC Color Space

Kodak decided to convert images to a YC_bC_r color space because the data in the chrominance channels of YC_bC_r images can be subsampled without greatly affecting image quality. This ability to subsample the image data was considered important because a single scanned-in image requires 18 Mbytes of storage. After subsampling, an image requires only half that amount.

The decision to develop the PhotoYCC color space instead of using an existing YC_bC_r color space (such as the one defined in CCIR Recommendation 709) came about because Kodak wanted a device-independent color space, one with a broad color gamut. CCIR Rec. 709 YC_bC_r, for example, limits its color gamut to colors that can be displayed on a high-definition television. Photo CD images may be displayed on televisions, but they may also be output to high-quality printers.

As an XIL programmer, you really only need to know (1) that when you read a Photo CD image into an XIL image, the image data is PhotoYCC data and (2) that the XIL library supports the conversion of images to and from the PhotoYCC color space. You must convert the data to another color space before you can display or print the image. See “Converting the Image’s Color Space” on page 83 for more information on this subject.

Image Resolutions Supported in the XIL Library

As mentioned in “The Photo CD Imaging Workstation” on page 75, Photo CD images are stored at multiple resolutions. Table 5-1 shows the resolutions supported by the XIL library and the names Kodak has given to them.

Table 5-1 Resolutions of Photo CD Images

Resolution	Name
192 by 128	Base/16
384 by 256	Base/4
768 by 512	Base
1536 by 1024	4Base
3072 by 2048	16Base
6144 by 4096	64Base

Reading Photo CD Images Using the XIL Library

Chapter 3, “XIL Images,” introduced the idea of a device image—an image that resides on a device such as a scanner or a printer and can be used as the source or destination image for an XIL operation. In general, support for these devices will be provided by third parties. However, the XIL library does include a loadable device handler that reads Photo CD images. This means that Photo CD images can be treated as device images.

The section “Reading and Writing Devices Other than Displays” on page 67 covered the general procedure for creating a device image, setting device-image attributes, and using a device image in an operation. The sections below detail the specific calls you use to read and operate on a Photo CD image.

Creating a Device Image

To create a device image that represents a Photo CD image, you use the function `xil_create_from_device()`, whose function prototype is shown below.

```
XilImage xil_create_from_device(XilSystemState state,
    char* devicename, XilDevice deviceObj);
```

The first parameter to this function is a handle to the system state. The second is a string that identifies the Photo CD device handler: `"ioSUNWPhotoCD"`. And the third specifies a device object that has been associated with the device type for initializing device attributes. The Photo CD reader doesn't use or recognize associated device objects, so this parameter must be `NULL`. Therefore, the code to create a device image associated with a Photo CD image will look something like this:

```
XilImage ycc_photocd_image;
XilSystemState state;

ycc_photocd_image = xil_create_from_device(state,
"ioSUNWPhotoCD", NULL);
```

The return value of the function, `ycc_photocd_image`, is a handle to the device image.

Setting Device-Image Attributes

Before you can read a Photo CD image, you must create the appropriate device image, then set the necessary attributes for that image. The attributes you may set for a Photo CD device image are shown in Table 5-2.

Table 5-2 Photo CD Image Attributes

Attribute	Meaning
FILEPATH	The full pathname of the Photo CD image you want to read
RESOLUTION	The resolution at which you want to read the image. The possible resolutions are listed in Table 5-1 on page 77.
MAX_RESOLUTION	A read-only attribute indicating the highest resolution at which an image is stored. (In general, each image is stored at all possible resolutions, but there are exceptions to this rule.)
ROTATION	A read-only attribute indicating the amount of rotation required to display the image in its proper orientation.

A Photo CD device images doesn't require any attribute values at creation time; therefore, you cannot associate a device object with it to initialize its attributes at creation time (see "Initializing a Device's Attributes" on page 68).

The `FILEPATH` Attribute

After creating a Photo CD image, you *must* set a `FILEPATH` attribute for it that indicates which image you want to read. To set this attribute, call the function `xil_set_device_attribute()`:

```
XilImage ycc_photocd_image;
char *pathname = "...";

xil_set_device_attribute(ycc_photocd_image, "FILEPATH",
    (void *)pathname);
```

Note – The last part of the path to the image usually has the form `.../PHOTO_CD/IMAGES/IMG`*nnnn*`.PCD`. This part of the pathname reflects the directory structure on the Photo CD disk.

You can also read the value of the FILEPATH attribute if you need to determine which image you'll be reading if you read the device image. You read this value using the function xil_get_device_attribute():

```
XilImage ycc_photocd_image;
char *pathname;

xil_get_device_attribute(ycc_photocd_image, "FILEPATH",
    (void **)pathname;
```

The RESOLUTION Attribute

By default, when you read a Photo CD image, you read the Base version of that image (the one whose resolution is 768 by 512). To read a higher or lower resolution version of the image, you must set the value of the RESOLUTION attribute to one of the following enumeration constants:

```
typedef enum{
    XIL_PHOTOCD_16TH_BASE,
    XIL_PHOTOCD_4TH_BASE,
    XIL_PHOTOCD_BASE,
    XIL_PHOTOCD_4X_BASE,
    XIL_PHOTOCD_16X_BASE,
    XIL_PHOTOCD_64X_BASE
} XilPhotoCDResolution;
```

These constants correspond to the resolutions listed in Table 5-1 on page 77.

For example, to read the 192-by-128 version of an image, you would set the RESOLUTION attribute as shown below.

```
XilImage ycc_photocd_image;

xil_set_device_attribute(ycc_photocd_image, "RESOLUTION",
    (void *)XIL_PHOTOCD_16TH_BASE);
```

You can also read the value of the attribute using code similar to that shown below.

```
XilImage ycc_photocd_image;
XilPhotoCDResolution resolution;

xil_get_device_attribute(ycc_photocd_image, “RESOLUTION”,
    (void **)&resolution);
```

The `MAX_RESOLUTION` Attribute

This is a read-only attribute that enables you to determine the highest resolution version of an image that is available for a particular Photo CD image. In general, each image is available at all possible resolutions; however, on some prerecorded Photo CD disks, this may not be the case.

To read the value of this attribute, you call the function `xil_get_device_attribute()` and request that it store the value in a variable of the enumerated type `XilPhotoCDResolution`.

```
XilImage ycc_photocd_image;
XilPhotoCDResolution resolution;

xil_get_device_attribute(ycc_photocd_image, “MAX_RESOLUTION”,
    (void **)&resolution);
```

The value returned in `resolution` will be one of the enumeration constants shown below.

```
typedef enum{
    XIL_PHOTOCD_16TH_BASE,
    XIL_PHOTOCD_4TH_BASE,
    XIL_PHOTOCD_BASE,
    XIL_PHOTOCD_4X_BASE,
    XIL_PHOTOCD_16X_BASE,
    XIL_PHOTOCD_64X_BASE
} XilPhotoCDResolution;
```

The `ROTATION` Attribute

This is a read-only attribute that enables you to determine the proper orientation for displaying a particular Photo CD image. The rotation attribute is analogous to the *portrait* or *landscape* modes of printed text: it merely indicates the orientation of the data.

To read the value of this attribute, you call the function `xil_get_device_attribute()` and request that it store the value in a variable of the enumerated type `XilPhotoCDRotate`.

```
XilImage ycc_photocd_image;
XilPhotoCDRotate rotation;

xil_get_device_attribute(ycc_photocd_image, "ROTATION",
    (void **)&rotation);
```

The value returned in `rotation` will be one of the enumeration constants shown below.

```
typedef enum{
    XIL_PHOTOCD_CCW0,
    XIL_PHOTOCD_CCW90,
    XIL_PHOTOCD_CCW180,
    XIL_PHOTOCD_CCW270
} XilPhotoCDRotate;
```

Based upon the constant returned by `xil_get_device_attribute()`, you can perform the appropriate rotation on the image (see "Rotating Images" on page 167). Before performing the rotation, be sure to use `xil_set_origin()` to set the image origin to the image's center; before displaying the image, you must call `xil_set_origin()` again to restore the origin to coordinate 0,0. Example code is shown on page 86.

Capturing an Image from a Photo CD Disk

Once you've created a Photo CD device image and set the appropriate device image attributes, capturing an image from the device is a simple two-step process. You need to:

1. Create an XIL image that has the same dimensions and data type as the image you're reading.

   The width and height of this image will be determined by the resolution of the Photo CD image and also by its rotation. The number of bands in the image should be three, since an image in PhotoYCC format has three channels. And the data type of the image should be `XIL_BYTE`, since each Y, C_b, and C_r value in a Photo CD image is an 8-bit value.

2. Perform an XIL operation using your device image as the source image and the XIL image mentioned above as the destination.

The code fragment below creates a destination image that will hold a Base Photo CD image and then copies an image from the Photo CD disk to the destination.

```
XilSystemState state;
XilImage ycc_photocd_image, dst_image;

dst_image = xil_create(state, 768, 512, 3, XIL_BYTE);
xil_copy(ycc_photocd_image, dst_image);
```

It is important to note that, after the copy operation, the data in `dst_image` is in the PhotoYCC color space. The next section tells you how to convert this data to another color space.

Converting the Image's Color Space

Before you can display or print an image you've read from a Photo CD image file, you must convert the data in the image from the PhotoYCC color space to another color space. For example, if you wanted to display the image on a 24-bit frame buffer, you would convert the data to an RGB space before

displaying it. Table 5-3 lists some of the different output devices to which you might want to send the image and the color-space conversion that is necessary for each case.

Table 5-3 Converting PhotoYCC Data to Another Color Space

Output Device	Color-Space Conversion
True-color display	Convert the data to an RGB color space, such as the one named `rgb709` in the XIL library.
Indexed-color display	Convert the data to an RGB color space, such as the one named `rgb709` in the XIL library. Then, dither the RGB image to an 8-bit image. For information about XIL dithering operations, see section "Dithering an Image" on page 105.
Grayscale display	Convert the data to a luminance-only color space, such as the XIL color space `y709` or `ylinear`.
Monochrome display	Convert the data to a luminance-only color space, such as the XIL color space `y709` or `ylinear`. Then, dither the 8-bit luminance data to a 1-bit image. For information about XIL dithering operations, see the section "Dithering an Image" on page 105.
Color printer	Convert the data to a CMYK color space. The XIL library supports a linear CMYK space named `cmyk`.

For detailed information about the color spaces that the library supports and about converting image data from one color space to another, see the section "Color-Space Conversion" on page 122.

The code fragment on this and the following page shows how the ROTATION attribute from a Photo CD image is read and used, and how the image's data is converted to the RGB color space specified in CCIR Recommendation 709.

```
XilSystemState state;
XilImage ycc_photocd_image;      /* image read from disk */
XilImage rgb_photocd_image;      /* converted to rgb color space
*/
XilImage photocd_image;          /* after rotation */
XilImage display;                /* image for display */
XilPhotoCDRotate rotation;

/*
 * This code assumes the photocd image has been created as shown
on
 * page 78, and that its width, height, number of bands, and
 * data type have been retrieved with xil_get_info()
 */

/* Read the Photo CD image's ROTATION attribute */
xil_get_device_attribute(ycc_photocd_image, "ROTATION",
    (void**)&rotation);

rgb_photocd_image = xil_create(state, width, height, nbands,
    datatype);

/* Set color spaces to prepare for color-space conversion */
xil_set_colorspace(ycc_photocd_image,
    xil_colorspace_get_by_name(state, "photoycc"));
xil_set_colorspace(rgb_photocd_image,
    xil_colorspace_get_by_name(state, "rgb709"));

/* Convert the image's color space so it can be displayed */
xil_color_convert(ycc_photocd_image, rgb_photocd_image);

/* Set the image origin in preparation for a rotation */
xil_set_origin(rgb_photocd_image, width/2.0, height/2.0);

/* Example continued on next page */
```

```
/* Based upon the Photo CD image’s ROTATION attribute, open a
 * window with appropriate dimensions; for 90 and 270 degree
 * rotation, the image is rotated. Also, set the Photo CD
 * image’s origin to its center.  */
if(rotation == XIL_PHOTOCD_CCW90 ||
   rotation == XIL_PHOTOCD_CCW270) {

/* ...code to open a window with dimensions height-by-width ...*/

      display = xil_create_from_window(state, xdisplay, xwindow);
       photocd_image = xil_create(state, height, width, nbands,
         datatype);
       xil_set_origin(photocd_image, height/2.0, width/2.0);
      }
else {

/* ...code to open a window with dimensions width-by-height ...*/

      display = xil_create_from_window(state, xdisplay, xwindow);
       photocd_image = xil_create(state, width, height, nbands,
         datatype);
       xil_set_origin(photocd_image, width/2.0, height/2.0);
     }

/* Set the radians for a rotation. Constant PI has be previously
 * defined and set equal to 3.14159 */
switch (rotation) {
  case XIL_PHOTOCD_CCW0:
    angle = 0; break;
  case XIL_PHOTOCD_CCW90:
    angle = PI * 0.5; break;
  case XIL_PHOTOCD_CCW180:
    angle = PI; break;
  case XIL_PHOTOCD_CCW270:
    angle = PI * 1.5; break;
}

/* Rotate the image, then return its origin to 0,0 */
xil_rotate(rgb_photocd_image, photocd_image, "nearest", angle);
xil_set_origin(photocd_image, 0.0, 0.0);

/* copy the image to the display */
xil_copy(photocd_image, display);
```

Preparing Images for Display

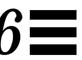

To introduce you to programming with the XIL library, Chapter 2, “Basic XIL Program,” presented an example program that read in an 8-bit grayscale image and displayed it in an X window created from one of four X visuals. This chapter looks at a revision of that example called `display`. The new example can display images of several data types and images of more than one band and was designed to illustrate how to handle some of the many display cases you may encounter.

The source files for the program can be found in the directory `$XILHOME/examples/display`. Table 6-1 list these source files and indicates what the code in each file does.

Table 6-1 Source Files for `display`

Source File	What the Code Does
`display.c`	Contains `main()`, which prepares different types of source images to be copied to various types of display images
`fileio.c`	Reads an image from a file and loads the image data into an XIL image
`window.c`	Creates the program’s X colormap and manages that colormap

The first part of this chapter explains briefly how to build and run the display example. The bulk of the chapter then discusses the most important tasks performed in the example. These include:

- Converting a single-band image to a multiband image
- Converting an XIL_SHORT image to an XIL_BYTE image
- Converting an RGB image to a pseudocolor image and a colormap
- Converting a 24-bit image to a 1-bit image
- Converting an 8-bit image to a 1-bit image
- Displaying a 1-bit image on a monochrome display

The chapter concludes with a list of all the cases handled in the program.

Running the Sample Display Program

To run this example, follow these instructions:

1. **Change your working directory to** `$XILHOME/examples/display`**.**
2. **Build the program using the makefile in that directory.**
3. **Execute the program from the directory in which you built it, using the command line**

```
% display toys.header
```

The program will display an RGB image of a group of toys in an X window. To stop the program, move your pointer into the program's window, and click any mouse button.

Converting a Single-Band Image to a Multiband Image

The display program approaches this problem in two ways: by passing the source image through a lookup table and by replicating the source image in the multiple bands of the destination image.

Passing the Source Image Through a Lookup Table

If there are only a few different values in the source image—for example, when the source image is a 1-bit image—passing the source image through a lookup table is an effective method of performing this conversion. For instance, if all the values in the source are 0 or 1, you can use a lookup table with only two

entries: one that specifies the values to appear in each band of the destination for pixels that have a value of 0 in the source, and another that specifies the values to appear in each band of the destination for pixels that have a value of 1.

For example, the program uses this method to prepare a bit image for display in a 24-bit window. To make the source image compatible with the display image, the program converts the source image to a 3-band, 8-bit image. To do this, the example:

1. Creates a temporary image, `retained_image`, that has the same width and height as the source image, but is 3 bands deep and will contain `XIL_BYTE` data. (This temporary image will later be copied to the display.)
2. Creates the lookup table (a data structure of type `XilLookup`) shown in Figure 6-1.

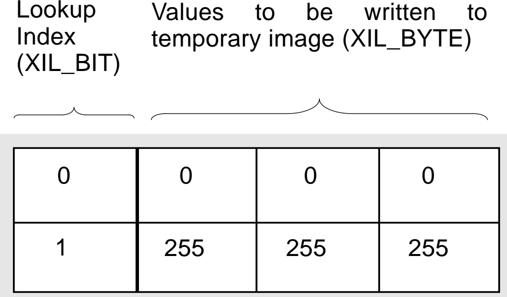

Figure 6-1 Lookup Table

3. For each pixel in the source image, looks up the proper value in the left column of the table and writes the three values to its right to the first, second, and third bands of the corresponding pixel in `retained_image`.

The code that implements these steps is shown below.

```
#define BITSIZE 0x2

XilImage src_image, retained_image;
XilLookup lookup;
Xil_unsigned8 lookupdata[] = {0, 0, 0, 255, 255, 255};

retained_image = xil_create(state, width, height, 3, XIL_BYTE);
lookup = xil_lookup_create(state, XIL_BIT, XIL_BYTE, 3,
    BITSIZE, 0, lookupdata);
xil_lookup(src_image, retained_image, lookup);
```

Replicating the Source Image in the Bands of the Destination

The second method the example uses to convert a single-band source image to a multiband destination image is to replicate the source image in each band of the destination. To perform this task using the XIL library, you must first create a child image representing each band of the destination and then copy the source to each child image. Because any changes made to an XIL child image affect the parent, these copies result in the source being copied to each band of the destination.

The display program converts a 1-band `XIL_BYTE` source image to a 3-band `XIL_BYTE` image using the code shown below.

```
XilImage image, retained_image, band0, band1, band2;

retained_image = xil_create(state, width, height, 3, XIL_BYTE);
band0 = xil_create_child(retained_image, 0, 0, width, height, 0,
     1);
band1 = xil_create_child(retained_image, 0, 0, width, height, 1,
     1);
band2 = xil_create_child(retained_image, 0, 0, width, height, 2,
     1);
xil_copy(image, band0);
xil_copy(image, band1);
xil_copy(image, band2);
```

Converting an XIL_SHORT Image to an XIL_BYTE Image

Converting an XIL_SHORT image to an XIL_BYTE image is a two-step process. First, you must rescale the values in the source image so that they fall in the range 0 to 255. Then, you must cast the values in the source image to values of type XIL_BYTE.

In the rescaling step, you could simply map the lowest possible value in the source (-32768) to the lowest possible value in the destination (0) and the highest possible value in the source (32767) to the highest possible value in the destination (255). However, you'll generally achieve better results by first finding the extrema in the source and then mapping the lowest value actually in the source to 0 and the highest value to 255.

The code that display uses to perform this task (in one instance) is shown below.

```
#define CMAPSIZE 256

float low[1], high[1];
float mult[1], offset[1];

xil_extrema(image, high, low);
mult[0] = (NCOLORS - 1) / (high[0] - low[0]);
offset[0] = -((low[0] * (NCOLORS - 1)) /
    (high[0] - low[0])) + (CMAPSIZE - NCOLORS);
xil_rescale(image, image, mult, offset);
xil_cast(image, retained_image);
```

The call to xil_extrema() stores the source image's highest value in high[0] and its lowest value in low[0]. Then, the call to xil_rescale() maps the lowest value in the image to 0 and the highest to 255. Once the values in image have been brought into the proper range, the example casts them to values of type XIL_BYTE. When casting a 16-bit value to an 8-bit value, xil_cast() preserves the 8 least significant bits of each input value, which is what the program requires.

Converting an RGB Image to an Indexed-Color Image and a Colormap

The display program prepares a true-color image to be shown on an indexed-color display by performing an ordered dither on the source image. The general procedure the example employs is to create a colorcube and a dither mask and then to call the function xil_ordered_dither(). After this operation, the program has a single-band version of the source image and a colorcube that defines the RGB values to be associated with each value in the indexed-color image. Before displaying the indexed-color image, the example creates a virtual X colormap and writes the values in the XIL colorcube to the X colormap.

Note – For a more detailed explanation of ordered dithering, see the section "Dithering an Image" on page 105.

The code the example uses to perform the ordered dither is shown below.

```
XilLookup cmap;
XilDitherMask mask;

cmap = xil_lookup_get_by_name(state, "cc496");
mask = xil_dithermask_get_by_name(state, "dm443");
set_colormap(xdisplay, xwindow, cmap);
retained_image = xil_create(state, width, height, 1, datatype);
xil_ordered_dither(image, retained_image, cmap, mask);
```

The colorcube cc496 and the dither mask dm443 are objects that the library creates when it is initialized. The colorcube is appropriate for dithering an RGB image to 216 colors (the indexes will range from 38 to 253), and dm443 is a 4-by-4 dither mask.

Note – These objects could also have been created with the functions xil_colorcube_create() and xil_dithermask_create().

Before the indexed-color image can be displayed, the color values in the colorcube cmap must be stored in colorcells 38 to 253 in the application's X colormap. This task is handled by the routine set_colormap(). Note that the colormap data that is being loaded into the X colormap is taken from the lookup table that was used as the colorcube for the dithering operation.

The XIL library also includes two other types of dithering operations that you can use in preparing a true-color image to be shown on an indexed-color display. The functions that perform these dithering operations are called `xil_nearest_color()` and `xil_error_diffusion()`. For further information about these functions, see the section "xil_nearest_color()" on page 108 and the section "xil_error_diffusion()" on page 114.

Converting a 24-Bit Image to a 1-Bit Image

The example uses a two-step process to convert an RGB image to a 1-bit image. The first step is to convert the 24-bit source image to an 8-bit grayscale image by extracting luminance information from the source image. Using the XIL library, you perform this task by converting the source image from the RGB color space to the Y color space. The destination image for this color-space conversion must be a single-band image.

The code `display` used to perform this color-space conversion is shown below.

```
image1 = xil_create(state, width, height, 1, XIL_BYTE);
xil_set_colorspace(image, xil_colorspace_get_by_name(state,
    "rgblinear"));
xil_set_colorspace(image1, xil_colorspace_get_by_name(state,
    "ylinear"));
xil_color_convert(image, image1);
```

The second step in the process is to dither the 8-bit image to a 1-bit image using either the function `xil_ordered_dither()` or `xil_error_diffusion()`. If execution speed is your primary concern, use the ordered-dither routine. To obtain the best quality image, use the error-diffusion function. For more information about these functions, see the sections "xil_ordered_dither()" on page 118 and "xil_error_diffusion()" on page 114.

The `display` example performs this task using `xil_error_diffusion()`.

```
#define BITSIZE 0x2

XilKernel distribution;
static int multipliers[1] = {-1};
static unsigned int dimensions[1] = {BITSIZE};

retained_image = xil_create(state, width, height, 1, XIL_BIT);
distribution = xil_kernel_get_by_name(state, "floyd-steinberg");
colormap = xil_colorcube_create(state, XIL_BIT, XIL_BYTE, 1, 0,
    multipliers, dimensions);
xil_error_diffusion(image1, retained_image, colormap,
    distribution);
```

Converting an 8-Bit Image to a 1-Bit Image

This task is the same as the second step described in the preceding section. You should dither the 8-bit image to a 1-bit image using either the function `xil_ordered_dither()` or `xil_error_diffusion()`. As noted above, the first of these functions is the fastest, but the second generally produces better looking results.

Displaying a 1-Bit Image on a Monochrome Display

The `display` program is set up to deal with `XIL_BIT` images in which 0's represent black pixels and 1's represent white pixels (the normal case). In the X colormap for monochrome displays supported by the XIL library, however, 0 represents white and a 1 represents black; therefore, the `display` program takes the one's complement of the source image (using `xil_not()`) before displaying it.

Types of Images Displayed

This section simply lists the cases handled by the `display` example so that you can quickly determine whether a case you're interested in is covered. In Table 6-2, the types of images the program can process are shown in the left

column, and the supported X visuals are listed in the top row. A check mark indicates that the program can display a particular type of image in an X window created using a particular X visual.

Table 6-2 Cases Handled by the `display` Program

	TrueColor 24 Bits	PseudoColor **8 Bits**	**GrayScale 8 Bits**	**StaticGray 8 Bits**	**StaticGray 1 Bit**
XIL_BIT 1 Band	✔		✔	✔	✔
XIL_BYTE 1 Band	✔			✔	✔
XIL_BYTE 3 Bands	✔	✔	✔		✔
XIL_SHORT 1 Band	✔		✔	✔	✔

Note – For an explanation of X visuals and their uses, see the *Xlib Programming Manual.*

Presentation Functions 7

This chapter covers a group of functions you are likely to use in displaying images, or printing them. These functions enable you to:

- Copy a source image to a destination image
- Copy specified image planes to a destination image
- Rescale the values in an image
- Cast an image from one data type to another
- Dither an image
- Convert an image from one color space to another
- Prepare CMY images for printing

These tasks are discussed in detail in the sections below.

Copying an Image to the Display

The `xil_copy()` function lets you copy a source image to a destination image. `xil_copy()` copies each ***plane*** (bit) in a source-image pixel to the destination image pixel, replacing the value in each corresponding plane in the destination pixel. To control the pixel planes copied from the source to the destination, you can define a ***plane mask*** to specify the bit planes to copy, then call `xil_copy_with_planemask()` to perform the copy.

The `xil_copy()` and `xil_copy_with_planemask()` functions are discussed separately in the following sections.

Copying All Bit Planes

The simplest way to display an image is to use `xil_copy()` to copy your source image to the display. The `xil_copy()` function copies each ***plane*** (bit) in a source-image pixel to the destination image pixel, replacing the value in each corresponding plane in the destination pixel. For example, an eight-bit-deep frame buffer has eight bit planes; if you use `xil_copy()` to copy an `XIL_BYTE` image to that frame buffer, all eight planes of each source pixel are copied to all eight planes in the frame buffer's destination pixel.

For `xil_copy()` to work, you must have created a display image that will serve as the destination for the copy. (For information about what display images are and how you create them, see the section "Display Images" on page 44.) In addition, the source image—the image to be displayed—and the display image must use the same number of bits to represent a pixel. Table 7-1 lists the combinations of source and display images that will match in this way.

Table 7-1 Matching Source and Display Images

Source Image Type	Required Display Image Depth
Single-band `XIL_BIT` image	One bit
Single-band `XIL_BYTE` image	Eight bits
Three-band `XIL_BYTE` image	Twenty-four bits

If your source and destination images match, you can display the source image using code similar to this.

```
XilImage src, display_image;

xil_copy(src, display_image);
```

Besides performing this kind of explicit copy using `xil_copy()`, you can also request that an image-processing function (like a rotate function) copy its output to the display. You do this by specifying a display image as the

destination for the operation. For example, the following call to `xil_rotate()` rotates the source image and copies the rotated image to the display.

```
XilImage src, display_image;

xil_rotate(src, display_image, "bilinear", -0.7854);
```

Again, before this type of operation will work, the source and display images must use the same number of bits to represent a pixel. If your source and display images do not match in this respect, you must perform one or more of the tasks discussed in this chapter before displaying your image.

If overlapping but not coincident sibling images (children of the same parent) are specified as the source and destination, `xil_copy()` detects the overlap and correctly generates the destination image. All other operations generate a warning message under these conditions and have undefined results, as mentioned on page 54 in the discussion on parent and child images.

Copying Only the Planes Defined in a Plane Mask

The `xil_copy()` function copies each ***plane*** (bit) in a source-image pixel to the destination image pixel, replacing the value in each corresponding plane in the destination pixel.To control the pixel planes copied from the source to the destination, you can define a ***plane mask*** to specify the bit planes to copy, then call the `xil_copy_with_planemask()` function to perform the copy. The prototype for `xil_copy_with_planemask()` is shown below.

```
void xil_copy_with_planemask(XilImage src, XilImage dst,
    unsigned int planemask[]);
```

A typical reason to copy with plane mask control is to overlay one image over another. Because you can copy specific planes from each source-image pixel, the underlying image isn’t visible where the overlay is drawn; however, all planes from the underlying image are still available and can be refreshed in the destination by rendering it with another plane mask. Overlays can improve performance by reducing the amount of graphic information that has to be redrawn; they can also be used to highlight graphics for selection.

Note – Because read/write color cells can be allocated only in PseudoColor and DirectColor visuals, you need to provide an alternative technique to overlays for other visuals; you may also want to provide an alternative in case the overlay fails. For more information on overlays and their advantages and disadvantages, see the *Xlib Programming Manual*.

When you use `xil_copy_with_planemask()`, each pixel in the destination image is defined by the following operation:

```
dst = (dst & ~mask) | (src & mask)
```

Here, `dst` is the destination image, `mask` is the plane mask, and `src` is the source image. Thus, if the plane-mask bit is "on," the copy overwrites the corresponding bit in the destination image; otherwise, the bit in the destination image is unchanged.

As an example, assume you have two single-band XIL_BYTE images you want to overlay. Table 7-2 shows the plane masks you might use.

Table 7-2 Plane Masks for an Overlay

Plane Mask	Binary Base	Hexadecimal Base	Decimal Base
Mask 1	00000001	0x1	1
Mask 2	11111110	0xfe	254

In Table 7-2, Mask 1 has only its low order bit turned on; thus, it ensures that only the source image's low order bit is copied to the destination. Mask 2, on the other hand, has its seven high-order bits turned on, so it ensures that the source image's seven high-order bits are copied to the destination.

The code below shows how you might define and use the plane masks shown in Table 7-2:

```
XilImage src1, src2, dst;
unsigned int planemask1[1], planemask2[1];

planemask1[0] = 0x1;
planemask2[0] = 0xfe;

xil_copy_with_planemask(src1, dst, planemask1);
xil_copy_with_planemask(src2, dst, planemask2);
```

The plane mask must be an array of unsigned integers. The number of array elements must match the number of image bands; each array element specifies the plane mask for the corresponding band in the destination image. Both the source and destination images must have the same type and number of bands, and in-place operations are supported.

When using a plane mask for copying an image to the display, the depth of the window is the upper limit on the number of meaningful bits you can set in the plane mask, and you must manipulate the colormap to get a reasonable display. For more information on allocating color cells for overlays, see the *Xlib Programming Manual*.

In addition to overlays, you can use a plane mask for double buffering on hardware that doesn't have separate memory buffers. Double buffering is useful for animation because you can render an image in a hidden memory buffer while you display another image in a second buffer. Once rendering is complete in the hidden buffer, you can display its contents, then hide the second buffer and render a new image into it. Rendering into the hidden buffer and quickly switching between buffers smooths the transition between images. As with overlays, double buffering requires you to toggle between two colormaps as you change plane masks between images. For more information on using double-buffering for animation, consult Foley, et al. *Computer Graphics: Principles and Practice* (see Appendix E, "Bibliography").

Rescaling an Image

The XIL function `xil_rescale()` maps the values in each band of an image from one range (for example, 0 to 2047) to another range (for example, 0 to 255). The function performs this mapping by multiplying each value in a band by one constant and then adding a constant to the result of the multiplication.

The prototype for `xil_rescale()` is shown below.

```
void xil_rescale(XilImage src, XilImage dst, float *scale,
    float *offset);
```

The parameters `src` and `dst` are handles to the source and destination images. The parameter `scale` is an array of floating-point numbers that will serve as the constants by which values in the bands of the source image are multiplied. The values in band 0 are multiplied by `scale[0]`, the values in band 1 are multiplied by `scale[1]`, and so on. Therefore, the number of elements in the array must match the number of bands in the image. The final parameter, `offset`, is an array of floating-point numbers, one of which is added to the scaled values in each band of the source image.

There are many cases in which this function will be useful as you prepare an image for display. A few of these are mentioned below.

One obvious case is that in which you want to display an image containing `XIL_SHORT` values. For example, say that you want to display a single-band `XIL_SHORT` image in an X window created using an 8-bit `GrayScale` visual. The values in your image may range from -32768 to 32767, but the values you write to the display image must fall in the range 0 to 255 (or some subset of that range). To handle this situation, you could perform these steps.

1. **Determine the minimum and maximum values in your source image.** You determine the minimum and maximum values using the function `xil_extrema()`. Your image will look better when displayed if you map the minimum value to 0 and the maximum value to 255, as opposed to mapping -32768 to 0 and 32767 to 255. In the latter case, you might wind up with very few gray levels in the image to be displayed.

2. **Rescale the values in your `XIL_SHORT` image so that they fall in the range 0 to 255.**
   The code you might use to accomplish this step is shown below.

```
XilImage src, dst;
float maximum[1], minimum[1];
float multiplier[1], offset[1];

multiplier[0] = 255 / (maximum[0] - minimum[0]);
offset[0] = -((minimum * 255) / (maximum[0] - minimum[0]));
xil_rescale(src, dst, multiplier, offset);
```

3. **Cast the values in the rescaled image to be of type `XIL_BYTE`.**
   Although you have rescaled your image so that its values fall in the range 0 to 255, those values are still 16-bit signed values. Before copying your image to the display, you must use the function `xil_cast()` to create an `XIL_BYTE` version of the image. For more information about `xil_cast()`, see the next section, "Casting an Image from One Data Type to Another."

Here's another case in which you may need to rescale an image before displaying it. Suppose you want to display an 8-bit grayscale image in an 8-bit `GrayScale` window, but that you do not want to use the currently installed X colormap in displaying it. You want to create a new virtual X colormap, store a grayscale ramp in that colormap, and have it installed when your application is active. If you write values to all 256 colorcells in the virtual colormap, you're almost certainly going to see colormap flashing when the colormap is installed, so you may decide not to write values to the first 16 colorcells—to write all the values needed to display your image in colorcells 17 through 255. This strategy requires that you rescale the values in your image so that they fall in the range 17 to 255.

Casting an Image from One Data Type to Another

The XIL function `xil_cast()` casts an XIL image from one data type to another. The function's prototype is shown below.

```
void xil_cast(XilImage src, XilImage dst);
```

The parameter `src` is a handle to your source image, and `dst` is a handle to a destination image. This destination must have the same width, height, and number of bands as the source image and must have the data type to which you want to cast the source image.

There are a number of instances in which you may need to use this function as you prepare an image for display. For example, you may have a single-band `XIL_SHORT` image that you want to display in an 8-bit window. To do this you need to follow these steps.

1. **Rescale the image.**
   Unless the values in your `XIL_SHORT` image already fall in the range 0 to 255, you must use the function `xil_rescale()` to map them to that range, or a subset of that range. For more information about rescaling images, see “Rescaling an Image” on page 102.

2. **Cast the `XIL_SHORT` image to an `XIL_BYTE` image.**
   Before you can display your image in an 8-bit window, you must cast the 16-bit values in the source image to 8-bit values. See the code fragment below.

```
XilImage short_image, byte_image;
unsigned int width, height, nbands; /* Dimensions of images */

byte_image = xil_create(state, width, height, nbands,
    XIL_BYTE);
xil_cast(short_image, byte_image);
```

Similarly, if you have an `XIL_BIT` image that you want to display in an 8-bit window, you must cast the source image to an `XIL_BYTE` image before displaying it. In this cast, the values 0 and 1 in the `XIL_BIT` image are cast to indices 0 and 1 in the `XIL_BYTE` image. If you want different indices, convert the image by passing it through a lookup table rather than by casting its data type. For information about lookup tables, see “Passing an Image Through a Lookup Table” on page 199.

Note – When casting the image so you can display it, you may want to use a display image as the destination image for the cast. However, if you intend to use the converted image again, you can cast the source image to an interim destination image, then use `xil_copy()` to copy the interim image to the display.

Dithering an Image

The XIL library provides several functions you can use to prepare images for display by dithering them. Before looking at these individual functions, though, the section explains what it means to dither an image in an XIL application.

What Is Dithering?

First, you need to know that the XIL library gives you the ability to pass a single-band image of any data type through a lookup table to produce a single-band or multiband image of the same or another data type. The lookup table used for this operation is a data structure of type `XilLookup` and has (among others) the following attributes:

- An input data type
- An output data type
- A number of bands on the output side

Figure 7-1 shows a single-band `XIL_BIT` image being passed through a lookup table to produce a three-band `XIL_BYTE` image.

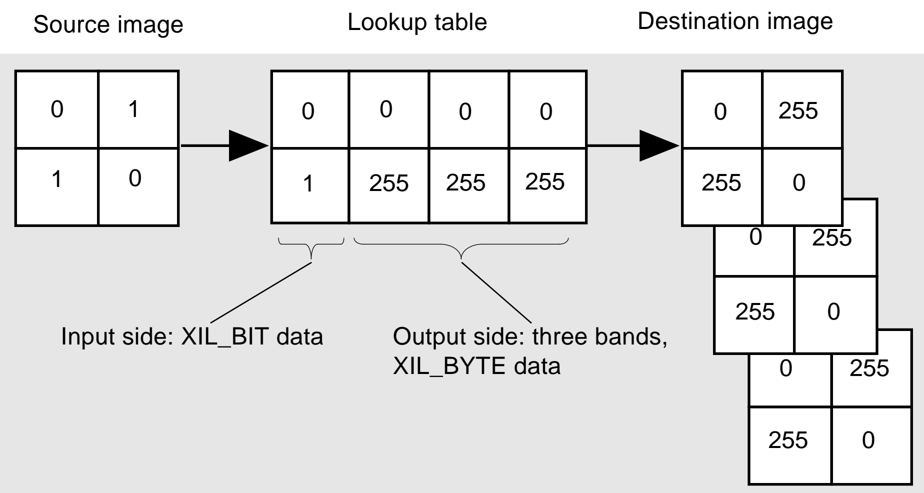

Figure 7-1 XIL Lookup Operation

A dither operation is an *inverse* lookup operation. As Figure 7-2 indicates, the dither operation matches a value (or values) in the source image—the image to be dithered—with a value from the *output* side of the lookup table being used, and then writes the corresponding value from the *input* side of the lookup table to a single-band destination image.

Note – Since the values in the source image don't actually match values in the output side of the table, each source-image value is paired with the value in the table closest to it.

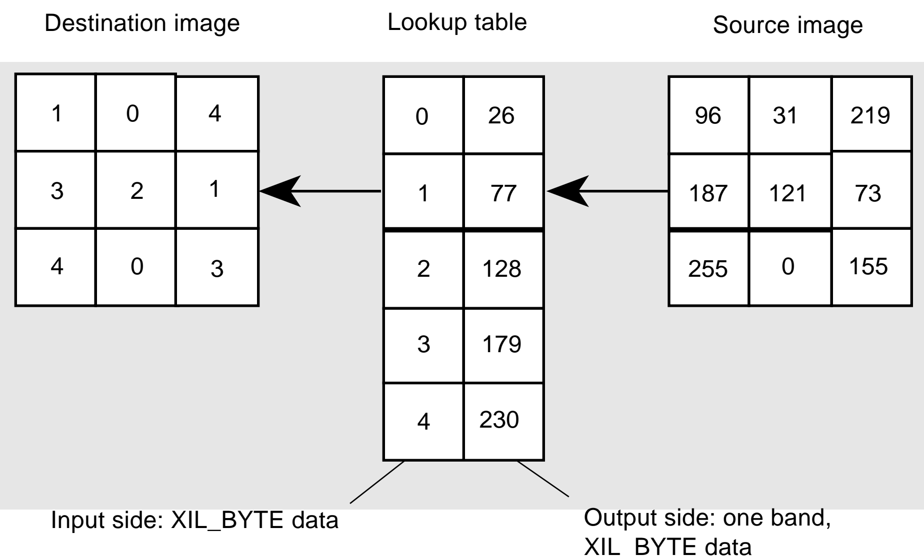

Figure 7-2 Dithering an Image

The purpose of this dithering operation is to produce an image that when mapped forward through the lookup table will produce an image as similar as possible to the original source image. Another way to state this is that the dithering operation produces an image that when displayed using the lookup table as its colormap will look as much like the original as possible.

Such dithering operations have many applications, but a couple of them are by far the most common. One common use of dithering is to convert a true-color (3-band `XIL_BYTE`) image to a pseudocolor (1-band `XIL_BYTE`) image. Another common use is to convert a grayscale (1-band `XIL_BYTE`) image to a monochrome (1-band `XIL_BIT`) image.

Methods of Dithering

This section contains information about the three functions the XIL library provides for performing dithering operations: `xil_nearest_color()`, `xil_error_diffusion()`, and `xil_ordered_dither()`.

xil_nearest_color()

The function `xil_nearest_color()` is the simplest of the dithering functions in that its algorithm for performing the inverse lookup described above is the most straightforward. This algorithm includes no provision for eliminating unwanted contours in the dithered image.

The function prototype for `xil_nearest_color()` is shown below.

```
void xil_nearest_color(XilImage src, XilImage dst,
    XilLookup lookup);
```

The parameters `src` and `dst` are handles to the source image—the image to be dithered—and the destination image. The parameter `lookup` is a handle to the lookup table through which the source image will be passed to produce the destination. There are actually two types of lookup tables that can be referred to here. One is the generic lookup table that is used for lookup operations, and the other is a special type of lookup table called a ***colorcube***. The basic difference between the two types of lookup tables is this: when you create a generic lookup table, you specify the dimensions of the table and the data to be stored in the table; when you create a colorcube, you specify the dimensions of the colorcube, and the function you use to create the colorcube fills in the data in the table for you. The sections below provide additional information about these two types of lookup tables.

Lookup Tables

You can create a generic lookup table in one of two ways: using the function `xil_lookup_create()` or using the function `xil_choose_colormap()`. In either case, `xil_nearest_color()` pairs the values in your source image with values on the output side of the lookup table by searching for nearest matches. This means that for each pixel in your source image, `xil_nearest_color()` must examine each entry in the lookup table. This is a time-consuming process. However, generic lookup tables do give you the best control over the contents of your lookup table and, therefore, over the quality of the dithered image.

The function prototype for `xil_lookup_create()` is shown below.

```
XilLookup xil_lookup_create(XilSystemState state,
    XilDataType input_datatype, XilDataType output_datatype,
    unsigned int output_nbands, unsigned int num_entries,
    short first_entry_offset, void *table_data);
```

Here's an example of a time when you might want to use this function to create your lookup table. You have a true-color image that you want to dither to a single-band `XIL_BYTE` image, and you want to be able to display the dithered image using the currently installed X colormap. The code you use to create this lookup table might look like this:

```
/* Read the color values stored in the X colormap and write them
   to an array of Xil_unsigned8. The values should be written
   in the following order: b0, g0, r0, b1, g1, r1, and so on.
   Let's say the array is called table_data */

xil_lookup_create(state, XIL_BYTE, XIL_BYTE, 3, 256, 0,
    table_data);
```

Using this table for your dither operation enables you to produce a dithered image that contains the maximum number of unique values (256) and to display the image without having to write anything to the X colormap. You won't have to worry about colormap flashing. The disadvantage of this method is that there may be few or no exact matches between the RGB values in your source image and the RGB values on the output side of the lookup table. Clearly, this can affect how accurately the dithered image, when displayed, will reflect the original source image.

The other function you might use to create a lookup table for dithering the true-color image mentioned above to a pseudocolor image is `xil_choose_colormap()`. This function returns the best lookup table of a particular size to use in dithering the image. The best lookup table is defined to be the one that contains as many as possible of the most frequently used colors in the source image.

The function prototype for `xil_choose_colormap()` is shown below.

```
XilLookup xil_choose_colormap(XilImage src, unsigned int size);
```

When you use this function to create your lookup table, there are two basic approaches you might take to dithering your true-color image. One approach is to use `xil_choose_colormap()` to create a lookup table that contains 256 entries. This strategy will produce a dithered image that looks very much like the original source image when displayed because the 256 most frequently used colors in the original will be reproduced exactly. However, you also have to get 256 specific colors into a hardware colormap to display your image. This will lead to colormap flashing on most displays.

A second approach is to use `xil_choose_colormap()` to create a lookup table with fewer than 256 entries, say 240 entries. You will still have most of the colors you need to display your image, and you can write those colors to your hardware colormap (if you only have one) without overwriting the first 16 entries in that colormap. This should prevent your application from contending with the OpenWindows™ tools for colors.

Note – `xil_choose_colormap()` accepts only 3-banded `XIL_BYTE` source images.

Colorcubes

You create a colorcube using the function `xil_colorcube_create()`. The prototype for this function is shown below.

```
XilLookup xil_colorcube_create(XilSystemState state,
    XilDataType input_type, XilDataType output_type,
    unsigned int nbands, short offset, int multipliers[],
    unsigned int dimensions[]);
```

Note that the parameters to this function are similar to those for `xil_lookup_create()`. However, you don't specify a number of entries in the lookup table or the data to be loaded into the table. Instead you provide arrays called `dimensions` and `multipliers`. The number of elements in both of these arrays must equal the number of bands in the image being dithered.

To understand the roles of dimensions and multipliers, consider the situation where you want to dither an XIL_BYTE RGB image to a 1-band XIL_BYTE image. To handle this case, you might declare and initialize dimensions and multipliers as follows:

```
unsigned int dimensions[] = {4, 9, 6};
int multipliers[] = {1, 4, 36};
```

These values would lead to the creation of a colorcube that would dither blue values in the source image to one of 4 blue levels, green values to one of 9 green levels, and red values to one of 6 red levels. You could picture this colorcube as a cube with dimensions of 4, 9, and 6, but it's probably more helpful to think of it as a lookup table with three bands on the output side. Figure 7-3 shows what the first 22 elements of the 216-element table would look like.

Values to be written to indexed-color image.

	Blue values	Green values	Red values
0	0	0	0
1	85	0	0
2	170	0	0
3	255	0	0
4	0	32	0
5	85	32	0
6	170	32	0
7	255	32	0
8	0	64	0
9	85	64	0
10	170	64	0
11	255	64	0
12	0	96	0
13	85	96	0
14	170	96	0
15	255	96	0
16	0	128	0
17	85	128	0
18	170	128	0
19	255	128	0
20	0	159	0
21	85	159	0

Figure 7-3 Colorcube for Dithering a True-Color Image to a Pseudocolor Image

This illustration should help clarify the significance of the three elements of multipliers. The element multipliers[0] indicates the distance between changes in blue levels on the output side of the table, multipliers[1] indicates the distance between changes in green levels, and multipliers[2] indicates the distance between changes in red levels. That is, the blue level changes with each entry, the green level changes every fourth entry, and the red level changes every thirty-sixth entry.

Note – The elements of multipliers can be negative numbers. Negative numbers indicate that the values in a color ramp should appear in decreasing as opposed to increasing order.

The values in the table are calculated by xil_colorcube_create() based on the dimensions and multipliers you supply. This means that when you dither an image using a colorcube, the dither function does not need to search the output side of a lookup table to find a set of values. Instead, it can calculate the value to be written to the dithered image using the values stored in the dimensions and multipliers arrays. As a result, dithering operations that use a colorcube are much faster than those that use a generic lookup table. What you give up is control over the exact contents of the colorcube.

One other note about colorcubes. Although they are called color*cubes*, colorcubes need not have three dimensions. As mentioned earlier, a colorcube has a number of dimensions equal to the number of bands in the image to be dithered. Therefore, to create a colorcube suitable for dithering a grayscale image to a monochrome image, you might define dimensions and multipliers as follows.

```
unsigned int dimensions[] = {2};
int multipliers[] = {-1};
```

Note – The XIL library creates two colorcubes when you initialize the library. One of these has the dimensions 4:9:6 and is useful for dithering RGB images to 216 colors. The other has the dimensions 8:5:5 and is useful for dithering YC_bC_r images to 200 colors. To get a handle to one of these colorcubes, use the function xil_lookup_get_by_name().

In addition to xil_colorcube_create(), the XIL library contains the colorcube-related functions shown in Table 7-3.

Table 7-3 Functions for Managing Colorcubes

Function Name	What the Function Does
xil_lookup_get_colorcube	Determines whether a lookup table is a colorcube or a generic lookup table
xil_lookup_get_colorcube_info	Determines whether a lookup table is a colorcube and, if it is, returns the dimensions and multipliers used to create the colorcube and the colorcube's origin

xil_error_diffusion()

The function xil_error_diffusion() is similar to xil_nearest_color(), but in addition to performing an inverse lookup, it uses a process called ***error diffusion*** to deal with any difference between a source-image pixel value and a value on the output side of a lookup table with which that pixel value is matched. The difference (or error) is distributed to the source-image pixels immediately to the right of and below the last pixel processed.

The function prototype for xil_error_diffusion() is shown below:

```
void xil_error_diffusion(XilImage src, XilImage dst,
    XilLookup lookup, XilKernel distribution);
```

The first three parameters are the same as the parameters to xil_nearest_color(). The parameter src is a handle to your source image, and dst is a handle to your destination image. The parameter lookup can be either a generic lookup table or a special type of lookup table called a colorcube. For information about these lookup tables, see the sections "Lookup Tables" on page 108 and "Colorcubes" on page 110. The final parameter, distribution, is an error-distribution kernel.

Note – XIL kernels (data structures of type `XilKernel`) are used primarily for convolution operations, which are discussed in the section "Filtering an Image" on page 187. The utility functions that affect kernels are also discussed in that section.

You create the error-distribution kernel using the function `xil_kernel_create()`, whose function prototype is shown below.

```
XilKernel xil_kernel_create(XilSystemState system_state,
    unsigned int width, unsigned int height, unsigned int keyx,
    unsigned int keyy, float *data);
```

The parameters `width` and `height` determine the size of the kernel, `keyx` and `keyy` determine the kernel's ***origin***, and `data` is a pointer to the floating-point values to be stored in the kernel. For example, the code below creates a 3-by-3 kernel with an origin of 1,1.

```
XilKernel kernel;
float kernel_data[] = {
    0/16.0, 0/16.0, 0/16.0,
    0/16.0, 0/16.0, 7/16.0,
    3/16.0, 5/16.0, 1/16.0
};

kernel = xil_kernel_create(state, 3, 3, 1, 1, kernel_data);
```

Note – The library provides a special shorthand method of creating the kernel shown above because it is used so commonly in error-diffusion operations. This method is to call the function `xil_kernel_get_by_name()` using a name of `"floyd-steinberg"`.

Figure 7-4 shows what this kernel looks like.

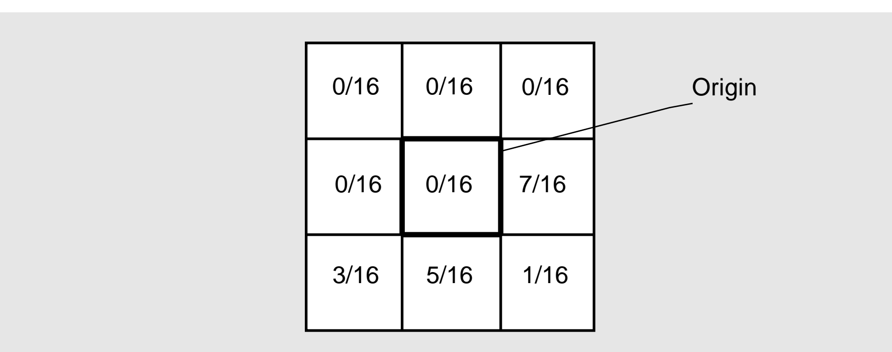

Figure 7-4 Error-Distribution Kernel

Note – The value of the kernel's origin and of the elements above and to the left of the origin must be 0 when you're calling `xil_error_diffusion()`.

Here's how `xil_error_diffusion()` uses the kernel. Say that you're dithering a grayscale image with values in the range 0 to 255 to a grayscale image with values in the range 0 to 4. This situation is shown in Figure 7-5.

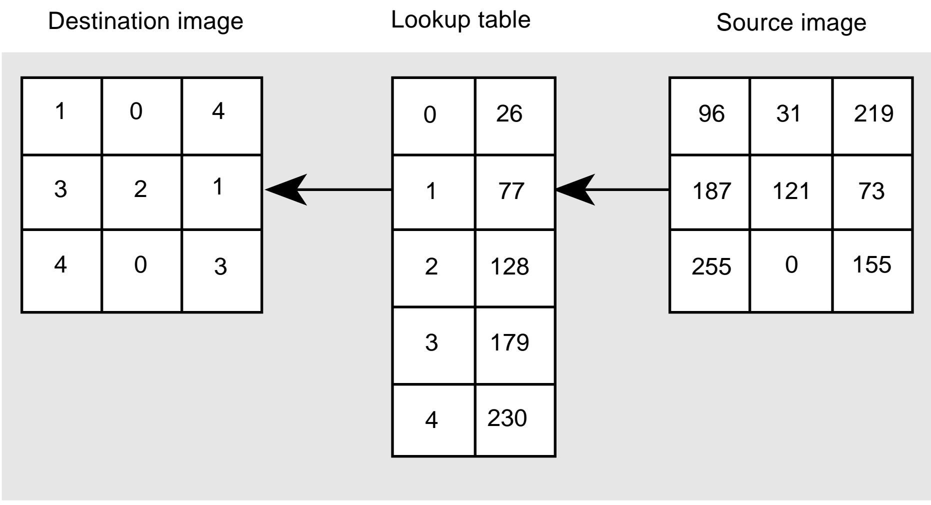

Figure 7-5 Using `xil_error_diffusion()` to Dither an Image

When the pixel in the middle of the top scanline of the source image (31) is passed through the lookup table, it is matched with a 26, so a 0 is written to the destination image. When the destination is displayed using the lookup table as its colormap, this pixel will have the value 26. Thus, in the dither-and-reconstruction process, there will be an error of 5 (31 - 26). To counteract the loss of data inherent in this process, `xil_error_diffusion()` distributes this error before passing the next pixel (219) through the lookup table. The way in which the error is distributed depends on the kernel created earlier. See Figure 7-6 for an example.

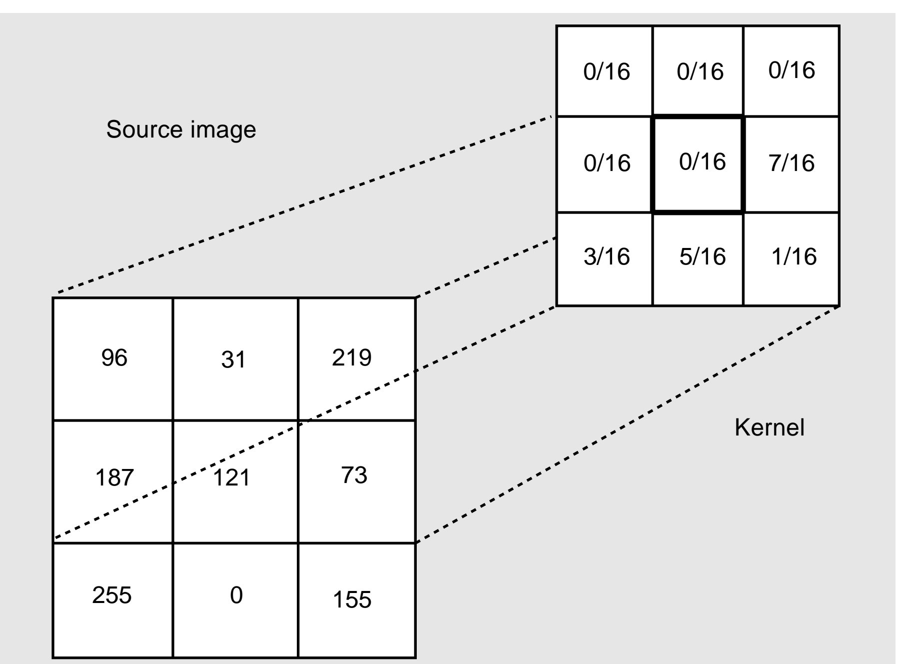

Figure 7-6 Error Diffusion

The kernel is laid on top of the source image so that its origin aligns with the last pixel to be passed through the lookup table. Before the next pixel is passed through, the following processing will take place:

- The pixel at 0,2 will be set to 219 + (5 * (7 / 16.0))
- The pixel at 1,0 will be set to 187 + (5 * (3 / 16.0))
- The pixel at 1,1 will be set to 121 + (5 * (5 / 16.0))
- The pixel at 1,2 will be set to 73 + (5 * (1 / 16.0))

This processing is fairly time-consuming, but can greatly reduce contouring in the dithered image.

xil_ordered_dither()

The XIL library's third dithering function is `xil_ordered_dither()`. Like `xil_nearest_color()` and `xil_error_diffusion()`, `xil_ordered_dither()` processes an image by performing an inverse lookup. However, two characteristics distinguish `xil_ordered_dither()` from the other dithering functions. One is that the lookup table the function uses to do its job must be a colorcube; it cannot be a generic lookup table. (For information about the differences between colorcubes and generic lookup tables, see the sections “Lookup Tables” on page 108 and “Colorcubes” on page 110.) Second, `xil_ordered_dither()` makes use of an XIL data structure called a dither mask to help eliminate contouring in the dithered image.

The function prototype for `xil_ordered_dither()` is shown below.

```
void xil_ordered_dither(XilImage src, XilImage dst,
    XilLookup lookup, XilDitherMask mask);
```

The parameters `src` and `dst` are handles to the source and destination images, and `lookup` is the lookup table to be used for the dither operation. As mentioned above, this lookup table must be a colorcube; that is, it must have been created with the function `xil_colorcube_create()`—or `xil_lookup_get_by_name()`. The final parameter, `mask`, is the dither mask, a data structure of type `XilDitherMask`.

You create the dither mask using the function `xil_dithermask_create()`, whose prototype is shown below.

```
XilDitherMask xil_dithermask_create(XilSystemState state,
        unsigned int width, unsigned int height,
        unsigned int nbands, float *data);
```

The parameters `width` and `height` define the width and height of the dither mask in pixels. These values are user defined. The next parameter, `nbands`, determines the number of bands in the dither mask. This number must match

the number of bands in the image being dithered. Finally, `data` is a pointer to the data to be stored in the dither mask. All the values in the mask must fall in the range 0.0 to 1.0. For example, the code below creates a 1-band, 4-by-4 dither mask.

```
XilDitherMask mask;
float mask_data[] = {
     0/16.0, 8/16.0, 2/16.0,10/16.0,
    12/16.0, 4/16.0,14/16.0, 6/16.0,
     3/16.0,11/16.0, 1/16.0, 9/16.0,
    15/16.0, 7/16.0,13/16.0, 5/16.0
}

mask = xil_dithermask_create(state, 4, 4, 1, mask_data);
```

Note – When you initialize the XIL library, four standard dither masks are created: 1- and 3-band 4-by-4 masks and 1- and 3-band 8-by-8 masks. To get a handle to one of these masks, you use the function `xil_dithermask_get_by_name()`.

Here's how the mask is used. Assume that you're dithering a grayscale image with values in the range 0 to 255 and that the destination is to contain values in the range 0 to 15. First, think of the 4-by-4 dither mask as having been replicated over the entire source image, as depicted in Figure 7-7.

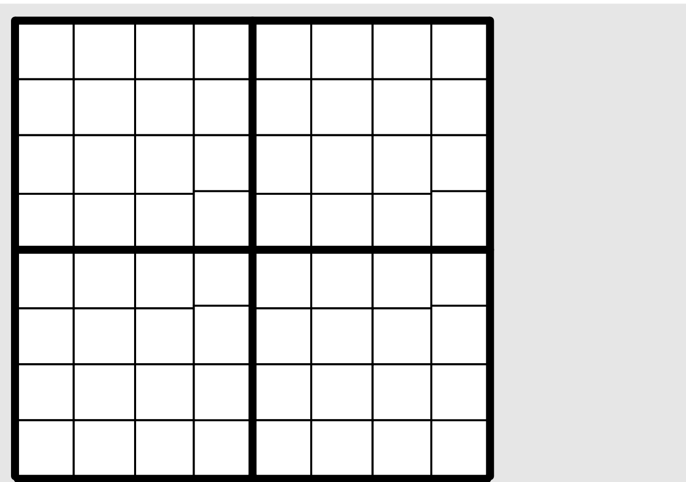

Figure 7-7 Dither Mask Replicated over a Source Image

At this point, each pixel in each 4-by-4 block of the image has been associated with (in this case) a unique dither-mask value. The `xil_ordered_dither()` function then performs the following steps for each pixel in the image:

1. **Passes the value through the lookup table.**
   In this step, the function is actually dividing values in the source image by 17 (the number of values that can be represented in 8 bits divided by the length of the colorcube minus one). That is, if no further processing were to take place, a 200 in the source image would map to an 11 since 200 divided by 17 equals 11.76.

2. **Considers the dither-mask value associated with the source-image pixel value.**
   The dither function now compares the fractional part of the quotient from the division operation with the source pixel's dither-mask value. If the fractional part is greater than the dither-mask value, the dividend shown above (11) is incremented by one, so a 12 is written to the destination image. Otherwise, an 11 is written to the image.

When you use the dither mask in this way, you're essentially performing averaging. This averaging helps prevent undesirable contours from appearing in the dithered image.

Table 7-4 below lists the XIL utility functions that affect dither masks and indicates what these functions do.

Table 7-4 Utility Functions for Dither Masks

Function Name	What It Does
`xil_dithermask_destroy`	Deallocates the memory used by a dither mask
`xil_dithermask_get_height`	Returns the height of a dither mask in pixels
`xil_dithermask_get_width`	Returns the width of a dither mask in pixels
`xil_dithermask_get_nbands`	Returns the number of bands in a dither mask
`xil_dithermask_create_copy`	Returns a copy of a dither mask
`xil_dithermask_set_name`	Sets the name of a dither mask
`xil_dithermask_get_name`	Returns a copy of a dither mask's name

When to Use Each Dithering Function

When you're trying to decide on the best way to perform a dither operation, it's helpful to think of the matrix shown in Table 7-5:

Table 7-5 Review of Dithering Operations

	Generic Lookup	Colorcube
xil_nearest_color()	✔	✔
xil_error_diffusion()	✔	✔
xil_ordered_dither()		✔

If you look at the right column—labeled "Colorcube"—you'll see that each of the dithering functions can perform its inverse lookup using a colorcube. If you compare the functions used with a colorcube, you'll see that choosing one over another involves a trade-off between speed and quality. The function `xil_nearest_color()` is the fastest, but produces the poorest quality image, and `xil_error_diffusion()` is the slowest, but produces the best quality image. The function `xil_ordered_dither()` is somewhat faster than `xil_error_diffusion()`, but does not produce quite as nice an image.

Similarly, if you compare operations that use generic lookup tables, the trade-off is one between speed and quality. The guidelines mentioned in the preceding paragraph apply, except that `xil_ordered_dither()` cannot work with a generic lookup table.

Now, if you read the table from left to right, you can compare operations that use generic lookup tables with corresponding operations that use colorcubes. Operations that use colorcubes are much faster than their counterparts because operations that use colorcubes do not need to search a lookup table. On the other hand, when you use a colorcube for dithering, you lose some of your control over the quality of the dithered image, because you cannot specify the exact values in the colorcube. You also lose the ability to use the colors in the currently installed colormap, again because you don't specify the values to be written to the lookup table.

Color-Space Conversion

The XIL library enables you to convert an image from any one to any other of the following color spaces:

- CCIR Rec. 709 RGB
- A linear version of CCIR Rec. 709 RGB
- CCIR Rec. 709 YC_bC_r
- A luminance-only space derived from CCIR Rec. 709 YC_bC_r
- A linear version of the luminance-only space mentioned above
- CCIR Rec. 601 YC_bC_r
- A luminance-only space derived from CCIR Rec. 601 YC_bC_r
- A YC_bC_r color space defined by Kodak for PhotoCD
- A linear CMY
- A linear CMYK

The procedure for converting an image from one color space to another is outlined below.

1. **Set the color-space attribute of your source image.**
   When you create an XIL image, it does not have a color space associated with it. To associate one with it, you must call the function `xil_colorspace_get_by_name()` to get the appropriate color-space data structure; then, you call `xil_set_colorspace()` to actually set the attribute.

The code fragment below shows a program getting a data structure of type `XilColorspace` that describes a linear RGB color space.

```
XilColorspace colorspace;

colorspace = xil_colorspace_get_by_name(state, "rgblinear");
```

The parameter `state` is a handle to the system state, and `rgblinear` is a string identifying the color space. The strings you use to identify the various XIL color spaces are listed in Table 7-6.

Table 7-6 Strings Used to Specify Color Spaces

Color Space	String
CCIR Rec. 709 RGB	`rgb709`
Linear version of CCIR Rec. 709 RGB	`rgblinear`
CCIR Rec. 709 YC_bC_r	`ycc709`
Luminance-only space derived from CCIR Rec. 709 YC_bC_r	`y709`
Linear version of `y709`	`ylinear`
CCIR Rec. 601 YC_bC_r	`ycc601`
Luminance-only space derived from CCIR Rec. 601 YC_bC_r	`y601`
YC_bC_r color space defined by Kodak for Photo CD	`photoycc`
Linear CMY	`cmy`
Linear CMYK	`cmyk`

After you've gotten this color-space data structure, you set an image's color-space attribute using the function `xil_set_colorspace()`.

```
XilImage image;
XilColorspace colorspace;

xil_set_colorspace(image, colorspace);
```

2. **Create a destination image, and assign it the color space to which you want to convert your source image.**
   The destination image must have the same width, height, and data type as the source image, and it must have a number of bands appropriate to the color space you will assign it. That is, if you want to assign the destination image the color space `ycc709`, it must have three bands. You assign a color space to the destination using the same method you used to assign a color space to the source image.

Note – It is possible to perform an in-place color conversion if the destination image is a child of the source and has the same width and height as the source. This procedure is not generally recommended, however, because the color-conversion operation will overwrite some or all of the data in the source image.

---

3. **Perform the color conversion.**
   You perform the color conversion by calling the function `xil_color_convert()`. This function takes handles to your source and destination images as its two parameters.

Note – Regions of interest are ignored when you perform a color conversion.

There are many applications for this type of color conversion. For example, you may want to:

- Convert an image from a YC_bC_r color space to an RGB color space in order to display the image
- Convert an image from an RGB color space to a YC_bC_r color space to prepare it as input to the JPEG or CellB compressor
- Convert an image to the CMY or CMYK color space to prepare it to be printed on a subtractive color printer
- Convert an RGB or YC_bC_r image to the `y709`, `ylinear`, or `y601` color space as a way of converting a 24-bit image to an 8-bit grayscale image
- Convert an `XIL_BYTE` image from a linear color space to a gamma-corrected color space to prevent contouring in low-intensity regions of the image

Black Generation

The XIL library also enables you to perform undercolor removal and black generation on CMYK images—images to be printed on a four-color printing system. Normally, the process of undercolor removal and black generation goes something like this. For each pixel in the image:

- Black (K) equals the minimum of C, M, and Y
- C = C - K
- M = M - K
- Y = Y - K

However, the XIL function that performs undercolor removal and black generation gives you more control over how much cyan, magenta, and yellow will be removed from your image and over how much black will be added to it. This function is called `xil_black_generation()`, and its function prototype is shown below.

```
void xil_black_generation(XilImage src, XilImage dst,
    float black, float undercolor);
```

The parameters `src` and `dst` are handles to your source and destination images. Both images must be CMYK images. Therefore, if you want to process a CMY image, you must first convert it to a CMYK image. For instructions on how to do this, see the section “Color-Space Conversion” on page 122. In-place operations are allowed.

The parameters `black` and `undercolor` are used as follows. For each pixel in your source image:

- K = (`black` * minimum of C, M, and Y)
- C = C - (`undercolor` * minimum of C, M, and Y)
- M = M - (`undercolor` * minimum of C, M, and Y)
- Y = Y - (`undercolor` * minimum of C, M, and Y)

Thus, if both `black` and `undercolor` are set to 1.0, you'll get standard undercolor removal. You'll frequently achieve better results, though, if the percentage of black you add to your image is slightly higher than the percentage of C, M, and Y you remove from it. For example, you might try using the following call:

```
XilImage src, dst;

xil_black_generation(src, dst, 0.7, 0.5);
```

Note – Regions of interest are ignored when you perform undercolor removal.

Error Handling 8

The list below summarizes the most important facts about XIL error handling.

- Error handling is asynchronous. For example, assume that your program includes a call to `xil_copy()` that attempts to copy a three-band image to a single-band image. This is an error, and at some point the library will report that the source and destination images do not contain the same number of bands. However, this error may not be reported when the line in your program containing the `xil_copy()` is called. This is true because the XIL library's deferred execution scheme enables it to store many operations before executing them. (See Chapter 21, "Acceleration in XIL Programs," for more information on this deferred execution feature.) Later, when a set of stored operations is performed, the error will be reported.

  The fact that error reporting is asynchronous does not really affect the way you handle errors in your XIL programs, but it is something you should be aware of.

- The amount of error handling that you can do by simply checking the return values of XIL functions is limited. About half of the functions in the XIL library return `void`, and for the remaining functions, while you may be able to tell from the return value whether an error occurred, you won't be able to determine exactly what the error was. To handle errors properly, your program needs to examine the error data structure (of type `XilError`) that the library creates whenever an error occurs. This structure contains such information as an error ID and a string containing a description of the error.
- Your application can deal with the error structure referred to above in one of two ways:

- One option is to do nothing. In this case, the library's default error handler will look at the error structure and print a message to `stderr`. As shown in "An Example" on page 135, this option isn't recommended because the default handler prints information that is intended for application developers and isn't useful to an end user.
- Your other option is to write your own error handler and to install it; this is the recommended option. For information on writing such an error handler, see the next section "Writing an Error Handler," and for information on installing the routine, see the section "Installing and Chaining Error Handlers" on page 136.

Note – You can actually have more than one error handler installed at a time. In addition, you can chain these error handlers so that if the first error handler in the chain does not deal with an error, it can pass the error structure containing information about the error to a second error handler, and so on. This subject is explained in "Installing and Chaining Error Handlers" on page 136.

Writing an Error Handler

Any error handler you write should take one parameter, a handle to the error structure the library created when the last error occurred. In addition, the function should return a value of type `Xil_boolean`. Thus, the header of your function should look like this:

```
Xil_boolean function-name(XilError error);
```

The function should return `TRUE` if it has handled the error and `FALSE` if it has not.

The XIL library contains a number of routines you can use in building the body of your error handler. These are discussed in the next section, "Functions You Can Call in Your Error Handler." The section after that—"An Example" on page 135—shows how you might use these functions to write an error handler.

Functions You Can Call in Your Error Handler

This section discusses the XIL functions you can use in writing your own error handler. These functions fall broadly into two groups: those that get information from the error structure that the library passes to your error handler and those that retrieve information from structures of type `XilObject`.

Getting Information from an Error Structure

The functions listed in this section all take as their only parameter the name of the error structure the library has passed to your error handler.

xil_error_get_string()

This function returns an internationalized string (`char *`) that describes the error that just occurred. For a complete list of possible messages, see Appendix B, "XIL Error Messages."

You should either use this string immediately or make a copy of it because the pointer to it will become invalid after the next call your program makes to `dgettext(3I)` or to one of the XIL error functions that returns a string. Also, you should not free or modify this string.

xil_error_get_id()

This function returns a string (`char *`) of the form `di-`*n*, where is *n* is a unique identifier. For a complete list of possible error IDs, see Appendix B, "XIL Error Messages."

You should not free or modify this string.

xil_error_get_category()

This function returns an enumeration constant of type XilErrorCategory. The declaration of this enumeration is shown below.

```
enum XilErrorCategory {
    XIL_ERROR_SYSTEM,
    XIL_ERROR_RESOURCE,
    XIL_ERROR_ARITHMETIC,
    XIL_ERROR_CIS_DATA,
    XIL_ERROR_USER,
    XIL_ERROR_CONFIGURATION,
    XIL_ERROR_OTHER};
```

Table 8-1 explains briefly what types of error fall into each category.

Table 8-1 XIL Error Categories

Category	Explanation
XIL_ERROR_SYSTEM	The library cannot perform an operation correctly. This type of error is usually a secondary error caused by a user, resource, or configuration error.
XIL_ERROR_RESOURCE	Usually means that there is insufficient memory for the library to create an image or other data structure.
XIL_ERROR_ARITHMETIC	Indicates an arithmetic error such as a divide by 0.
XIL_ERROR_CIS_DATA	Means that the datastream being decompressed does not conform to the appropriate datastream definition.

Table 8-1 XIL Error Categories

Category	Explanation
XIL_ERROR_USER	The programmer has specified an invalid parameter to an XIL function. For example, for a lookup operation, the data type of the source image and the input data type of the lookup table might not match.
XIL_ERROR_CONFIGURATION	Usually indicates that the XIL software has not been installed properly or that an environment variable has not been set properly. The specific error may be that the runtime system is unable to find a loadable driver it needs.
XIL_ERROR_OTHER	Miscellaneous errors, such as Xlib and DGA errors.

xil_error_get_category_string()

This function returns a string (`char *`) that indicates which category of error has occurred. The strings for the possible categories are

- `System`
- `Resource`
- `Arithmetic`
- `Cis Data`
- `User`
- `Configuration`
- `Other`

As with the string returned by `xil_error_get_string()`, you should either use this string immediately or make a copy of it because the pointer to it will become invalid after the next call your program makes to `dgettext(3I)` or to one of the XIL error functions that returns a string. Also, you should not free this string.

xil_error_get_primary()

This function returns a value of type `Xil_boolean`, either `TRUE` or `FALSE`. A value of `TRUE` indicates that the error being processed was the primary cause of a problem you're seeing. For example, if you run out of memory and

xil_create() is unable to create an image, the primary error produced will be a resource error. The XIL library may also generate secondary errors as the NULL image is used internally.

xil_error_get_location()

This function returns a string (char *) that indicates where in the XIL library the error was generated. This location will help Customer Support to resolve any problems you report to them.

You should not free this string.

xil_error_get_object()

This function returns a handle of type XilObject to the object affected by the error. This object may be:

- An image
- An image type
- A lookup table
- A compressed image sequence (CIS)
- A dither mask
- A kernel
- A structuring element
- A region of interest
- A histogram

Once you have this handle, you can use the functions discussed in the next section to obtain further information about the object.

Note – If the error did not affect a particular object, you will receive a NULL handle.

Getting Information About the Object Affected by the Error

If your error handler uses the function xil_error_get_object() to retrieve a handle to the object affected by the last error, you can use the functions discussed below to request further information about the object.

xil_object_get_error_string()

This function attempts to get from an object a string containing additional information about the object. For instance, if you query a CIS object, you will get a string containing information about the values of the start frame, read frame, and write frame at the time of the error.

The prototype for this function is shown below.

```
void xil_object_get_error_string(XilObject object,
    char *string, int string_size);
```

The parameter `object` is the object handle returned by `xil_error_get_object()`; `string` is an array of `char`s in which the string will be returned; and `string_size` is the maximum number of characters the array can contain.

In the code fragment below, an error handler uses `xil_object_get_error_string()` to request information about the object affected by an error and prints any information that is available.

```
#define MAX 1024

Xil_boolean UserDefinedErrorFunc(XilError error)
{
    XilObject object;
    char buffer[MAX];

    object = xil_error_get_object(error);
    if (object) {
        xil_object_get_error_string(object, buffer, MAX);
        if (buffer[0] != 0)
            fprintf(stderr, "object info: %s\n", buffer);
```

xil_object_get_type()

You use this function to determine what kind of object the last error affected. The prototype for the function is shown below.

```
XilObjectType xil_object_get_type(XilObject object);
```

The parameter to the function is the object handle returned by `xil_get_error_object()`. The return value is an enumeration constant of type `XilObjectType`. The declaration of this enumeration is written as follows:

```
enum XilObjectType {
   XIL_IMAGE,
   XIL_IMAGE_TYPE,
   XIL_LOOKUP,
   XIL_CIS,
   XIL_DITHER_MASK,
   XIL_KERNEL,
   XIL_SEL,
   XIL_ROI,
   XIL_ROI_LIST,
   XIL_HISTOGRAM,
};
```

Once you know the type of object that was affected by the last error, you can cast the object handle returned by `xil_get_error_object()` to a handle to the particular type of object you have. For example, if the object is an image, you can cast the handle of type `XilObject` to one of type `XilImage`. Doing so gives you the ability to write code like the following:

```
XilObject object;

object = xil_error_get_object(error);
if (object)
    if (xil_object_get_type(object) == XIL_IMAGE)
        fprintf(stderr, "image bands: %d\n",
            xil_get_nbands((XilImage)object));
```

Because the error handler has a handle to the image, it can use any of the XIL functions that read image attributes to get additional information about the image. In this case, the error handler is finding out how many bands are in the image.

An Example

The library's default error handler prints a message like the one shown below each time an error occurs.

```
XilDefaultErrorFunc:
   error category: User
     error string: Destination colorspace not specified
         error id: di-203
   primary error detected at location colorConvert79 in XIL
```

As you can see, it indicates that the XIL default error handler handled the error, and it prints the category of the error, an error message, an error ID, and the location at which the error occurred, indicating whether the error was a primary or secondary error.

This information should be useful to the application developer because it verifies that the XIL handler handled this error, and also prints the error location, which is useful when contacting Customer Support. However, you wouldn't want your application to identify the error handler and error location because this information isn't useful to an end user and, moreover, would be confusing to that user. The code below shows how you could implement an error handler that doesn't identify the error handler and doesn't print the error location.

```
Xil_boolean UserDefinedErrorFunc(XilError error)
{
    if (xil_error_get_primary(error) == TRUE)
        fprintf(stderr, "\nPrimary Error:\n");
    else
        fprintf(stderr, "\nSecondary Error:\n");
    fprintf(stderr, "   error category: %s\n",
        xil_error_get_category_string(error));
    fprintf(stderr, "     error string: %s\n",
        xil_error_get_string(error));
    fprintf(stderr, "         error id: %s\n",
        xil_error_get_id(error));
    return TRUE;
}
```

Installing and Chaining Error Handlers

By default, when an XIL error occurs, the library creates an error structure of type `XilError` and passes that structure to the default error handler. To override this behavior, you can write your own error handler and install it. Once you install your error handler, the library will call it instead of the default error handler when an error occurs. If you write and install more than one error handler, the library always calls the most recently installed error routine. For more information on installing error handlers, see the next section, "Installing Error Handlers."

As mentioned above, the XIL library allows you to have more than one user-defined error routine installed at the same time. At first glance, this seems a waste because the library can only call the most recently installed routine when an error occurs. However, having more than one user-defined error handler installed is often useful because the library allows you to call one error handler from another, an action referred to as *chaining*. For an explanation of why you might want to chain error routines and how you go about chaining them, see the section "Chaining Error Handlers" on page 139.

Installing Error Handlers

To install an error handler you've written, you call the function `xil_install_error_handler()`, whose prototype is shown below.

```
int xil_install_error_handler(XilSystemState state,
    XilErrorFunc func);
```

The function's return value indicates whether the installation was successful. A value of `XIL_SUCCESS` means that it was successful, and a value of `XIL_FAILURE` means that it wasn't.

The parameter `state` is the system-state data structure you received when you initialized the library, and `func` is a pointer to your error handler. (`XilErrorFunc` is a defined type representing a pointer to a function.) Thus, the call you would use to install the error handler shown in the section "An Example" on page 135 would look like this.

```
XilSystemState state;

xil_install_error_handler(state, UserDefinedErrorFunc);
```

When you install an error handler, the library adds it to a list of error handlers that it recognizes. Figure 8-1 shows what the list would look like at this point.

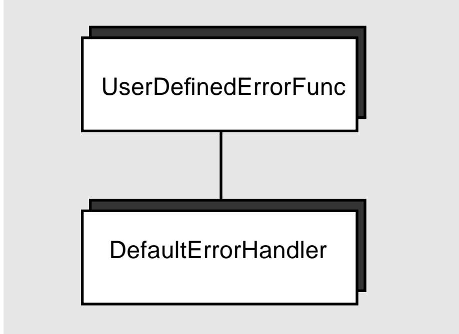

Figure 8-1 List of Error Handlers

If you then installed a third error handler, called `UserDefinedErrorFunc2`, the list would grow as shown in Figure 8-2.

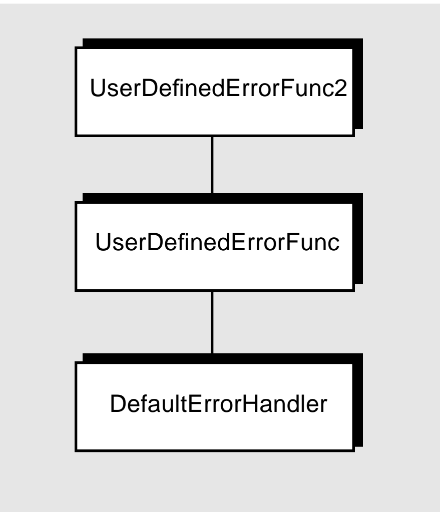

Figure 8-2 Adding to the List of Error Handlers

When an XIL error occurs, the library always calls the topmost routine (the last routine to be installed) in the list.

The library also includes a routine for removing error handlers from this list. Its name is `xil_remove_error_handler()`.

```
void xil_remove_error_handler(XilSystemState state,
    XilErrorFunc func);
```

The parameters to this function are the same as those for `xil_install_error_handler()`.

Note – If the same error handler has been installed twice, only the most recent version of the error handler will be removed.

Chaining Error Handlers

Again, having more than one error handler installed is useful because the library enables you to call one error handler from another. From the topmost error handler in a list of handlers, you can call either the next-to-last error handler to be installed or the default error handler. Thus, in your topmost error routine, you might handle only resource errors. If the error you're passed is a resource error, you take whatever action is appropriate and return `TRUE`; otherwise, you call the next error handler in the list. This next error handler might also check for one special type of error. Like the previous routine, if it is passed the proper type of error, it handles the error and returns `TRUE`; if it receives another type of error, it might call the error handler below it on the list. This process can continue until the error handler being called is the default error handler. This routine will handle any XIL error and always returns `TRUE`.

The functions you use to chain error handlers are `xil_call_next_error_handler()` and `xil_default_error_handler()`.

```
Xil_boolean xil_call_next_error_handler(XilError error);

Xil_boolean xil_default_error_handler(XilError error);
```

The parameter `error` is the error structure that the library passed to the topmost error handler in its list of error handlers (the last one to be installed). The return value of both functions indicates whether the error routine being called handled the error. A return value of `TRUE` indicates that the routine did handle the error, and a return value of `FALSE` indicates that it didn't.

The brief example below shows the routine `xil_call_next_error_handler()` being used in an error-handling routine. This routine simply ensures that if the error is a resource error, the program will exit after handling the error in the normal way.

```
Xil_boolean resource_errors(XilError error)
{
    int return_value;

    return_value = xil_call_next_error_handler(error);
    if (xil_error_get_category(error) == XIL_ERROR_RESOURCE)
        exit(1);
    return return_value;
}
```

Arithmetic, Relational, and Logical Functions 9

This chapter discusses the XIL library's arithmetic, relational, and logical functions. These functions can be used to lighten or darken an image, to increase or decrease contrast in an image, or to produce the "negative" of an image.

Note – All the functions discussed in this chapter can be performed in place; that is, the source and destination images can be the same image.

Arithmetic Functions

If you have two source images and a destination image, the XIL library enables you to:

- Add the two source images and store the results in the destination (`xil_add()`)
- Subtract one source image from the other and store the results in the destination (`xil_subtract()`)
- Multiply the two source images and store the results in the destination (`xil_multiply()`)
- Divide one source image into the other and store the results in the destination (`xil_divide()`)

For a single source image, the library lets you find the absolute value of pixels in an image and store the result in a destination image (`xil_absolute()`).

For each of these operations to work, its source and destination images must have the same data type: XIL_BIT, XIL_BYTE, or XIL_SHORT. Also, the images must have the same number of bands. The images do *not* have to have the same width and height. For an explanation of which pixels an operation will affect if the images have different sizes, see "Region of Interest" on page 48.

When you add two images, you take the value at location 0,0 in one source image, add it to the value at location 0,0 in a second source image, and write the sum at location 0,0 in a destination image. You then follow the same procedure for all other points in the images. Subtraction, multiplication, and division are handled similarly. When multiband images are involved, the arithmetic operation is performed on corresponding bands in the source images; that is, band 0 in the first source image is added to band 0 in the second source image, and so on.

This all seems very straightforward. But there are a few points you should bear in mind when using these functions:

- **If the result of an operation is out of range for a particular data type, the result is not truncated, but is clamped to the minimum or maximum value for the data type.** Thus, if you're working with XIL_BYTE images, adding a 200 and a 200 gives a result of 255 because 400 cannot be represented in 8 bits and the greatest valid value for an XIL_BYTE image is 255. Likewise, if you subtract 200 from 100, the result is 0. Table 9-1 indicates the valid range of values for each XIL data type.

  *Table 9-1*Valid Values for Each XIL Data Type

Data Type	Lowest Value	Greatest Value
XIL_BIT	0	1
XIL_BYTE	0	255
XIL_SHORT	-32768	32767

- **Division by 0 is permitted.** If the image serving as the divisor contains 0's, you will receive *one* error message indicating that division by 0 occurred. However, the division operation will be performed. If a 0 in an image is divided by 0, the result is 0, and if any other value is divided by 0, the result is the greatest valid value for the data type (or possibly the lowest valid value if the numerator is a negative value of type XIL_SHORT).

- **All division is performed using floating-point operands.** The `xil_divide()` function casts each dividend and divisor to a `float` before doing the division. It then rounds off each quotient and casts it to the type of the images involved in the division.

The `xil_absolute()` function finds the absolute value of pixels in a source image and stores the result in a destination image. Since `XIL_BIT` and `XIL_BYTE` images don't have negative values, `xil_absolute()` is useful only for `XIL_SHORT` images. For each band in the image, the absolute value at location 0,0 in the source is written to location 0,0 in the destination; the same procedure is then followed for all other points in the image.

Note – On `XIL_BIT` and `XIL_BYTE` images, `xil_absolute` is effectively `xil_copy`.

Relational Functions

If you have two source images and a destination image, the XIL library enables you to:

- Find the larger of pixels in the two source images and store the results in the destination (`xil_max()`)
- Find the lesser of pixels in the two source images and store the results in the destination (`xil_min()`)

As with the arithmetic operations, `xil_max()` and `xil_min()` require the source and destination images to have the same data type and the same number of bands. The images do not have to have the same width and height.

Taking the maximum or minimum values from two images performs a band by band comparison. For example, if `src1` and `src2` images are 4-banded images, comparing the value at location 0,0 writes to location 0,0 in the destination image the maximum (or minimum) for each source band; thus, the first band at 0,0 in the destination might receive its value from `src1`, whereas the remaining three bands might receive their values from `src2`. This band by band comparison is repeated for all other points in the source images.

Note – Don't confuse these relational functions with finding the maximum or minimum pixel value within each band of a single image. To find those values, use `xil_extrema()`, discussed in "Finding the Minimum and Maximum Values in an Image" on page 178.

Logical Functions

If you have two source images and a destination image, the XIL library also enables you to:

- Take the bitwise AND of the two source images and store the results in the destination (`xil_and()`)
- Take the bitwise OR of the two source images and store the results in the destination (`xil_or()`)
- Take the bitwise XOR of the two source images and store the results in the destination (`xil_xor()`)

As with the arithmetic and relational operations, for these operations to work, all three images must have the same data type and the same number of bands. The images do not have to have the same width and height.

The XIL library also contains a bitwise NOT operator (`xil_not`). This function works on a single source image. It looks at the values in that image as binary values and changes all the 1's in those values to 0's, and all the 0's to 1's. The function then writes this one's complement version of the source image to the destination.

Operations with Constants

For each of the arithmetic and logical functions that operate on two source images—addition and so on—the XIL library includes a similar function that takes as input one source image and a constant. For instance, the library contains the function `xil_add()`, whose prototype is shown below.

```
void xil_add (XilImage src1, XilImage src2, XilImage dst);
```

It also contains the function `xil_add_const()`, whose prototype is shown below.

```
void xil_add_const (XilImage src1, float *constants,
    XilImage dst);
```

To add a constant to a single-band image, you declare a one-element array of type `float`, assign that element a value, and then call `xil_add_const()`. For instance, the code to add 8.0 to each value in a single-band image might look like this:

```
XilImage src1, dst;
float constants[1] = {8.0};

xil_add_const(src1, constants, dst);
```

This operation is roughly the equivalent of adding to `src1` a second source image, all of whose values are 8. The only difference between these two operations is that using a constant instead of a second source image enables you to use noninteger values in the operation. When you add a constant to an image, each sum is rounded and then cast to the data type of the source and destination images.

Note – As stated on page 142, if the result of an operation is out of range, the value is clamped and assigned the minimum or maximum value for the data type. It's important to realize that the operation is performed first, then the result is clamped. For example, if the operation is division and you divide a negative constant into the values for an `XIL_BYTE` image, the division is first performed. Thus, for a pixel value of 10 divided by the constant -2, the result is

10 / -2 = -5, which is clamped to 0.

Notice that clamping the -2 to 0 before the division is performed would yield 10 / 0, which would result in a value of 255 in the destination image.

In operations that involve a multiband source image and a constant, the constant must actually be an array of constants, and the number of constants in the array must equal the number of bands in the image. The following code adds a different constant to each band of a 3-band image.

```
XilImage src1, dst;
float constants[3] = {8.0, 12.0, 2.0};

xil_add_const(src1, constants, dst);
```

This operation adds 8.0 to each value in band 0 of the image, 12.0 to each value in band 1, and 2.0 to each value in band 2.

For operations that are not associative—subtraction and division—the XIL library enables you to specify the constant as the first operand or the second. That is, for a subtraction, the constant can be either the minuend or the subtrahend, and for a division, it can be either the dividend or the divisor.

Table 9-2 lists the arithmetic functions that can take constants.

Table 9-2 Arithmetic Operations Using a Source Image and a Constant

Function Name	What the Function Does
`xil_add_const`	Adds an image and a constant
`xil_subtract_const`	Subtracts a constant from an image
`xil_subtract_from_const`	Subtracts an image from a constant
`xil_multiply_const`	Multiplies an image by a constant
`xil_divide_by_const`	Divides an image by a constant
`xil_divide_into_const`	Divides an image into a constant

You can also perform logical operations using an image and a constant. For instance, besides the function `xil_and()`, the library contains the function `xil_and_const()`, whose prototype is shown below.

```
void xil_and_const (XilImage src1, unsigned int *constants,
    XilImage dst);
```

To take the logical AND of a constant and a single-band image, you declare a one-element array of type `unsigned int`, assign that element a value, and then call `xil_and_const()`. For instance, the code to find the logical AND of each value in an image and 8 might look like this:

```
XilImage src1, dst;
unsigned int constants[1] = {8};

xil_and_const(src1, constants, dst);
```

This operation is the equivalent of taking the logical AND of `src1` and a second source image, all of whose values are 8.

Table 9-3 lists the logical functions that can take constants.

Table 9-3 Logical Operations Using a Source Image and a Constant

Function Name	What the Function Does
`xil_and_const`	ANDs an image and a constant
`xil_or_const`	ORs an image and a constant
`xil_xor_const`	XORs an image and a constant

Note – The relational operations `xil_max()` and `xil_min()` always require two source images and cannot operate on one source image and one constant.

Arithmetic and Logical Operations with Bit Images

The arithmetic and logical operations discussed in this chapter *are* defined for `XIL_BIT` images. If you bear in mind that the values in the destination image produced by such an operation cannot fall below 0 or go above 1, the results to expect are usually obvious. For example, when you add two `XIL_BIT` images, the following calculations are used:

- 0 + 0 = 0
- 0 + 1 = 1
- 1 + 0 = 1
- 1 + 1 = 1

Note that 1 plus 1 equals 1. Two is not a valid value for a 1-bit data element, so the result has been clamped at 1. Similarly, when you subtract one `XIL_BIT` image from another, 0 minus 1 equals 0 because -1 is not a valid value.

When performing division, bear in mind that 0 divided by 0 equals 0 and that a positive nonzero value divided by 0 equals the maximum valid value for the data type, in this case 1.

When you perform an arithmetic or logical operation that involves a bit image and a constant, similar rules apply. For arithmetic operations, think of the following sequence of steps taking place:

1. Addition, subtraction, multiplication, or division is performed using a 1-bit value (0 or 1) and a floating-point value.
2. The result is rounded and clamped to 0 or 1.
3. The 0 or 1 is cast to the data type `XIL_BIT`.

For logical operations involving a constant, the constant (an `unsigned int`) is set to 1 if it is greater than 0; then the operation is performed.

Geometric Functions 10

This chapter discusses the XIL library's geometric functions. These functions enable you to perform such operations as:

- Translating an image (moving an image up, down, left, or right)
- Scaling an image (changing an image's width or height)
- Rotating an image
- Warping an image
- Transposing an image (flipping an image across a horizontal or vertical axis, or across a diagonal)

The following sections discuss the XIL functions that perform the geometric functions; because you must pass an interpolation option as an argument on most of them, this chapter begins with a discussion of the available interpolation options.

Note – The functions discussed in this chapter cannot be performed in place. That is, the source and destination images for a geometric operation must be different images.

Interpolation Options

Geometrically transforming images centers around the notion of *point sampling*. In point sampling, each pixel in a destination image is located with integer coordinates at a distinct point *D* in the image plane. The geometric transform *T* identifies each destination pixel with a corresponding point *S* in the source image; thus, *S* is the point that *T* maps to *D*. In general, *S* doesn't correspond to a single source pixel; that is, it doesn't have integer coordinates. Therefore, the value assigned to the pixel *D* must be computed as an interpolated combination of the pixel values closest to *S* in the source image.

When performing most geometric transformations, you must specify the interpolation method to be used in calculating the destination pixel values. Table 10-1 shows the interpolation options the XIL library provides, and the strings you use to request them.

Table 10-1 Types of Interpolation

Interpolation Type	String	What it Assigns to *D*
Nearest neighbor	`nearest`	The value of the pixel nearest *S*
Bilinear	`bilinear`	A value that's a bilinear function of the four pixels nearest *S*
Bicubic	`bicubic`	A value that's a bicubic function of the sixteen pixels nearest *S*
General	`general`	A value that's a separable user-defined function of the pixels in a rectangular region surrounding *S*

To see how the interpolation type can affect the result of a geometric operation, see Color Plate 1.

Nearest Neighbor Interpolation

Nearest-neighbor interpolation, sometimes called zero-order interpolation, is the fastest interpolation method because it simply assigns to point *D* in the destination image the value of the pixel nearest S in the source image (see "Interpolation Options" on page 150). Though it's a good choice when speed is important, nearest-neighbor interpolation can produce undesirable artifacts

in the destination image, especially near edges where there may be a big change in color or gray level between two adjacent pixels. For example, smooth lines in the source image may show up as jagged lines in the destination.

Bilinear Interpolation

A routine that performs bilinear, or first-order, interpolation assigns to point *D* in the destination image a value that's a bilinear function of the four pixels nearest *S* in the source image (see "Interpolation Options" on page 150). This interpolation type yields better results than nearest-neighbor interpolation, but can itself have an undesirable smoothing effect on an image. To alleviate this problem, you can use bicubic interpolation.

Bicubic Interpolation

A routine that performs bicubic interpolation assigns to point *D* in the destination image a value that's a bicubic function of the 16 pixels nearest S in the source image (see "Interpolation Options" on page 150). Using bicubic interpolation preserves fine detail present in a source image, but it takes more time than the nearest-neighbor or bilinear interpolation methods.

General Interpolation

The interpolation options discussed so far base the interpolated values on relatively few pixels: nearest neighbor uses 1 pixel, bilinear uses 4 pixels, and bicubic uses 16. If these options don't provide the quality you need, you can use the general interpolation option. For example, if you're subsampling by a factor of four, bicubic interpolation would result in aliasing artifacts that can be improved by using more source pixels in the interpolation. General interpolation lets you:

- Determine how many pixels nearest point *S* in the source image are used to calculate the interpolated pixel value of point D in the destination image (see "Interpolation Options" on page 150). If desired, you could use every pixel in the source image.

- Weight the pixels used in the calculation. By ensuring that the pixels closest to point *S* in the source image have more influence on the value assigned to *D* than pixels that are further from *S*, you can reduce the contrast between adjacent pixels in the destination image, thus providing smoother line and color transitions.
- Designate the number of pixel subsamples to use for interpolating pixel values in the destination. This effectively lets you divide into fractional locations the space between adjacent source-image pixels so you can interpolate a destination pixel's value differently, depending on which location point *S* falls in.

To use general interpolation, you must:

1. Create horizontal and vertical interpolation tables; these form the filters or *kernels* for the interpolation. The tables determine the kernel sizes, weighted pixel values, and number of subsamples used for a general interpolation.
2. Set the interpolation tables on the system-state object. The tables affect all general interpolation operations using images created from this system-state object.
3. Pass `general` as the interpolation string on any XIL function that requires an interpolation argument. This causes the interpolation tables to be used as the interpolation method.
4. Destroy the interpolation tables when they are no longer needed. This releases the memory that was allocated for them.

The sections that follow discuss these steps in more detail. Before proceeding to them, however, it's useful to consider a conceptual model of a general interpolation.

Figure 10-1 represents an interpolation kernel as a 5-by-5 matrix of weighted values; for simplicity, the values are labelled `a` through `y`.

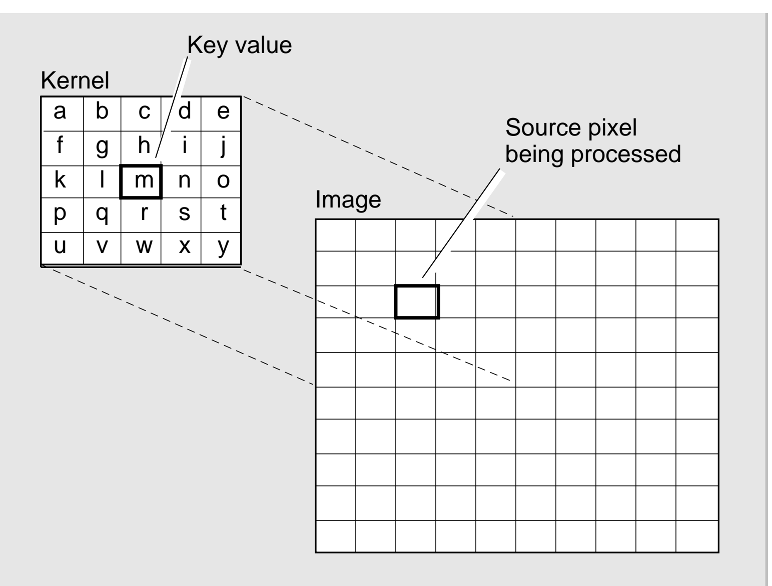

Figure 10-1 Conceptual Model of a General Interpolation

During a general interpolation, the kernel's key value—generally its center value—is laid over the source-image pixel to be processed; this means the other kernel values lie over neighboring pixels. Each source-image pixel that is covered by the kernel is then multiplied by the kernel value that lies over it. In Figure 10-1, the source-image pixel 2,2 is multiplied by kernel value `m`; pixel 0,0 is multiplied by kernel value `a`; pixel 1,0 is multiplied by `b`; and so on. The multiplication products are then summed together, and this sum becomes the pixel value in the destination.

Though Figure 10-1 provides a useful conceptual model, the general interpolation method doesn't actually construct the two-dimensional kernel shown in the figure because it's computationally expensive. To improve efficiency, general interpolation uses separate horizontal and vertical vector arrays—the kernels in the interpolation tables you create—to calculate the same values a two-dimensional kernel would calculate. The vector arrays require you to provide fewer data elements for the kernel values; this reduction is particularly significant for large tables with many subsamples. Nonetheless, for a horizontal table with M elements and a vertical table with N elements, the number of pixels that contribute to the interpolated value is still given by M * N.

Note – To translate the horizontal and vertical kernels into a two-dimensional matrix as shown in Figure 10-1, the 1xn kernel values would have to be matrix-multiplied to obtain the corresponding two-dimensional kernel values.

When interpolating pixels that are at, or near, the source image edges, the general interpolation method temporarily reduces the kernel height and width so the edges fit within the defined kernel size. This strategy lets the method interpolate pixels that originally formed the image edges. Pixels outside of the image don't contribute to any interpolation values.

Creating Vertical and Horizontal Interpolation Tables

To support general interpolation, XIL has an `XilInterpolationTable` object, which is an array of 1xn kernels representing either a horizontal or vertical interpolation filter. General interpolation can be performed using two interpolation tables, one to represent the horizontal filter and one to represent the vertical filter. Or it can be performed using a single interpolation table that represents either the horizontal or the vertical filter, in which case the missing matrix dimension is 1. For example, if the horizontal table defines 1-by-7 kernels and the vertical table defines 1-by-5 kernels, the effective matrix is 7-by-5; if only the 1-by-7 horizontal table is defined, the effective matrix is 7-by-1. If both the horizontal and vertical interpolation tables are `NULL`, nearest-neighbor interpolation is performed.

To create an `XilInterpolationTable` object, you call the `xil_interpolation_table_create()` function, whose prototype is shown below.

```
XilInterpolationTable xil_interpolation_table_create(
     XilSystemState state, unsigned int kernel_size,
     unsigned int subsamples, float* data);
```

The table's data type is `XIL_FLOAT`. The parameter `state` is the handle to the system-state structure created when you initialized the XIL library. The parameter `kernel_size` specifies the number of elements in each kernel. The parameter `subsamples` indicates the number of subsamples or fractional locations between source-image pixels; each subsample requires its own kernel data. The parameter `data` specifies the data values for each kernel. There is no limit or restriction on the `kernel_size` or number of `subsamples`.

The following code fragment creates a horizontal and a vertical interpolation table. The horizontal table has seven kernel elements, and the vertical table has three; each table has only one subsample.

```
XilSystemState state;
XilInterpolationTable horizTable;
XilInterpolationTable vertTable;

int subsamples = 1,
    horizSize  = 7,
    vertSize   = 3;

float horizData[] = {.005, .005, .04, .9, .04, .005, .005};
float  vertData[] = {.12, .8, .08};

horizTable = xil_interpolation_table_create(state,
                 horizSize, subsamples, horizData);

vertTable = xil_interpolation_table_create(state,
                 vertSize, subsamples, vertData);
```

Kernel Size

The `xil_interpolation_table_create()`function's `kernel_size` parameter determines the number of data elements in each subsample of the kernel defined by the interpolation table. There is no limit on the kernel size. The horizontal and vertical tables can have different kernel sizes, as well as different subsampling values.

The key element in a kernel is its center element; for an even-numbered kernel size, the key element is offset to the next-lowest index element. Thus, the key element for a 5-element table is the center element, which has array index 2; the key element for an 8-element table is the fourth element, which has index 3. The key element's array index can be computed as an integer calculation:

```
int array_index = (kernel_size - 1) / 2
```

Subsamples

The `xil_interpolation_table_create()`function's `subsamples` parameter determines the number of subsamples used for a general interpolation. Increasing the number of subsamples lets you specify different data values to be used for the interpolation kernel, depending on where point

S in the source image falls between pixels. There is still only one point sample taken to determine each destination pixel value, but the data used to interpolate the value depends on which subsample's kernel is used, and that, in turn, depends on the exact location of point *S*.

There is no limit on the number of subsamples you can use. For each subsample, you must define separate kernel data. Thus, for *n* subsamples, an interpolation table must have *n * kernel_size* data elements. For example, if you create a 7-element horizontal interpolation table with 4 subsamples, you must define 28 data elements. The first 7 elements define the first subsample's kernel, the second 7 elements define the second subsample's kernel, and so on.

Figure 10-2 shows how the interpolation tables are used to determine which kernel applies to a particular subsample location. In the figure, the subsampling is 4 in both the horizontal and vertical directions.

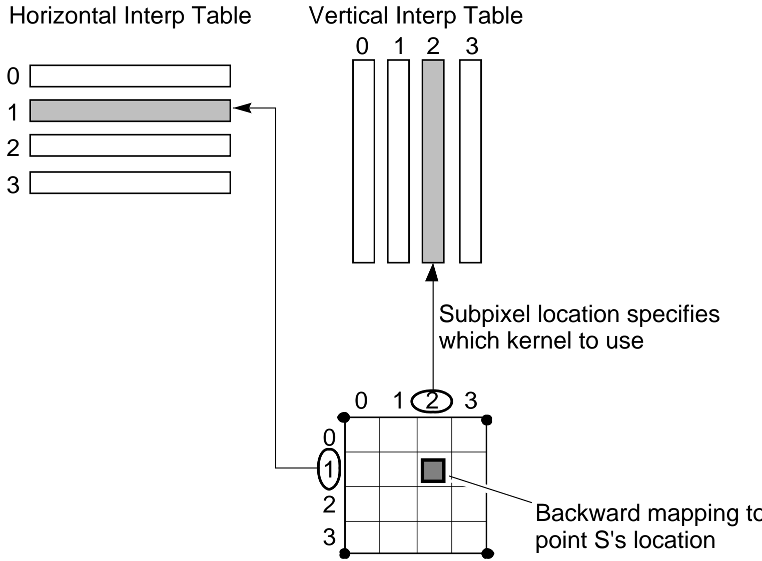

Figure 10-2 Determining the Kernel to Use for a General Interpolation

Typically, the kernel values for each subsample are weighted according to the subsample location's proximity to the pixels used in the calculation: the closer a pixel is to the subsample location, the more weight it carries in the kernel.

Kernel Data

The kernel data for each table is an array of floating point numbers. As mentioned in “Subsamples” on page 155, for *n* subsamples, there must be *n* * *kernel_size* data elements. Thus, for a 3-element kernel size with 2 subsamples, you must define an array of 6 floating point numbers; the first 3 numbers form the first subsample’s kernel, and the second 3 numbers form the second subsample’s kernel.

Caution – Providing too few kernel values to complete an interpolation table results in bad values.

To preserve the source image’s intensity in the destination image, the sum of the data values in each interpolation kernel should equal one. Kernel values whose sum is greater than one tend to increase the destination image’s intensity, and those whose sum is less than one tend to diminish the intensity. The results of a given kernel depend on the image it is used on.

To sharpen an image, the kernel values should be weighted heavily toward the center, with approximately 90% of the kernel value being concentrated there. To blur an image, the kernel values can be weighted away from the center.

Setting the Interpolation Tables on the System-State Object

After creating the vertical and horizontal tables, you must set them on the system-state object by calling the `xil_set_interpolation_tables()` function, whose prototype is shown below.

```
void xil_state_set_interpolation_tables(
     XilSystemState state,
     XilInterpolationTable horiz_kernel,
     XilInterpolationTable vertical_kernel);
```

The parameter `state` is a system-state data structure; the parameters `horiz_kernel` and `vertical_kernel` are the horizontal and vertical interpolation tables you created for general interpolation (see “Creating Vertical and Horizontal Interpolation Tables” on page 154).

The following code fragment creates horizontal and vertical interpolation tables, sets them on the system-state object, and uses them to perform a general interpolation for translating an image (see “Translating Images” on page 160).

```
XilSystemState state;
XilInterpolationTable horizTable;
XilInterpolationTable vertTable;
XilImage src, dst;

int subsamples = 8,   /* subsample and kernel sizes */
    horizSize  = 2,
    vertSize   = 2;

float kernelData[] = {1.0,  0.0,
                       .875, .125,    /* 7/8,  1/8 */
                       .75,  .25,     /* 6/8,  2/8 */
                       .625, .375,    /* 5/8,  3/8 */
                       .5,   .5,      /* 4/8,  4/8 */
                       .375, .625,    /* 3/8,  5/8 */
                       .25,  .75,     /* 2/8,  6/8 */
                       .125, .875 };  /* 1/8,  7/8 */

horizTable = xil_interpolation_table_create(state,
                 horizSize, subsamples, kernelData);

vertTable = xil_interpolation_table_create(state,
                 vertSize, subsamples, kernelData);

xil_state_set_interpolation_tables(
    state, horizTable, vertTable);

xil_translate(src, dst, "general", 50.0, 50.0);
```

In this example, the general interpolation effectively performs a bilinear interpolation with 8 subsamples. The kernel values indicate how much influence the source-image pixels have on the destination value; the kernel value 1 indicates that a source pixel completely determines the value for the destination pixel, and 0 indicates a source pixel has no influence on the destination value. If *S* is a source-image point that maps to a given destination pixel *D* (*S* and *D* are described in “Interpolation Options” on page 150), the

closer *S* is to a source pixel used to interpolate *D*'s value, the more influence that source pixel has on the computation. Conversely, the further *S* is from the source pixel, the less influence the source pixel has on *D*'s computed value.

Additional Kernel-Related Functions

The following XIL functions are available for getting an interpolation table's kernel and subsample sizes, and for getting the kernel data values.

```
unsigned int xil_interpolation_table_get_kernel_size(
     XilInterpolationTable table);
```

```
unsigned int xil_interpolation_table_get_subsamples(
     XilInterpolationTable table);
```

```
float* xil_interpolation_table_get_data(
     XilInterpolationTable table);
```

To retrieve an interpolation table, call the `xil_state_get_interpolation_tables()` function, whose prototype is shown below.

```
void xil_state_get_interpolation_tables(
    XilSystemState state,
    XilInterpolationTable* horiz_kernel,
    XilInterpolationTable* vertical_kernel);
```

Destroying an Interpolation Table

After performing a general interpolation, you should destroy the vertical and horizontal interpolation tables if you aren't going to use them again. Destroying the tables releases the memory that was allocated to them.

To destroy an interpolation table, call the `xil_interpolation_table_destroy()` function, whose prototype is shown below.

```
void xil_interpolation_table_destroy(
     XilInterpolationTable table);
```

The only parameter to this function is a handle to the kernel you want to destroy.

Translating Images

Translating an image means moving it up, down, left, or right. The XIL function you use to translate an image is called `xil_translate()`. The prototype for this function is shown below.

```
void xil_translate(XilImage src, XilImage dst,
    char *interpolation, float xoffset, float yoffset);
```

The parameters `src` and `dst` are handles to the source and destination images. The parameter `interpolation` is a string that specifies an interpolation type. For information on this parameter, see “Interpolation Options” on page 150. The parameters `xoffset` and `yoffset` are floating-point numbers that represent the number of pixels the image should be moved horizontally and vertically. If `xoffset` is positive, the image is moved to the right, and if it is negative, the image is moved to the left. If `yoffset` is positive, the image is moved down, and if it is negative, the image is moved up.

If you have already read the section “Origin” on page 46, which discusses image origins, you may think that translating an image is the equivalent of setting the destination image’s origin and then copying the source image to the destination. For the most part, this is true. That is, the two code fragments shown below have the same effect.

```
XilImage src, dst;

xil_set_origin(dst, 50.0, 50.0);
xil_copy(src, dst);
```

```
XilImage src, dst;

xil_translate(src, dst, "nearest", 50.0, 50.0);
```

Both fragments move the source image 50 pixels to the right and 50 pixels down.

However, what if the 50.0's in the fragments above are changed to 49.5's? In the first case, the coordinates that make up the destination image's origin (49.5, 49.5) will be rounded to integers (50, 50) before the copy occurs, so the result will not change. In the second case, though, `xil_translate()` actually moves the image 49.5 pixels down and 49.5 pixels to the right so that pixels in the destination image map to noninteger coordinates in the source image. Since there is no longer a one-to-one correspondence between the pixels in the source and destination images, the values of the pixels in the destination image must be estimated or ***interpolated***. The different types of interpolation the XIL library provides are discussed in "Interpolation Options" on page 150.

Scaling and Subsampling Images

The XIL library includes three functions that enable you to change the width and height of an image:

- `xil_scale()`
- `xil_subsample_adaptive()`
- `xil_subsample_binary_to_gray()`

The first function, `xil_scale()`, is the most general of the routines: it increases or decreases either the width or height of an image. The second, `xil_subsample_adaptive()`, is useful only for reducing an image's size; it is not useful for enlarging an image. It is available for reducing images because it often produces better results than using `xil_scale()` to reduce images. The last function, `xil_subsample_binary_to_gray()` is a special function used only for reducing the size of `XIL_BIT` images.

xil_scale()

The XIL library's general scaling routine is `xil_scale()`, whose function prototype is shown below.

```
void xil_scale(XilImage src, XilImage dst, char *interpolation,
    float xscale, float yscale);
```

The parameters `src` and `dst` are handles to the image to be scaled and the destination image. The parameter `interpolation` is a string that specifies the type of interpolation to be used for the operation: nearest neighbor, bilinear, bicubic, or general. (For information about these interpolation methods and how to request one of them, see "Interpolation Options" on page 150.) The parameters `xscale` and `yscale` are the horizontal and vertical scale factors. If `xscale` is greater than 1, the width of the source image will be increased, and if it is between 0 and 1 (exclusive), the width will decrease. Similarly, if `yscale` is greater than 1, the height of the source image will be increased, and if it is between 0 and 1 (exclusive), the height will decrease.

One thing to bear in mind when scaling images is the way image origins can affect the scaling operation. (For information about image origins, see "Origin" on page 46.) For instance, assume the source and destination images are the same size and that you want to scale the source image by a factor of 2 both horizontally and vertically. Figure 10-3 shows what will happen if, when you scale the image, the origins of both the source and destination images are set to 0.0, 0.0 (the default).

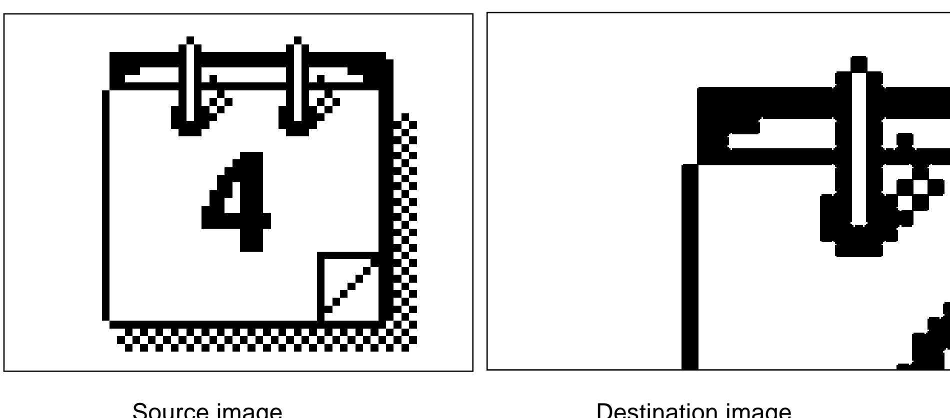

Figure 10-3 Zooming the Upper-Left Corner of an Image

The image on the left is the original image, and the image on the right is the scaled image. Because the origins are in the upper-left corner of the images, the destination image contains a zoomed version of the upper-left quadrant of the source image.

If, on the other hand, the origins of the source and destination images are set to the pixels at the centers of the images, the scale operation will produce the results shown in Figure 10-4.

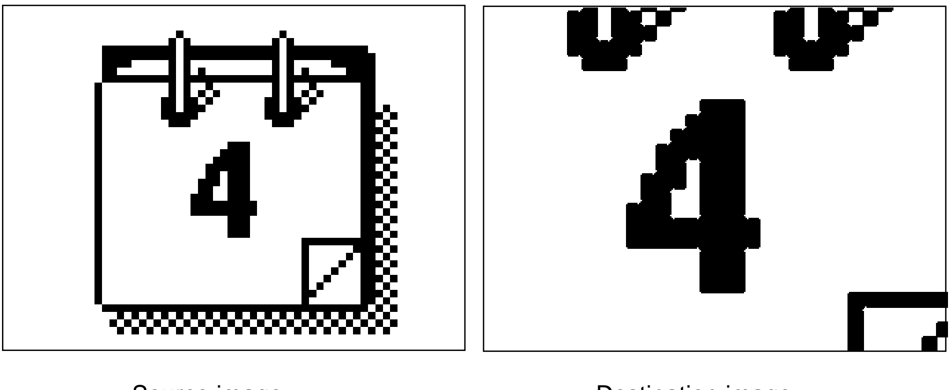

Figure 10-4 Zooming the Center of an Image

This time the destination image contains a zoomed version of a block taken from the center of the source image.

Note – The function `xil_scale()` has been optimized for the case where the *x* and *y* scale factors are .5.

---

xil_subsample_adaptive()

The `xil_scale()` function discussed in the last section is the only function XIL provides for enlarging images. However, you can reduce an image's size using either `xil_scale()` or `xil_subsample_adaptive()`.

The library includes the second function because when you scale down an image using `xil_scale()`, you must request one of the interpolation methods discussed earlier: nearest neighbor, bilinear, bicubic, or general. The first three of these interpolation methods look at relatively few source-image pixels when calculating the value of a destination pixel. Therefore, if you use one of these options on `xil_scale()` to reduce the size of an image, there are going to be many cases where some pixels in the source image make no contribution to the value of a pixel in the destination. For example, say that you scale down a 512-by-512 image using an *x* scale factor of .25, a *y* scale factor of .25, and bilinear interpolation. The source image contains 262,144 pixels, and the part of the destination image that represents the scaled-down source image contains 16,384 pixels. Since the operation used bilinear interpolation, the maximum number of source-image pixels that could have contributed to the scaled-down image is 4 times 16,384, or 65,536. In other words, at least three-fourths of the pixels in the source image had no effect on the values in the scaled-down image.

To increase the number of source-image pixels used in the calculation with `xil_scale()`, you could specify the general interpolation method and define your own interpolation tables. However, the function `xil_subsample_adaptive()` circumvents this problem for you because it guarantees that no matter what scale factors you use, every pixel in the source image will contribute to the value of one (and only one) pixel in the scaled-down image. For this reason, scaling down images using `xil_subsample_adaptive()` is generally preferable to using `xil_scale()`.

The function prototype for `xil_subsample_adaptive()` is shown below.

```
void xil_subsample_adaptive(XilImage src, XilImage dst,
    float xscale, float yscale);
```

The parameters to this function look much like those used by the `xil_scale()` routine. Note, however, that you don't specify an interpolation type. The reason for this is that `xil_subsample_adaptive()` has an interpolation scheme built in.

Because `xil_subsample_adaptive()` is useful only for reducing the size of images, `xscale` and `yscale` must be less than or equal to 1.

xil_subsample_binary_to_gray()

The function `xil_subsample_binary_to_gray()` is similar to `xil_subsample_adaptive()` in that you use it to scale down an image and it guarantees that all the pixels in the source image will make a contribution to the value of one pixel in the scaled-down image. However, `xil_subsample_binary_to_gray()` is designed specifically for scaling down `XIL_BIT` images.

This specialized function is necessary because if you scale down a 1-bit image using `xil_subsample_adaptive()` or `xil_scale()`, no matter how many source pixels you consider in determining the value of a destination-image pixel, the destination pixel can only have one of two values: 0 or 1. Obviously, a great deal of information about the image can be lost in this way.

The function `xil_subsample_binary_to_gray()` helps solve this problem by allowing up to 256 gray levels in the destination image, which must have the data type `XIL_BYTE`. For defining gray levels, one colormap index is used for each possible gray level that could result from the *x* and *y* scaling factors, and the indexes are consecutive values, with 0 representing all 0's in the source image. You must modify your colormap to define a gray level for each resulting index.

For example, Figure 10-5 shows a source image being scaled down by a factor of .5 in both the *x* and *y* dimensions. This means that a block of four pixels in the source image will determine the value of each pixel in the destination.

Since there are five possible combinations of values that can appear in a 2-by-2 block of pixels in the source, there will be five possible gray levels in the destination. The equations below indicate how the subsampling is performed.

- Zero 1's in the source (all 0's) = 0 in the destination
- One 1 in the source = 1 in the destination
- Two 1's in the source = 2 in the destination
- Three 1's in the source = 3 in the destination
- Four 1's in the source = 4 in the destination

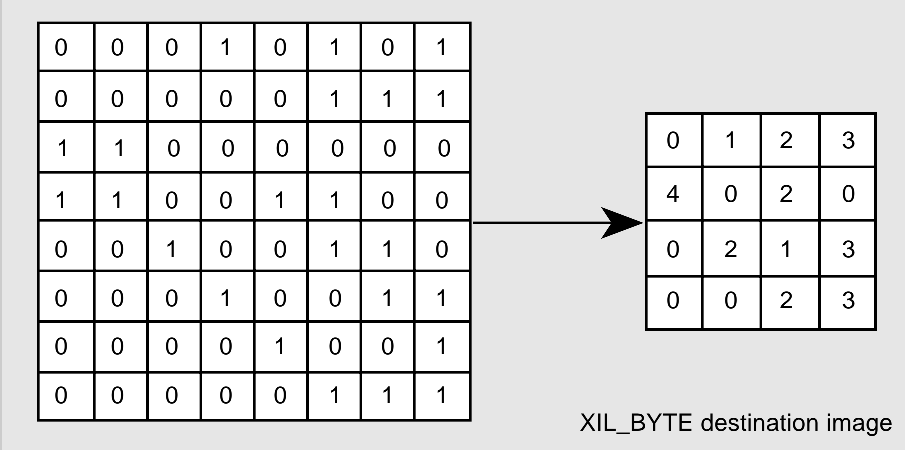

Figure 10-5 Subsampling Bit Images

If the scaling factors require a fractional block of pixels in the source to determine the destination pixel values, the block size is rounded up. For example, if a 2.2-by-2.2 block of pixels would be required to determine pixel values in the destination, a 3-by-3 block is used, resulting in 10 possible gray levels and therefore 10 colormap indexes whose values are 0 through 9.

The result of this type of operation is much clearer than the image obtained by scaling down a 1-bit image using `xil_subsample_adaptive()` or `xil_scale()`.

Note – The function `xil_subsample_binary_to_gray()` has been optimized for the case where the *x* and *y* scale factors are .33 and the width of the destination image is a multiple of 8.

Rotating Images

The XIL library's rotation function is called `xil_rotate()`. This function's prototype is shown below.

```
void xil_rotate(XilImage src, XilImage dst, char *interpolation,
    float angle);
```

The parameter `src` is a handle to the image to be rotated, and `dst` is a handle to the destination image. The parameter `interpolation` is a string that specifies the type of interpolation to be used for the operation: nearest neighbor, bilinear, bicubic, or general. For information about these interpolation methods and how to request one of them, see "Interpolation Options" on page 150. The last parameter, `angle`, indicates the desired angle of rotation in radians. (One radian is equal to 57.2958 degrees.) A positive angle results in a counterclockwise rotation, and a negative angle results in a clockwise rotation.

One thing to keep in mind when you're rotating an image is that `xil_rotate()` rotates an image around its origin. (For information about image origins, see "Origin" on page 46.) For example, consider the situation where you have not set the origin of the source or destination image. Figure 10-6 shows what will happen when you rotate the source image by 90 degrees counterclockwise.

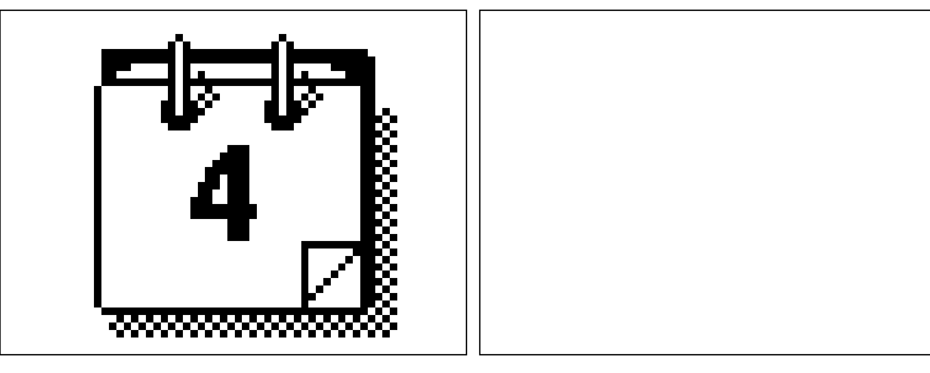

Figure 10-6 Rotating an Image Around Its Default Origin

Because you're rotating the image about its default origin, which is in the upper-left corner, the image is basically rotated out of sight. Only the pixels in the leftmost column of the source image are written to the destination. They become part of the topmost scanline in that image.

To rotate the source image in the usual sense, you must first set the origins of the source and destination images to their centers.

```
unsigned int width, height;

width = xil_get_width(src);
height = xil_get_height(src);
xil_set_origin(src, width/2.0, height/2.0);
xil_set_origin(dst, width/2.0, height/2.0);
xil_rotate(src, dst, "nearest", 1.5708);
```

This code produces the results shown in Figure 10-7.

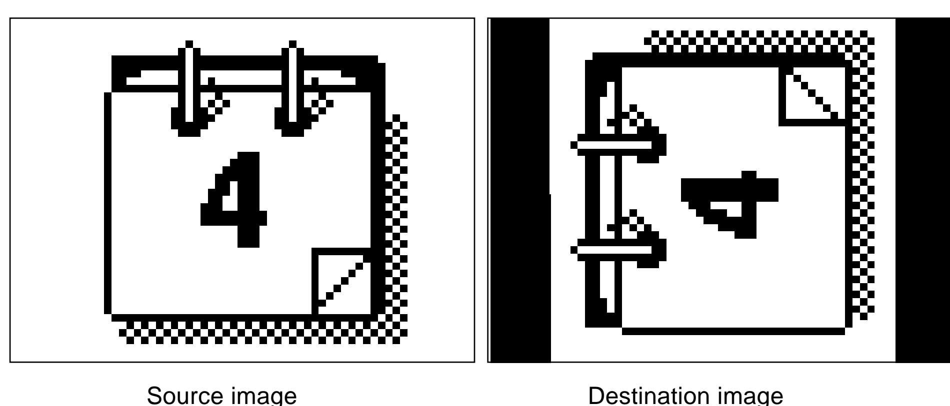

Figure 10-7 Rotating an Image Around Its Center

Performing General Affine Transforms

An *affine transform* is a transformation of an image in which straight lines remain straight and parallel lines remain parallel, but the distance between lines and the angles between lines may change. Thus, translation, scaling, and rotation are all affine transforms.

The XIL functions that perform *one* of these transforms are discussed in "Translating Images," "Scaling and Subsampling Images," and "Rotating Images." However, the XIL library also contains a function that can perform *any* of these transforms, or a combination of them, such as a rotate and a scale. This function is called `xil_affine()`, and its prototype is shown below.

```
void xil_affine(XilImage src, XilImage dst, char *interpolation,
    float *matrix);
```

The parameter `src` is a handle to the image to be transformed, and `dst` is a handle to the destination image. The parameter `interpolation` is a string that specifies the type of interpolation to be used for the operation. For information about the available methods of interpolation and how to request one of them, see "Interpolation Options" on page 150. The last parameter, `matrix`, must be a six-element array of floating-point numbers.

If we call the six elements in the array a, b, c, d, e, and f, the equations used in performing the transform look like this:

x‘ = (a * x) + (c * y) + e
y’ = (b * x) + (d * y) + f

Or written as a matrix operation:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & c \\ b & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} e \\ f \end{bmatrix}$$

To illustrate how these equations relate to particular transforms, look at how you would perform the translation, scale, and rotation considered earlier in the chapter using `xil_affine()`.

The following code fragment illustrates an image being translated 50 pixels to the right and 50 pixels down:

```
XilImage src, dst;
float matrix[6] = {1.0, 0.0, 0.0, 1.0, 50.0, 50.0};

xil_affine(src, dst, "nearest", matrix);
```

The fifth value in the matrix is used to control the horizontal translation of the image, and the sixth is used to control the amount of vertical translation.

The following code fragment shows an image being zoomed by a factor of 2 in both the *x* and *y* directions:

```
XilImage src, dst;
float matrix[6] = {2.0, 0.0, 0.0, 2.0, 0.0, 0.0};

xil_affine(src, dst, "nearest", matrix);
```

And this code fragment shows an image being rotated 90 degrees counterclockwise.

```
XilImage src, dst;
float matrix[6];

matrix[0] = (float)cos(1.5708);
matrix[1] = (float)-sin(1.5708);
matrix[2] = (float)sin(1.5708);
matrix[3] = (float)cos(1.5708);
matrix[4] = 0.0;
matrix[5] = 0.0;
xil_affine(src, dst, "nearest", matrix);
```

You can also use `xil_affine()` to perform many other affine transforms. For example, the code below shears a source image along the *x* axis and then translates the image to the left so that it is centered in the destination.

```
XilImage src, dst;
float matrix[6] = {1.0, 0.0, -0.4, 1.0, 50.0, 0.0};

xil_affine(src, dst, "nearest", matrix);
```

Note – In this case, the *y* coordinate of a pixel in the source image affects the *x* coordinate of a pixel in the destination. Figure 10-8 shows an image before and after this shearing operation.

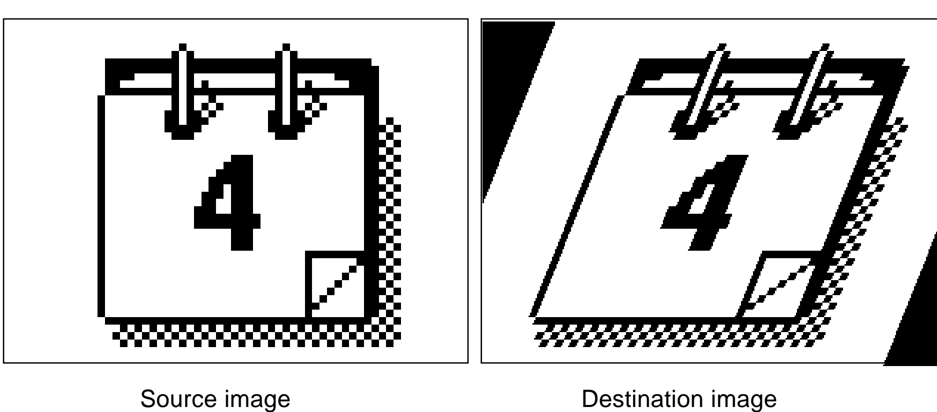

Figure 10-8 Shearing an Image Along Its *x* Axis

Warping Images

Each of the geometric functions discussed so far in this chapter performs a specific type of affine transformation: `xil_translate()` moves an image up, down, left, or right; `xil_scale()` changes an image's width or height; and so on. The XIL library also contains a set of general functions that let you customize a transformation to meet an individual set of needs by creating *warp* tables that perform pixelwise displacement horizontally, vertically, or in both directions. The warp tables are XIL images whose pixel values define the backward mapping from a pixel in the destination to a pixel in the source.

Warp tables are typically used to stretch an image according to predefined rules and are most useful for performing nonlinear transformations. For example, they might be used to correct distortions that were imposed on an image by the equipment used to capture it. Or they might be used for cartographic projection of an image.

To displace pixels in a single dimension, use either the `xil_tablewarp_horizontal()` or `xil_tablewarp_vertical()` function.

```
void xil_tablewarp_horizontal(XilImage src, XilImage dst,
    char* interpolation, XilImage warp_table);
```

```
void xil_tablewarp_vertical(XilImage src, XilImage dst,
    char* interpolation, XilImage warp_table);
```

To displace pixels in two dimensions, use the xil_tablewarp() function.

```
void xil_tablewarp(XilImage src, XilImage dst,
    char* interpolation, XilImage warp_table);
```

xil_tablewarp_horizontal() specifies a horizontal displacement, xil_tablewarp_vertical() specifies a vertical displacement, and xil_tablewarp() specifies a displacement in both directions. For these functions, the parameters src and dst are handles to the source and destination images. The parameter interpolation is a string that specifies an interpolation type for the operation: nearest neighbor, bilinear, bicubic, or general (see “Interpolation Options” on page 150). The parameter warp_table is a handle to a defined warp table. For a one-dimensional displacement, the warp table is a 1-banded image that specifies the displacement in one direction. For a two-dimensional displacement, the warp table is a 2-banded image; the first band specifies the horizontal displacement and the second band specifies the vertical displacement.

A warp table is applied to the destination image’s origin; the source image’s origin is then added to the backward mapping position specified by the warp table. By default, an image’s origin is 0,0; to change the origin, use the xil_set_origin() function.

A warp table must have the data type XIL_SHORT, though it can be used to warp images of any data type; the XIL_SHORT value is interpreted as fixed point data with a 12-bit value and 4 bits of precision. Both the source and destination images must have the same type and number of bands. Pixels that map to positions outside the source image are considered zeros. The table warp operations cannot be performed in place.

Warp tables can be used to perform any type of geometric transformation. For example, as discussed in “Translating Images” on page 160, you can move a source image up, down, left, or right with xil_translate(). You can also

use a warp table to move the image. The following code fragment shows how you could define a two-dimensional warp table to move an image to the left 100 pixels and down 500 pixels.

```
XilSystemState state;
XilImage src, dst, warp_table;
float values[2];
int width, height;

warp_table = xil_create(state, width, height, 2, XIL_SHORT);
/* multiply offsets by 16 because of 12 bit values with
 * 4 bit precision */
values[0] = 100.0 * 16;
values[1] = 500.0 * 16;
xil_set_value(warp_table, values);

xil_tablewarp(src, dst, “bilinear”, warp_table);
```

In this example, variable `warp_table` stores the warp table for the operation; it is assumed that the source image's width and height have already been stored in the `width` and `height` variables. The `warp_table` image has two bands; the first band is needed to store the horizontal displacement defined by `values[0]`, and the second is needed to store the vertical displacement defined by `values[1]`. The call to `xil_set_values()` sets all pixels in the table-warp image to the values defined by the `values` array; thus, the displacement is identical for all pixels. The call to `xil_tablewarp()` then applies the warp table to the source image, using a bilinear interpolation.

Of course, you wouldn't use a warp table to perform a simple translation; warp tables are better used to define algorithms that implement nonlinear displacements. The algorithm would typically define a pixelwise displacement for source-image pixels, in which case you would set the displacements on the warp table by calling `xil_set_pixel()` rather than `xil_set_value()`. For a discussion of warping transformations, see Gonzalez and Wintz, *Digital Image Processing* (see Appendix E, “Bibliography”).

Transposing Images

The XIL library's transposition function, `xil_transpose()`, enables you to:

- Flip an image across an imaginary horizontal line that runs through the center of the image
- Flip an image across an imaginary vertical line that runs through the center of the image
- Flip an image across its main diagonal
- Flip an image across its antidiagonal
- Rotate an image counterclockwise about its center by 90, 180, or 270 degrees

The function's prototype is shown below:

```
void xil_transpose(XilImage src, XilImage dst,
    XilFlipType fliptype);
```

The parameters `src` and `dst` are handles to the source and destination images. The parameter `fliptype` is an enumeration constant of type `XilFlipType`. The constants in the enumeration and their meanings are shown in Table 10-2.

Table 10-2 Constants in the Enumeration `XilFlipType`

Constant	Meaning
XIL_FLIP_X_AXIS	Flips the image across a horizontal line running through its center
XIL_FLIP_Y_AXIS	Flips the image across a vertical line running through its center
XIL_FLIP_MAIN_DIAGONAL	Flips the image across its main diagonal
XIL_FLIP_ANTIDIAGONAL	Flips the image across its antidiagonal
XIL_FLIP_90	Rotates the image 90 degrees counterclockwise about its center
XIL_FLIP_180	Rotates the image 180 degrees counterclockwise about its center
XIL_FLIP_270	Rotates the image 270 degrees counterclockwise about its center

---

Note – You don't specify an interpolation type when you call `xil_transpose()`. This is not necessary because `xil_transpose()` differs slightly from the other geometric operators in that it maps pixels in the source

image directly to pixels in the destination. Another important thing to note about `xil_transpose()` is that it ignores image origins when it flips or rotates images. The function always flips an image across a line that passes through the center of the image and always rotates an image around its center.

Figure 10-9 illustrates several of the operations that `xil_transpose()` can perform.

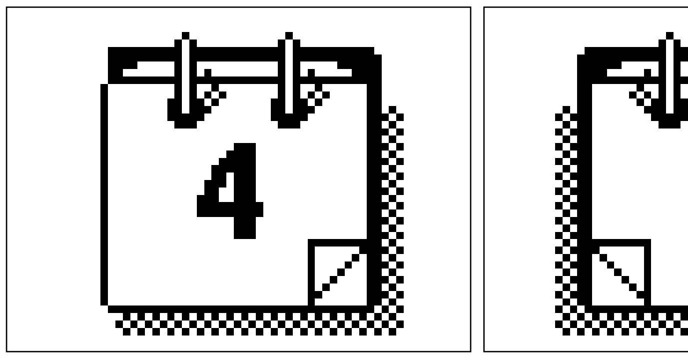

Original image   Flipped across vertical axis

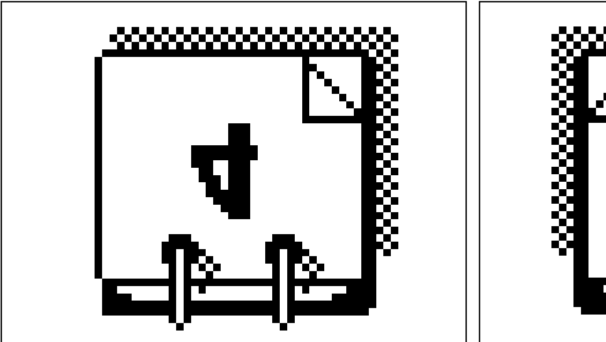

Flipped across horizontal axis   Rotated 180 degrees counterclockwise

Figure 10-9 Flipping and Rotating Images Using `xil_transpose()`

Miscellaneous Image Processing Functions 11

Earlier chapters in this book have discussed several classes of XIL image-processing functions. Chapter 7, “Presentation Functions,” talked about the XIL functions you’re most likely to use in preparing an image for display. Chapter 9, “Arithmetic, Relational, and Logical Functions,” considered the XIL library’s arithmetic and logical functions. And Chapter 10, “Geometric Functions,” discussed the library’s geometric functions. This chapter covers the library’s remaining image-processing functions. These functions enable you to:

- Find the minimum and maximum values in an image
- Produce a histogram for an image
- Threshold an image
- Fill an area in an image
- Filter an image
- Detect an image’s edges
- Dilate or erode an image
- Pass an image through a lookup table
- Perform a linear combination of the values in the different bands of an image
- Blend images
- Paint on an image
- Set and get the values of pixels in an image
- Copy a pattern to an image

These operations are discussed in the sections below.

Finding the Minimum and Maximum Values in an Image

The function `xil_extrema()` finds the minimum and maximum values in a single-band image or the minimum and maximum values in each band of a multiband image. The prototype for this function is shown below:

```
void xil_extrema(XilImage src, float *max, float *min);
```

The parameter `src` is a handle to a source image. The parameter `max` is an array of floating-point numbers in which the function will return the maximum value found in each band of the source image. The parameter `min` is an array of `float`s in which the function will return the minimum value found in each band of the image. The number of elements in the arrays `max` and `min` must match the number of bands in the image.

Consider the following example:

```
XilImage src;
float max[3], min[3];

xil_extrema(src, max, min);
```

After this call, `max[0]` will contain the maximum value in band 0 of `src`, `max[1]` will contain the maximum value in band 1, and `max[2]` will contain the maximum value in band 2. Similarly, `min` will contain the minimum values in the three bands of the image.

Note – `xil_extrema()` returns a minimum or maximum value, depending on which one you specify. Don't confuse this with finding the larger or lesser of pixels between two source images and storing the results in a destination image. To perform that relational function, use `xil_max()` or `xil_min()`, discussed on page 143.

Producing a Histogram for an Image

The XIL library defines a data type called `XilHistogram`. A data structure of this type is designed to hold gray-level (or color-level) information about an image. This section discusses how to produce a histogram for an image and how to read intensity-level information from an XIL histogram structure. The primary tasks you will need to perform are listed below:

1. Create a histogram data structure.
2. Write gray- or color-level information about an image to the data structure.
3. Read the histogram data stored in the structure.
4. Destroy the histogram.

These tasks are discussed in detail in the sections that follow.

Creating a Histogram

An XIL histogram structure contains a set of *bins* for each band of the image whose histogram you want to take. These bins are used to hold information about gray or color levels. For instance, if you want to take the histogram of an 8-bit grayscale image, you might create a histogram structure that contains 256 bins. When you actually take the histogram, the number of 0's in the image will be stored in bin 0, the number of 1's in the image will be stored in bin 1, and so on.

Your histogram need not contain a bin for each possible value in the image. You can specify the lowest and highest values that will result in a bin count being incremented. In addition, you can specify that the histogram structure contain a number of bins less than the number of levels you will be checking for. In this case, each bin holds a count for a range of values. For example, if you create a histogram for an 8-bit grayscale image and specify four bins, occurrences of the values 0 to 63 will be stored in bin 0, occurrences of 64 to 127 will be stored in bin 1, and so on.

You create a histogram data structure by calling the function `xil_histogram_create()`, whose function prototype is shown below.

```
XilHistogram xil_histogram_create(XilSystemState state,
    unsigned int nbands, unsigned int *nbins, float *low_value,
    float *high_value);
```

Table 11-1 discusses the purpose of each of these parameters.

Table 11-1 Parameters to `xil_histogram_create()`

Parameter	What It Represents
`state`	The system-state data structure returned when you initialized the library.
`nbands`	The number of bands in the histogram. This number must match the number of bands in the image whose histogram you want to take.
`nbins`	An array of `unsigned int`s, each element of which specifies the number of bins to be used for one band of the image. The number of elements in the array must match the number of bands in the image.
`low_value`	An array of `float`s, each element of which specifies the lowest gray or color level that will be checked for in one band of the image. The number of elements in the array must match the number of bands in the image.
`high_value`	An array of `float`s, each element of which specifies the highest gray or color level that will be checked for in one band of the image. The number of elements in the array must match the number of bands in the image.

Writing Level Information to the Histogram Structure

Once you have created an appropriate histogram structure, you generate the histogram for an image using the function `xil_histogram()`. The prototype for this function is shown below.

```
void xil_histogram(XilImage src, XilHistogram histogram,
    unsigned int skip_x, unsigned int skip_y);
```

The parameter `src` is a handle to the image whose histogram you want to take. The parameter `histogram` is a handle to the histogram structure you created earlier. The parameters `skip_x` and `skip_y` tell `xil_histogram()` whether it should count every pixel in the source image, or whether it can skip some pixels, either horizontally or vertically.

If `skip_x` is set to 1, `xil_histogram()` counts every pixel on a scanline; if it is set to 2, the function counts every other pixel; and so on. The value of `skip_y` has an analogous effect on whether `xil_histogram()` counts every pixel in the vertical direction.

Reading Data from a Histogram

Once you've generated histogram data for an image, you'll want to read the counts stored in the bins of the histogram structure. You read this data using the function `xil_histogram_get_values()`, whose prototype is shown below.

```
void xil_histogram_get_values(XilHistogram histogram,
    unsigned int *data);
```

The parameter `histogram` is a handle to your histogram structure, and `data` is a pointer to an array of `unsigned int`s in which `xil_histogram_get_values()` can place the histogram data.

The code fragment below shows how you might retrieve the histogram data for a 3-band `XIL_BYTE` image. This code assumes that 32 bins were used for each band of the image.

```
#define BINS 32

XilHistogram histogram;
unsigned int *data;

data = (unsigned int *)malloc(BINS * BINS * BINS *
    sizeof(unsigned int));
xil_histogram_get_values(histogram, data);
```

The XIL library also contains functions that enable you to retrieve other information from histogram structures and to name histograms. These are listed in Table 11-2.

Table 11-2 Additional Histogram Functions

Function Name	What The Function Does
`xil_histogram_get_nbands`	Returns the number of bands in the histogram
`xil_histogram_get_nbins`	Fills a user-supplied array with values representing the number of histogram bins for each histogram band
`xil_histogram_get_limits`	Fills one user-supplied array with floating-point numbers that represent the low values for the histogram's bands and a second array with numbers representing the high values
`xil_histogram_get_info`	Returns information about the number of bands, the number of bins, and the low and high values in the histogram
`xil_histogram_set_name`	Sets the name of a histogram
`xil_histogram_get_name`	Returns a copy of a histogram's name
`xil_histogram_get_by_name`	Returns a handle to a histogram that has the name you specify

Destroying a Histogram

After reading the histogram data for your image, you should destroy the histogram structure if you will not be using it again. Destroying the histogram frees the memory that was allocated to store it.

You destroy a histogram structure by calling the function `xil_histogram_destroy()`, whose prototype is shown below.

```
void xil_histogram_destroy(XilHistogram histogram);
```

The only parameter to this function is a handle to the histogram you want to destroy.

Thresholding an Image

The XIL library's thresholding function, `xil_threshold()`, provides a simple mechanism for defining the boundaries of objects that appear on a contrasting background. The function's prototype is shown below.

```
void xil_threshold(XilImage src, XilImage dst, float *lowvalue,
    float *highvalue, float *mapvalue);
```

The parameters `src` and `dst` are handles to the source and destination images for the operation. The parameter `lowvalue` is a pointer to an array of floating-point numbers that define the lower bound for the threshold operation for each band of `src`. That is, the operation will affect only values greater than or equal to `lowvalue[0]` in band 0, only values greater than or equal to `lowvalue[1]` in band 1, and so on. The parameter `highvalue` is a pointer to an array of `float`s that define the upper bound for the threshold operation for each band of `src`. The final parameter, `mapvalue`, is also a pointer to an array of `float`s and is used as follows. For band *n* of an image, all values in the range `lowvalue[`*n*`]` to `highvalue[`*n*`]` inclusive are set to `mapvalue[`*n*`]`.

A standard way of arriving at the optimal values for the elements of the arrays `lowvalue` and `highvalue` is to create a histogram for the image. For information about XIL histograms, see the section "Producing a Histogram for an Image" on page 179.

Filling an Area in an Image

The XIL library contains two fill functions: `xil_fill()` and `xil_soft_fill()`. The first function performs boundary fills, and the second performs soft fills.

xil_fill()

The function `xil_fill()` fills a 4-connected region of pixels with a specified color. (One pixel is 4-connected to another if it is located directly above, below, to the left of, or to the right of that pixel.) The region begins at a starting pixel (the *seed*) and grows until it encounters a boundary color or the edge of the image.

Figure 11-1 shows a fill operation in which the boundary color is black, the fill color is dark gray, and the starting point for the fill is 3,2.

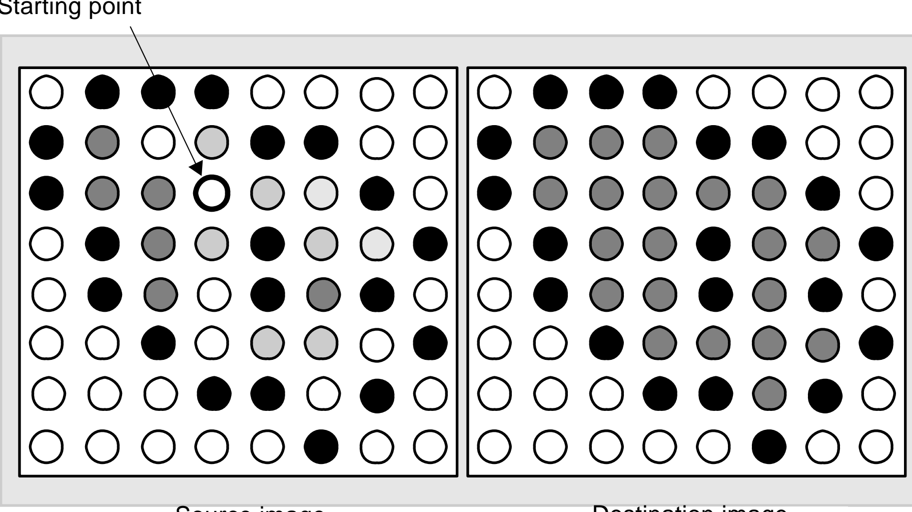

Figure 11-1 Boundary Fill

The function prototype for `xil_fill()` is shown below.

```
void xil_fill(XilImage src, XilImage dst, float xseed,
    float yseed, float *boundary, float *fill_color);
```

The parameters `src` and `dst` are handles to the source and destination images for the operation. These images must have the same number of bands and the same data type. The parameters `xseed` and `yseed` represent the *x* and *y* coordinates of the starting point. The parameters `boundary` and `fill_color` are arrays of floating-point numbers. Both arrays must contain a number of

elements equal to the number of bands in the source and destination images. The elements of `boundary` represent the boundary color, and those in `fill_color` represent the fill color.

xil_soft_fill()

The function `xil_soft_fill()` is designed to solve the problem of filling a region that does not have a distinct boundary, usually because the outline of the region has been antialiased. It also fills regions that are shaded or contain specular reflections without losing the special effects.

Here's how `xil_soft_fill()` works. Assume for the purpose of illustration that you want to fill a red object that appears on a blue background and that the object has been antialiased so that there's a transition at the edge of the object from red to magenta to blue. You want to fill the object with green. To get started, you specify:

- The color of the *foreground*, the region to be filled: red
- The color of the background: blue. (There may be more than one background color.)
- A fill color: green
- A starting point (seed) known to be inside the region to be filled

The soft-fill function then determines which pixels are inside the region and sets those pixels appropriately. To be considered part of the region to be filled, a pixel must be 4-connected to a pixel inside the region and must contain some fraction of the foreground color. That is, the value of each pixel in the region can be expressed as:

Pixel value = (fraction * foreground-color) + ((1 - fraction) * background-color)

When setting the values of pixels in the destination image, `xil_soft_fill()` uses the equation:

Pixel value = (fraction * fill-color) + ((1 - fraction) * background-color)

As a consequence, where there was a transition from red to magenta to blue in the source, there will be a transition from green to cyan to blue in the destination.

The prototype for xil_soft_fill() is shown below.

```
void xil_soft_fill(XilImage src, XilImage dst, float xseed,
    float yseed, float *fgcolor, unsigned int num_bgcolor,
    float *bgcolor, float *fill_color);
```

Table 11-3 below explains the purpose of each of these parameters.

Table 11-3 Parameters to xil_soft_fill()

Parameter	What It Is Used For
src	A handle to the source image for the operation.
dst	A handle to the destination image. This image must have the same number of bands and the same data type as the source image.
xseed	The *x* coordinate of the starting point for the operation.
yseed	The *y* coordinate of the starting point for the operation.
fgcolor	An array of floats containing a number of elements equal to the number of bands in the source and destination images. These elements specify the current color of the region to be filled.
num_bgcolor	The number of background colors in the image.
bgcolor	An array of floats containing a number of elements equal to the number of background colors times the number of bands in the source and destination images. If your RGB source image contained two background colors—black and red—your declaration of bgcolor might look like this: `float bgcolor[6] = {0.0,0.0,0.0,0.0,0.0,255.0};`
fill_color	An array of floats containing a number of elements equal to the number of bands in the source and destination images. These elements specify the color to be used in filling the region of interest.

Filtering an Image

The XIL function for filtering is called `xil_convolve()`. Depending on the convolution filter, or ***kernel***, you specify as a parameter to this function, the routine can perform a variety of tasks, including:

- Sharpening images
- Blurring images
- Highlighting the edges in images

See Color Plate 2 for a few examples of how different kernels can affect a source image.

Note – To detect edges, you can also use `xil_edge_detection()`, which uses the Sobel algorithm. For information about `xil_edge_detection()`, see "Detecting Edges in an Image" on page 192.

The kernel you supply as a parameter is a two-dimensional array of weighted values that the `xil_convolve()` function uses as follows. To calculate the value of each pixel in the destination image, the function takes the following steps:

1. It lays the kernel over the corresponding pixel in the source image.

   To be precise, the function lays the key value of the filter—usually the value at the center of the filter—over this source-image pixel. This means that the other values in the filter lie on top of neighbors of this pixel. Figure 11-2 shows a 3-by-3 kernel being laid on top of the source-image pixel at 1,1 and its neighbors.

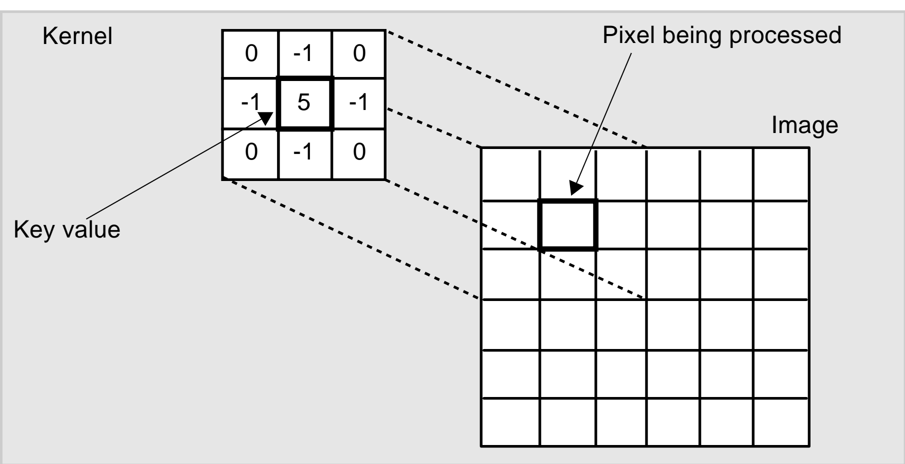

Figure 11-2 Convolution Operation

2. The convolution function multiplies each pixel in the neighborhood defined by the filter by the appropriate value in the filter. Then, it sums the products of these multiplications. This sum becomes the value of the pixel in the destination.

Three basic steps are involved in performing this type of convolution operation using the XIL library. These are:

- Creating the convolution kernel
- Filtering your image
- Destroying the kernel

These tasks are discussed in detail in the next several sections.

Creating a Convolution Kernel

Convolution kernels are XIL data structures of type `XilKernel` and are created with the function `xil_kernel_create()`. The prototype for this function is shown below.

```
XilKernel xil_kernel_create(XilSystemState state,
    unsigned int width, unsigned int height, unsigned int keyx,
    unsigned int keyy, float *data);
```

The parameter `state` is a system-state data structure. You received a handle to this structure when you initialized the library. The parameters `width` and `height` are the width and height of the kernel in pixels. These values are usually odd numbers so that the kernel will have a center value, but need not be. Common sizes for kernels are 3-by-3 and 5-by-5. The parameters `keyx` and `keyy` define the coordinates of the key value in the kernel. These coordinates are specified with respect to the upper-left value in the kernel (0,0). As mentioned earlier, the key value is usually in the center of the kernel. The last parameter, `data`, is a pointer to the floating-point values that will be written to the kernel.

The code needed to create the kernel shown in Figure 11-2 would look something like this.

```
XilKernel high_pass_filter;
XilSystemState state;
unsigned int width, height, keyx, keyy;
float data[] = {0.0, -1.0, 0.0,
                -1.0, 5.0, -1.0,
                0.0, -1.0, 0.0};

width = height = 3;
keyx = keyy = 1;
high_pass_filter = xil_kernel_create(state, width, height, keyx,
    keyy, data);
```

The kernel created here is a high-pass filter, used for sharpening an image. Two other common high-pass filters are shown in Figure 11-3.

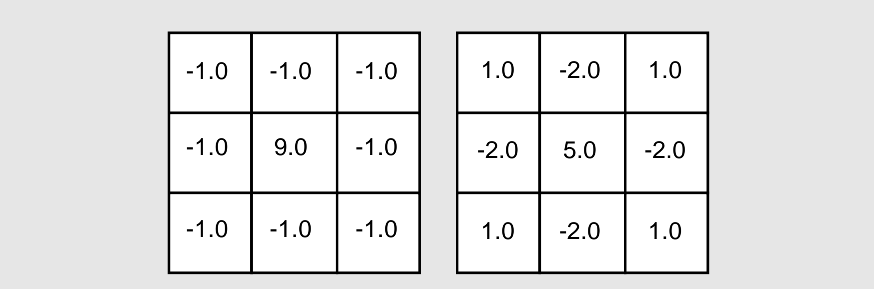

-1.0	-1.0	-1.0
-1.0	9.0	-1.0
-1.0	-1.0	-1.0

1.0	-2.0	1.0
-2.0	5.0	-2.0
1.0	-2.0	1.0

Figure 11-3 High-Pass Filters

Figure 11-4 shows several low-pass filters. These are useful for smoothing or blurring images.

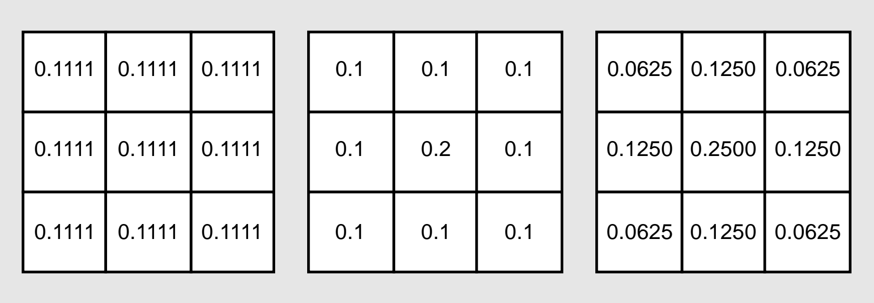

0.1111	0.1111	0.1111
0.1111	0.1111	0.1111
0.1111	0.1111	0.1111

0.1	0.1	0.1
0.1	0.2	0.1
0.1	0.1	0.1

0.0625	0.1250	0.0625
0.1250	0.2500	0.1250
0.0625	0.1250	0.0625

Figure 11-4 Low-Pass Filters

Filtering an Image

Once you have created the convolution kernel you want to use for your filtering operation, you perform the actual filtering by calling `xil_convolve()`. The prototype for this function is shown below.

```
void xil_convolve(XilImage src, XilImage dst, XilKernel kernel,
    XilEdgeCondition edge_condition);
```

The parameters `src` and `dst` are handles to a source and destination image. These two images must have the same number of bands and contain the same type of data. The parameter `kernel` is a handle to the convolution kernel you created earlier. The final parameter, `edge_condition`, is an enumeration

constant that indicates how `xil_convolve()` should handle pixels at, or near, the edge of the source image. Table 11-4 lists the three constants you can use here and explains what effect they have.

Table 11-4 Handling Edges in a Convolution Operation

Value for `edge_condition`	Effect on Convolution Operation
XIL_EDGE_NO_WRITE	The values of pixels at the edge of the source image are copied to the corresponding pixels in the destination. Thus, the edge of the source image is not filtered.
XIL_EDGE_ZERO_FILL	Pixels at the edge of the destination image are set to zero.
XIL_EDGE_EXTEND	The convolution operator temporarily extends the width and height of the source image by replicating the pixels at the edge of the image. This strategy enables the operator to filter the pixels that were originally at the edge of the image.

Destroying a Convolution Kernel

After filtering your image, you should destroy the convolution kernel you used for the filtering if you will not be using it again. Destroying this kernel frees the memory that was allocated to store it.

You destroy a convolution kernel by calling the function `xil_kernel_destroy()`, whose prototype is shown below.

```
void xil_kernel_destroy(XilKernel kernel);
```

The only parameter to this function is a handle to the kernel you want to destroy.

Additional Kernel-Related Functions

The preceding sections have discussed the most frequently used XIL functions that affect kernels. However, the library also contains the kernel-related functions shown in Table 11-5. These functions enable you to make a copy of a kernel, read the values of kernel attributes, and assign a name to a kernel, among other things.

Table 11-5 Utility Functions for Convolution Kernels

Function Name	What the Function Does
xil_kernel_create_copy	Creates a copy of a kernel
xil_kernel_get_width	Gets the width of a kernel in pixels
xil_kernel_get_height	Gets the height of a kernel in pixels
xil_kernel_get_key_x	Gets the *x* coordinate of the kernel's key value
xil_kernel_get_key_y	Gets the *y* coordinate of the kernel's key value
xil_kernel_set_name	Assigns a name to a kernel
xil_kernel_get_name	Returns a copy of a kernel's name
xil_kernel_get_by_name	Gets a handle to a kernel by specifying the name of the kernel

Detecting Edges in an Image

The `xil_edge_detection()` function detects edges in a source image and writes the result to a destination image. The prototype for this function is shown below.

```
void xil_edge_detection(XilImage src, XilImage dst,
    XilEdgeDetection edge_detection_method);
```

The parameters `src` and `dst` are handles to a source and destination image. The two images must have the same number of bands and contain the same type of data. The parameter `edge_detection_method` is an enumeration constant describing the type of edge detection to perform.

Currently, XIL_EDGE_DETECT_SOBEL is the only edge detection method supported; to detect vertical and horizontal edges in an image, it uses the filtering kernels shown in Figure 11-5.

Vertical

-0.5	0.0	0.5
-1.0	0.0	1.0
-0.5	0.0	0.5

Horizontal

-0.5	-1.0	-0.5
0.0	0.0	0.0
0.5	1.0	0.5

Figure 11-5 Filters Used by the XIL_EDGE_DETECT_SOBEL Algorithm

The XIL_EDGE_DETECT_SOBEL method detects edges by finding the gradient of an image as follows:

- It performs two convolution operations on the source image; one operation detects vertical edges using the vertical filter shown in Figure 11-5, the other detects horizontal edges using the horizontal filter. This yields two intermediate images: *a* and *b*.
- It squares all the pixel values in image *a* and then in image *b*, yielding the intermediate images a^2 and b^2.
- It forms the final destination image by taking the square root of $a^2 + b^2$.

Note – The convolution operations duplicate the source-image edges during convolution, similar to using the XIL_EDGE_EXTEND edge condition on xil_convolve(). For information on xil_convolve(), see "Filtering an Image" on page 187.

Note – For a discussion of the Sobel edge condition algorithm, consult Gonzalez and Wintz, *Digital Image Processing* (see Appendix E, "Bibliography").

Dilating or Eroding an Image

The XIL library includes two functions—`xil_dilate()` and `xil_erode()`—that enable you to dilate and erode regions within images. One of the main uses of these functions is to perform *trapping* on color images that will be printed. This trapping involves creating some overlap where regions of different colors meet. Once a trap is in place, slight registration problems during the printing process should not result in gaps appearing between the edge of an object and its background.

Dilation and erosion are similar to convolution in that both calculate destination-image pixels by looking at a neighborhood of source-image pixels and the values in a matrix (called a *structuring element*). Like a convolution kernel, this structuring element contains a key value. As each pixel in the destination is calculated, this key value is laid over the corresponding pixel in the source. See Figure 11-6.

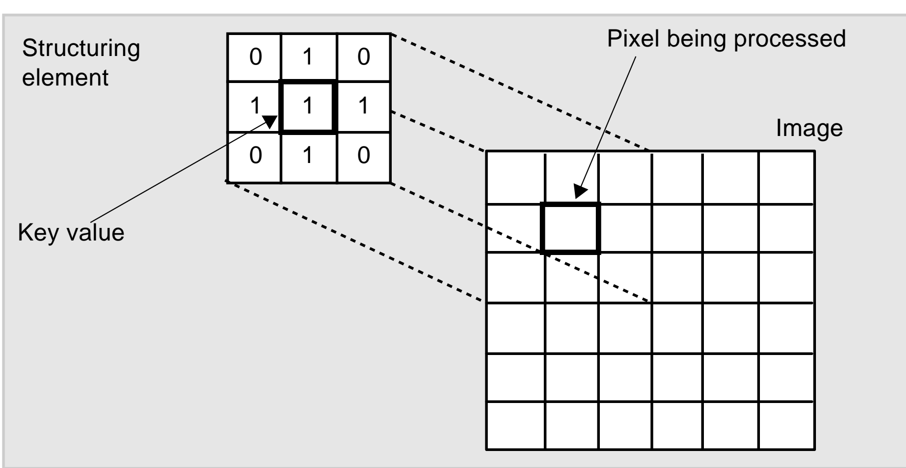

Figure 11-6 Dilating and Eroding Images

Note that a structuring element contains only 0's and 1's.

For each placement of the structuring element, the dilation function calculates a destination pixel value as follows. It looks at the values of the source-image pixels that correspond to a 1 in the structuring element and assigns the maximum of these values to the appropriate destination pixel. In the case of an `XIL_BIT` image, this means that if any of the source-image pixels associated

with a 1 in the structuring element is set, the destination pixel will be set. Figure 11-7 shows a binary image before and after dilation of the image's white region. The dilation was performed using the structuring element shown in Figure 11-6.

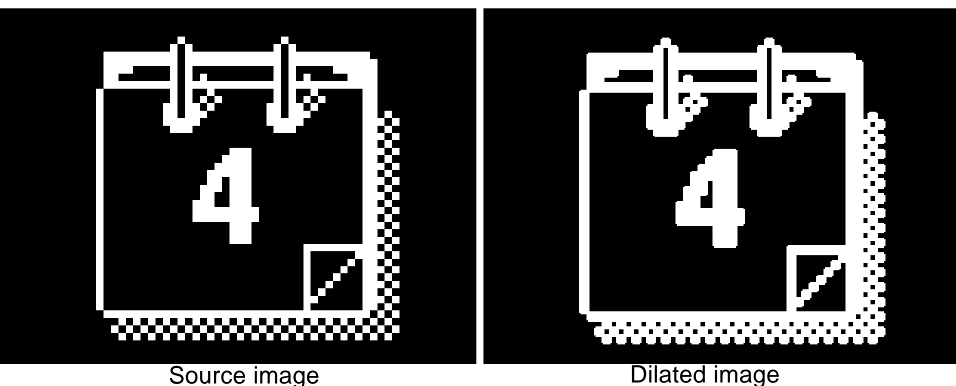

Figure 11-7 Dilating an Image

The erosion function does the opposite of this. It looks at the values of the source-image pixels that correspond to a 1 in the structuring element and assigns the minimum of these values to the appropriate destination pixel. Figure 11-8 shows a binary image before and after erosion of the image's white region.

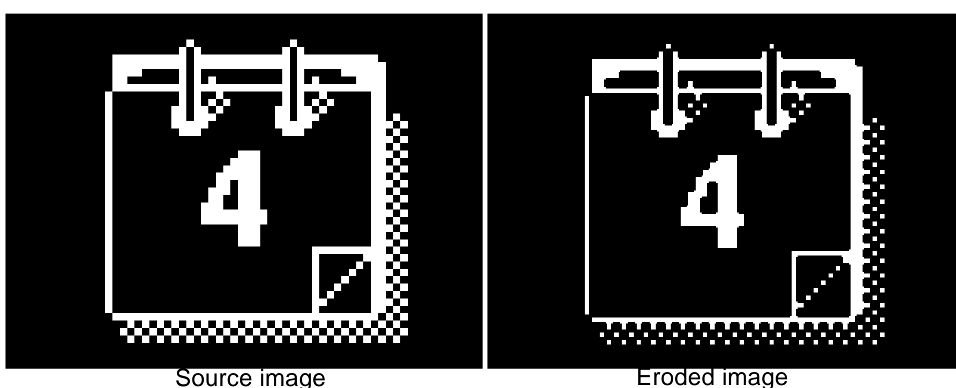

Figure 11-8 Eroding an Image

Performing either a dilation or an erosion in an XIL program is a three-step process:

1. Create a structuring element.
2. Perform the dilation or erosion.
3. Destroy the structuring element.

These tasks are discussed in detail in the sections below.

Creating a Structuring Element

Structuring elements, which are XIL data structures of type `XilSel`, are created with the function `xil_sel_create()`. The prototype for this function is shown below.

```
XilSel xil_sel_create(XilSystemState state,
    unsigned int width, unsigned int height, unsigned int keyx,
    unsigned int keyy, unsigned int *data);
```

The parameter `state` is a system-state data structure. You received a handle to this structure when you initialized the library. The parameters `width` and `height` are the width and height of the structuring element in pixels. These values are usually odd numbers so that the structuring element will have a

center value, but need not always be odd. Common sizes for structuring elements are 3-by-3 and 5-by-5. The parameters `keyx` and `keyy` define the coordinates of the key value in the kernel. These coordinates are specified with respect to the upper-left value in the structuring element (0,0). The key value is usually in the center of the structuring element. The last parameter, `data`, is a pointer to the Boolean values that will be written to the kernel.

The code needed to create the structuring element shown in Figure 11-6 would look something like this.

```
XilSel cross_sel;
XilSystemState state;
unsigned int width, height, keyx, keyy;
unsigned int data[] = {0, 1, 0,
                       1, 1, 1,
                       0, 1, 0};

width = height = 3;
keyx = keyy = 1;
cross_sel = xil_kernel_create(state, width, height, keyx, keyy,
    data);
```

You can also assign names to structuring elements that you've created and then get handles to those structuring elements using the function `xil_sel_get_by_name()`. To assign a name to a custom structuring element, you use the function `xil_sel_set_name()`, whose prototype is shown below.

```
void xil_sel_set_name(XilSel sel, char *sel_name);
```

You can determine whether a structuring element has a name and, if it does, what that name is using the function `xil_sel_get_name()`.

```
char *xil_sel_get_name(XilSel sel);
```

Dilating or Eroding an Image

Once you have created a structuring element, you perform a dilation by calling `xil_dilate()`, whose prototype is shown below.

```
void xil_dilate(XilImage src, XilImage dst, XilSel sel);
```

The parameters `src` and `dst` are handles to the source and destination images for the operation. These images must have the same number of bands and the same data type. The parameter `sel` is a handle to the structuring element you created earlier.

To erode a region in an image, you call `xil_erode()`, which takes the same parameters as `xil_dilate()`.

```
void xil_erode(XilImage src, XilImage dst, XilSel sel);
```

Destroying a Structuring Element

After performing your dilation or erosion, you should destroy your structuring element if you will not be using it again. Destroying the structuring element frees the memory that was allocated to store that structure.

You destroy a structuring element by calling `xil_sel_destroy()`.

```
void xil_sel_destroy(XilSel sel);
```

The only parameter to this function is a handle to the structuring element you want to destroy.

Additional Structuring-Element Functions

The preceding sections have discussed the most frequently used XIL functions that affect structuring elements. However, the library also contains the structuring-element-related functions shown in Table 11-6. These functions enable you to make a copy of a structuring element and read the values of structuring-element attributes.

Table 11-6 Utility Functions for Structuring Elements

Function Name	What the Function Does
`xil_sel_create_copy`	Returns a copy of a structuring element
`xil_sel_get_width`	Gets the width of a structuring element
`xil_sel_get_height`	Gets the height of a structuring element
`xil_sel_get_key_x`	Gets the *x* coordinate of the key value of a structuring element
`xil_sel_get_key_y`	Gets the *y* coordinate of the key value of a structuring element

Passing an Image Through a Lookup Table

XIL lookup tables provide a very general mechanism for modifying images. A lookup table enables you to 1) convert a single-band image of any data type to a single-band or multiband image of any data type, or 2) convert a multiband image of any data type to a multiband image of any data type. Each lookup table allows you to specify precisely the correlation between the values of pixels in the source image and values in the destination image.

The general procedure for performing a lookup operation involves three steps:

1. Create a lookup table.
2. Pass a source image through the lookup table.
3. Destroy the table.

These tasks are discussed in detail in the sections that follow.

Creating a Lookup Table

The steps for creating a lookup table vary slightly, depending on whether the input image is a single-band or multiband image. For a single-band input image, you create a single lookup table; for a multiband input image, you create a single lookup table for each band in the input image, then combine those tables into a *combined* lookup table. Regardless of whether you create a single lookup table or a combined lookup table, you pass the image through the table the same way.

Lookup tables for single-band and multiband input images are discussed separately in the following sections.

Creating Lookup Tables for Single-Band Input Images

To create a lookup table for a single-band input image, you make a single call to the function `xil_lookup_create()`, whose prototype is shown below.

```
XilLookup xil_lookup_create(XilSystemState state,
    XilDataType input_data_type, XilDataType output_data_type,
    unsigned int output_nbands, unsigned int num_entries,
    short first_entry_offset, void *data);
```

This function returns a handle to a data structure of type `XilLookup`, which is the lookup table. Table 11-7 lists the parameters to `xil_lookup_create()` and explains what each parameter represents.

Table 11-7 Parameters to `xil_lookup_create()`

Parameter	What It Represents
`state`	A handle to the system-state data structure that was created when you initialized the XIL library.
`input_data_type`	The data type of the pixel values in the source image. The three possible data types are `XIL_BIT`, `XIL_BYTE`, and `XIL_SHORT`.
`output_data_type`	The data type of the pixel values in the destination image.
`output_nbands`	The number of bands in the destination image.

Table 11-7 Parameters to `xil_lookup_create()`

Parameter	What It Represents
`num_entries`	The number of entries in the lookup table. Each entry specifies the correspondence between a source-image pixel value and the value, or values, that will define the corresponding pixel in the destination image.
`first_entry_offset`	The source-image pixel value for the first entry in the table. The source-image pixel value for the last entry in the table is `first_entry_offset` + `num_entries` - 1.
`data`	A pointer to the data to be stored in the table.

As an example of how this function might be used, say that you have an `XIL_BIT` image that you want to convert to a three-band `XIL_BYTE` image so that you can display it on a 24-bit display. Further, you want black pixels in the source to appear dark gray in the destination image, and white pixels to appear light blue. The code to create a lookup table for this operation is shown below.

```
XilLookup lookup_table;
XilSystemState state;
Xil_unsigned8 data[] = {79, 79, 47, 216, 216, 191};

lookup_table = xil_lookup_create(state, XIL_BIT, XIL_BYTE, 3, 2,
    0, data);
```

The lookup table this code would produce is shown in Figure 11-9.

0	79	79	47
1	216	216	191

Figure 11-9 Single Lookup Table

The numbers to the left of the bold line are XIL_BIT values to be looked up in the source image. The numbers to the right of the line are XIL_BYTE values to be written to the destination image. Because the destination has three bands, the table contains three values for each pixel.

A lookup table can also be used to convert an XIL_BIT image to an 8-bit XIL_BYTE image; however, it may be easier to use xil_cast() to cast the XIL_BIT data type to XIL_BYTE. Generally, the xil_cast() function casts the values 0 and 1 in the XIL_BIT image to indices 0 and 1 in the XIL_BYTE image. If you need different indices, convert the image by passing it through a lookup table.

After creating the lookup table, you must pass the image through it, as discussed in "Passing an Image Through the Table" on page 204.

Creating Lookup Tables for Multiband Input Images

To create a lookup table for a multiband input image, you call xil_lookup_create() once for each band in the input image, then combine these single lookup tables into a *combined* lookup table. To create the combined lookup table, you call the function xil_lookup_create_combined(), whose prototype is shown below:

```
XilLookup xil_lookup_create_combined(XilSystemState state,
    XilLookup lookup_list[], unsigned int num_lookups);
```

This function returns a handle to a data structure of type XilLookup, which is the combined lookup table. The parameter state is a handle to the system-state data structure that was created when you initialized the XIL library, lookup_list[] is an array of type XilLookup that stores the single lookup tables created for each of the input image's bands, and num_lookups indicates how many lookup tables are stored in the lookup_list[] array.

To create the combined lookup table:

1. **Create an array variable of data type** XilLookup**.**
   The variable should have as many array elements as there are bands in the input image so you can store each band's values in a separate element. Store the values for band 0 in element 0, those for band 1 in element 1, and so on.

2. **For each band in the input image, create a single lookup table by calling** `xil_lookup_create()`**, as discussed in "Creating Lookup Tables for Single-Band Input Images" on page 200.**
   Each single lookup table defines values for only one band in the destination image; thus, you must pass a `1` for the `xil_lookup_create()` function's `output_nbands` argument. The single lookup tables must all have the same data type, but each can use a different offset.

3. **Combine the single tables by calling the** `xil_lookup_create_combined()` **function, passing it the array variable you created in Step 1.**

4. **Pass the input image through the combined lookup table, as discussed in the next section.**

The following code fragment shows how you might alter a 3-band XIL_BYTE input image whose green band is accented but whose red and blue bands are subdued:

```
XilImage image;
XilLookup lookup_tables[3]; /* var to store 3 lookup components
*/
XilLookup combined_lookup_table;
Xil_unsigned8 red[256];     /* red component of lookup */
Xil_unsigned8 green[256];   /* green component of lookup */
Xil_unsigned8 blue[256];    /* blue component of lookup */
int i;

for(i=0; i<256; i++) {
   green[i] = (i + 20) < 255 ? i + 20 : 255;
    blue[i] = red[i] = (i - 10) < 0 ? 0 : i - 10;
}
/* create single lookup tables for each input band */
lookup_tables[0] = xil_lookup_create(state, XIL_BYTE, XIL_BYTE,
                                     1, 256, 0, red);
lookup_tables[1] = xil_lookup_create(state, XIL_BYTE, XIL_BYTE,
                                     1, 256, 0, green);
lookup_tables[2] = xil_lookup_create(state, XIL_BYTE, XIL_BYTE,
                                     1, 256, 0, blue);
/*combine the tables and pass input image through combined
table*/
combined_lookup_table = xil_lookup_create_combined(state,
                                              lookup_tables, 3);
xil_lookup(image, image, combined_lookup_table);
```

Passing an Image Through the Table

Once you've created a lookup table, you call xil_lookup() to perform the actual lookup operation. The prototype for this function is shown below.

```
void xil_lookup(XilImage src, XilImage dst, XilLookup lookup);
```

The parameter src is a handle to your source image. The data type of this image must match the data type of the values on the input side of your lookup table. The parameter dst is a handle to your destination image. The data type of this image must match the data type of the values on the output side of the

table. For a single-band input image, the number of bands in the destination image must match the number of values per pixel on the output side of the table; for a multiband input image, the number of bands in the destination image must match the number of lookup tables that compose the combined lookup table. The final parameter, `lookup`, is the handle to the lookup table you received when you called `xil_lookup_create()` for a single-band input image, or `xil_lookup_create_combined()` for a multiband input image.

Destroying a Lookup Table

After performing your lookup operation, you should destroy your lookup table if you won't be using it in a subsequent operation. Destroying the table frees the memory that was allocated to store it. For multiband input images, be sure to destroy the tables created for each input band, as well as the combined table.

You destroy a lookup table by calling the function `xil_lookup_destroy()`, whose prototype is shown below.

```
void xil_lookup_destroy(XilLookup lookup);
```

The only parameter to this function is a handle to the lookup table you want to destroy.

Additional Lookup-Table Functions

The preceding sections have discussed the most frequently used XIL functions that affect lookup tables. However, the library also contains the lookup-table-related functions shown in Table 11-8. These functions enable you to make a copy of a lookup table, read the values of lookup-table attributes, assign a name to a lookup table, and so on.

Note – Some of the functions cannot be used on a combined lookup table. However, you can extract a band from a combined lookup table and perform those functions on the extracted copy. You can then use the altered table for creating another combined lookup table.

Table 11-8 Additional Functions for Lookup Tables

Function Name	What the Function Does
`xil_lookup_create_copy`	Creates and returns a copy of an existing lookup table (LUT).
`xil_lookup_get_input_datatype`	Gets the data type of an LUT's input.
`xil_lookup_get_output_datatype`	Gets the data type of an LUT's output.
`xil_lookup_get_input_nbands`	Gets the number of bands in the LUT's input.
`xil_lookup_get_output_nbands`	Gets the number of bands defined for the LUT's output.
`xil_lookup_get_num_entries`	Gets the number of entries in the LUT. Cannot be used on a combined LUT.
`xil_lookup_get_offset`	Gets the input value for the first entry in the LUT. Cannot be used on a combined LUT.
`xil_lookup_set_offset`	Sets the input value for the first entry in the LUT. Cannot be used on a combined LUT.
`xil_lookup_get_band_lookup`	Gets a particular lookup table out of a combined LUT.
`xil_lookup_get_values`	Retrieves the values in an LUT. Cannot be used on a combined LUT.

Table 11-8 Additional Functions for Lookup Tables

Function Name	What the Function Does
xil_lookup_set_values	Sets the values in an LUT. Cannot be used on a combined LUT
xil_lookup_get_version	Gets a unique identifier associated with an LUT.
xil_lookup_convert	Creates an LUT that includes input data from one lookup table and output data from another. Cannot be used on a combined LUT.
xil_squeeze_range	Creates an LUT that will map a single-band image to a single-band image with contiguous values. Cannot be used on a combined LUT.
xil_lookup_get_name	Reads the name of an LUT.
xil_lookup_set_name	Assigns a name to an LUT.
xil_lookup_get_by_name	Gets a handle to an LUT by specifying the name of the table.

Linear Combination

The function `xil_band_combine()` calculates each value in a destination image by performing a linear combination (matrix multiply) of the values of all the bands of a pixel in the source image. The nonimage values used in this calculation are stored in a matrix. For example, assume that `xil_band_combine()` is operating on the three-band single-pixel image shown in Figure 11-10 using the matrix shown in the figure and is writing its output to a one-band single-pixel image.

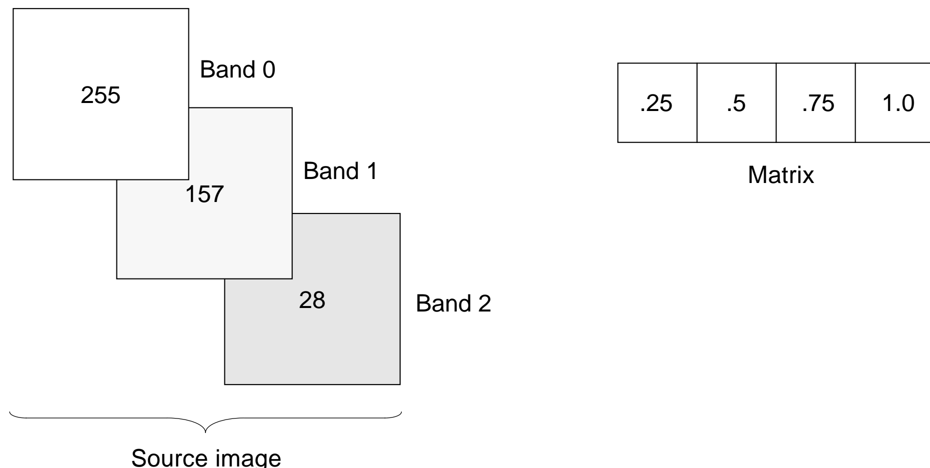

Figure 11-10 Interband Linear Combination

The equation used to calculate the value of the destination-image pixel would be

dst = (255 * .25) + (157 * .5) + (28 * .75) + 1.0

As you can see, the number of columns in the matrix is equal to the number of bands in the source image plus one. This number of columns provides a multiplier for each band in the source image plus a constant that is to be added to the sum of the products obtained using the multipliers and the values in the source image. The number of rows in the matrix must equal the number of bands in the destination image. If the destination image had three bands, the values in the second row of the matrix would be used in calculating the values in the second band of the destination, and the values in the third row would be used in calculating the values in the third band.

Thus, the matrix shown in Figure 11-11 would operate on a three-band source image and would produce a three-band image identical to the source.

1.0	0.0	0.0	0.0
0.0	1.0	0.0	0.0
0.0	0.0	1.0	0.0

Figure 11-11 Linear Combination Matrix

Performing a Linear Combination

Performing a linear combination is a two-step process. You first create an XIL kernel (a object of type `XilKernel`) in which to store a matrix. Then, you call `xil_band_combine()` to perform the operation.

To create the kernel that holds the matrix, you call the function `xil_kernel_create()`. For instance, to create a kernel that describes the matrix shown in Figure 11-11, you could use the code shown below.

```
XilSystemState state;
XilKernel matrix;
unsigned int width = 4, height = 3, keyx = 0, keyy = 0;
float data[] = {1.0, 0.0, 0.0, 0.0,
                0.0, 1.0, 0.0, 0.0,
                0.0, 0.0, 1.0, 0.0};

matrix = xil_kernel_create(state, width, height, keyx, keyy,
    data);
```

The values of `keyx` and `keyy`, which define the key value of the kernel, are irrelevant if the kernel will be used in a linear-combination operation because `xil_band_combine()` ignores these values. For more information about key values, see the section “Filtering an Image” on page 187.

Note – When you’ve finished using this kernel, you should destroy it using the function `xil_kernel_destroy()`.

Once you've created this kernel, you can call `xil_band_combine()`, whose function prototype is shown below.

```
void xil_band_combine(XilImage src, XilImage dst, XilKernel
    matrix);
```

The parameter `src` is a handle to the source image, `dst` is a handle to the destination image, and `matrix` is a handle to the XIL kernel created earlier.

How to Use Linear Combinations

The function `xil_band_combine()` is a very general-purpose function, but let's look at a couple of ways in which it could be used.

One possibility is to use the function to convert an image from one color space to another. For example, you could use the matrix shown in Figure 11-12 to convert a three-band `XIL_BYTE` image from the RGB to the CMY color space.

0.0	0.0	-1.0	255.0
0.0	-1.0	0.0	255.0
-1.0	0.0	0.0	255.0

Figure 11-12 RGB-to-CMY Conversion Using `xil_band_combine()`

Similarly, you could use the matrix shown in Figure 11-13 to convert a three-band `XIL_BYTE` RGB image to a one-band Y image.

0.114	0.587	0.299	0.0

Figure 11-13 RGB-to-Y Conversion Using `xil_band_combine()`

Note, however, that in most cases using `xil_band_combine()` is not the best way of performing color-space conversions using the XIL library. See the section "Color-Space Conversion" on page 122 for information about the preferred method of performing these conversions.

Another possibility would be to use the matrix shown in Figure 11-14 to convert a three-band source image to a single-band destination of the same data type.

0.333	0.333	0.333	0.0

Figure 11-14 Calculating the Normalized Sum of an Image

Each pixel in the destination would equal the normalized sum of the three values of the corresponding pixel in the source.

Blending Images

The XIL library's blending function, `xil_blend()`, blends two images using an alpha mask. The function's prototype is shown below.

```
void xil_blend(XilImage src1, XilImage src2, XilImage dst,
    XilImage alpha);
```

The parameters `src1` and `src2` are handles to the two images to be blended, and `dst` is a handle to the destination image. The last parameter, `alpha`, is a handle to an XIL image that serves as an alpha mask for the blending operation. This alpha image must be a single-band image.

The equation below shows how `xil_blend()` uses the values in the mask as it calculates the values of pixels in the destination image:

dst = (((1 - normalized-alpha) * src1) + (normalized-alpha * src2))

Note – A normalized value is one that has been scaled into the range 0 to 1, where 0 corresponds to the minimum value for the image's data type and 1 corresponds to the maximum value. For example, in an XIL_BYTE image, a value of 127 would be normalized to approximately .5 since it falls about half way between 0 and 255.

Figure 11-15 illustrates the blending of two XIL_BYTE images, labeled src1 and src2. The pixels in the dark areas of these images have values of 100, and the white pixels have values of 255. The alpha mask, alpha, has all its pixels set to 127, so normalized alpha is approximately .5. Thus, the equation used to calculate destination-image values in this example is

dst = ((.5 * src1) + (.5 * src2))

There will only be three distinct values in the destination:

- Where src1 and src2 were both white, the destination value is calculated as follows:

  (.5 * 255) + (.5 * 255) = 255

  These pixels appear white in the destination.

- Where src1 and src2 were both dark (pixel value of 100), the destination value is calculated as follows:

  (.5 * 100) + (.5 * 100) = 100

  These pixels appear dark gray in the destination.

- Where src1 was white and src2 was dark, or vice versa, the destination value is calculated as follows:

  (.5 * 255) + (.5 * 100) = 178

  These pixels appear light gray in the destination.

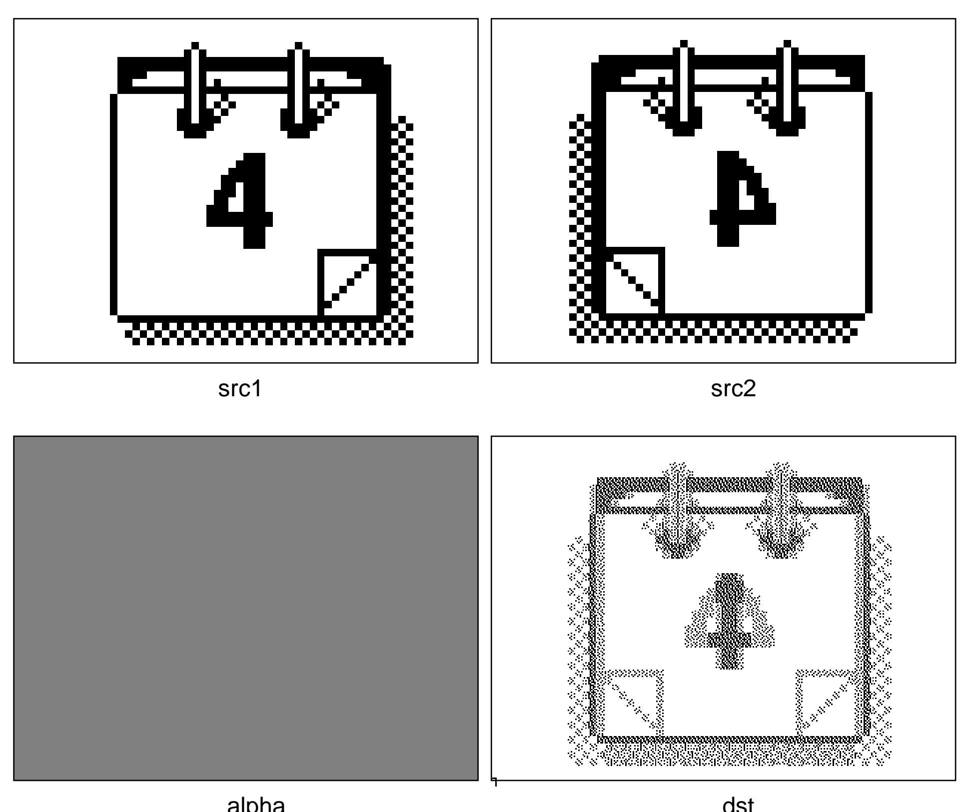

Figure 11-15 Blending Images

To see an example of the blending of two color images, see Color Plate 3.

Painting on an Image

The library's paint operation requires you to provide a list of the pixels in the source image you want to paint by specifying their coordinates. You also specify a kernel that the operation will use as a brush. This kernel is like the kernel you use with `xil_convolve()`, except that all the values in the kernel must fall in the range 0.0 to 1.0.

Note – For information on creating an XIL kernel, see the section “Creating a Convolution Kernel” on page 188.

Finally, you specify a color that you want to use for your painting.

In one sense, the paint function uses the kernel you supply much as the convolution operation does. For each pixel that you want to paint, the paint function lays the key value of the kernel over the pixel. See Figure 11-16.

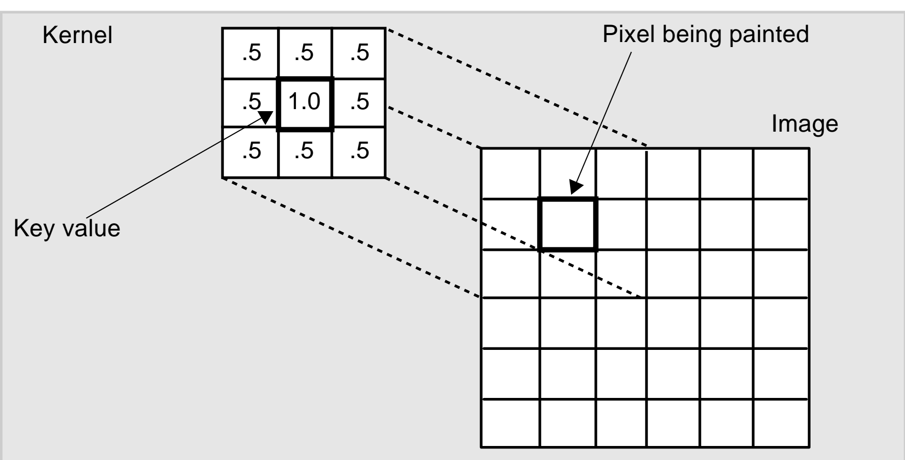

Figure 11-16 Painting on an Image

In another sense, however, the paint function uses the kernel differently. In a paint operation, a neighborhood of pixels in the source does not contribute to the value of a singe pixel in the destination. Instead, for each placement of the kernel, each source-image pixel that lies under the kernel affects the value of the corresponding pixel in the destination. This part of the painting operation is similar to a blending operation. The pixels under the kernel are blended with the paint color, and the kernel serves as the equivalent of an alpha mask. Destination-image pixel values are determined by the following formula:

dst = (brush-value * color) + ((1.0 - brush-value) * src)

You perform a paint operation in an XIL program by calling `xil_paint()`, whose prototype is shown below.

```
void xil_paint(XilImage src, XilImage dst, float *color,
    XilKernel brush, unsigned int count, float *coord_list);
```

Table 11-9 explains the purpose of each of the parameters shown above.

Table 11-9 Parameters to `xil_paint()`

Parameter	What It Represents
`src`	A handle to the source image for the operation.
`dst`	A handle to the destination image. This image must have the same number of bands and the same data type as the source image.
`color`	An array of `float`s containing a number of elements equal to the number of bands in the source and destination images. The elements in this array define the color to be used for painting.
`brush`	A handle to the kernel that will be used as the brush.
`count`	The number of pixels to be painted.
`coord_list`	An array of `float`s containing `2 * count` elements. These elements represent the *x* and *y* coordinates of the pixels to be painted.

Setting and Getting the Values of Pixels in an Image

The XIL library contains three functions that set or retrieve the values of pixels in an image. Two of these functions—`xil_set_pixel()` and `xil_get_pixel()`—set or get the value of individual pixels. The third, `xil_set_value()`, sets all the pixels in an image.

xil_set_pixel() and xil_get_pixel()

The prototypes for these two functions are shown below.

```
void xil_set_pixel(XilImage image, unsigned int x,
    unsigned int y, float *values);

void xil_get_pixel(XilImage image, unsigned int x,
    unsigned int y, float *values);
```

In both prototypes, the parameter `image` is a handle to the image you're working with. The parameters `x` and `y` are the *x* and *y* coordinates of the pixel you want to read or write. These coordinates should describe the location of the pixel of interest with respect to the image's origin, which by default is in the upper-left corner of the image (0,0). The final parameter, `values`, is an array of floating-point numbers. The number of elements in this array must match the number of bands in the image.

When you read a pixel, the value of band 0 of that pixel is written to `values[0]`, the value of band 1 is written to `values[1]`, and so on. Similarly, when you write a pixel, `values[0]` is written to band 0 of the pixel, and so forth. Before the writing actually takes place, the floating-point numbers in `values` are converted to integers.

Note – If you attempt to write to a pixel a value that is out of range for the source image's data type, the value will be clamped to the low or high limit for the data type. For example, an `XIL_BYTE` image can only accommodate values in the range 0 to 255. If you try to write a 400 to such an image, the function will actually write a 255.

---

xil_set_value()

The function `xil_set_value()` is useful for clearing an image by setting all the pixels in the image to the same color. The prototype for this function is shown below.

```
void xil_set_value(XilImage dst, float *values);
```

The parameter dst is a handle to the image you are working with, and values is an array of floating-point numbers. The number of elements in this array must match the number of bands in the image dst. The first element in the array is rounded to an integer and is then written to all of band 0 of your image; the second element is rounded to an integer and written to band 1 of the image; and so on.

Note – If you attempt to write to a band a value that is out of range for the image's data type, the value will be clamped to the low or high limit for the data type. For example, an XIL_BYTE image can only accommodate values in the range 0 to 255. If you try to write a 400 to a band, the function will actually write a 255.

Copying a Pattern to an Image

Besides xil_copy(), which copies a source image to a destination image, the XIL library includes a function, xil_copy_pattern(), that writes as many copies of the source image as possible to the destination image. The function prototype for xil_copy_pattern() is shown below.

```
void xil_copy_pattern(XilImage src, XilImage dst);
```

The parameter src is a handle to the source image, and dst is a handle to the destination.

If your destination image is the same size as, or smaller than, your source image, using xil_copy() and xil_copy_pattern() will produce the same result. However, if your destination image is larger than your source, xil_copy() will write one copy of the source image to the destination, while xil_copy_pattern() will write as many copies as possible. Some of these copies will be partial copies if the dimensions of the destination are not multiples of the dimensions of the source. Figure 11-17 below illustrates the difference between using xil_copy() and xil_copy_pattern() to copy a source image to a destination that is twice as high and wide as the source.

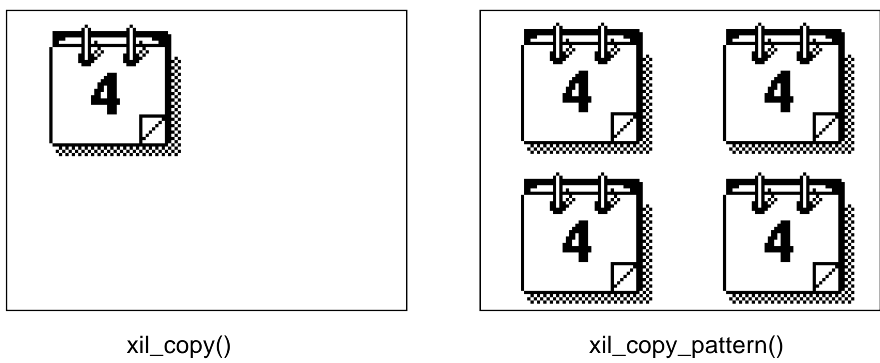

Figure 11-17 Replicating a Source Image

Compressing and Decompressing Sequences of Images 12

Besides enabling you to perform image-processing operations on XIL images, the XIL library also gives you the ability to compress one or more XIL images and to store the compressed images in a data structure called a compressed image sequence (CIS). The XIL images you compress may represent frames of video from a movie, pages from a document, or any number of things. You can also decompress the data stored in a CIS and write your output to one or more XIL images. This compression-decompression process is shown in Figure 12-1.

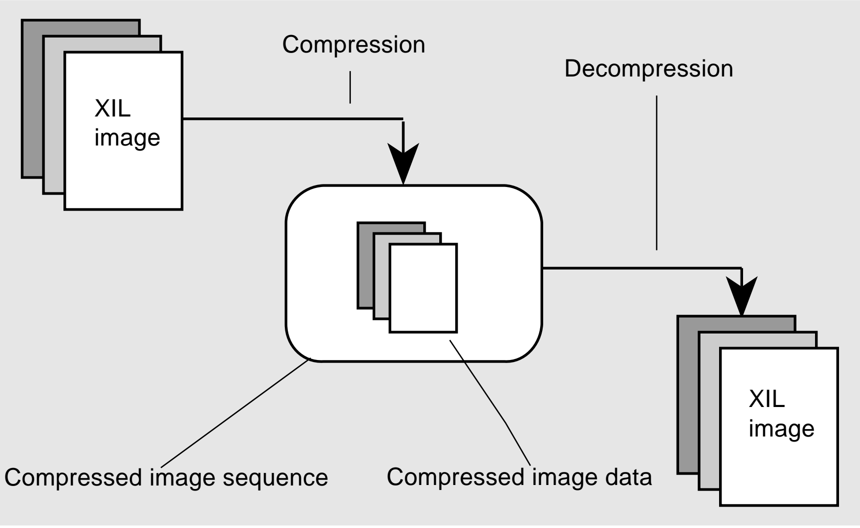

Figure 12-1 Compressing and Decompressing XIL Images

As you probably know, uncompressed images can take up a lot of disk space and take a long time to move over the network. For example, a 512-by-512 3-band image containing XIL_BYTE data takes up .75 Mbytes. A compressed version of this image might be anywhere from one-half to one-fiftieth this size, depending on the type of compressor used.

The XIL library contains several compression and decompression modules:

- Cell and CellB codecs
- JPEG baseline sequential and lossless codecs
- An H.261 decompressor
- An MPEG-1 decompressor
- CCITT Group 3 and Group 4 codecs

The Cell image compression technology, which was developed by Sun, has been optimized for the rapid decompression and display of images on simple hardware. Therefore, the Cell codec is able to achieve reasonable display quality on indexed-color frame buffers. The initial focus of the Cell technology is on Sun-to-Sun communications, where the benefits of fast decoding outweigh the benefits of standards. Possible areas of application include media distributions on CD-ROM and multimedia mail.

The CellB codec, which derives from its Cell counterpart, is intended for use primarily in videoconferencing applications. It features greater balance between the time spent compressing and decompressing images than the Cell codec. The CellB codec's strengths include

- Software compression at interactive rates
- Very fast decoding and display, especially on indexed-color frame buffers
- Low rates of CPU use
- Good quality output

The JPEG compression standards were developed by the Joint Photographic Experts Group to support the compression of still images, both grayscale and color. Although not specifically designed for the compression of sequences of images, or movies, JPEG compressors are also used frequently for that purpose. A *lossy* JPEG compressor compresses an image in such a way that when the compressed data is decompressed, the decompressed image and the original image may not match exactly. On the other hand, a *lossless* JPEG compressor processes an image so that the decompressed image does match the original image pixel for pixel.

The H.261 compression-decompression scheme was developed by the International Telegraph and Telephone Consultative Committee (CCITT). The H.261 video encoder is intended to be used to compress video data that will be sent over Integrated Services Digital Network (ISDN) lines. The H.261 codec is intended primarily for use in video telephony and videoconferencing applications.

The MPEG-1 video compression standard was developed by the Moving Picture Experts Group. The group's goal was to compress full-motion video and the associated audio at the rate of about 1.5 Mbits/s. This is approximately the rate at which data can be read from a CD-ROM, so MPEG-1 compressed video is a good choice for use in interactive multimedia applications.

The CCITT Group 3 and Group 4 compression standards were developed by the International Telegraph and Telephone Consultative Committee to enable facsimile machines to compress and decompress digitized documents. Now, Group 3 and Group 4 compressors and decompressors are used for general document storage and retrieval.

The remainder of this chapter introduces you to the subject of XIL compression and decompression by examining two example programs: `encode` and `xilcis_example`. The first program builds a Cell, CellB, or JPEG movie. It reads a series of video frames from disk files, compresses each frame, and writes the compressed image data to a CIS. As frames are compressed, the program writes the compressed data in the CIS to an output file. This output file can serve as input to the second example program, a movie player.

The second example program plays back a Cell, CellB, JPEG, H.261, or MPEG-1 datastream on an 8- or 24-bit frame buffer. It reads a datastream from a disk file and stores the compressed data in a CIS. Then, for each compressed image in the CIS, it takes the following actions. First, it decompresses a frame of video from the CIS and stores the resulting image in a 3-band `XIL_BYTE` image. Then, it prepares each image for display on a particular frame buffer and writes the image to a display image. These steps result in a frame of video being displayed.

Creating a JPEG Datastream

The source files for the movie-making example—encode.c, load_file.c, and image_transform.c—can be found in the directory $XILHOME/examples/movie_maker_example. This directory also contains the image files that the program should use as input and a makefile you can use to build the program.

After briefly discussing how to build and run the example, this section walks you through the program's code to clarify how the program produces a JPEG datastream. (Cell and CellB bytestreams are produced in a very similar way.) The basic algorithm is shown below:

1. Create a CIS.
2. While all the video frames have not been compressed, perform the following steps:
   a. Load data from an image file into an XIL image.
   b. Convert the format of the XIL image if that's necessary or desirable. It will be necessary if the image does not contain XIL_BYTE data because the JPEG compressor expects input of that type. If your image is not a YC_bC_r image, it is desirable to convert the image to that color space because YC_bC_r images are optimized for compression and decompression. The JPEG compressor will work with images of any color space, but it typically produces superior compression with YC_bC_r images.
   c. Compress the image and write the compressed data to the CIS.
   d. Destroy the image in order to free the memory allocated to hold the image structure and the image data.
   e. Write any compressed data in the CIS to a disk file.
3. Perform any outstanding compression operations, and write any resulting compressed data to the output file.

Building and Running the Example

To build and run the example `encode`, perform the following steps:

1. **Change your working directory to** `$XILHOME/examples/movie_maker_example` **and type** `make`**. The makefile in that directory will build the program.**

2. **While in the same directory, execute the program using the following syntax:**

   `%` **`encode`** `[-c | -cb]` *file-list* `[`*output-file*`]`

   By default, the program creates a JPEG bitstream. To produce a Cell bytestream instead, use the `-c` option, and to produce a CellB bytestream, use the `-cb` option.

   The argument *file-list* must be the name of a file that contains a list of files (images) to be processed. The list of files supplied with this example is called `mifkin.list`.

   The argument *output-file* is optional. If you supply a file name here, the example will write its output to a file of that name; otherwise, it will write its output to `out.jpg`, `out.cell`, or `out.cellb`, depending on which compressor you've used.

Creating a CIS

Before a program can compress an image or sequence of images, it must create a CIS in which to store compressed data. The example program creates this data structure using code similar to this.

```
XilSystemState state;
XilCis cis;

cis = xil_cis_create(state, "Jpeg");
xil_cis_set_attribute(cis, "ENCODE_411_INTERLEAVED",
    (void *)TRUE);
xil_cis_set_max_frames(cis, 100);
```

You use the function `xil_cis_create()` to create the CIS, which is a data structure of type `XilCis`. The two arguments to this function are a system-state structure (this was returned by an earlier call to `xil_open()`) and the name of a compressor/decompressor. The example requests the JPEG baseline sequential codec (`Jpeg`), but could have requested another compressor/decompressor. (For information about the strings you use to create CISs associated with other compressors/decompressors, see the section "Creating and Destroying a CIS" on page 250.)

The next line sets a JPEG-specific attribute called `ENCODE_411_INTERLEAVED` to `TRUE`. If the images being compressed are 3-band YC_bC_r images, setting this attribute causes the codec to subsample the data in the color bands so that only one color value will be encoded for each four color values in the original (4:1:1). This subsampling enables the compressor to achieve a much higher level of compression than it could otherwise.

The last line sets a CIS attribute that limits the number of frames the CIS will buffer. In this case, the CIS will maintain 100 frames in its buffer. If more than 100 frames are compressed, the oldest frames will become inaccessible. This attribute is set to keep the program from using too much memory in instances when there are many frames to be compressed.

CIS attributes are discussed further in the section "General CIS Attributes" on page 257.

Compressing Video Frames and Writing Compressed Data to a File

After creating a CIS, the example begins compressing frames of video. For each frame, the program performs the steps described in the sections below.

Loading Data from a File into an XIL Image

The example calls a routine called `load_file()` to read an image from a disk file into an XIL image. For a complete explanation of what is happening in this subroutine, see the section "Acquiring an Input Image" on page 13. The only thing new in this example is that `load_file()` reads a color space from the

image file header and, after creating the XIL image into which it will read pixel values, sets that XIL image's color-space attribute appropriately. See the code fragment below.

```
FILE *in_file;
char colorspace[80];
XilSystemState state;
unsigned int width, height, nbands;
XilDataType datatype;
XilImage image;

fscanf(in_file, "%s", colorspace);
image = xil_create(state, width, height, nbands, datatype);
xil_set_colorspace(image, xil_colorspace_get_by_name(state,
    colorspace));
```

The reason for checking the color space of the input image is that the JPEG compressor is optimized for compressing YC_bC_r images. The images supplied with the example are RGB images, so the example converts each image to the YC_bC_r color space before compressing it. This task is described in the next section.

Converting an XIL Image to the Proper Format

Compressors have certain requirements for the images they compress. (See the section "Putting Compressed Data into a CIS" on page 251 for a list of these requirements.) For example, the JPEG compressor used in this example requires `XIL_BYTE` images. In addition, the compressor will produce a bitstream that can be decompressed most rapidly if the input is YC_bC_r images.

The images provided for you to use as input to the example are RGB XIL_BYTE images. Thus, the example does not have to convert the images in any way, but can produce a bitstream that can be played back the fastest if it converts the color space of the images. This conversion is handled in the routine image_transform().

```
XilImage
image_transform(XilSystemState state, XilImage image,
    int desired_nbands, char *desired_colorspace)
{
    XilImage new_image;
    unsigned int width, height, nbands;
    XilDataType datatype;

    xil_get_info(image, &width, &height, &nbands, &datatype);
    new_image = xil_create(state, width, height, desired_nbands,
        datatype);
    if (new_image == NULL) {
      /* XIL sends an error msg to stderr if image create fails */
        return(NULL);
    }
    xil_set_colorspace(new_image,
        xil_colorspace_get_by_name(state, desired_colorspace));
    xil_color_convert(image, new_image);
    xil_destroy(image);
    return(new_image);
}
```

This routine is passed a source image and a desired color space (ycc601). Given this data, the routine:

- Creates a second XIL image, new_image, that has the same width, height, and data type as image and a number of bands that fits the desired color space. Because a YC_bC_r image must have three bands, new_image is a 3-band image.
- Uses the function xil_set_colorspace() to set new_image's color space attribute to ycc601.
- Calls xil_color_convert() to convert the data in image from the RGB color space to the YC_bC_r color space. The converted data is stored in new_image.
- Destroys image since it is no longer needed.

A JPEG playback program will also run fastest if the images it decompresses are multiples of 16 in width and height. For this reason, the example uses the following code to clip each source image to a multiple of 16 in width and height before compressing it.

```
XilImage src, child_clip;

child_clip = xil_create_child(src, 0, 0,
                              xil_get_width(src) & ~0xf,
                              xil_get_height(src) & ~0xf, 0, 3);
```

Compressing an XIL Image

Compressing the clipped YC_bC_r image requires a single function call.

```
XilImage src;
XilCis cis;

xil_compress(child_clip, cis);
```

The first argument to this function is the child YC_bC_r image `child_clip`, and the second is a handle to the program's CIS. The call both compresses the image and writes the compressed data to the CIS.

Destroying an XIL Image

Once an image has been compressed, the example destroys that image to free the resources associated with it.

```
xil_destroy(child_clip);
xil_destroy(src);
```

Writing the Compressed Data to a File

After compressing each frame of video, the example calls a routine `write_file()` to write any compressed data in the CIS to the program's output file. The basic algorithm for this routine is shown below.

- While the CIS contains data that has not been read,
  a. Get a pointer to the first byte of compressed data that has not been read yet. After the first call to the routine that gets this pointer, the pointer will point to the first byte of data in the CIS.
  b. Using the pointer retrieved in the previous step, read data from the CIS and write it to a file.

Note – The loop is necessary because all of the compressed data in the CIS may not be in one continuous buffer.

Here is code that implements this algorithm:

```
int nbytes, nframes;
XilCis cis;
Xil_unsigned8 *data;
FILE *out;

while (xil_cis_has_frame(cis) == TRUE) {
    data = (Xil_unsigned8 *)xil_cis_get_bits_ptr(cis, &nbytes,
        &nframes);
    fwrite((char *)data, sizeof(Xil_unsigned8), nbytes, out);
}
```

The function `xil_cis_has_frame()` determines whether there is at least one complete frame's worth of compressed data in the CIS that has not yet been read. If there is remaining data, the function returns `TRUE`, and the statements inside the loop are executed.

Note – The routine must check for more than one frame's worth of data in the CIS because compressors are not required to write data to the CIS each time `xil_compress()` is called. Instead, they may read a number of frames of video and store them internally before doing any compression. Then, they may write a number of frames' worth of compressed data to the CIS at once. This type of strategy is necessary, for instance, when a compressor wants to do interframe compression.

The first call to `xil_cis_get_bits_ptr()` returns a generic pointer to the beginning of the compressed data in the CIS. It also returns the number of bytes the pointer points to (`nbytes`) and the number of video frames these bytes represent (`nframes`). In addition to returning these values, the function changes an attribute of the CIS called its ***read frame*** from 0 to `nframes`. (Frames are numbered beginning with 0.) Thus, a second call to `xil_cis_get_bits_ptr()` will return a pointer to the first byte of compressed data that is part of the frame numbered `nframes`. The arguments to this function are a handle to your CIS and the addresses of the variables `nbytes` and `nframes`.

The only other action taking place in the loop is that the example is using the function `fwrite()` to write the compressed data pointed to by `data` to the output file `out`.

Performing Any Outstanding Compression Operations

When the loop described in the preceding section ends, the compressor may have read frames of video that it has not yet compressed. To deal with this possibility, the example makes the following two calls.

```
XilCis cis;
FILE *out;

xil_cis_flush(cis);
write_file(cis, &total_nbytes, &total_nframes, out);
```

The function `xil_cis_flush()` instructs the compressor to compress any images it has read but not compressed and to write the last of its output to the CIS. The call to `write_file()` is the same as the one used in the loop discussed in the last section. It writes to the output file any compressed data written to the CIS as a result of the call to `xil_cis_flush()`.

Playing a JPEG Movie

This section discusses a movie-player example supplied with the library. This program can display Cell, CellB, JPEG, H.261, and MPEG-1 movies on 8- and 24-bit displays.

The source files for the movie-player example—`xilcis_example.c`, `memmap.c`, `memmap.h`, and `xilcis_color.c`—can be found in the directory `$XILHOME/examples/movie_player_example`. This directory also contains a one-frame JPEG bitstream that the example can used as input and a makefile you can use to build the program. In addition, you can use as input to this program the output of the movie-maker example discussed earlier in this chapter.

Note – The JPEG code in the program is optimized to handle movies containing YC_bC_r images. The JPEG movie supplied with the example contains this type of image.

After briefly discussing how to build and run the example, the section takes a look at the code the example uses to play back a JPEG movie on an 8-bit display. (If you look at the program, you'll notice that the CellB, H.261, and MPEG-1 cases are handled very similarly to the JPEG case.) The code specific to the Cell case for 8-bit displays is covered in the section "Playing Cell Movies" on page 241.

The basic algorithm the example uses to play a JPEG movie on an indexed-color display is shown below.

1. Memory map the contents of the movie file (the compressed data) into your process's address space.

2. Create a CIS.

3. Give the CIS a pointer to the compressed data you memory mapped earlier. At this point, the CIS is all set up.

4. Create a display image in which to show the movie by following these steps:
   a. Determine the dimensions of the frames in the movie.
   b. Create an X window that is equal in width and height to the frames in the movie.
   c. Create an XIL display image from the X window.
5. Create an XIL image that will hold frames as they are decompressed. This image must have the same width, height, number of bands, and data type as the frames in the movie.
6. Initialize the parameters the example will use when dithering each 24-bit XIL image (frame of video) so that it can be displayed on an 8-bit frame buffer.
7. Install an X colormap.
8. Play back the movie. While the CIS contains unread frames of compressed data:
   a. Decompress a frame of video and store the resulting image in the XIL image created above.
   b. Dither the XIL image, and write the dithered image to the display.

Running the Movie Player

To run the movie player, change your working directory to `$XILHOME/examples/movie_player_example`, and then build the program in that directory by typing `make`. You should execute the example from this same directory, using a command line of the following form:

```
% xilcis_example [-i filename] [-c | -cb | -h | -m] [-s width height]
```

Table 12-1 describes the example's command-line options.

Table 12-1 Command-Line Options for `xilcis_example`

Option	What It Does
`-i` *filename*	Specifies the name of the movie (a file containing compressed movie frames) that you want to play back. If you omit this option, the program will play back a one-frame movie called `earth.jpg`.
`-c`	Instructs the movie player to play back a Cell movie. By default, the player is set up to play a JPEG movie.
`-cb`	Instructs the movie player to play back a CellB bytestream.
`-h`	Instructs the movie player to play back an H.261 bitstream.
`-m`	Instructs the movie player to play back an MPEG-1 bitstream.
`-s` *`width height`*	If you're decoding a CellB bytestream, you must use this option. The arguments *width* and *height* are the width and height in pixels of the images being decoded.

When the program runs, it shows the movie you have selected in an X window. To terminate the program, move your cursor into the program's window and press any mouse button.

Memory Mapping the Movie

When you run the example program, you use the `-i` *`filename`* option to pass it the name of a file containing a movie. Before the example can attach the compressed image data in the file to a CIS, it must map the contents of the file to system memory. You can find the code that handles this task in the source file `memmap.c`.

As you'll see in that source file, after opening the movie file, the example uses the system call `fstat(2V)` to determine the length of the movie in bytes and the system call `mmap(2)` to memory map the file and get a pointer to the beginning of the datastream. Following these calls, the structure member `memfile->mstart` is a pointer to the beginning of the datastream, and `memfile->mlen` is the number of bytes in the datastream.

Creating a CIS

After mapping the compressed movie data to memory, the example creates a CIS to hold that data. The code used to create this `XilCis` data structure is shown below.

```
XilCis cis;
XilSystemState state;
char *cis_type = "Jpeg";

cis = xil_cis_create(state, cis_type);
if (!cis) {
    /* XIL sends error message to stderr if xil_cis_create fails */
     exit(1);
}
```

The arguments to `xil_cis_create()` are a handle to a system-state data structure returned by an earlier call to `xil_open()` and the name of the codec or decompressor that will be used to decompress frames from the datastream, in this case `"Jpeg"`. The function returns a handle to the newly created CIS.

Putting Compressed Data in a CIS

Now that the CIS exists and the example has a pointer to the JPEG bitstream, the example can set the data-pointer member of the CIS to point to the beginning of the JPEG bitstream. It does this using a call to `xil_cis_put_bits_ptr()`:

```
XilCis cis;
int frame_count = -1;

xil_cis_put_bits_ptr(cis, memfile->mlen, frame_count,
     memfile->mstart, NULL);
```

Table 12-2 explains what the several arguments to the function represent.

Table 12-2 Parameters to `xil_cis_put_bits_ptr()`

Argument	What It Represents
`cis`	A handle to the CIS whose data pointer is being set.
`memfile->mlen`	The length in bytes of the JPEG bitstream.
`frame_count`	Indicates that the number of video frames the JPEG bitstream represents is unknown (because `frame_count` is set to -1). If the number of frames in the movie were known, `frame_count` would be set to that number.
`memfile->mstart`	A pointer to the beginning of the JPEG bitstream.
`NULL`	Indicates that the example has assumed responsibility for freeing the memory in which the JPEG bitstream is stored. Alternatively, the example could have passed a pointer to a function to `xil_cis_put_bits_ptr()`. In the latter case, the function passed to `xil_cis_put_bits_ptr()` would be called if the CIS were destroyed or reset.

After the call to `xil_cis_put_bits_ptr()` has been made, the CIS is ready for use. Before the example can begin decompressing frames of video, however, it must create two XIL images: an 8-bit display image in which to display the video and a 24-bit memory image to serve as a destination image for the decompression function.

Creating a Display Image

There are three steps involved in creating the 8-bit display image in which the movie will be shown.

1. **Determine the dimensions of the images stored in the CIS.**
   The code used to get these dimensions is shown below.

```
XilImageType outputtype;
XilCis cis;
unsigned int cis_xsize, cis_ysize, cis_nbands;
XilDataType cis_datatype;

outputtype = xil_cis_get_output_type(cis);
xil_imagetype_get_info(outputtype, &cis_xsize, &cis_ysize,
    &cis_nbands, &cis_datatype);
```

   The first function called, `xil_cis_get_output_type()`, takes a handle to the CIS as its only argument and returns an image type. This image type, a data structure of type `XilImageType`, contains information about the images stored in the CIS, such as their width, height, number of bands, and data type.

   The following call to `xil_imagetype_get_info()` takes the image type as its first argument and returns the width, height, number of bands, and data type of the image type in its remaining arguments. The data type, stored in `cis_datatype`, will be one of the following enumeration constants: `XIL_BIT`, `XIL_BYTE`, or `XIL_SHORT`. (In this case, the value will be `XIL_BYTE` because the movie supplied with the XIL release was made from 3-band `XIL_BYTE` images.)

2. **Create an X window that matches the width and height of the images in the CIS.**
   In this example, the X window is created with a call to `XCreateSimpleWindow()`. When the example is run on a system with an 8-bit display, this window will be 8 bits deep by default.

```
int screen_num;
Display *display;
Window window;

screen_num = DefaultScreen(display);
window = XCreateSimpleWindow(display,
    RootWindow(display, screen_num), 0, 0, cis_xsize,
    cis_ysize, 0, BlackPixel(display, screen_num),
    WhitePixel(display, screen_num));
```

3. **Create a display image based on the X window created above.**
   This display image will be the destination image to which frames are written in order to display them. To create the display image, the example calls `xil_create_from_window()`.

```
XilImage displayimage = NULL;
XilSystemState state;
Display *display;
Window window;

displayimage = xil_create_from_window(state, display, window);
if (!displayimage) {
    /* XIL err msg to stderr if xil_create_from_window fails */
    exit(1);
}
```

   For a complete discussion of display images, see the section “Display Images” on page 44.

Creating an Image to Hold Decompressed Frames

If you have loaded the JPEG movie supplied with the XIL release (`earth.jpg`) or a JPEG movie you created with the example movie maker into your CIS, the images in the CIS are YC_bC_r images (24 bits deep). Because they cannot be decompressed directly into the program’s display image (8 bits deep), the example must create an XIL memory image whose purpose is to hold those decompressed frames. That is, for each frame in the movie, the example must perform the two-step sequence shown in Figure 12-2.

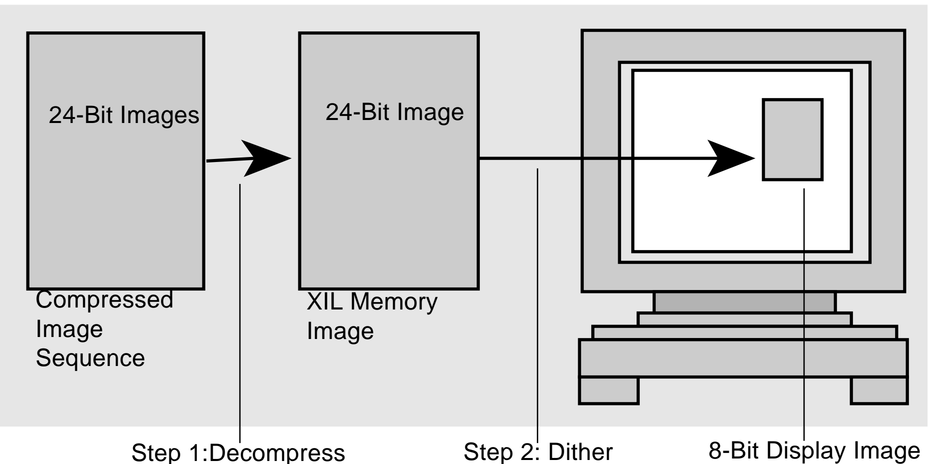

Figure 12-2 Decompressing and Dithering a Frame of Video

The code used to create the XIL memory image is shown below:.

```
XilSystemState state;
XilImage imageYCC = NULL;
unsigned int cis_xsize, cis_ysize, cis_nbands;
XilDataType cis_datatype;

imageYCC = xil_create(state, cis_xsize, cis_ysize, cis_nbands,
    cis_datatype);
```

The variables `cis_xsize`, `cis_ysize`, `cis_nbands`, and `cis_datatype` were set earlier by a call to `xil_imagetype_get_info()`, which returned the width, height, number of bands, and data type of the compressed images in the CIS. Thus, the image returned here, `imageYCC`, will be identical to the images stored in the CIS in both dimensions and data type.

Initializing Parameters to Be Used with the Dither Function

Next, the example program produces the colorcube and dither mask that it will need later to dither images to the display. To get the colorcube, the example calls `xil_lookup_get_by_name()`.

```
XilLookup colorcube;

colorcube = xil_lookup_get_by_name(state, "cc855");
```

This call returns a handle to a special colorcube that the library creates when it is initialized. This colorcube has dimensions of 8, 5, and 5 and is designed for dithering the values in a YC_bC_r image to 200 colors.

The program then creates a dither mask:

```
XilDitherMask dmask;

dmask = xil_dithermask_get_by_name(state, "dm443");
```

This call returns a handle to a special dither mask the library creates when it is initialized. The mask is 4 pixels high and wide and 3 bands deep.

Installing an X Colormap

At this point, the program calls the function `create_cmap()`, from the source file `xilcis_color.c`, to set up the X colormap that will be used in displaying the dithered frames. This function contains primarily Xlib code, but does get some information from the XIL library to get started.

First, the function is passed an XIL lookup table called `yuv_to_rgb`. This is a standard lookup table that is useful for converting the values in a dithered YC_bC_r image to RGB values.

Note – The YC_bC_r images must have been dithered using the standard colorcube `cc855`.

This lookup table specifies which color values should be written to which colorcells in the X colormap.

The function create_cmap() then calls xil_lookup_get_num_entries() to determine the number of entries in the colormap.

```
int cmapsize;

cmapsize = xil_lookup_get_num_entries(yuvtorgb);
```

Once the function knows how many colorcells it needs to allocate, it creates an X colormap and allocates the colorcells it needs to hold the information stored in yuvtorgb. Note that to avoid colormap flashing, the function avoids allocating the last two colorcells in the X colormap if possible and also leaves free as many colorcells as possible at the beginning of the X colormap.

At this point, create_cmap() allocates a buffer called data and then uses the function xil_lookup_get_values() to copy colormap entries from the lookup table yuvtorgb to the buffer data.

```
Xil_unsigned8 *data;

data = (Xil_unsigned8 *)malloc(sizeof(Xil_unsigned8) * cmapsize
     * 3);
xil_lookup_get_values(yuvtorgb,
    xil_lookup_get_offset(yuvtorgb), cmapsize, data);
```

The second argument to xil_lookup_get_values() enables you to start copying entries from the lookup table at an entry other than 0, and here is set to the offset of yuvtorgb.

Once create_cmap() has copied the colormap information from yuvtorgb to the buffer data, the function has direct access to the color values it needs to store in the X colormap, and it goes on to store those values. The only remaining XIL call in the function create_cmap() changes the offset of the colorcube that main uses to dither decompressed images to match the offset at which create_cmap() began writing color values to the X colormap.

```
xil_lookup_set_offset(colorcube, (unsigned int)pixels[0]);
```

The first argument to xil_lookup_set_offset() is a handle to the colorcube—main passed this handle to cmap_create()—and the second represents the pixel value associated with the first colorcell in which

cmap_create() stored color values. This call ensures that all the pixel values in the dithered images created in main will map to the correct spot in the X colormap.

Playing the Movie

After returning from create_cmap(), the program is ready to play the movie. The calls it uses to do this are shown below.

```
while (xil_cis_has_frame(cis)) {
    xil_decompress(cis, imageYCC);
    xil_rescale(imageYCC, imageYCC, scale, offset);
    xil_ordered_dither(imageYCC, displayimage, colorcube,
        dmask);
}
```

The loop is controlled by the return value of xil_cis_has_frame(). As long as the CIS contains at least one complete frame of video, the function returns 1, and the statements in the loop are performed.

The call to xil_decompress() decompresses an image from the CIS and stores it in the intermediate image imageYCC. Since this is a CCIR Rec. 601 YC_bC_r image, the values in the Y band of the image can range from 16 to 235, and those in the color bands can range from 16 to 240. Before such an image is dithered, the values in each band should be scaled to fall in the range 0 to 255.

This rescaling is handled by the call to xil_rescale(). Since imageYCC is both the source and destination image for the operation, the rescale is performed in place. The scale and offset arguments were defined earlier in this way:

```
float scale[3], offset[3];

scale[0] = 255.0 / (235.0 - 16.0);
scale[1] = 255.0 / (240.0 - 16.0);
scale[2] = 255.0 / (240.0 - 16.0);
offset[0] = -16.0 * scale[0];
offset[1] = -16.0 * scale[1];
offset[2] = -16.0 * scale[2];
```

Finally, `xil_ordered_dither()` dithers the 3-band YC_bC_r image in `imageYCC` using the colorcube and dither mask created earlier and writes the result to the display image `displayimage`. This results in a frame of the movie being shown in the X window.

Note – When a JPEG movie made from YC_bC_r images is displayed on an 8-bit frame buffer, as in this example, the exact sequence of calls you use to play the movie has a dramatic effect on the speed with which the movie is shown. The reason for this is that by using the deferred-execution scheme explained in Chapter 21, "Acceleration in XIL Programs," the XIL library can look for a certain sequence of functions at runtime and, if it finds that sequence, replace all the functions in the sequence with a single, highly optimized routine. In this example the sequence `xil_decompress()`, `xil_rescale()`, `xil_ordered_dither()` is such a sequence. Thus, when you play a JPEG movie using this program, these functions are not executed, but are replaced by an optimized function (***molecule***) that performs the jobs of all three functions. See the section "Video Decompression Molecules" on page 399 for a complete list of decompression molecules.

Playing Cell Movies

The example movie player performs the same steps regardless of whether its input is a JPEG movie or a Cell movie up through the point where it creates a display image. From that point on, however, it treats the JPEG and Cell cases differently to get the maximum playback speed for each case. The algorithm for the Cell case is summarized below:

1. Install an X colormap.
2. Create an XIL image to hold images as they are decompressed.
3. Decompress frames from the movie and display them.

Each of these steps is considered in more detail in the following sections.

Installing an X Colormap

As you look at this part of the example program, keep in mind that the Cell playback code may update the X colormap many times as it plays back a movie. In fact, the decompressor may use a different colormap for each frame of a movie. This need to modify the X colormap makes colormap handling a bit more complex than it is in the JPEG case.

The first thing the example does in this section is to call the function `create_cmap()`, which is defined in the file `xilcis_color.c`:

```
XilLookup xil_cmap;
Colormap x_cmap;
XilIndexList *ilist;

xil_cmap = create_cmap(state, cis, display, window,
  DefaultScreen(display), &x_cmap, CELL, ilist, NULL, NULL);
```

As you'll see below, this routine returns an XIL lookup table that serves as the colormap for the Cell decompressor. The argument `&x_cmap` is the address of an X `Colormap` and `ilist` is a structure that is defined (in an XIL header file) as follows:

```
typedef struct
{
    Xil_unsigned32 *pixels;
    Xil_unsigned16 ncolors;
} XilIndexList;
```

The array `pixels` can hold up to `ncolors` colormap index values, and these values will later determine which X colormap entries the Cell decompressor can modify.

The first thing the `create_cmap()` function does is to look at the CIS's `DECOMPRESSOR_MAX_CMAP_SIZE` attribute to determine how many colorcells to allocate in the X colormap.

```
int cmapsize;

xil_cis_get_attribute(cis, “DECOMPRESSOR_MAX_CMAP_SIZE”,
    (void **)&cmapsize);
```

The function then creates an X colormap and allocates `cmapsize` colorcells in that colormap.

The next section of code sets up the `ilist` structure that will be used to determine which X colormap entries the Cell decompressor can modify. In this example, `ilist` is being set up so that the decompressor can modify any of the colorcells allocated by `create_cmap()`.

```
ilist->pixels = (Xil_unsigned32 *)malloc(sizeof(Xil_unsigned32)
    * cmapsize);
ilist->ncolors = cmapsize;
for (i = 0; i < cmapsize; i++)
    ilist->pixels[i] = (Xil_unsigned32)pixels[i];
```

Next, the `cmap_create()` function creates the XIL colormap that it will pass back to `main`. Because the XIL routine that creates this colormap requires as one of its arguments an array of bytes containing the color values to be stored

in the colormap, cmap_create() first reads the color values in the default X colormap into a buffer called data. It then uses xil_lookup_create() to create the XIL colormap.

```
data = (Xil_unsigned8 *)malloc(sizeof(Xil_unsigned8) * cmapsize
    * 3);
for (i = 0; i < cmapsize; i++)
    cdefs[i].pixel = i + pixels[0];
XQueryColors(display, DefaultColormap(display, screen), cdefs,
    cmapsize);
for (i = 0, j = 0; i < cmapsize; i++, j += 3) {
    data[j] = cdefs[i].blue >> 8;
    data[j + 1] = cdefs[i].green >> 8;
    data[j + 2] = cdefs[i].red >> 8;
}
lut = xil_lookup_create(state, XIL_BYTE, XIL_BYTE, 3, cmapsize,
    (int)pixels[0], data);
```

The lookup table created contains three bands on the output side and cmapsize entries. Also, note that color values are entered in the table in BGR order.

After create_cmap() returns the XIL lookup table to main, the example gets the version number of this colormap (xil_cmap):

```
XilVersionNumber lu_version;

lu_version = xil_lookup_get_version(xil_cmap);
```

This version number, lu_version, changes whenever the lookup table's contents are changed by the Cell decompressor. Therefore, after decompressing each movie frame, the program can test the version number to determine whether it needs to update the X colormap.

The example then sets two important attributes of the Cell decompressor: DECOMPRESSOR_COLORMAP and RDWR_INDICES.

```
xil_cis_set_attribute(cis, "DECOMPRESSOR_COLORMAP", xil_cmap);
xil_cis_set_attribute(cis, "RDWR_INDICES", ilist);
```

The DECOMPRESSOR_COLORMAP attribute specifies the XIL lookup table that will serve as the Cell decompressor's XIL colormap. Because colormap information is stored in the Cell bytestream along with pixel values, any time a frame is decompressed, the contents of the decompressor's colormap can change. As you'll see later, when the contents of the colormap change, a subset of its entries are written to the X colormap before the next video frame is displayed.

By default, however, the decompressor's colormap is read only. To make an entry writable, you register its index in RDWR_INDICES (in this case, ilist->pixels). Only the entries whose indexes are stored in this array can change. For example, in this program, only the entries whose indexes correspond to those of the X colorcells allocated earlier can be written to.

Note – If you want to play Cell movies using a fixed colormap (for example, to avoid colormap flashing between windows), don't set the RDWR_INDICES attribute. In this case, the Cell compressor will do its best using the colors stored in the lookup table xil_cmap.

Creating an Image to Hold Decompressed Frames

The program now creates a 3-band XIL_BYTE image to hold frames as they are decompressed. This task is handled with a call to xil_create():

```
XilImage image24 = NULL;

image24  = xil_create(state, cis_xsize, cis_ysize, 3, XIL_BYTE);
```

Later in the program, these images will be converted to 8-bit images for display.

Playing the Movie

Like the JPEG player, the Cell player plays a movie by executing a few statements in a loop that terminates when the function call `xil_cis_has_frame(cis)` returns 0.

First, the example decompresses a frame of data using the function `xil_decompress()` and checks the version number of the decompressor's XIL colormap:

```
xil_decompress(cis, image24);
if (lu_version != xil_lookup_get_version(xil_cmap))
{
    cell_install_cmap(x_cmap, displayimage, xil_cmap, ilist);
    lu_version = xil_lookup_get_version(xil_cmap);
}
```

The version number of the original XIL colormap was stored in `lu_version`. If this colormap has changed, then `xil_cmap`'s version number will have changed, and the program calls `cell_install_cmap()` to update the X colormap. The program then saves `xil_cmap`'s new version number in `lu_version` so that it can check again to see if the colormap has changed after the next call to `xil_decompress()`.

Finally, the example converts the RGB image stored in `image24` to an 8-bit image and copies the 8-bit image to `displayimage`. This is done using the function `xil_nearest_color()`.

```
xil_nearest_color(image24, displayimage, xil_cmap);
```

For each set of RGB values in `image24`, this routine performs a pixel-by-pixel search for the nearest matching color in the supplied colormap (`xil_cmap`) and sets the destination image pixel value to the appropriate colormap index.

Note – When a Cell movie is displayed on an 8-bit frame buffer, as in this example, the exact sequence of calls you use to play the movie has a dramatic effect on the speed with which the movie is shown. The reason for this is that by using the deferred-execution scheme explained in Chapter 21, “Acceleration in XIL Programs,” the XIL library can look for a certain sequence of functions and, if it finds that sequence, replace all the functions in the sequence with a single highly optimized routine. In this example the sequence `xil_decompress()`, `xil_nearest_color()` is such a sequence. Thus, when you play a Cell movie using this program, these functions are not executed, but are replaced by an optimized function (*molecule*) that performs the jobs of both functions. See the section “Cell Molecules” on page 288 for a complete list of Cell-related molecules.

Compressed Image Sequences 13

An XIL compressed image sequence (CIS) is a container for compressed images. These compressed images may represent a series of still images, frames from a movie, or pages from a document, and they may have been compressed in one of several formats:

- Cell or CellB
- JPEG baseline sequential or lossless
- H.261
- MPEG-1
- CCITT Group 3 or Group 4

For more information about these compression formats—what a particular type of compression is useful for, how a certain type of compressor is implemented, and so on—see chapters 13 through 19.

This chapter has three main sections. The first section discusses basic CIS operations, such as creating a CIS and writing data to it. The second explains what attributes CISs have and how you set and read the values of those attributes. The third talks about recovering from errors in a CIS's datastream.

Basic CIS Management

This section explains several basic CIS-related tasks:

- Creating and destroying a CIS
- Putting compressed data in a CIS
- Reading data from a CIS

If you have looked at the example programs presented in Chapter 12, "Compressing and Decompressing Sequences of Images," you already know something about performing these tasks. Don't skip the following sections, however, because they introduce a number of functions that were not used in the examples.

Creating and Destroying a CIS

You create a CIS by calling the function `xil_cis_create()`. This function takes two arguments: a handle to the system-state data structure that was created when you initialized the XIL library and a string that identifies the compressor/decompressor that will be used to write compressed data to, or read compressed data from, the CIS. For example, the call to create a CIS might look like this.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, "Jpeg");
```

The string you pass to the function—`Jpeg` in the code fragment above—can be any one of the following strings:

- `Cell` (a Cell codec)
- `CellB` (a CellB codec)
- `Jpeg` (a JPEG baseline sequential codec)
- `JpegLL` (a JPEG lossless codec)
- `H261` (an H.261 decompressor)
- `Mpeg1` (an MPEG-1 decompressor)
- `faxG3` (a CCITT Group 3 codec)
- `faxG4` (a CCITT Group 4 codec)

The return value of `xil_cis_create()` is a handle to the newly created CIS. You use this handle as an argument to all subsequent functions that affect the CIS.

A function related to `xil_cis_create()` is `xil_cis_reset()`. The latter function takes an existing CIS, which may contain compressed data, and returns it to its initial state. That is, it clears out all existing compressed data

and frees any memory the CIS has allocated to hold that data. The only argument to this function is a handle to the CIS whose compressed data you want to clear.

To destroy a CIS, you call the function `xil_cis_destroy()`. This function deallocates any buffers allocated to hold compressed data and frees the memory used to hold other information about the CIS. The function's only argument is a handle to the CIS.

Putting Compressed Data into a CIS

There are two primary ways to get compressed data into a CIS. One way is to call the function `xil_compress()` to compress one or more XIL images and write the compressed data to the CIS. The second way is to move an existing compressed data stream from system memory to a CIS. This compressed data may have been compressed originally by an XIL or a non-XIL compressor.

Both of these methods change the value of an internal CIS index so that the next operation that writes compressed data to the CIS will append its data to existing data in the CIS. For more information about this index, see the section "Start Frame, Read Frame, Write Frame" on page 260.

Using `xil_compress()`

The function `xil_compress()` reads an XIL image (or possibly more than one image), compresses it, and writes the compressed data to a CIS. The function takes two arguments. The first argument is a handle to a source image, and the second is a handle to a CIS. You do not have to specify a compressor or codec because each CIS is associated with a single compressor/decompressor when the CIS is created.

If the image handle passed to `xil_compress()` is a handle to a memory image or a display image, the function compresses a single image at a time. If the handle is a handle to a device image such as a video-capture card, `xil_compress()` may compress one or a number of frames per invocation.

Note – This description of what `xil_compress()` does is somewhat oversimplified. Some compressors may prefer to compress groups of frames. Therefore, they read an image each time your application calls `xil_compress()`, but they may not actually compress the images and write

data to the CIS until they have read *n* frames. If you need to make sure that a compressor has compressed all the frames it has read so far, call the function `xil_cis_flush()`.

Each compressor has certain requirements regarding the type of XIL image you can compress using `xil_compress()`. These requirements are summarized in Table 13-1. An N/A in a table box indicates that there is no requirement in a particular area.

Table 13-1 Types of Images Supported by XIL Compressors

	Width and Height	**Number of Bands**	**Data Type**	**Color Space**
Cell	Multiple of 4	3	`XIL_BYTE`	`rgb709`, `ycc601`, `ycc709`
CellB	Multiple of 4	3	`XIL_BYTE`	`ycc601`
JPEG baseline sequential	N/A	1 to 255	`XIL_BYTE`	N/A
JPEG lossless	N/A	1 to 255	`XIL_BYTE` `XIL_SHORT`	N/A
H.261	176 x 144 (QCIF) or 352 x 288 (CIF)	3	`XIL_BYTE`	`ycc601`
MPEG-1	N/A	3	`XIL_BYTE`	`ycc601`
CCITT Group 3	N/A	1	`XIL_BIT`	N/A
CCITT Group 4	N/A	1	`XIL_BIT`	N/A

You can also obtain information about the type of XIL image that can be written to a particular CIS using the function `xil_cis_get_input_type()`. See “Input and Output Image Type” on page 259 for more information about this function.

Putting a Compressed Datastream into a CIS

If you have an existing stream of compressed data in system memory, you can put that data into a CIS buffer in one of two ways. In both cases, the compressed data must have been compressed with a compressor of the same type as the compressor/decompressor associated with the CIS.

First, you can copy the datastream into the CIS buffer using the function `xil_cis_put_bits()`, whose function prototype is shown below.

```
void xil_cis_put_bits(XilCis cis, int nbytes, int nframes,
    void *data);
```

This routine copies `nbytes` of compressed data representing `nframes` of uncompressed data into the CIS `cis`. The argument `data` is a generic pointer to the data to be copied into the CIS.

The argument `nframes` normally indicates the number of frames represented by `nbytes` of compressed data. If you do not know the exact number of frames in the datastream, but know that there are no partial frames (frames that are not completely represented in the datastream), you should set `nframes` to -1. This value informs the CIS that the data being placed into it contains one or more complete frames and no partial frames. If you know that the first or last frame is not represented completely by the data in the stream, or you're not sure whether the datastream contains a partial frame, you should set `nframes` to 0. This value alerts the compressor that the datastream may contain a partial frame.

Note – Don't set `nframes` to 0 if you know that the datastream contains only complete frames. Doing so will slow down the copy.

You can also put a compressed datastream into a CIS using the function `xil_cis_put_bits_ptr()`. Its function prototype is shown below:

```
void xil_cis_put_bits_ptr(XilCis cis, int nbytes, int nframes,
    void *data, XIL_FUNCPTR_DONE_WITH_DATA done_with_data);
```

This function does not copy data into the CIS buffer, but gives the CIS a pointer to the datastream. If you use this function, your application is responsible for ensuring that the datastream remains valid.

You must also free the memory where the datastream is stored. One way to do this is to pass `xil_cis_put_bits_ptr()` a pointer to a function (***done_with_data***) that frees the memory. The `xil_cis_put_bits_ptr()` routine will call this function if the CIS is destroyed, is reset, or no longer needs the data. You should define this ***done_with_data*** routine to return `void` and to take as its only argument a generic pointer to `data`. Thus, the prototype for the function should look like this:

```
void free_memory(void *data)
```

When you call `xil_cis_put_bits_ptr()`, put the function's name in the spot where ***done_with_data*** is shown in the function prototype.

```
xil_cis_put_bits_ptr(cis, nbytes, nframes, data, free_memory);
```

If you don't pass `xil_cis_put_bits_ptr()` a pointer to a function, pass `NULL` as the final argument. In this case, your program must determine when the memory holding the datastream is no longer needed.

The value of `nframes` has the same meanings when you're using `xil_cis_put_bits_ptr()` as it does when you're using `xil_cis_put_bits()`.

Reading Data from a CIS

As with putting data into a CIS, there are two basic ways to get data out of a CIS. First, you can use the function `xil_decompress()` to read one frame's worth of compressed data from the CIS, decompress the data, and write the decompressed data to an XIL image. Second, you can get a pointer to the compressed data in the CIS and write the data to a file using a function like `fwrite()`.

Both of these methods change the value of an internal CIS index so that the next operation that reads compressed data from the CIS will begin reading with the first frame in the CIS following the last frame read during the just-completed operation. For more information about this index, see the section "Start Frame, Read Frame, Write Frame" on page 260.

Using `xil_decompress()`

The function prototype for `xil_decompress()` is shown below.

```
void xil_decompress(XilCis cis, XilImage dst);
```

The first argument to the function is a handle to your CIS, and the second is a handle to an XIL image into which a frame of video can be decompressed. The decompressor used for the operation will be the one you specified when you created your CIS.

The XIL image into which you decompress a frame of data must have the same width, height, number of bands, and data type as the images stored in the CIS. If you do not know these attributes of the images in the CIS, you can determine them using the function `xil_cis_get_output_type()`. See the section "Input and Output Image Type" on page 259 for more information on this function.

Copying Compressed Data from a CIS

If you want to copy data from a CIS without decompressing it, you should follow this procedure:

1. Use the function `xil_cis_get_bits_ptr()` to get a generic pointer to the compressed data in the CIS.
2. Use a function like `fwrite()` to copy the data from the CIS buffer to a file.

The function prototype for `xil_cis_get_bits_ptr()` is shown below.

```
void *xil_cis_get_bits_ptr(XilCis cis, int *nbytes,
    int *nframes);
```

The first argument to this function is a handle to the CIS. The second argument is the address of a variable in which the function stores the number of bytes to which its return value (a pointer to `void`) points. The third argument is the address of a variable in which the function stores the number of frames represented by `nbtyes`.

Note that nbytes is frequently not the total number of bytes of compressed data stored in the CIS. This is true because the CIS's "buffer" may actually be a list of buffers. As a result, you may have to call xil_cis_get_bits_ptr() multiple times to read all the compressed data in a CIS. See the next section for information about controlling loops in which you're reading data from a CIS.

Reading Data from a CIS in a Loop

If you want to decompress all the frames in a CIS or copy all the compressed data in a CIS to a file, you need to decompress data or copy data repeatedly, until all the data in the CIS has been read. A good way to control the loop that decompresses or copies data is to use the function xil_cis_has_frame(), as shown below.

```
XilCis cis;
XilImage dst;

while (xil_cis_has_frame(cis) == TRUE)
    xil_decompress(cis, dst);
```

The function xil_cis_has_frame() returns TRUE as long as one complete frame of compressed data remains in the CIS. Its single argument is a handle to your CIS.

Note – A loop like the one shown above works because a call to xil_decompress() or xil_cis_get_bits_ptr() changes an attribute of the CIS called its *read frame*. (For complete information about this attribute, see the section "Start Frame, Read Frame, Write Frame" on page 260.) If you decompress a frame, the read frame is incremented by 1, and if you copy 50 frames, the read frame is increased by 50.

You can also control this type of loop using the functions xil_cis_has_data() and xil_cis_number_of_frames(). The first of these functions returns the number of unread bytes in the CIS, and the second returns the number of unread frames. Neither function, however, is as effective as xil_cis_has_frame() for normal loop control: xil_cis_has_data() can return a nonzero value even after the last complete frame has been read (if the CIS contains a partial frame), and xil_cis_number_of_frames() may take longer to execute than xil_cis_has_frame().

Nonsequential Reads

Normally, when you read data from a CIS, you begin reading at the last frame read plus 1. However, you can also begin reading at a different point. To do this, you call the routine `xil_cis_seek()` to indicate the frame you would like to read next. The function prototype for this routine is shown below.

```
void xil_cis_seek(XilCis cis, int framenumber, int relative_to);
```

The first argument to this function is a handle to your CIS. The second and third arguments, taken together, determine which frame will become the CIS's read frame (the frame that will be read first the next time you read data from the CIS). The second argument, `framenumber`, is a number of frames and is to be construed as an offset from one of three locations. This location is specified by the value of `relative_to`. A value of 0 indicates that `framenumber` is an offset from the beginning of the CIS (frame 0), a value of 1 that `framenumber` is an offset from the current read frame, and a value of 2 that `framenumber` is an offset from the last frame in the CIS.

It is an error to seek for a frame prior to the first frame currently in the CIS or for a frame beyond the end of the CIS. See the section "Start Frame, Read Frame, Write Frame" on page 260 for information about determining the frame numbers of the first and last frames in the CIS. It is also an error to seek a frame prior to the current read frame if the CIS's random-access attribute is set to 0. See the section "Random Access Flag" on page 260 for a discussion of this attribute.

General CIS Attributes

Up to this point, this chapter has presented the CIS as a buffer or database in which you store compressed data. However, you can also think of a CIS as a structure with members that contain information about the compressor/decompressor associated with the CIS, the type of the images in a CIS, and so on. These members are referred to as *attributes* of the CIS. The following list shows you the general CIS attributes:

Note – The term *general attribute* is used here because CISs also have codec-specific attributes. These are covered in chapters 13 to 19, where the different XIL compressors/decompressors are discussed.

- Compressor (or decompressor) associated with the CIS
- Type of the compressor
- Type of image that can be compressed and stored in the CIS
- Type of image that will result when compressed data in the CIS is decompressed
- Flag indicating whether the data in the CIS can be accessed randomly
- Index to the first image currently in the CIS
- Index to the image that will be read next
- Index to the image that will be written next
- Maximum number of images the CIS can contain
- Number of already-read images the CIS should try to keep in its buffer
- Flag indicating whether the compressor should automatically recover from an error it knows how to handle
- Name of the CIS

These attributes are discussed in detail in the following sections.

Compressor and Compression Type

A CIS's compressor attribute is a string that identifies the compressor/decompressor that will be used to compress or decompress data for the CIS. This attribute is set when you create the CIS. For example, if you create your CIS using the statement

```
cis = xil_cis_create(state, "Cell");
```

the attribute will be set to `Cell`. You cannot change the value of this attribute after creating the CIS; however, you can read its value using the function `xil_cis_get_compressor()`.

The compression-type attribute is a string that identifies a compressor's class. In this release of the XIL library, each class contains only one compressor/decompressor. However, as Sun and third parties write new XIL compressors/decompressors, this situation will change. For instance, if a third party writes a JPEG baseline sequential compressor, that compressor will have a unique compressor name, but a compression type of `JPEG`, which is the compression type of the XIL JPEG baseline sequential codec. Like the

compressor attribute, this attribute is set when you create your CIS and cannot be changed afterwards. You can retrieve the value, however, using the function `xil_cis_get_compression_type()`.

Table 13-2 shows the currently available compressors/decompressors and their types.

Table 13-2 Compressors and Compressor Types

Compressor	**Compression Type**
`Jpeg`	`JPEG`
`JpegLL`	`JPEGLL`
`Cell`	`CELL`
`CellB`	`CELLB`
`H261`	`H261`
`Mpeg1`	`MPEG1`
`faxG3`	`FAXG3`
`faxG4`	`FAXG4`

Input and Output Image Type

A CIS's input-image-type attribute is a data structure of type `XilImageType` that defines the type of XIL image that may be compressed and written to the CIS. This structure contains information about an image's width, height, number of bands, and data type.

This information is available as soon as the CIS is created and may be retrieved with the function `xil_cis_get_input_type()`. For example, assume you use the following statement to create a CIS:

```
cis = xil_cis_create(state, "Cell");
```

This statement creates a CIS that will be written to by the Cell compressor. If you then call `xil_cis_get_input_type()` to determine the type of image that can be compressed and written to this CIS, you will find that the image must have three bands and contain `XIL_BYTE` data. The Cell compressor can only work with images that have these characteristics. At this point, the values of the image type's width and height members will be 0, which indicates that there are currently no specific requirements for these characteristics.

After you have used `xil_compress()` to write compressed data to your CIS, additional information about the input image type will become available. For instance, a call to `xil_cis_get_input_type()` will return nonzero values for the image type's width and height. These will correspond to the width and height of the images that have already been compressed.

A CIS's output-data-type attribute is also a structure of type `XilImageType`, but indicates the type of image that will be produced when data stored in the CIS is decompressed. You get this image-type structure by calling `xil_cis_get_output_type()`.

It's an error to call this function before writing data to your CIS. However, once the CIS contains data, `xil_cis_get_output_type()` will return information about the width, height, number of bands, and data type of the images stored in the CIS. In general, programs that decompress data must call this function because `xil_decompress()` expects its second argument to be an XIL image to which it can write its output. Before you can create this XIL image, you must know the output image type.

Random Access Flag

You can always seek forward in a CIS using the function `xil_cis_seek()`. However, only certain compressors/decompressors allow you to seek backwards. To determine whether your compressor/decompressor allows this type of seek, call the function `xil_cis_get_random_access()` to determine the value of the CIS's random-access attribute. A return value of 1 indicates that backward seeks are supported, and a return value of 0 indicates that such seeks are not allowed.

Start Frame, Read Frame, Write Frame

These attributes are integer indexes to certain important images or frames stored in a CIS.

The start-frame attribute is an index to the first frame in the CIS that is currently accessible. For now, just think of the start frame as the first frame written to the CIS (frame 0). To obtain the number of the current start frame, you call the function `xil_cis_get_start_frame()`.

The read-frame attribute is an index to the frame that will be read the next time you read data from the CIS. If you use `xil_decompress()` to decompress a frame, the function decompresses the current read frame and then increments the read frame by 1. If you get a pointer to data you want to read using the function `xil_cis_get_bits_ptr()`, the read frame is incremented by the number of frames that you read from the CIS. You can determine the current read frame using the function `xil_cis_get_read_frame()`.

The write-frame attribute is an index to the frame that will be written to the next time you put data into the CIS. (Thus, the index to the last frame currently in the CIS is the write-frame index minus 1.) When you add data to the CIS using `xil_compress()`, `xil_cis_put_bits()`, or `xil_cis_put_bits_ptr()`, this write-frame index is incremented by the number of frames added to the CIS. To read the current value of the write-frame attribute, use the function `xil_cis_get_write_frame()`.

The diagrams below may shed further light on how these indexes relate to one another. In each diagram, each row of rectangles represents a CIS buffer, and each rectangle in the row represents a frame of compressed data. Start frames, read frames, and write frames are labeled S, R, and W respectively.

The first diagram below shows a CIS buffer before and after a call to `xil_compress()`.

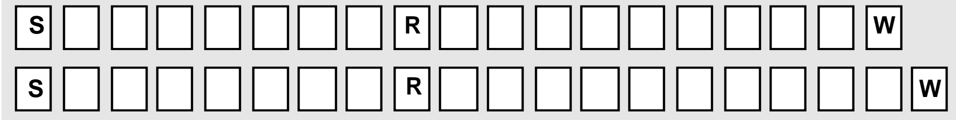

The diagram below illustrates the buffer before and after a call to `xil_decompress()`.

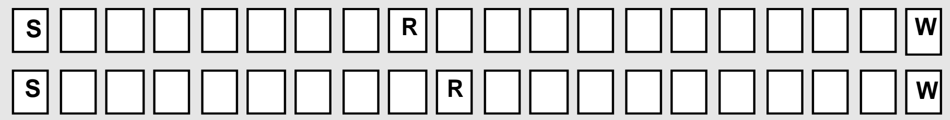

This illustration shows the buffer before and after a call to `xil_cis_get_bits_ptr()` that returns a pointer to the data from the read frame to the write frame minus 1 inclusive. After obtaining this pointer, you can use a function like `fwrite()` to write the data to a file.

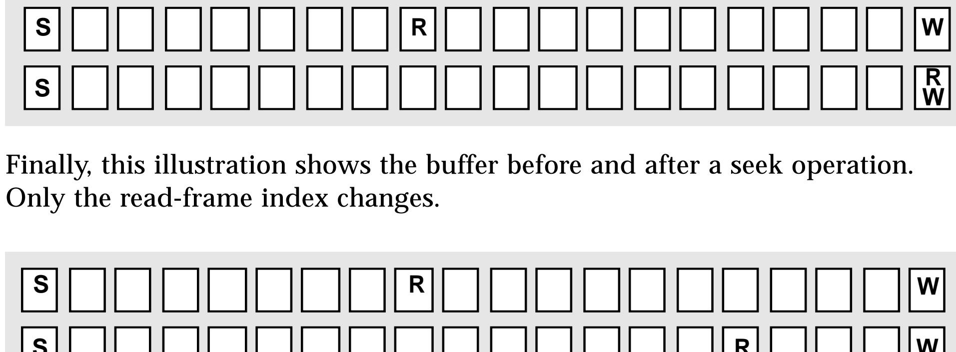

Finally, this illustration shows the buffer before and after a seek operation. Only the read-frame index changes.

Maximum Frames and Keep Frames

The maximum-number-of-frames attribute specifies the maximum number of frames a CIS should buffer at any one time. You set this attribute using the function `xil_cis_set_max_frames()`, and you can read the value of the attribute using `xil_cis_get_max_frames()`.

Each CIS has a maximum-frames attribute whether you set it or not. If you don't set the value of the attribute, the value is a default value that depends on the compressor/decompressor that is associated with the CIS. You can set the value of the attribute by passing an integer greater than 0 to `xil_cis_set_max_frames()`. You can also pass a -1 to this function: this value indicates that there should be no limit on the number of frames that the CIS can buffer.

Note – In the case where you set the maximum-frames attribute to an integer greater than 0, this setting is actually a suggestion rather than a requirement. Some compressors/decompressors cannot function properly if a CIS's buffer is too small.

The keep-frames attribute specifies the number of frames *prior to the current read frame* that a CIS should try to keep in its buffer. As with the maximum-frames attribute, the keep-frames attribute will have a compressor-specific value if you do not set it. It you do set the attribute, you can pass to `xil_cis_set_keep_frames()` an integer representing the number of frames to buffer or a -1, which means that there is no limit on the number of keep frames. You read the value of the keep-frames attribute using the function `xil_cis_get_keep_frames()`.

Note – Like the maximum-frames attribute, the keep-frames attribute is only a hint to the compressor/decompressor. Some decompression algorithms will not work if certain already-decompressed frames, such as key frames, aren't available.

The diagrams below illustrate how the maximum-frames and keep-frames attributes affect the state of the CIS buffer.

The first diagram shows how compressing frames affects a buffer that already contains the maximum number of frames. The first line in the illustration shows the initial state of the buffer. The second line shows how the buffer looks after a call to `xil_compress()`, the third line how the buffer looks after a second call to `xil_compress()`, and so on. Note that an error occurs if, in adding new data to the buffer, you lose the current read frame. An error also occurs if R minus S becomes less than the value of the keep-frames attribute. This error is reported only once, after the operation that first causes this condition.

Note – An error does not occur if you cause R minus S to become less than the value of the keep-frames attribute by seeking backwards in the buffer.

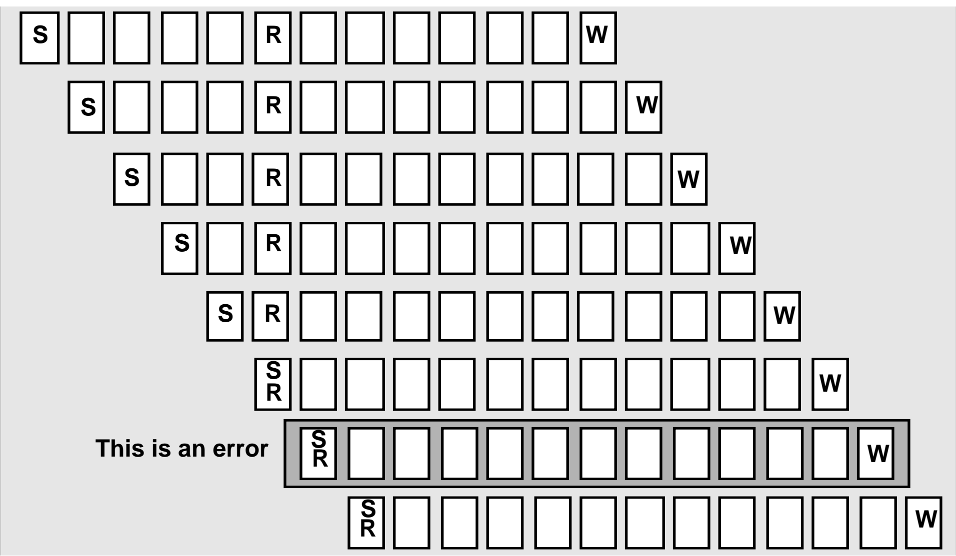

In the next two illustrations, R minus S equals the value of the keep-frames attribute in the initial view of the buffer. The first diagram below shows the effect of decompressing a frame from the CIS. Note that S is incremented to prevent R minus S from exceeding the value of the keep-frames attribute. Frame S - 1 is no longer accessible.

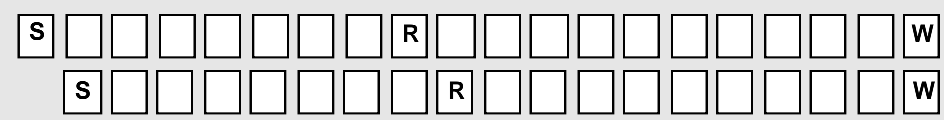

This picture shows the CIS buffer before and after an operation that reads all the compressed data between the read frame and the write frame minus 1 inclusive. Note that R minus S equals the value of the keep-frames attributes in both depictions of the buffer.

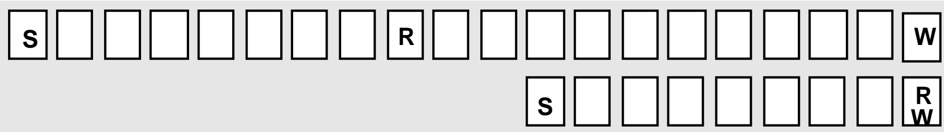

Error-Recovery Flag

There are two types of XIL datastream errors: those that a codec or decompressor knows how to recover from and those that it can't recover from without help from the programmer. Here's an example of the former type of error. Some datastreams contain end-of-line markers. If a decompressor, while it is decoding data, finds a line that is longer or shorter than expected, it can (but doesn't have to) recover from this error by truncating or zero-filling the line. The value of a CIS's error-recovery flag determines whether a codec or decompressor automatically recovers from errors that it knows how to handle or doesn't. For information on recovering from errors that the module doesn't know how to handle, see "CIS Error Recovery" on page 267".

By default, a CIS's error-recovery flag is set to `FALSE`. This setting indicates that a codec or decompressor should not automatically recover from the type of error mentioned above.

Note – The library will generate an error structure when any type of datastream error occurs and pass that structure to the currently installed error handler. The error handler can then deal with the error as it sees fit.

To set this flag to `TRUE` in order to enable automatic error recovery, you call the function `xil_cis_set_autorecover()`.

```
void xil_cis_set_autorecover(XilCis cis, Xil_boolean on_off);
```

The parameter cis is a handle to the CIS whose datastream is being accessed, and on_off should be the enumeration constant TRUE. (You would use the constant FALSE to turn automatic recovery back off.)

The XIL library also includes a function that determines whether the error-recovery flag is currently set. This function is called xil_cis_get_autorecover().

```
Xil_boolean xil_cis_get_autorecover(XilCis cis);
```

Note – If you set the error-recovery flag for a codec that cannot automatically recover from any datastream errors—like the Cell compressor—the flag has no effect.

Name

The library enables you to assign a name (char *) to a CIS. This type of naming is useful because it enables you to get a handle to a CIS later in your program by specifying the name of the CIS. The functions that allow for the naming of CISs are shown in Table 13-3.

Table 13-3 CIS Naming Functions

Function Name	What the Function Does
xil_cis_set_name	Sets the name of a CIS
xil_cis_get_name	Returns a copy of a CIS's name
xil_cis_get_by_name	Returns a handle to the CIS that has the name you specify

CIS Error Recovery

This section discusses how you recover from datastream errors that a codec or decompressor does not know how to handle. For information on handling errors that such a module can deal with, see the section “Error-Recovery Flag” on page 265.

When a datastream error occurs that a codec or decompressor cannot recover from, two things happen:

- The CIS is marked invalid. If the datastream error occurs while data is being decompressed, the CIS is marked invalid for further reading. If the error occurs while data is being compressed, the CIS is marked invalid for further writing. To determine whether a CIS is invalid, you can call either `xil_cis_get_read_invalid()` or `xil_cis_get_write_invalid()`.
- The XIL library generates an error structure and passes it to the currently installed error handler. If this is a user-installed error handler, it may take some action to deal with the datastream error.

Once a CIS has been marked invalid, there are three ways to make it valid again.

- Reset the CIS by calling `xil_cis_reset()`. This function clears all existing compressed data from the CIS.
- Seek a valid frame using `xil_cis_seek()`. If this seek fails, the library will generate a seek error.
- Ask the library to recover from the error by calling `xil_cis_attempt_recovery()`. The remainder of this section explains how to use this function.

You use the function `xil_cis_attempt_recovery()` to try to recover from a datastream error that (1) occurred while you were decompressing a CIS and (2) made the CIS invalid for further reading. The prototype for this function is shown below.

```
void xil_cis_attempt_recovery(XilCis cis, unsigned int nframes,
    unsigned int nbytes);
```

The parameter `cis` is a handle to the CIS you're working with. The parameter `nframes` indicates the maximum number of frames the function should scan in its attempt to recover from the error. The parameter `nbytes` is the maximum number of bytes the function should scan. If both `nframes` and `nbytes` are set to 0, the function can search as far forward as necessary in order to recover.

If one of the last two parameters is zero and the other is nonzero, the function behaves as follows. If `nframes` is nonzero and `nbytes` is zero, the error recovery mechanism attempts to search `nframes` frames ahead, using its best guess as to exactly how many bytes those frames would contain. If `nframes` is zero, and `nbytes` is nonzero, the recovery routine will scan `nbytes` bytes, regardless of how many frames those bytes represent.

When `xil_cis_attempt_recovery()` returns, you can check to see whether it was successful by calling the routine `xil_cis_get_read_invalid()`. If you asked the recovery routine to scan a relatively few frames or bytes, it may be necessary to call the routine several times to recover from the error.

The example below shows `xil_cis_attempt_recovery()` being used inside a user-defined error handler. (For a complete discussion of writing error handlers, see Chapter 8, "Error Handling.")

```
Xil_boolean cis_error_handler(XilError error)
{
    XilCis cis;
    XilObject object;

    if ((xil_error_get_category(error) == XIL_ERROR_CIS_DATA) &&
        ((object = xil_error_get_object(error)) != NULL) &&
        (xil_object_get_type(object) == XIL_CIS)) {

        cis = (XilCis)object;
        if (xil_cis_get_read_invalid(cis)) {
            xil_cis_attempt_recovery(cis, 0, 0);
            if (!xil_cis_get_read_invalid(cis))
                return TRUE;
        }
    }
    return xil_call_next_error_handler(error);
}
```

If the error affects a CIS and the CIS has been marked read invalid, the error handler uses `xil_cis_attempt_recovery()` to try to recover from the error. The recovery routine has permission to scan the entire datastream if necessary. The error handler returns `TRUE` if the recovery is successful. Otherwise, it calls the next error handler.

Note – When `xil_cis_attempt_recovery()` is called from within an error handler, the call to `xil_decompress()` that produced the error will fail even if the recovery is successful. The next call to `xil_decompress()` will succeed.

Cell Codec 14

The Cell image compression technology, which was developed by Sun, has been optimized for the rapid decompression and display of images on simple hardware. Therefore, Cell compression is able to achieve reasonable display quality on indexed-color frame buffers. The initial focus of the Cell technology is on Sun-to-Sun communications, where the benefits of fast decoding outweigh the benefits of standards. Some possible areas of application include media distributions on CD-ROM and multimedia mail applications.

Note – To get an idea of the quality of an image that has been compressed and decompressed using the Cell compressor, see Color Plate 4.

This chapter is divided into four sections. First, it explains how the Cell compressor works and the type of applications it was designed for. Second, it explains briefly how to create a Cell CIS. Third, it discusses CIS attributes that apply specifically to the Cell codec (as opposed to the general CIS attributes covered in the section “General CIS Attributes” on page 257). Fourth, the chapter discusses some accelerated playback routines (called *molecules*) that you may be able to take advantage of when you’re playing back movies.

14

How the Cell Codec Works

The Cell encoding process transforms individual video frames into a bytestream that is displayable with the Cell decompressor. Normally, the encoder works with RGB images; however, it can also handle XIL images whose color-space attribute is set to `ycc601` or `ycc709`. The decompressor always produces RGB images.

In the first step of the encoding process, video images are analyzed to produce an appropriate colormap to represent the frames to be encoded (unless the programmer has already specified one). This step allows the specification of the colormap size in order to leave colors unused. This strategy enhances cooperation with the window manager and other applications. Cell also provides for Adaptive Colormap Selection, in which a new colormap is generated when the current colormap becomes unsuitable. This colormap can be used in subsequent frames.

Given the images and the colormap, the second step is to encode the individual frames into a Cell bytestream. The basic coding scheme used in Cell encoding is based on an image coding method called Block Truncation Coding. A 4-by-4 region of pixels (a *cell*) from an image is represented by 2 colors and a 16-bit mask. See Figure 14-1.

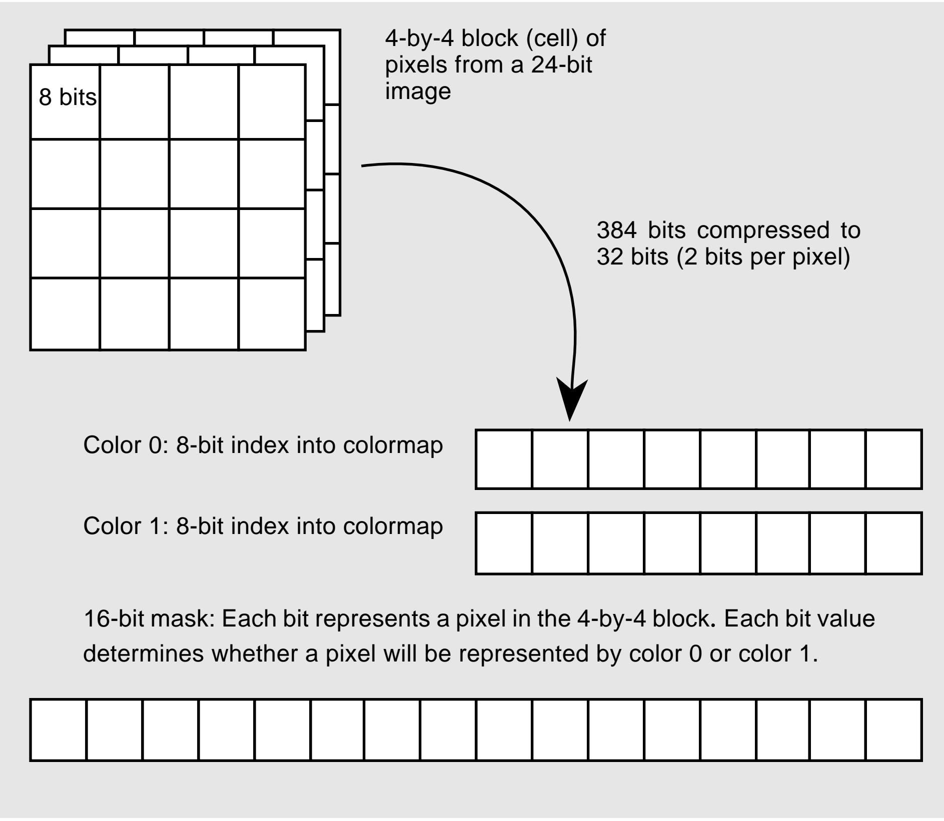

Figure 14-1 Cell Compression

The mask indicates which color to place at each of the pixel positions. The mask and colors may be chosen to maintain certain statistics of the cell, or they may be chosen to reduce contouring in a manner similar to ordered dithering.

The primary advantage of this coding method lies in the similarity of the decoding process to the operation of character fonting. The character display process for a color frame buffer takes as input a foreground color, a background color, and a mask that specifies which color to use at each pixel. Because this function is so important to the window system, it is often implemented as a display primitive in graphics accelerators. The Cell

compression technique uses this existing capability to provide full-motion video decoding with no special hardware or modifications to the window system.

There are actually two different encoding methods used to generate the encoded bytestream. One encoding method, known as Block Truncation Coding (BTC), chooses a 4-by-4 binary pixel mask and two colors (a foreground and a background color), while attempting to maintain the mean and variance of the luminance within the block. A second method, called the Dither technique, determines the pair of colors from the colormap that produce the least error when dithered over the 4-by-4 region. The BTC method is fastest and usually produces good results; the Dither method is slower, but is less likely to produce noticeable contours in regions of slow color variation.

Finally, encoded frames are combined into a frame sequence in the interframe compression step, where the compression ratio can be increased by anywhere from a factor of two to a factor of five. The colors and mask in each cell are compared to those used in the previous frames of the movie. If the colors and mask match to within a certain tolerance, a special skip code is generated. Runs of skip codes are combined to further reduce the bytestream. If either a single color changes or the mask changes, special escape codes are sent to update the changed data. Changes in the colormap are also detected by the interframe encoder, which causes special codes to be inserted into the stream to update the colormap.

Note – Details about the makeup of a Cell bytestream can be found in Appendix D, "Cell and CellB Bytestream Definitions."

Choosing a Colormap

The compressor chooses which colormap to use in encoding the current image in one of three ways. If Adaptive Colormap Selection (ACS) is enabled and a new colormap has not been associated with the compressor since the last call to `xil_compress()`, the compressor uses a colormap adapted to the current frame. When ACS is disabled, the compressor will use the colormap associated with the `COMPRESSOR_COLORMAP` attribute, if that attribute has been set. If ACS is disabled and the `COMPRESSOR_COLORMAP` attribute has not been set, the compressor calls `xil_choose_colormap()` to generate an optimal

colormap for the image. When this optimal colormap is created, it is associated with the COMPRESSOR_COLORMAP attribute and will be used in encoding subsequent frames.

Cell Compression Ratios

The basic Cell compression technique achieves a compression ratio of 2 bits per pixel. Additional encodings further reduce the data required per pixel. First, the compressor can use code words to specify a run of constant intensity cells. This type of coding is very effective in synthetic imagery, often leading to rates of less than 1 bit per pixel. The decompressor can also optimize the decompression of these constant colored cells by using rectangle-fill hardware. A second coding technique involves the introduction of frame-to-frame coherence (interframe coding), by instructing the decoder to skip over cells that are identical to those in the previous frame. These skip codes are implemented by simply updating the current writing position to skip over the number of cells specified.

The combination of the basic Cell encoding with run codes and skip codes can lead to compression ratios of about half a bit per pixel.

Image Types

The Cell compressor is designed to work with 3-band RGB and YC_bC_r images. The width and height of these images must be divisible by four. The Cell decompressor always produces RGB images.

Creating a Cell CIS

Before you can use the Cell codec to compress images or decompress a Cell bytestream, you must create a Cell CIS. You do this by passing the compressor name Cell to the function xil_cis_create(). See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, “CellB”);
```

Cell Codec Attributes

As discussed in the section "General CIS Attributes" on page 257, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for CISs attached to a Cell codec. You set compressor-specific attributes using the function `xil_cis_set_attribute()`, and you read compressor-specific attributes using the function `xil_cis_get_attribute()`.

Note – Some attributes can only be read, others can only be set, and others can be both read and set. The example, or examples, that conclude the discussion of each attribute indicate how the attribute can be used.

The Cell attributes can be broadly grouped into those that affect compression and those that affect decompression. The attributes are discussed under these headings.

Compression Attributes

Setting any of the following attributes affects how the Cell compressor compresses images.

BITS_PER_SECOND

The `BITS_PER_SECOND` attribute controls the size of the bytestream the Cell compressor produces. You specify the maximum number of bits the compressor can use to encode one second's worth of video, and the compressor guarantees that it will meet this bit-rate requirement on a frame-group basis (where a frame group is a key frame and all the ensuing frames up to the next key frame). That is, if every sixth frame is a key frame and your video was captured at 30 frames per second, the compressor will encode each frame group in a maximum of `BITS_PER_SECOND` divided by 5.

Note – You use the attribute `COMPRESSOR_FRAME_RATE` to indicate the rate at which your images were captured.

Setting BITS_PER_SECOND to 0, the default value, disables bit-rate control. In addition, if you set BITS_PER_SECOND to a rate lower than the compressor can achieve, an error is generated, and bit-rate control is disabled.

The code below shows the bit-rate being set to 1152000 bits per second. This is the rate necessary to encode 30 320-by-240 frames at half a bit per pixel.

```
XilCis cis;
int bit_rate;

bit_rate = 1152000;
xil_cis_set_attribute(cis, "BITS_PER_SECOND",
    (void *)bit_rate);
```

This code reads the value of the BITS_PER_SECOND attribute.

```
XilCis cis;
int bit_rate;

xil_cis_get_attribute(cis, "BITS_PER_SECOND",
    (void **)&bit_rate);
```

COLORMAP_ADAPTION

The value of the COLORMAP_ADAPTION attribute determines whether the Cell compressor's Adaptive Colormap Selection (ACS) feature is enabled. When enabled, ACS generates a colormap for the next image to be compressed by looking at the colors in the current image. Thus, each frame in a movie may have its own colormap.

The possible values for this attribute are TRUE and FALSE, which are values of type Xil_boolean. A setting of TRUE, the default value, enables ACS, and FALSE disables it.

The code below shows the COLORMAP_ADAPTION attribute being set to FALSE.

```
XilCis cis;

xil_cis_set_attribute(cis, "COLORMAP_ADAPTION", (void *)FALSE);
```

This code reads the value of the attribute.

```
Xil_boolean acs_enabled;

xil_cis_get_attribute(cis, "COLORMAP_ADAPTION",
    (void **)&acs_enabled);
```

COMPRESSOR_COLORMAP

The COMPRESSOR_COLORMAP attribute specifies the colormap the Cell compressor should use as it encodes images. You set this attribute by passing to xil_cis_set_attribute() an XIL lookup table (a data structure of type XilLookup) that contains 8-bit indexes on the input side and 24-bit RGB values on the output side. The default value of this attribute is NULL.

The code below shows the attribute being set.

```
XilCis cis;
XilLookup colormap;

xil_cis_set_attribute(cis, "COMPRESSOR_COLORMAP",
    (void *)colormap);
```

COMPRESSOR_FRAME_RATE

Use the COMPRESSOR_FRAME_RATE attribute to let the compressor know the rate at which the images to be compressed were captured. Express this rate in microseconds per frame (a microsecond is one millionth of a second). The default value is 33333, which indicates that the frames were captured at 30 frames per second.

The frame rate you supply, or the default frame rate, is encoded in the Cell bytestream. When you play your movie back, you can read the value of the DECOMPRESSOR_FRAME_RATE attribute to determine the rate at which frames should be decompressed.

The code fragment below shows the frame rate being set to 66666 microseconds per frame (15 frames per second).

```
XilCis cis;
Xil_unsigned32 microseconds;

microseconds = 66666;
xil_cis_set_attribute(cis, "COMPRESSOR_FRAME_RATE",
    (void *)microseconds);
```

COMPRESSOR_MAX_CMAP_SIZE

The `COMPRESSOR_MAX_CMAP_SIZE` attribute is an integer that defines the maximum number of entries in the colormap, or colormaps, encoded in the Cell bytestream. If Adaptive Colormap Selection is enabled, this attribute limits the size of the colormaps produced by the compressor. If ACS is disabled, the attribute limits the size of the colormap you can pass to the compressor using the `COMPRESSOR_COLORMAP` attribute. If you pass a colormap with more than the maximum number of entries, the colormap will be truncated.

When you first create a Cell CIS, the `COMPRESSOR_MAX_CMAP_SIZE` attribute is set to -1, which indicates that the attribute is settable at this point. To set the attribute to a value other than 256, you must set it prior to your application's first call to `xil_compress()`. If you have not set the attribute by that point, it defaults to 256. Also, note that the colormap size can only be set once during the life of a CIS.

Note – When decompressing a Cell bytestream, you can determine the maximum colormap size by reading the CIS attribute `DECOMPRESSOR_MAX_CMAP_SIZE`. Knowing this maximum will enable you to allocate the appropriate number of X colorcells to hold the colormap.

The following code shows the COMPRESSOR_MAX_CMAP_SIZE attribute being set to 240.

```
XilCis cis;
int cmap_size = 240;

xil_cis_set_attribute(cis, "COMPRESSOR_MAX_CMAP_SIZE",
    (void *)cmap_size);
```

This code reads the value of the attribute.

```
XilCis cis;
int cmap_size;

xil_cis_get_attribute(cis, "COMPRESSOR_MAX_CMAP_SIZE",
    (void **)&cmap_size);
```

COMPRESSOR_USER_DATA

This attribute writes *user data* to a Cell bytestream. This user data can be anything you choose as long as it does not exceed 8 Kbytes. It can be an XGL rendering or an executable. The attribute's main purpose, however, is to enable you to store audio data in the bytestream.

Before setting this attribute, you must assign values to the members of a structure of type XilCellUserData. The definition of this structure is shown below.

```
typedef struct {
    Xil_unsigned8 *data;
    Xil_unsigned32 length;
} XilCellUserData;
```

The member data is a pointer to the user data, and the member length specifies the length of the data in bytes.

Once you've set up this structure, you can set the attribute using code similar to this fragment.

```
XilCis cis;
XilCellUserData user_data;

xil_cis_set_attribute(cis, "COMPRESSOR_USER_DATA",
    (void *)&user_data);
```

After you set the attribute, the next time your application calls `xil_compress()`, the Cell compressor:

1. Writes the user data to the bytestream just before the compressed image data.
2. Clears the setting of the attribute.

Thus, each time you set `COMPRESSOR_USER_DATA`, the compressor writes data to the bytestream only once.

To read user data from a Cell bytestream, you must read the attribute `DECOMPRESSOR_USER_DATA`.

ENCODING_TYPE

The `ENCODING_TYPE` attribute indicates whether a Cell compressor will use Block Truncation Coding or the Dither method. Block Truncation Coding chooses a 4-by-4 binary pixel mask and two colors (a foreground and a background color), while attempting to maintain the mean and variance of the luminance within the block. The Dither technique determines the pair of colors from the colormap that produce the least error when dithered over the 4-by-4 region. The BTC method is fastest, and usually produces good results; the Dither method is slower, but is less likely to produce noticeable contours in regions of slow color variation.

The value of the attribute can be either of the enumeration constants shown in the definition below.

```
typedef enum {
    BTC, DITHER
} XilCellEncodingType;
```

The default value is BTC.

The call below shows the ENCODING_TYPE attribute being set to DITHER.

```
XilCis cis;

xil_cis_set_attribute(cis, "ENCODING_TYPE", (void *)DITHER);
```

This code shows the attribute being read.

```
XilCis cis;
XilCellEncodingType encode_type;

xil_cis_get_attribute(cis, "ENCODING_TYPE",
    (void **)&encode_type);
```

KEYFRAME_INTERVAL

A key frame in a Cell bytestream is one that contains a bytestream information header and a colormap and uses no interframe escape codes. The KEYFRAME_INTERVAL attribute is an integer (type int) that specifies how frequently the compressor should encode key frames in the bytestream. That is, if the value of the attribute is 6, every sixth frame will be a key frame. The default value of the attribute is 6.

If you set the KEYFRAME_INTERVAL attribute to 0, no key frames are encoded in the resulting Cell bytestream. In this case, bit-rate control is disabled. (For further information about bit-rate control, see the section "BITS_PER_SECOND" on page 276.)

The code below shows KEYFRAME_INTERVAL being set to 10.

```
XilCis cis;
int key_frame = 10;

xil_cis_set_attribute(cis, "KEYFRAME_INTERVAL",
    (void *)key_frame);
```

This code reads the value of the KEYFRAME_INTERVAL attribute.

```
XilCis cis;
int key_frame;

xil_cis_get_attribute(cis, “KEYFRAME_INTERVAL”,
    (void **)&key_frame);
```

TEMPORAL_FILTERING

This attribute turns on or off a form of temporal filtering. When this filtering is turned on, the compressor does not encode a new value for a pixel in frame *n* if the value of that pixel is within a certain tolerance of the same pixel in frame *n* - 1. Having the filter on generally reduces noise in an image sequence and also significantly reduces the size of the bytestream the compressor produces. The primary reason to turn the filter off would be to eliminate ghosting artifacts.

By default, the attribute is set to TRUE.

The code below shows TEMPORAL_FILTERING being set to FALSE.

```
XilCis cis;

xil_cis_set_attribute(cis, “TEMPORAL_FILTERING”, (void *)FALSE);
```

This code reads the value of the TEMPORAL_FILTERING attribute.

```
XilCis cis;
Xil_boolean status;

xil_cis_get_attribute(cis, “TEMPORAL_FILTERING”,
    (void **)&status);
```

Decompression Attributes

Setting any of the following attributes affects how the Cell compressor decompresses images.

DECOMPRESSOR_COLORMAP

When playing a Cell movie by decompressing frames and then using `xil_nearest_color()` to convert the resulting RGB images to 8-bit images, you must set this attribute if you want your playback code to be accelerated. The attribute specifies the XIL lookup table that `xil_nearest_color()` should use when doing its 24-bit to 8-bit conversion. Of course, your application also needs to write the values stored in this lookup table to the X colormap your application is using.

As the Cell decompressor decompresses frames, it may change the values in the lookup table. (By default, the lookup table is read-only, but you can make it writable by setting the attribute `RDWR_INDICES`, which is discussed in the section "RDWR_INDICES" on page 287.) Therefore, your application may need to call the function `xil_lookup_get_version()` to check the version number of the lookup table after each call to `xil_decompress()`. A change in version number means that the values in the lookup table have changed. If the values have changed, your application must ensure that corresponding changes are made in its X colormap before displaying the most recently decompressed frame.

Note – The section "Playing Cell Movies" on page 241 discusses an example program that sets this attribute and checks the version number of the lookup table.

You can also read the value of the `DECOMPRESSOR_COLORMAP` attribute. If it has been set, you will get back a handle to an XIL lookup table (the decompressor's colormap). If it has not been set, `xil_cis_get_attribute()` will return `NULL`.

The code fragment below shows DECOMPRESSOR_COLORMAP being set.

```
XilCis cis;
XilLookup colormap;

xil_cis_set_attribute(cis, “DECOMPRESSOR_COLORMAP”,
    (void *)colormap);
```

This code reads the value of the attribute.

```
XilCis cis;
XilLookup colormap;

xil_cis_get_attribute(cis, “DECOMPRESSOR_COLORMAP”,
    (void **)&colormap);
```

DECOMPRESSOR_FRAME_RATE

If the COMPRESSOR_FRAME_RATE attribute was set when the Cell compressor encoded your movie, you can use the DECOMPRESSOR_FRAME_RATE attribute to determine the rate, in microseconds per frame, at which the frames in the movie were captured. The attribute may contain different values at different points in the Cell bytestream.

If the COMPRESSOR_FRAME_RATE attribute was not set when the Cell compressor encoded your movie, the DECOMPRESSOR_FRAME_RATE attribute will contain the default value 33333 (30 frames per second).

The code fragment below shows the DECOMPRESSOR_FRAME_RATE attribute being read.

```
XilCis cis;
Xil_unsigned32 frame_rate;

xil_cis_get_attribute(cis, “DECOMPRESSOR_FRAME_RATE”,
    (void **)&frame_rate);
```

DECOMPRESSOR_MAX_CMAP_SIZE

This is a read-only attribute that indicates the size of the colormap you should use when playing back your Cell movie. It's important to use the smallest colormap possible because your application must create an X colormap that is the same size as the decompressor's XIL colormap. If the X colormap is too large, colormap flashing may occur when it is installed.

Information about a movie's maximum colormap size is encoded in the Cell bytestream that represents the movie. Therefore, it an error to read this attribute of an empty CIS.

The example code below shows the `DECOMPRESSOR_MAX_CMAP_SIZE` attribute being read.

```
XilCis cis;
int cmapsize;

xil_cis_get_attribute(cis, "DECOMPRESSOR_MAX_CMAP_SIZE",
    (void **)&cmapsize);
```

DECOMPRESSOR_USER_DATA

This attribute holds a pointer to any user data that was encoded with the most recently decompressed image. To get this pointer, you might use code similar to this.

```
XilCis cis;
XilCellUserData user_data;

xil_cis_get_attribute(cis, "DECOMPRESSOR_USER_DATA",
    (void **)&user_data);
```

If user data was encoded with the last image that was decompressed, `user_data.data` will be a pointer to that data, and `user_data.length` will be the length of the data in bytes. If no user data was encoded with that image, `user_data.data` will be `NULL`, and `user_data.length` will be 0.

Any data pointer you get by using this attribute is only valid until the next call to `xil_decompress()`.

RDWR_INDICES

If your Cell movie-playback application uses the function xil_nearest_color() to convert 24-bit RGB images to 8-bit pseudocolor images, xil_nearest_color() uses the decompressor's colormap (XIL lookup table) to perform this conversion. For your application to display images using the best colors possible, the Cell decompressor must be able to alter the color values in this lookup table as the colors in the frames it is decompressing change. By default, however, the colormap is read only.

The RDWR_INDICES attribute enables you to make the colormap writable for a subset of its entries. Before you actually set the attribute, you must set up a structure of type XilIndexList. The definition of this structure is shown below.

```
typedef struct {
    Xil_unsigned32 *pixels;
    Xil_unsigned16 ncolors;
} XilIndexList;
```

You set the ncolors member of the structure to indicate how many of the elements in the colormap are being made writable. The array pixels lists the pixel values in the colormap whose RGB values the decompressor may change.

Note – If you allow the decompressor to alter the contents of its colormap, you must check the version number of the colormap after each call to xil_decompress(). You check this version number using the function xil_lookup_get_version(). If the version number has changed—that is, the contents of the colormap have changed—you must make corresponding changes in your application's X colormap before displaying the most recently decompressed frame.

Calls to RDWR_INDICES are not cumulative. Only the RGB values associated with the pixel values specified in your most recent setting of the attribute are writable.

The code fragment below shows the RDWR_INDICES attribute being set.

```
XilCis cis;
XilIndexList indexes;

xil_cis_set_attribute(cis, "RDWR_INDICES", (void *)&indexes);
```

To see how the structure indexes might be filled out, see the function xilcis_color.c in the directory $XILHOME/examples/movie_player_example.

Cell Molecules

The XIL library includes a series of *molecules* that accelerate the playback of Cell movies. These molecules are optimized routines that perform the jobs of two or more functions from the XIL API. You do not call such an optimized routine directly; rather, the library calls a molecule when your program calls a predefined sequence of XIL functions, sometimes with specific arguments. (These sequences of functions are discussed later in this section.)

For example, if your program calls xil_decompress() to decode an image stored in a Cell CIS and then calls xil_nearest_color() to convert the decoded image from a 24- to an 8-bit image, the library may not call these two functions. Instead, it may call a molecule that performs the decompression and the conversion in an optimized way.

Note – This replacement of two or more atomic functions by a molecule is made possible by the XIL library's deferred-execution scheme, which is discussed in Chapter 21, "Acceleration in XIL Programs." This chapter also contains a more detailed definition of molecules than the one presented here and explains how to determine whether a molecule you want to call is actually being executed.

The library's Cell-decompression molecules enable you to accelerate the playback of a Cell bytestream on

- A one-bit destination image
- A display image associated with a GX frame buffer (available only on local GX frame buffer screens)
- An 8-bit destination image other than a GX display image

The procedures you follow to call these molecules are documented later in this chapter, in the section "Calling Cell Molecules" on page 290. Before moving on to this subject, however, the chapter lists some general rules you must follow to execute any decompression molecule.

Rules for Calling Decompression Molecules

This section goes over the global rules for calling decompression molecules.

- As mentioned earlier, you must call a predefined sequence of XIL functions, sometimes with specific arguments. This sequence will be replaced by the molecule.

Note - The calls to the functions in the sequence do not necessarily have to be consecutive statements in your program. For instance, notice how the example shown in "Other 8-Bit Destination Images," which begins on page 294, calls one function in the sequence, `xil_decompress()`, and then immediately afterward calls an XIL function that is not part of the sequence and checks that function's return value. Depending on this return value, the example may then call a routine that loads an X colormap before it gets around to calling the second function in the sequence. The key here is that between calls to the functions in the sequence, the program cannot call any XIL functions whose execution the library will defer. For more information about the XIL library's deferred-execution scheme, read the section "What Is Deferred Execution?" on page 395.

- The images being decompressed and the molecule's destination image must have the same width and height, except when the molecule performs a scaling operation. In that case, the destination image must have the same dimensions as the scaled version of the source image.
- The destination image must not have a region of interest. This image will not have a region of interest unless you have explicitly set its region-of-interest attribute.
- The destination image must have an origin of 0.0, 0.0. This is the destination image's default origin.
- If a molecule uses an intermediate image, it too must have an origin of 0.0, 0.0 and no region of interest.

If any of these conditions is not met, the XIL functions in your playback code will be executed individually, and will not be replaced by a molecule. This will have a significant impact on the speed at which your movie is played.

Note – There are certain decompression molecules to which some of these global rules do not apply. These exceptions will be noted as they arise.

Calling Cell Molecules

Cell decompression molecules have been implemented for applications that play back movies using 1- and 8-bit destination images (usually display images). These molecules and the series of XIL functions you must call for these molecules to be executed are discussed below.

One-Bit Destination Image

One molecule has been defined for playing back Cell movies using a 1-bit destination image. This molecule is appropriate for playing movies on monochrome displays.

The molecule performs several tasks:

- Decompresses an image
- Converts the RGB image to a 1-band image by extracting the luminance of the RGB image
- Rescales the values in the 8-bit grayscale image so that they fall in the range `0` to `255`
- *Optionally* zooms the 8-bit image by a factor of `2` in both the *x* and *y* dimensions
- Performs an ordered dither on the 8-bit image to produce a 1-bit image

For this molecule to be called, your application must include code similar to that shown below.

```
XilColorspace rgblinear, ylinear;
float scale[1], offset[1];
XilLookup colorcube;
int mult[1] = {-1};
unsigned short dims[1] = 2;
XilDitherMask dmask;

rgblinear = xil_colorspace_get_by_name(state, "rgblinear");
ylinear = xil_colorspace_get_by_name(state, "ylinear");
xil_set_colorspace(imageRGB_24, rgblinear);
xil_set_colorspace(imageY_8, ylinear);
scale[0] = 255.0 / (235.0 - 16.0);
offset[0] = -16.0 * scale[0];

/* Basically, the colorcube to be used for the ordered dither
   must be defined as shown here. The molecule will still execute
   if mult[0] is 1 instead of -1; however, this change will cause
   images to be displayed in reverse video. */
colorcube = xil_colorcube_create(state, XIL_BIT, XIL_BYTE, 1, 0,
    mult, dims);

/* The dither mask to be used in the ordered dither must be an
   8-by-8 mask. You can change the values in the mask if you
   want to. */
dmask = xil_dithermask_get_by_name(state, "dm881");

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageRGB_24);
xil_color_convert(imageRGB_24, imageY_8);
xil_rescale(imageY_8, imageY_8, scale, offset);
if (ZOOM) {
    xil_scale(imageY_8, zoom_imageY_8, "nearest", 2.0, 2.0);
    xil_ordered_dither(zoom_imageY_8, zoom_image_1, colorcube,
    dmask);
}
else
    xil_ordered_dither(imageY_8, image_1, colorcube, dmask);
/* MOLECULE ENDS HERE */
```

Note – Arguments shown in boldface must be typed as shown for the molecule to execute correctly.

SPARC *Eight-Bit Display Image (GX)*

Another molecule has been defined specifically for playing back Cell movies on *local* GX frame-buffer screens. This molecule performs these tasks:

- Decompresses an image
- Converts the RGB image to an 8-bit image by finding the closest match for the RGB values in the source image in a lookup table
- *Optionally* zooms the 8-bit image by a factor of 2 in both the *x* and *y* dimensions

Note – Several of the molecule rules presented in the section "Rules for Calling Decompression Molecules" on page 289 do not apply to the version of this molecule that does not perform the zooming operation. If no zooming is requested, the destination image for the molecule does not need to be the same size as the images stored in the Cell CIS. In addition, the destination image can have a region of interest and an origin other than 0.0, 0.0.

For this molecule to be called, your application must use the code shown below.

```
XilLookup cmap;
XilIndexList indexlist;

/* You must set this attribute, or the molecule will not
   execute. cmap is the lookup table that will be passed
   to xil_nearest_color(). */
xil_cis_set_attribute(cis, "DECOMPRESSOR_COLORMAP",
    (void *)cmap);

/* To get the best results, you should set this attribute
   so that the decompressor can change any of the entries in
   cmap. */
xil_cis_set_attribute(cis, "RDWR_INDICES", (void *)&indexlist);

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageRGB_24);

/* This checking of the lookup table's version number is
   necessary, but is not part of the molecule */
if (lu_version != xil_lookup_get_version(cmap)) {
    /* Include code to write the appropriate values from cmap to
       the application's X colormap */
    ...
    lu_version = xil_lookup_get_version(cmap);
}

if (ZOOM) {
    xil_nearest_color(imageRGB_24, image_8, cmap);
    xil_scale(image_8, zoom_image_GX, "nearest", 2.0, 2.0);
}
else
    xil_nearest_color(imageRGB_24, image_GX, cmap);
/* MOLECULE ENDS HERE */
```

Other 8-Bit Destination Images

If you want to play back your Cell movie using an 8-bit destination image other than a local GX display image, you can use one of three molecules. The first two molecules use a nearest-color strategy to convert images from 24 bits to 8 bits, and the third molecule uses an ordered dither. All three molecules require that you set the CIS attribute DECOMPRESSOR_COLORMAP before you execute them.

Nearest-Color Approach

Functionally the first molecule in this class is identical to the basic molecule (no zoom) designed for playing movies on a local GX frame buffer. (See the section “Eight-Bit Display Image (GX)” on page 292.) However, it is not as fast as the GX version, so displaying a movie on an indexed-color frame buffer other than the GX will not be quite as fast as on a GX.

For this molecule to execute, your application must use the code shown in the section “Eight-Bit Display Image (GX).” In addition:

- The destination image must be stored on a word boundary. It will be stored on such a boundary unless it is a child image created with an *x* offset from its parent that is not a multiple of four.
- The pixel stride in the destination image must be one. This will always be true unless the destination image is a single-band child of a multiband parent.

The second molecule in this class rapidly decompresses Cell images and displays them in a small window. It performs the following tasks:

- Decompresses an image
- Converts the RGB image to an 8-bit image by finding the closest match for the RGB values in the source image in a lookup table
- Scales the image down using *x* and *y* scale factors of $1 / (4 * n)$, where *n* is a positive integer

The destination image must be at least 4-by-4 pixels in size.

For this molecule to be called, your application must use code similar to that shown below.

```
/* MOLECULE BEGINS HERE */
xil_decompress(cis, imageRGB_24);
if (lu_version != xil_lookup_get_version(cmap)) {
    /* Include code to write the appropriate values from cmap to
       the application's X colormap */
    ...
    lu_version = xil_lookup_get_version(cmap);
}
xil_nearest_color(imageRGB_24, image_8, cmap);
xil_scale(image_8, small_image_8, "nearest", .25, .25);
/* MOLECULE ENDS HERE */
```

Ordered-Dither Approach

This molecule performs the following tasks:

- Decompresses an image
- *Optionally* zooms the RGB image by a factor of 2 in the *x* and *y* dimensions
- Dithers the 3-band image to a 1-band image using a colorcube

For the molecule to execute, the destination image must be word aligned and must have a pixel stride of one.

To call this molecule, your application must use the code similar to that shown below.

```
/* Create a colorcube and dither mask. All legal XIL
   colorcubes are supported, including those that have
   decreasing ramps in one or more bands. The dither mask
   must be a 4-by-4 mask, but there are no restrictions
   regarding its contents. */
...

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageRGB_24);
if (ZOOM) {
    xil_scale(imageRGB_24, zoom_imageRGB_24, "nearest", 2.0,
        2.0);
   xil_ordered_dither(zoom_imageRGB_24, zoom_image_8, colorcube,
        dithermask);
}
else
    xil_ordered_dither(imageRGB_24, image_8, colorcube,
        dithermask);
/* MOLECULE ENDS HERE */
```

CellB Codec 15

The Cell codec discussed in the preceding chapter is useful primarily in authoring applications. Its advantages are its fast decoding and the high quality of the images it produces, particularly on indexed-color frame buffers. The CellB codec, which derives from its Cell counterpart, is intended for use primarily in videoconferencing applications. It features a greater balance between the time spent compressing and decompressing images than the Cell codec and employs a fixed colormap. The CellB codec's strengths include

- Software compression at interactive rates
- Very fast decoding and display, especially on indexed-color frame buffers
- Low rates of CPU use
- Good quality output

The remainder of this chapter is divided into four sections. First, the chapter explains how the CellB compressor/decompressor works. Second, it explains briefly how to create a CellB CIS. Third, it discusses the CIS attributes that apply specifically to the CellB codec (as opposed to the general CIS attributes covered in the section "General CIS Attributes" on page 257.) Fourth, the chapter introduces the subject of accelerating the playback of CellB bytestreams. For further information on this subject, see Chapter 21, "Acceleration in XIL Programs."

How the Codec Works

The CellB compressor works on YC_bC_r images that conform to the guidelines set forth in CCIR Recommendation 601. The compressor performs *intraframe* compression by representing 4-by-4 blocks of pixels using cell codes. It performs *interframe* encoding using skip codes.

Cell Codes

The images you compress using the CellB compressor must have a width and height that are multiples of four because the compressor works with 4-by-4 cells of pixels. For each frame, the compressor begins with the cell in the upper-left corner and then proceeds from left to right. The compressor processes rows of cells in this way, moving from the top of the frame to the bottom.

When the compressor encodes a cell without reference to a cell in a preceding frame, it uses a four-byte cell code to represent the content of that cell. This cell code specifies two colors and includes a 16-bit bit mask that indicates which of the two colors should be used to represent each pixel in the cell. See Figure 15-1.

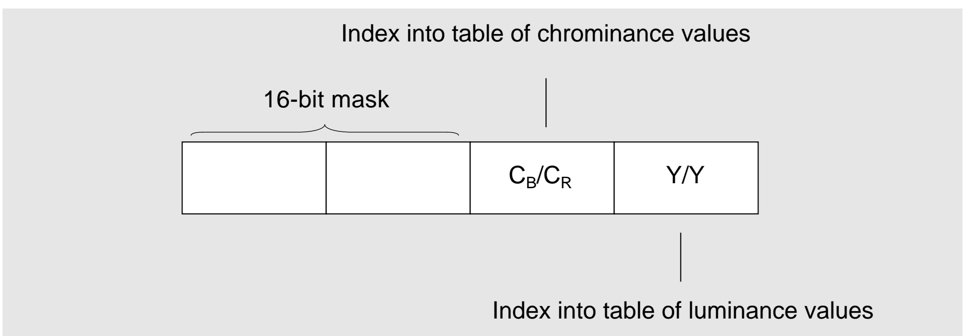

Figure 15-1 Cell Code

The two colors are encoded as follows. The compressor calculates the average C_b and C_r values for the 4-by-4 block. Then, it calculates an index into a table of 256 vectors in which each vector looks like the one shown in Figure 15-2.

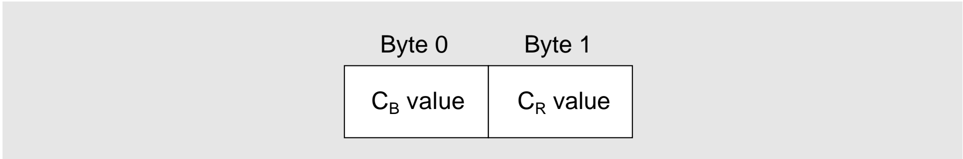

Figure 15-2 Vectors in Chrominance Table

The values in the vector pointed to by the index are the pair of values in the table nearest to the mean C_b and C_r values for the cell. The compressor writes the index to this vector to the third byte of the cell code.

The compressor also analyzes the luminance values in the cell. First, it calculates the mean luminance for the cell. Second, it partitions the 16 luminance values in the cell into those values that fall below the mean and those that fall above the mean. Then, it calculates the average luminance in the two partitions.

After arriving at the average luminance values for the two partitions, the compressor calculates an index into a table of vectors of the form shown in Figure 15-3.

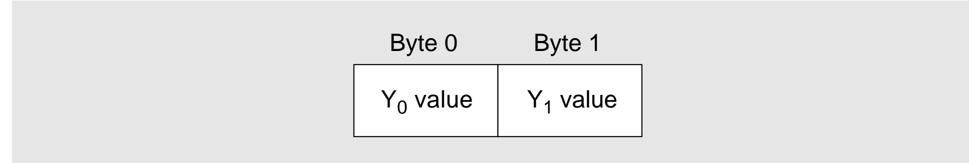

Figure 15-3 Vectors in Luminance Table

The values in the vector pointed to by the index are the pair of values in the table closest to the average luminance values for the two partitions. The compressor then writes the index to this vector to the fourth byte of the cell code.

The first color for a cell consists of the first byte of the cell's luminance vector and the chrominance values in the cell's chrominance vector. The second color consists of the second byte of the cell's luminance vector and the same chrominance values.

The bit mask shown in Figure 15-1 is filled out in this way. Each bit in the mask is associated with a pixel in the cell. If the luminance value for a pixel is below the mean luminance value for the cell, its bit is set to 0. This pixel will be represented by the first color when the cell is decompressed. If a pixel's luminance value is above the mean, its bit is set to 1.

Because each cell code represents the values of 16 pixels using 32 bits, compression using cell codes alone leads to a compression rate of 2 bits per pixel. To better this compression rate, the CellB compressor uses skip codes to achieve interframe compression.

Skip Codes

The CellB compressor encodes the first image in a sequence using cell codes exclusively. But after the first image, the compressor begins looking for cells in the current image that match—within a certain tolerance—the corresponding cells in the preceding image. Anywhere from 1 to 32 consecutive cells that match their counterparts in the previous image can be represented by a single 1-byte skip code.

The only restriction on the use of skip codes is that the cell at a particular set of coordinates can not be skipped indefinitely. There are a couple of reasons for this restriction. First, a user might join a videoconference that is already in progress. Cells represented by skip codes in the first image he receives will not be displayed correctly, and this problem must be corrected within a certain period of time. Similarly, if the CellB bytestream is being sent over an unreliable transport and a packet of data is lost, the period of the resulting error should be limited.

The policy concerning skip codes is that a particular cell must be updated at least every *n* frames, where *n* is an implementation-specific maximum. The exact value used is selected randomly for each cell each time that cell is encoded using a cell code. A random number is used to prevent the periodic bit-rate increase that might result were each cell to be updated at a fixed interval.

In a typical videoconference, about 80 percent of the cells in the average frame are represented by skip codes. This ratio leads to an average compression rate of about .8 bits per pixel.

Note – Details about the makeup of a CellB bytestream can be found in Appendix D, “Cell and CellB Bytestream Definitions.”

Creating a CellB CIS

Before you can use the CellB codec to compress images or decompress a CellB bytestream, you must create a CellB CIS. You do this by passing the compressor name `CellB` to the function `xil_cis_create()`. See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, “CellB”);
```

CellB Decompression Attributes

As discussed in the section “General CIS Attributes” on page 219, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for a CIS attached to a CellB codec. You set these codec-specific attributes using the function `xil_cis_set_attribute()` and read them using `xil_cis_get_attribute()`.

WIDTH and *HEIGHT*

If you have put compressed data into your CIS using `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()`, you *must* set the values of these attributes to the width and height in pixels of the images to be decompressed. If you do not set these attributes, their values will be 0, and an error will occur if you:

- Create an XIL image into which to decompress the compressed images in the CIS by calling the functions `xil_cis_get_output_type()` and `xil_create_from_type()`
- Decompress an image stored in the CIS by calling `xil_decompress()`

The legal values for both attributes are integers in the range 4 to 32764 (short int). The code fragment below shows the WIDTH attribute being set to 320 and the HEIGHT attribute being set to 240.

```
XilCis cis;
short width = 320;
short height = 240;

xil_cis_set_attribute(cis, "WIDTH", (void *)width);
xil_cis_set_attribute(cis, "HEIGHT", (void *)height);
```

IGNORE_HISTORY

The IGNORE_HISTORY attribute affects your ability to seek forward and backward in a CellB bytestream. These seeks are somewhat problematic in CellB because the codec relies so heavily on interframe encoding and does not require periodic key frames.

By default, IGNORE_HISTORY is set to FALSE. In this case, backward seeks are illegal because no past frame contains all the information necessary to reproduce an entire image. Forward seeks are possible, but to ensure that it can decode the frame you seek to properly, the decoder must actually decode all the frames you "skip."

If you set IGNORE_HISTORY to TRUE, you're telling the decoder that it's acceptable if, after a seek, it does not produce a correct image. You're willing to let it decode only the cells described with cell codes in the frame you seek to, and to fill in bad cells as it decodes subsequent frames. In this case, backward seeks are legal, and forward seeks are faster than they would be otherwise. The drawback to setting IGNORE_HISTORY to TRUE is that you may have to decode a number of frames after a seek before you get a properly reconstructed picture. (The exact number of frames is implementation dependent.)

The code below shows IGNORE_HISTORY being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "IGNORE_HISTORY", (void *)TRUE);
```

You can also read the value of this attribute using code similar to that shown below.

```
XilCis cis;
Xil_boolean history_status;

xil_cis_get_attribute(cis, "IGNORE_HISTORY",
    (void **)&history_status);
```

CellB Molecules

The XIL library includes a series of *molecules* that accelerate the playback of CellB bytestreams. These molecules are optimized routines that perform the jobs of two or more functions from the XIL API. You do not call such an optimized routine directly; rather, the library calls a molecule when your program calls a predefined sequence of XIL functions, sometimes with specific arguments.

For example, if your program calls `xil_decompress()` to decode an image stored in a CellB CIS and then calls `xil_ordered_dither()` to convert the decoded image from a 24- to an 8-bit image, the library may not call these two functions. Instead, it may call a molecule that performs the decompression and the dithering in an optimized way.

For information about the CellB molecules that are available and information about how to call those molecules, see the section "XIL Molecules" on page 398.

JPEG Baseline Sequential Codec 16

The JPEG baseline sequential coder-decoder is one of the digital-image codecs that has been specified by the Joint Photographic Experts Group, a joint ISO and CCITT technical committee. The JPEG lossless compressor discussed in Chapter 17, “JPEG Lossless Codec,” is another of these codecs. Taken together, the JPEG codecs are meant to provide a standard means of compressing still continuous-tone (grayscale and color) images.

The JPEG standard was originally developed for use in areas such as desktop publishing, graphic arts, medical imaging, and document imaging, where the archiving of still images is important. However, the introduction of high-performance hardware capable of coding and decoding JPEG images in real-time has enabled the development of full-motion video applications based on JPEG.

This chapter is divided into four sections. The first section explains how the compressor works and the type of applications it was designed for. The second explains briefly how to create a JPEG baseline sequential CIS. The third discusses CIS attributes that apply specifically to the baseline sequential codec (as opposed to the general CIS attributes covered in the section “General CIS Attributes” on page 257). The fourth introduces the subject of accelerating the playback of JPEG bitstreams. For further information on this subject, see Chapter 21, “Acceleration in XIL Programs.”

How the JPEG Baseline Sequential Codec Works

The JPEG baseline sequential compressor is one of the DCT-based compressors, which also include the MPEG-1 and H.261 compressors. Figure 16-1 shows the basic steps the compressor uses to compress an image.

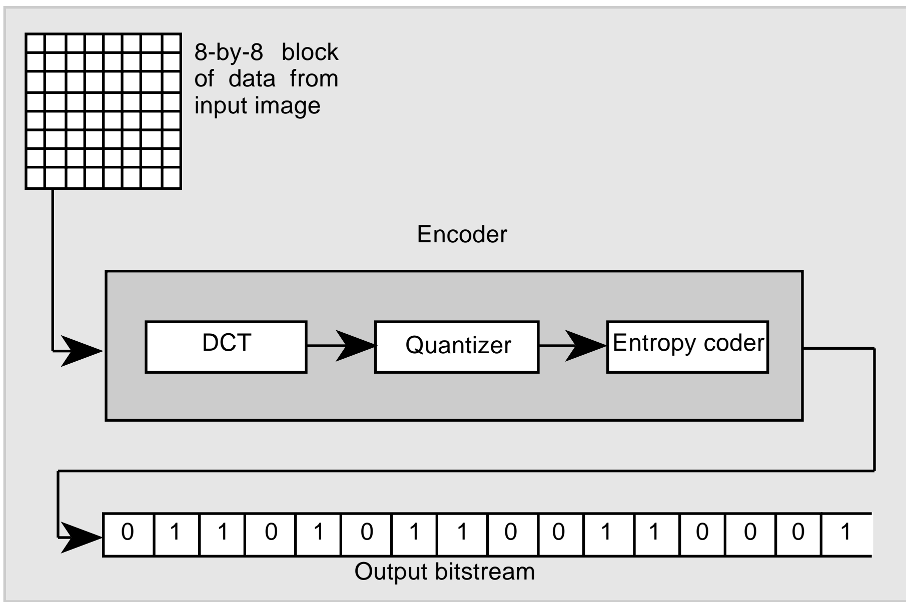

Figure 16-1 JPEG Baseline Sequential Compressor

As the figure indicates, the input to the encoder is an 8-by-8 block of samples from the image being compressed. The compressor encodes each block of data in an image by:

- Performing a Discrete Cosine Transform (DCT) on the 8-by-8 block of data
- Quantizing the results of the DCT
- Entropy coding the results of the quantization step

Each of these steps is considered in more detail in the sections below.

Discrete Cosine Transform

The Discrete Cosine Transform is a mathematical operation that takes a block of image samples as its input and converts the information in that input from the spatial domain to the frequency domain. For example, in JPEG, the input to the DCT is an 8-by-8 matrix whose values represent brightness levels at particular *x*, *y* coordinates, and the output is an 8-by-8 matrix whose values represent relative amounts of the 64 spatial frequencies that make up the input data's spectrum. In the output matrix, information about the lowest frequencies is stored in the upper-left corner, and information about the highest frequencies is stored in the lower-right corner. See Figure 16-2.

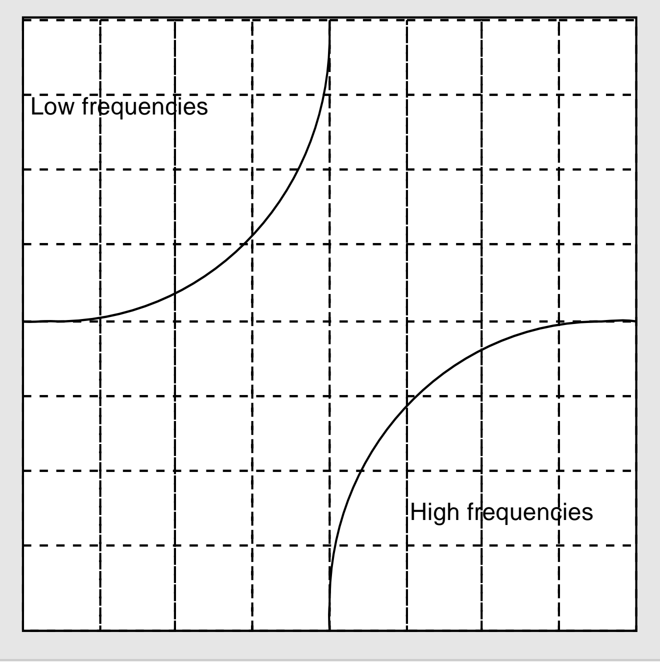

Figure 16-2 Output of the Discrete Cosine Transform

This transformation provides a strong basis for compression because in a typical block of input, low spatial frequencies far outweigh high spatial frequencies. As a result, most of the values in the output matrix, outside of those in the upper-left corner, will have values close to 0 and will end up not being encoded.

Quantization

Quantization is the simplest step in the encoder's algorithm. It simply involves dividing each value in the matrix output by the DCT by the corresponding value in a quantization table.

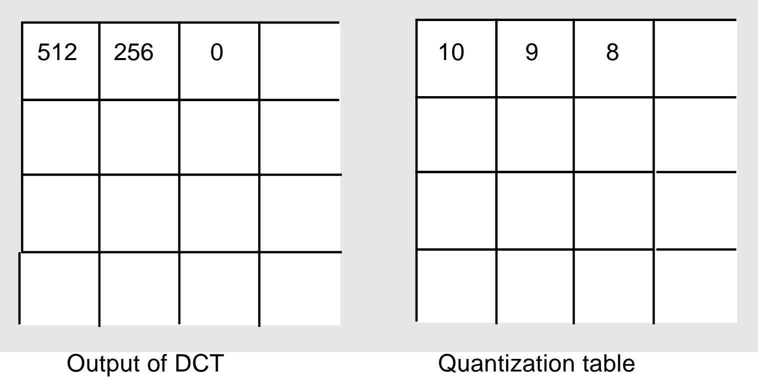

Figure 16-3 Quantization in the JPEG Encoder

Figure 16-3 shows part of an 8-by-8 block of values output by the Discrete Cosine Transform and part of a table to be used in quantizing this data. In this example, the quantizer will divide 512 by 10 and then round 51.2 off to 51. Likewise, it will divide 256 by 9 and 0 by 8. During quantization, any value in the matrix on the left that is divided by a number greater than itself times 2 will go to 0 and will not be encoded in the JPEG bitstream. The quantizer is the part of the JPEG baseline sequential encoder that causes the encoder to be a lossy one.

Entropy Coding

Once a block of data has been quantized, it enters an entropy coder, which creates the actual JPEG bitstream. The entropy coder first looks at the values in the quantized block of data in the zigzag sequence shown in Figure 16-4. Using this sequence ensures that the encoder will encounter all nonzero values in the block as early as possible.

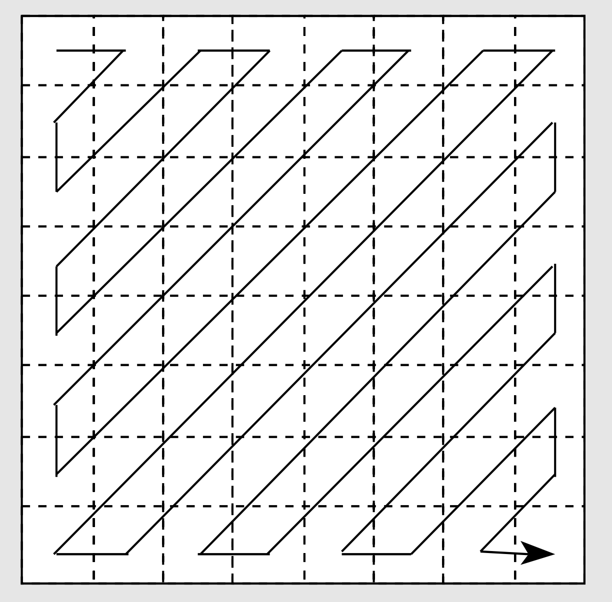

Figure 16-4 Zigzag Sequencing in JPEG Encoder

As the entropy encoder moves through the values in this order, it records three pieces of information each time it encounters a nonzero value: the number of 0's it passed over before finding the nonzero value, the number of bits it will take to encode the nonzero value, and the value itself. The first two pieces of information are considered a pair. For example, if the encoder skips over 7 zeros and then finds a 3, the pair will be 7/2 (2 is number of bits required to encode the value 3). The encoder then consults a Huffman table to find the bit sequence that represents the pair 7/2 and writes this bit sequence to the JPEG bitstream. It then encodes the value 3 using a variable-length code. The

encoder continues this process until all the remaining values in the block are 0's, at which point, it writes a special end-of-block bit sequence to the bitstream.

Creating a JPEG Baseline Sequential CIS

Before you can use the JPEG codec to compress images or decompress a JPEG bitstream, you must create a JPEG CIS. You do this by passing the compressor name `Jpeg` to the function `xil_cis_create()`. See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, "Jpeg");
```

JPEG Baseline Sequential Codec Attributes

As discussed in the section "General CIS Attributes" on page 257, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for CISs attached to a JPEG baseline sequential codec. You set codec-specific attributes using the function `xil_cis_set_attribute()`, and you read them using the function `xil_cis_get_attribute()`.

The JPEG baseline sequential attributes can be broadly grouped into those that affect compression and those that affect decompression. The attributes are discussed under these headings below.

Compression Attributes

Setting any of the following attributes affects how the JPEG baseline sequential compressor compresses images.

BAND_HUFFMAN_TABLE

This attribute enables you to associate a particular Huffman table with a particular band of your image. As shown in Table 16-2 on page 318, the JPEG compressor supports four tables (0 to 3) for both types of DCT coefficients (DC and AC). Table 16-1 shows the default relationship between image bands and Huffman tables.

Table 16-1 Image Bands and Huffman Tables

Table	Type	Band
0	DC	Used to encode the DC coefficient in band 0
1	DC	Used to encode the DC coefficients in all bands except band 0
2	DC	Not used
3	DC	Not used
0	AC	Used to encode the AC coefficients in band 0
1	AC	Used to encode the AC coefficients in all bands except band 0
2	AC	Not used
3	AC	Not used

Note – Tables 2 and 3, for both DC and AC coefficients, can be used for JPEG extended-baseline-sequential mode encoding. Bear in mind, though, that many JPEG decompressors will not be able to decompress the extended-mode bitstream.

Before calling `xil_cis_set_attribute()` to change one of these relationships, you must declare and assign values to the members of a structure of type `XilJpegBandHTable`.

```
typedef struct {
    int band;
    int table;
    XilJpegHTableType type;
} XilJpegBandHTable;
```

The member `band` can have a value in the range 0 to 255 and represents the image band to be associated with a different Huffman table. (The first band in an image is band 0.) The member `table` can be 0 to 3, and the member `type` can be set to `DC` or `AC` (which are enumeration constants of type `XilJpegHTableType`). The setting of `type` determines whether the table you specify will be used to encode DC or AC coefficients in the band you selected.

The code fragment below associates table 0 (type `DC`) with band 1 of the image to be encoded.

```
XilCis cis;
XilJpegBandHTable table_for_band = {1, 0, DC};

xil_cis_set_attribute(cis, "BAND_HUFFMAN_TABLE",
    (void *)&table_for_band);
```

BAND_QUANTIZER

The `BAND_QUANTIZER` attribute enables you to specify the quantization table that will be used in encoding a particular band of an image. Before setting this attribute, you must declare and assign values to the members of a structure of type `XilJpegBandQTable`.

```
typedef struct {
    int band;
    int table;
} xilJpegBandQTable;
```

The integer `band` can have a value in the range 0 to 255. In a YC_bC_r image, the Y band is band 0, the C_b band is band 1, and the C_r band is band 2.

The integer `table` can have a value in the range 0 to 3. By default, table 0 contains the values shown in Table K.1 of Annex K of the ISO JPEG specification, table 1 contains the values shown in Table K.2 of Annex K, and tables 2 and 3 are not loaded. You can load tables 2 and 3 or change the values in tables 0 and 1 using the attribute `QUANTIZATION_TABLE`. See the section "QUANTIZATION_TABLE" on page 322.

By default, band 0 is associated with table 0, and all other bands are associated with table 1. The example code below shows band 2 being associated with table 2.

```
XilCis cis;
XilJpegBandQTable band_to_table;

band_to_table.band = 2;
band_to_table.table = 2;
xil_cis_set_attribute(cis, "BAND_QUANTIZER",
    (void *)&band_to_table);
```

BYTES_PER_FRAME

`BYTES_PER_FRAME` is a read-only attribute that tells you the number of bytes in the last compressed image written to the CIS. The value of this attribute is 0 if no data has been written to the CIS.

`BYTES_PER_FRAME` is useful for helping you select an appropriate setting for the `COMPRESSION_QUALITY` attribute. (See the section "COMPRESSION_QUALITY" on page 314.) If the number of bytes being used to store a compressed image is too high, you can lower it by lowering the value of `COMPRESSION_QUALITY`, and vice versa.

The code fragment below reads the value of the `BYTES_PER_FRAME` attribute.

```
XilCis cis;
int bytes_per_frames;

xil_cis_get_attribute(cis, "BYTES_PER_FRAME",
    (void **)&bytes_per_frame);
```

COMPRESSED_DATA_FORMAT

`COMPRESSED_DATA_FORMAT` is a set-only attribute that defines the format of the JPEG compressor's output. It can be set to either `INTERCHANGE` or `ABBREVIATED_FORMAT`.

Setting the attribute to INTERCHANGE causes the compressor to produce output in JPEG interchange format. In this format, the quantization tables and Huffman tables required by the decompressor are included in each compressed frame.

The default value, ABBREVIATED_FORMAT, causes the compressor to produce output in JPEG abbreviated format. In this format, quantization tables and Huffman tables are not included in a compressed frame if the tables needed to decompress that frame have been defined in a previous frame in the sequence. If a table changes during the course of the sequence, the new table definition is included in the first compressed frame that uses the new table values.

Note – The compressor does not enable you to produce abbreviated-table output in which frames contain only table specifications. However, the decompressor will accept bitstreams in this format.

The code fragment below shows COMPRESSED_DATA_FORMAT being set to INTERCHANGE.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSED_DATA_FORMAT",
    (void *)INTERCHANGE);
```

COMPRESSION_QUALITY

The COMPRESSION_QUALITY attribute tells the compressor how it should handle the trade-off between image quality and compression ratio. The attribute can be set to any value between 1 and 100. Setting the attribute to 100 is a request that the compressor produce very high quality images, even though this will mean a lower compression ratio. A setting of 1 tells the compressor to increase its compression ratio, even though the result will be lower image quality. By default, this attribute is set to 50.

Studies by the JPEG committee indicate that for color images of average complexity, the following relationships exist between level of compression and image quality:

- 0.25-0.5 bits/pixel: moderate to good quality, sufficient for some applications;
- 0.5-0.75 bits/pixel: good to very good quality, sufficient for many applications;
- 0.75-1.5 bits/pixel: excellent quality, sufficient for most applications;
- 1.5-2.0 bits/pixel: usually indistinguishable from the original, sufficient for the most demanding applications.[1]

Also, see Color Plate 4. It shows an image compressed at a high quality setting (95) and one compressed at a low quality setting (5).

The setting of the COMPRESSION_QUALITY attribute controls image quality and the compression ratio by determining a scaling factor the compressor will use in creating scaled versions of its quantization tables. These scaled tables are used during compression. A COMPRESSION_QUALITY value of 50 results in a scaling factor of 1; that is, the scaled tables are identical to the original tables. A COMPRESSION_QUALITY value greater than 50 results in a scaling factor that is less than 1, and a value less than 50 results in a scaling factor greater than 1.

Note – The trade-off between image quality and compression ratio can also be affected by the values stored in the compressor's quantization tables. See the section "QUANTIZATION_TABLE" on page 322.

The code fragment below shows the COMPRESSION_QUALITY attribute being set to 100.

```
XilCis cis
int quality = 100;

xil_cis_set_attribute(cis, "COMPRESSION_QUALITY",
    (void *)quality);
```

1. G. K. Wallace, "The JPEG Still Picture Compression Standard," *Communications of the ACM*, April 1991, p. 35.

ENCODE_INTERLEAVED

ENCODE_INTERLEAVED is a set-only attribute that determines whether the compressor produces an interleaved bitstream when compressing multiband images. If the attribute is set to TRUE, the default value, the compressor produces an interleaved bitstream, and if the attribute is set to FALSE, the compressor produces a noninterleaved bitstream.

This attribute is ignored if the image being compressed has more than four bands because the bitstream for such images is never interleaved. Also, note that the attribute ENCODE_411_INTERLEAVED takes precedence over this attribute if the image being compressed has three bands.

In the example below, ENCODE_INTERLEAVED is being set to FALSE.

```
XilCis cis;

xil_cis_set_attribute(cis, "ENCODE_INTERLEAVED", (void *)FALSE);
```

ENCODE_411_INTERLEAVED

ENCODE_411_INTERLEAVED is a set-only attribute that affects how the compressor handles 3-band images and is intended specifically for use with YC_bC_r images (not RGB images). If the input to the compressor is an image of a different type, this attribute should be set to FALSE (its default value).

Note – If the input images are not 3-band images and this attribute is set to TRUE, the compressor will operate as if the attribute were set to FALSE.

If the input is a YC_bC_r image and ENCODE_411_INTERLEAVED is set to TRUE, the compressor:

- Subsamples the data in the color bands so that only one color value will be encoded for each four color values in the original (4:1:1)
- Interleaves the data for the three image bands on a macroblock basis, where a macroblock is defined as a 16-by-16 block of pixels

Note – This is the only method the XIL library provides for producing JPEG bitstreams with subsampled bands.

If you are making a movie from YC_bC_r images, ENCODE_411_INTERLEAVED should be set to TRUE because the XIL library's accelerated routines for playing back JPEG movies require input in the format described above. The ENCODE_411_INTERLEAVED attribute tells the JPEG compressor to subsample the u and v data instead of using the evenly sampled bands that the XIL software hands it. The output of a decompression (and later export) is still an evenly sampled 3-banded image with the u and v duplicated.

If the attribute is set to FALSE, the setting of the ENCODE_INTERLEAVED attribute determines the format of the JPEG bitstream.

In the example below, ENCODE_411_INTERLEAVED is being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "ENCODE_411_INTERLEAVED",
    (void *)TRUE);
```

HUFFMAN_TABLE

This attribute enables you to supply the compressor with a Huffman table to use in encoding quantized DC or AC coefficients. Before calling xil_cis_set_attribute() to set this attribute, you must store information about the Huffman table in a structure of type XilJpegHTable.

```
typedef struct {
    int table;
    XilJpegHTableType type;
    XilJpegHTableValue *value;
} XilJpegHTable;
```

The member table must be set to 0 to 3, and the member type can be set to DC or AC. In a Huffman table of type DC, the array value contains 16 elements. In a table of type AC, this array contains 256 elements.

The default contents of the eight Huffman tables are shown in Table 16-2.

Table 16-2 Default Huffman Tables

Table	Type	Default Contents
0	DC	Contains the values specified in Table K.3 of the ISO JPEG specification. These values are useful for encoding the DC coefficients of the luminance band of 8-bit YC_bC_r images.
1	DC	Contains the values specified in Table K.4 of the ISO JPEG specification. These values are useful for encoding the DC coefficients of the chrominance bands of 8-bit YC_bC_r images.
2	DC	Empty.
3	DC	Empty.
0	AC	Contains the values specified in Table K.5 of the ISO JPEG specification. These values are useful for encoding the AC coefficients of the luminance band of 8-bit YC_bC_r images.
1	AC	Contains the values specified in Table K.6 of the ISO JPEG specification. These values are useful for encoding the AC coefficients of the chrominance bands of 8-bit YC_bC_r images.
2	AC	Empty.
3	AC	Empty.

Note – Tables 2 and 3, for both DC and AC coefficients, can be used for JPEG extended-baseline-sequential mode encoding. Bear in mind, though, that many JPEG decompressors will not be able to decompress the extended-mode bitstream.

The final member of the `XilJpegHTable` structure, `value`, is an array of structures of type `XilJpegHTableValue`.

```
typedef struct {
    int bits;
    int pattern;
} XilJpegHTableValue;
```

Each structure of this type defines a pair of values to be written to the Huffman table. The member `bits` is the length of a Huffman code in bits, and `pattern` contains the code itself (the `bits` least significant bits of `pattern` are the code).

The examples below indicate how the structures referred to above are used in setting up a Huffman table. This first example shows table 0 (type `DC`) being loaded with values suitable for encoding DC coefficients.

```
XilCis cis;
XilJpegHTable huffman_table;

XilJpegHTableValue huffman_values[16] = {
    {2, 0x0000}, {3, 0x0002}, {3, 0x0003}, {3, 0x0004},
    {3, 0x0005}, {3, 0x0006}, {4, 0x000e}, {5, 0x001e},
    {6, 0x003e}, {7, 0x007e}, {8, 0x00fe}, {9, 0x01fe},
    {0, 0x0000}, {0, 0x0000}, {0, 0x0000}, {0, 0x0000}
};

huffman_table.table = 0;
huffman_table.type = DC;
huffman_table.value = huffman_values;
xil_cis_set_attribute(cis, "HUFFMAN_TABLE",
    (void *)&huffman_table);
```

Note that there are only twelve meaningful value pairs being loaded into the table. The last four are there to fill out `huffman_table.value`, which is an array of sixteen structures.

This next example shows table 0 (type `AC`) being loaded with values suitable for encoding AC coefficients.

```
XilCis cis;
XilJpegHTable huffman_table;

XilJpegHTableValue huffman_values[256] = {
    {4, 0x000a}, {2, 0x0000}, {2, 0x0001}, {3, 0x0004},
    {4, 0x000b}, {5, 0x001a}, {7, 0x0078}, {8, 0x00f8},
    {10, 0x03f6}, {16, 0xff82}, {16, 0xff83}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {4, 0x000c}, {5, 0x001b}, {7, 0x0079},
    {9, 0x01f6}, {11, 0x07f6}, {16, 0xff84}, {16, 0xff85},
```

```
    {16, 0xff86}, {16, 0xff87}, {16, 0xff88}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {5, 0x001c}, {8, 0x00f9}, {10, 0x03f7},
    {12, 0x0ff4}, {16, 0xff89}, {16, 0xff8a}, {16, 0xff8b},
    {16, 0xff8c}, {16, 0xff8d}, {16, 0xff8e}, {0, 0x0000},
    {0, 0x0000},{0, 0x000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {6, 0x003a}, {9, 0x01f7}, {12, 0x0ff5},
    {16, 0xff8f}, {16, 0xff90}, {16, 0xff91}, {16, 0xff92},
    {16, 0xff93}, {16, 0xff94}, {16, 0xff95}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {6, 0x003b}, {10, 0x03f8}, {16, 0xff96},
    {16, 0xff97}, {16, 0xff98}, {16, 0xff99}, {16, 0xff9a},
    {16, 0xff9b}, {16, 0xff9c}, {16, 0xff9d}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},


    {0, 0x0000}, {7, 0x007a}, {11, 0x07f7}, {16, 0xff9e},
    {16, 0xff9f}, {16, 0xffa0}, {16, 0xffa1}, {16, 0xffa2},
    {16, 0xffa3}, {16, 0xffa4}, {16, 0xffa5}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {7, 0x007b}, {12, 0x0ff6}, {16, 0xffa6},
    {16, 0xffa7}, {16, 0xffa8}, {16, 0xffa9}, {16, 0xffaa},
    {16, 0xffab}, {16, 0xffac}, {16, 0xffad}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {8, 0x00fa}, {12, 0x0ff7}, {16, 0xffae},
    {16, 0xffaf}, {16, 0xffb0}, {16, 0xffb1}, {16, 0xffb2},
    {16, 0xffb3}, {16, 0xffb4}, {16, 0xffb5}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {9, 0x01f8}, {15, 0x7fc0}, {16, 0xffb6},
    {16, 0xffb7}, {16, 0xffb8}, {16, 0xffb9}, {16, 0xffba},
    {16, 0xffbb}, {16, 0xffbc}, {16, 0xffbd}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {9, 0x01f9}, {16, 0xffbe}, {16, 0xffbf},
    {16, 0xffc0}, {16, 0xffc1}, {16, 0xffc2}, {16, 0xffc3},
    {16, 0xffc4}, {16, 0xffc5}, {16, 0xffc6}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {9, 0x01fa}, {16, 0xffc7}, {16, 0xffc8},
```

```
    {16, 0xffc9}, {16, 0xffca}, {16, 0xffcb}, {16, 0xffcc},
    {16, 0xffcd}, {16, 0xffce}, {16, 0xffcf}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {10, 0x03f9}, {16, 0xffd0}, {16, 0xffd1},
    {16, 0xffd2}, {16, 0xffd3}, {16, 0xffd4}, {16, 0xffd5},
    {16, 0xffd6}, {16, 0xffd7}, {16, 0xffd8}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {10, 0x03fa}, {16, 0xffd9}, {16, 0xffda},
    {16, 0xffdb}, {16, 0xffdc}, {16, 0xffdd}, {16, 0xffde},
    {16, 0xffdf}, {16, 0xffe0}, {16, 0xffe1}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {11, 0x07f8}, {16, 0xffe2}, {16, 0xffe3},
    {16, 0xffe4}, {16, 0xffe5}, {16, 0xffe6}, {16, 0xffe7},
    {16, 0xffe8}, {16, 0xffe9}, {16, 0xffea}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {0, 0x0000}, {16, 0xffeb}, {16, 0xffec}, {16, 0xffed},
    {16, 0xffee}, {16, 0xffef}, {16, 0xfff0}, {16, 0xfff1},
    {16, 0xfff2}, {16, 0xfff3}, {16, 0xfff4}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000},

    {11, 0x07f9}, {16, 0xfff5}, {16, 0xfff6}, {16, 0xfff7},
    {16, 0xfff8}, {16, 0xfff9}, {16, 0xfffa}, {16, 0xfffb},
    {16, 0xfffc}, {16, 0xfffd}, {16, 0xfffe}, {0, 0x0000},
    {0, 0x0000},{0, 0x0000},{0, 0x0000},{0, 0x0000}
};

huffman_table.table = 0;
huffman_table.type = AC;
huffman_table.value = huffman_values;
xil_cis_set_attribute(cis, "HUFFMAN_TABLE",
    (void *)&huffman_table);
```

Note that `huffman_values` is an array of 256 structures of type `XilJpegHTableValue`, and that the structures are organized as 16 groups of 16 structures. Because there are not 16 bit lengths (`bits`) for each of the 16 run lengths, each group of 16 structures contains 5 meaningless structures at the end.

OPTIMIZE_HUFFMAN_TABLES

You can use this attribute to tell the compressor to generate optimal Huffman tables instead of using the default tables included in the ISO specification or tables you loaded earlier. When the compressor creates optimal Huffman tables, its Huffman tables can vary from image to image, and you should see higher rates of compression.

If OPTIMIZE_HUFFMAN_TABLES is set to FALSE, its default value, the compressor uses fixed Huffman tables for each image in a sequence. If the attribute is set to TRUE, the compressor uses optimal Huffman tables.

Note – If an application sets OPTIMIZE_HUFFMAN_TABLES to TRUE and then later sets it to FALSE, the compressor loads a default set of tables. It does not restore the tables it was using when the attribute was first set to TRUE.

The following code fragment shows OPTIMIZE_HUFFMAN_TABLES being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, “OPTIMIZE_HUFFMAN_TABLES”,
    (void *)TRUE);
```

QUANTIZATION_TABLE

The QUANTIZATION_TABLE attribute enables you to load values into one of four quantization tables that the compressor may use to quantize DCT coefficients. Prior to actually setting the attribute you must declare and assign values to the members of an XilJpegQTable structure.

```
typedef struct {
    int table;
    int (*value)[8];
} XilJpegQTable;
```

The integer table must have a value between 0 and 3. This value indicates which of the four quantization tables you want to load. By default, table 0 contains the values shown in Table K.1 of Annex K of the ISO JPEG

specification. These values are designed to be used with the luminance band of 8-bit YCC images. Likewise, table 1 contains the values shown in Table K.2 of Annex K. The values in this table are designed to be used with the chrominance bands of 8-bit YCC images. By default, tables 2 and 3 are not loaded.

You use the pointer `value` (a pointer to an array of 8 integers) to point to the 64 quantization values that you want to load into the table.

The code fragment below illustrates 64 quantization values being loaded into table 2.

```
Xilcis cis;
XilJpegQTable quantization;
int quantization_array[8][8] = {
    {16, 11, 10, 16, 24, 40, 51, 61},
    {12, 12, 14, 19, 26, 58, 60, 55),
    {14, 13, 16, 24, 40, 57, 69, 56},
    {14, 17, 22, 29, 51, 87, 80, 62},
    {18, 22, 37, 56, 68, 109, 103, 77},
    {24, 35, 55, 64, 81, 104, 113, 92},
    {49, 64, 78, 87, 103, 121, 120, 101},
    {72, 92, 95, 98, 112, 100, 103, 99}
};

quantization.table = 2;
quantization.value = quantization_array;
xil_cis_set_attribute(cis, "QUANTIZATION_TABLE",
    (void *)&quantization);
```

TEMPORAL_FILTERING

This attribute controls whether the JPEG compressor filters 3-band images before it encodes them. If the attribute is set to `FALSE`, its default value, the compressor does not filter the images, and if it is set to `TRUE`, the compressor does filter them. This filtering can help in cases where the images being encoded contain a lot of noise; however, the filtering can also have the undesirable side effect of producing ghosting artifacts.

The example below shows the attribute being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "TEMPORAL_FILTERING", (void *)TRUE);
```

And this fragment demonstrates how to read the attribute's value.

```
XilCis cis;
Xil_boolean filtering_enabled;

xil_cis_get_attribute(cis, "TEMPORAL_FILTERING",
    (void **)&filtering_enabled);
```

Decompression Attributes

Setting the following attribute affects how the JPEG baseline sequential codec decompresses images.

DECOMPRESSION_QUALITY

You use the DECOMPRESSION_QUALITY attribute to provide a hint to the compressor concerning how it should handle the trade-off between the quality of decompressed images and speed of decompression. The decompressor is free to ignore this hint.

The attribute's value must be an integer in the range of 1 to 100. A setting of 100 is a request for a high level of image quality, and a setting of 1 is a request from a high playback speed. By default, the attribute is set to 100.

Note – The JPEG compressor increases playback speed by decreasing the number of quantized coefficients it uses in reconstructing an image. It drops high-frequency coefficients first.

The example code below shows DECOMPRESSION_QUALITY being set to 75.

```
XilCis cis;
int quality = 75;

xil_cis_set_attribute(cis, "DECOMPRESSION_QUALITY",
    (void *)quality);
```

IGNORE_HISTORY

The setting of this attribute affects what happens when you seek backward or forward in a JPEG CIS whose random-access flag is not set. If your CIS's random-access flag is set, the value of IGNORE_HISTORY is irrelevant: you can always seek backward or forward and get correct results. To determine whether the random-access flag is set, call the function xil_cis_get_random_access().

Note – When you create a JPEG CIS, its random-access flag is always set. This flag will remain set if the quantization and Huffman tables needed to decode the entire JPEG bitstream are encoded with the first image or if those tables are encoded with each image in the sequence. Otherwise, the decoder will clear the flag as soon as it discovers that the bitstream cannot be accessed randomly.

If your CIS's random-access bit is not set and IGNORE_HISTORY is FALSE (its default value), the following rules apply:

- Backward seeds are illegal.
- A forward seek will be successful; that is, the image you decode after the seek will be reconstructed properly. However, a forward seek with IGNORE_HISTORY set to FALSE may be slower than one with the attribute set to TRUE.

If your CIS's random-access bit is not set and IGNORE_HISTORY is TRUE, these rules apply:

- Backward seeks are legal. However, it is the responsibility of your application to seek to an image that can be decoded correctly. That is, the image you seek to must either define its own tables or depend on the tables that were most recently loaded into the decoder.

- You can also seek forward. However, your application should not seek forward past images that contain table definitions if those definitions will be needed to decode the image you're seeking to. The decoder does not ensure that these table definitions are loaded.

The code below shows IGNORE_HISTORY being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "IGNORE_HISTORY", (void *)TRUE);
```

JPEG Molecules

The XIL library includes a series of *molecules* that accelerate the playback of JPEG baseline sequential bitstreams. These molecules are optimized routines that perform the jobs of two or more functions from the XIL API. You do not call such an optimized routine directly; rather, the library calls a molecule when your program calls a predefined sequence of XIL functions, sometimes with specific arguments.

For example, if your program calls xil_decompress() to decode an image stored in a JPEG CIS and then calls xil_ordered_dither() to convert the decoded image from a 24- to an 8-bit image, the library may not call these two functions. Instead, it may call a molecule that performs the decompression and the dithering in an optimized way.

For information about the JPEG molecules that are available and information about how to call those molecules, see the section "XIL Molecules" on page 398.

JPEG Lossless Codec 17 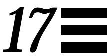

Like the JPEG baseline sequential compressor, the JPEG lossless compressor is designed to compress still continuous-tone images. The difference between the two is that the baseline sequential compressor is a lossy compressor: once an image has been compressed, you cannot recover the original image samples. The JPEG lossless compressor does allow you to recover these samples; however, the lossless compressor produces much lower compression ratios than the lossy compressor. The compression ratio for the lossless compressor is in the neighborhood of 2:1. One other difference between the two compressors is that while the baseline sequential compressor works only with `XIL_BYTE` images, the lossless compressor can work with both `XIL_BYTE` and `XIL_SHORT` images.

The remainder of this chapter is divided into three sections. The first section provides an overview of how the JPEG lossless compressor works. The second explains briefly how to create a JPEG lossless CIS. The third discusses a set of CIS attributes that apply specifically to the JPEG lossless codec (as opposed to the general CIS attributes covered in the section “General CIS Attributes” on page 257).

How the JPEG Lossless Codec Works

The JPEG lossless compressor is not based on the DCT like the baseline sequential encoder. Instead, it uses a predictive method, as shown in Figure 17-1.

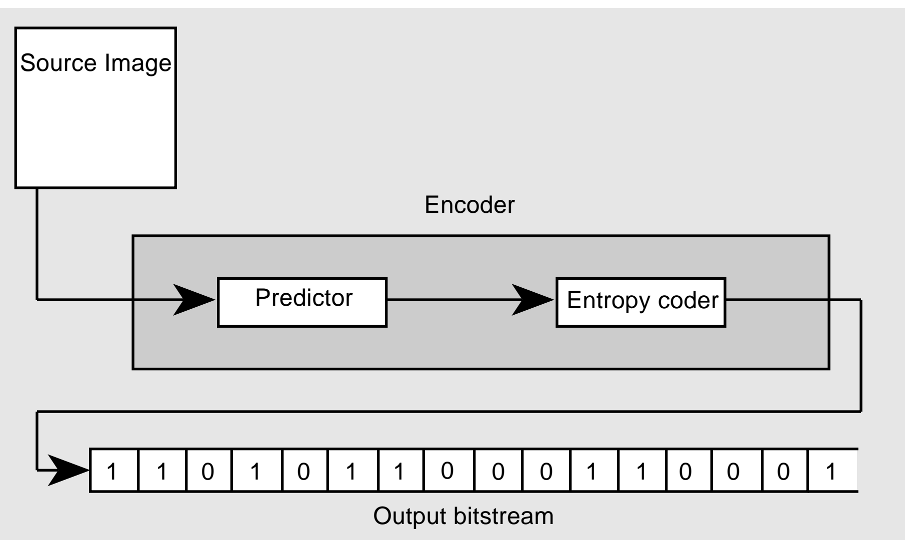

Figure 17-1 JPEG Lossless Compressor

Each sample in the source image is encoded as follows: The predictor makes a guess at the sample's value based on its knowledge of the values of neighboring samples and then subtracts the predicted value of the sample from its actual value. The difference calculated by the predictor is then passed on to the entropy coder, which does a lookup in a Huffman table and writes encoded data to the bitstream.

The prediction and entropy-coding steps in the encoding process are discussed in more detail in the following sections.

Note – The XIL library's JPEG lossless encoder actually lets you perform an operation called a point transformation on your source image before the encoding summarized above begins. The point transformation can be done on a band-by-band basis. What happens is this. If you supply a point-transform value other than 0 for a band of an image, each sample in the band is divided by 2 raised to the power *x*, where *x* is the point-transform value. This operation leads to a higher compression ratio because smaller values will need to be encoded; however, it is not actually part of the lossless-encoding process because it can result in the loss of some data.

Prediction

As mentioned above, the predictor predicts the value of an image sample based on its knowledge of the values of neighboring samples. The neighboring samples it can take into consideration are shown in Figure 17-2 and are labeled A, B, and C. The sample whose value is being predicted is labeled P.

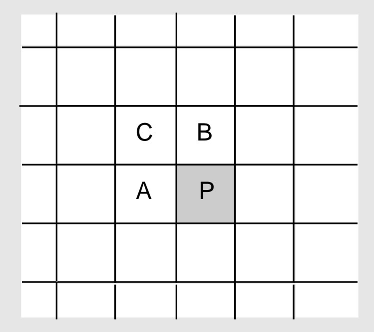

Figure 17-2 Predicting Values in the JPEG Lossless Compressor

Seven prediction methods are defined. These are shown below.

1. P = A
2. P = B
3. P = C
4. P = A + B - C
5. P = A + ((B - C) / 2)
6. P = B + ((A - C) / 2)
7. P = (A + B) / 2

By default, the compressor uses prediction method 1 for all bands in an image. However, you can select a different prediction method for a particular band using the CIS attribute `LOSSLESS_BAND_SELECTOR`. See the section "LOSSLESS_BAND_SELECTOR" on page 334.

After the predictor has predicted the value of a sample, it subtracts the predicted value of the sample from its actual value. It is this difference that is encoded by the entropy coder.

Entropy Coding

The entropy coder takes as input a difference, the difference between a predicted and actual sample value. On the output side, the entropy coder writes a sequence of bits to the JPEG bitstream.

To produce its output, the entropy coder uses the number of bits required to represent a difference as an index into a Huffman table. At the proper location in the table, the encoder reads a code word—a sequence of bits—and then writes this code word to the bitstream. Immediately after the code word, the encoder writes to the bitstream the difference itself. The encoder performs these steps for each sample in the image.

JPEG Lossless Compressor Attributes

The attributes discussed below affect the way that the JPEG lossless compressor works.

BAND_HUFFMAN_TABLE

This attribute enables you to specify that a particular Huffman table be used in encoding a particular band of your image. Before calling `xil_cis_set_attribute()` to set this attribute, you must declare and assign values to the members of a structure of type `XilJpegBandHTable`.

```
typedef struct {
    int band;
    int table;
    XilJpegHTableType type;
} XilJpegBandHTable;
```

The member `band` can have a value in the range 0 to 255 and represents the image band to be associated with a particular Huffman table. The member `table` can have a value between 0 and 3 and specifies one of four possible Huffman tables, and the member `type` must be set to `DC`.

The code fragment below associates table 0 (type `DC`) with band 1 of the image to be encoded.

```
XilCis cis;
XilJpegBandHTable table_for_band = {1, 0, DC};

xil_cis_set_attribute(cis, “BAND_HUFFMAN_TABLE”,
    (void *)&table_for_band);
```

By default, table 0 is used in encoding band 0 of an image, and table 1 is used in encoding all other bands (1 to 255).

COMPRESSED_DATA_FORMAT

COMPRESSED_DATA_FORMAT is a set-only attribute that defines the format of the JPEG lossless compressor's output. It can be set to either INTERCHANGE or ABBREVIATED_IMAGE.

Setting the attribute to INTERCHANGE causes the compressor to produce output in JPEG interchange format. In this format, the Huffman tables required by the decompressor are included in each compressed frame.

The default value, ABBREVIATED_IMAGE, causes the compressor to produce output in JPEG abbreviated format. In this format, Huffman tables are not included in a compressed frame if the tables needed to decompress that frame have been defined in a previous frame in the sequence. If a table changes during the course of the sequence, the new table definition is included in the first compressed frame that uses the new table values.

The code fragment below shows COMPRESSED_DATA_FORMAT being set to INTERCHANGE.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSED_DATA_FORMAT",
    (void *)INTERCHANGE);
```

ENCODE_INTERLEAVED

ENCODE_INTERLEAVED is a set-only attribute that determines whether the compressor produces an interleaved bitstream when compressing multiband images. If the attribute is set to TRUE, the default value, the compressor produces an interleaved bitstream, and if the attribute is set to FALSE, the compressor produces a noninterleaved bitstream.

This attribute is ignored if the image being compressed has more than four bands, or if different point transformations or different predictors are being used in encoding different bands. In these instances, the compressor never produces an interleaved bitstream.

In the example below, ENCODE_INTERLEAVED is being set to FALSE.

```
XilCis cis;

xil_cis_set_attribute(cis, "ENCODE_INTERLEAVED", (void *)FALSE);
```

HUFFMAN_TABLE

This attribute enables you to supply the compressor with a Huffman table to use in encoding difference values. By default, the compressor's four Huffman tables are loaded as follows: Tables 0 and 2 contain the values specified in Table K.3 of the ISO JPEG specification (extended to 17 entries). These values are useful for encoding pixel differences for the luminance band of 8-bit YC_bC_r images. Tables 1 and 3 contain the values specified in Table K.4 of the ISO JPEG specification (again extended to 17 entries). These values are useful for encoding pixel differences for the chrominance bands of 8-bit YC_bC_r images.

Before calling xil_cis_set_attribute() to set this attribute, you must store information about the Huffman table in a structure of type XilJpegHTable.

```
typedef struct {
    int table;
    XilJpegHTableType type;
    XilJpegHTableValue *value;
} XilJpegHTable;
```

The member table must be set to a value between 0 and 3, and the member type must be set to DC. The final member of the XilJpegHTable structure, value, is an array of 17 structures of type XilJpegHTableValue.

```
typedef struct {
    int bits;
    int pattern;
} XilJpegHTableValue;
```

Each structure of this type defines a pair of values to be written to the Huffman table. The member `bits` is the length of a Huffman code in bits, and `pattern` contains the code itself (the `bits` least significant bits of `pattern` are the code).

The example below shows table 0 (type `DC`) being loaded with values suitable for encoding difference values.

```
XilCis cis;
XilJpegHTable huffman_table;

XilJpegHTableValue huffman_values[17] = {
    {2, 0x0000}, {3, 0x0002}, {3, 0x0003}, {3, 0x0004},
    {3, 0x0005}, {3, 0x0006}, {4, 0x000e}, {5, 0x001e},
    {6, 0x003e}, {7, 0x007e}, {8, 0x00fe}, {9, 0x01fe},
    {10, 0x03fe},{11, 0x07fe}, {12, 0x0ffe}, {13, 0x1ffe},
    {14, 0x3ffe}
};

huffman_table.table = 0;
huffman_table.type = DC;
huffman_table.value = huffman_values;
xil_cis_set_attribute(cis, "HUFFMAN_TABLE",
    (void *)&huffman_table);
```

LOSSLESS_BAND_SELECTOR

As mentioned in the section "Prediction" on page 329, there are seven methods the predictor can use to predict the values of an image sample. This attribute enables you to specify which of these methods the compressor should use in encoding a particular band of an image.

Before calling `xil_cis_set_attribute()` to set this attribute, you must declare and assign values to the members of a structure of type `XilJpegLLBandSelector`.

```
typedef struct {
    int band;
    XilJpegLLBandSelectorType selector;
} XilJpegLLBandSelector;
```

The structure member `band` must be set to a value between 0 and 255 (where band 0 is the first band in an image). The member `selector` can have any of the values shown in the following enumeration.

```
typedef enum {
    ONE_D1=1, ONE_D2, ONE_D3, TWO_D1, TWO_D2, TWO_D3, TWO_D4
} XilJpegLLBandSelectorType;
```

The correspondence between these enumeration constants and the prediction methods described earlier is shown in Table 17-1.

Table 17-1 JPEG Lossless Prediction Methods

Constant	Prediction Method
ONE_D1	P = A
ONE_D2	P = B
ONE_D3	P = C
TWO_D1	P = A + B - C
TWO_D2	P = A + ((B - C) / 2)
TWO_D3	P = B + ((A - C) / 2)
TWO_D4	P = (A + B) / 2

The following code fragment shows the prediction formula P = A + ((B - C) / 2) being selected for band 2 of an image.

```
XilCis cis;
XilJpegLLBandSelector predictor_band = {2, TWO_D2};

xil_cis_set_attribute(cis, "LOSSLESS_BAND_SELECTOR",
    (void *)&predictor_band);
```

LOSSLESS_BAND_PT_TRANSFORM

This attribute enables you to request that a point transformation be performed on a particular band of an image before the image is encoded. If you request this operation, each sample in the band is divided by 2 raised to power of *x*, where *x* is the prediction-transformation value. The default prediction-transformation value for all bands is 0.

Before calling `xil_cis_set_attribute()` to set this attribute, you must declare and assign values to the members of a structure of type `XilJpegLLBandPtTransform`.

```
typedef struct {
    int band;
    int PtTransform;
} XilJpegLLBandPtTransform;
```

The structure member `band` must be set to a value between 0 and 255 (where band 0 is the first band in an image), and the `PtTransform` member must be set to an integer in the range 0 to 15 that represents the power of two by which you want to divide all the samples in a band.

The code below requests that the samples in band 2 of an image be divided by 16 (2^4) before the lossless encoding process takes place.

```
XilCis cis;
XilJpegLLBandPtTransform pt_transform = {2, 4};

xil_cis_set_attribute(cis, "LOSSLESS_BAND_PT_TRANSFORM",
    (void *)&pt_transform);
```

H.261 Codec 18

The H.261 compression-decompression scheme takes its name from the title of the recommendation in which the H.261 codec is specified: Recommendation H.261 published by the International Telegraph and Telephone Consultative Committee (CCITT). This recommendation defines a video encoder that is intended to be used to compress video data that will be sent over Integrated Services Digital Network (ISDN) lines.

The H.261 codec is intended primarily for use in video telephony and videoconferencing applications. Video telephony, in which generally a picture of the speaker's face against a stationary background is transmitted, is possible when only one or two ISDN channels (each capable of carrying 64Kb/s of information) are available. If more channels are available, more complex images can be sent.

The remainder of this chapter is divided into four sections. The first section provides an overview of how an H.261 codec works. The second explains briefly how to create an H.261 CIS. The third discusses CIS attributes that apply specifically to a CIS associated with an XIL H.261 compressor or decompressor (as opposed to the general CIS attributes covered in the section "General CIS Attributes" on page 257). And the fourth section introduces the subject of accelerating the playback of H.261 bitstreams. For further information on this subject, see Chapter 21, "Acceleration in XIL Programs.".

Note – This chapter discusses both H.261 compression and decompression. However, the current release of the XIL library includes only an H.261 decompressor. The compressor interface is defined for third parties who want to implement XIL H.261 compressors.

How an H.261 Codec Works

This section presents an overview of how an H.261 codec works. It discusses

- The format of the images that can be used as input to the encoder
- The basic encoding scheme
- Methods of controlling the size of the bitstream produced by the encoder
- How the codec supports multipoint conferencing

Source Images

The images supplied as input to an H.261 compressor must meet both color space and size (width and height) requirements. In terms of color space, the images must be YC_bC_r images that conform to the standard set forth in CCIR Recommendation 601. In terms of size, the images must adhere to either the Common Interchange Format (CIF) or the Quarter-CIF (QCIF) format. Table 18-1 below indicates the widths and heights defined by these formats.

Table 18-1 Sizes of CIF- and QCIF-Format Images

	Width	Height
CIF images	352	288
QCIF images	176	144

All H.261 encoders must be able to compress QCIF images. The ability to compress CIF images is optional.

Given an image of the appropriate format, the H.261 compressor subsamples the chrominance values so that there is one C_b value and one C_r value for each two-by-two block of luminance values. It then processes the image in segments called macroblocks. Each macroblock consists of a 16-by-16 block of luminance values and the chrominance values associated with those luminance values. See Figure 18-1 below.

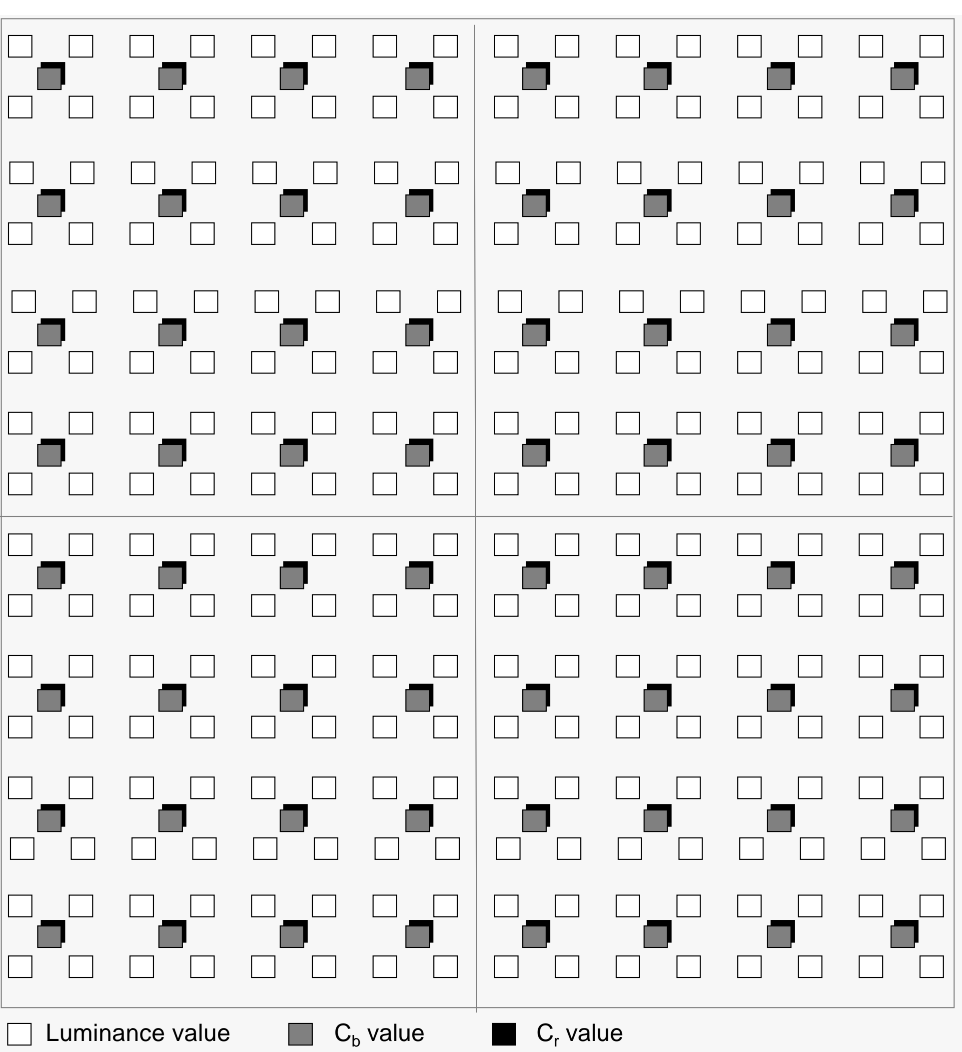

Figure 18-1 Macroblock

As you'll see shortly, the encoder performs some operations on 8-by-8 blocks of values. Each macroblock contains six blocks: four blocks of luminance values, one block of C_b values, and one block of C_r values.

Basic Encoding Scheme

The flow chart in Figure 18-2 illustrates the procedure the H.261 encoder uses to encode a macroblock.

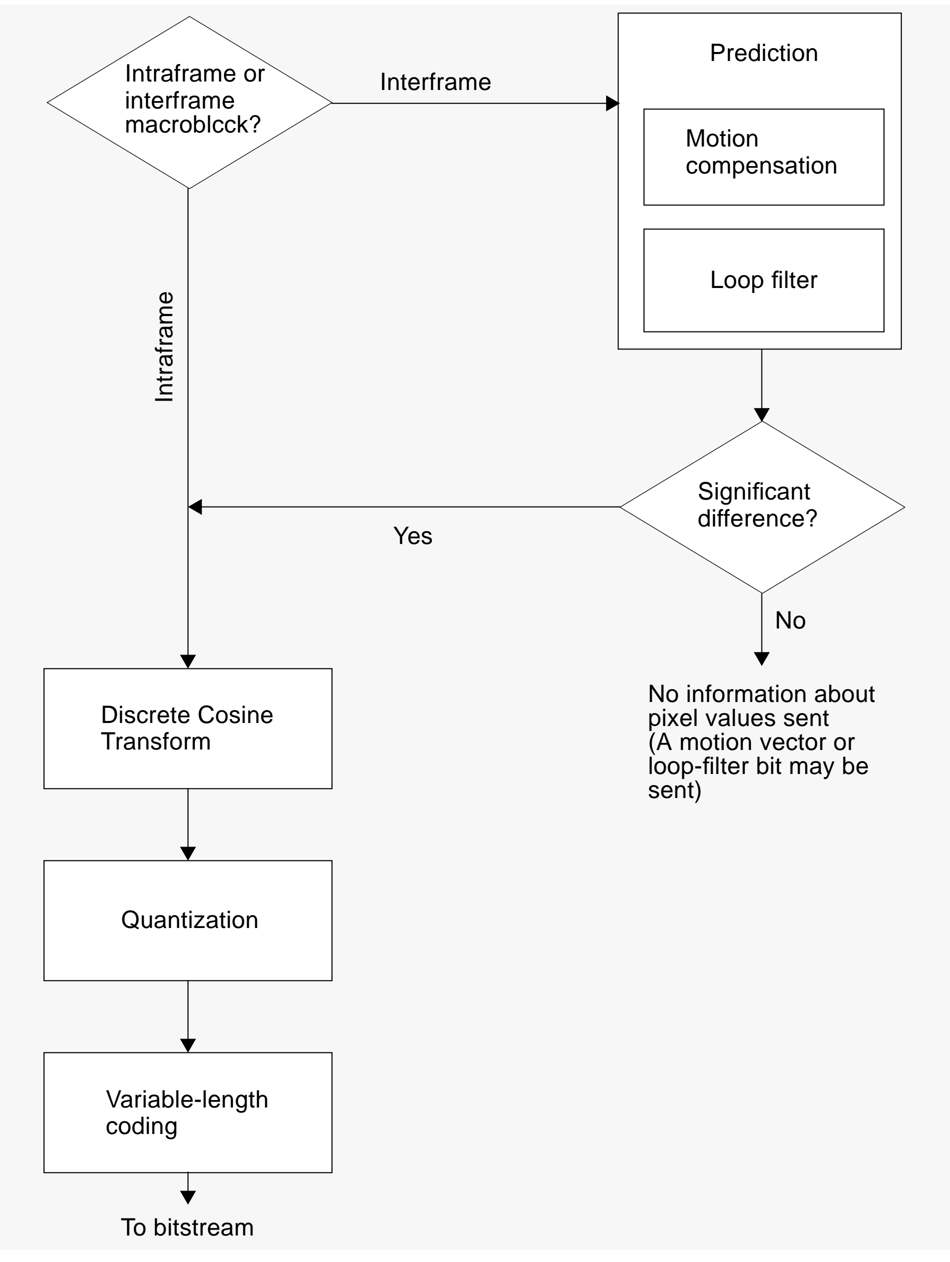

Figure 18-2 Flow Diagram for H.261 Encoding

Intraframe Versus Interframe Encoding

For each macroblock that it encodes, the H.261 encoder can perform intraframe or interframe compression. In intraframe mode, the compressor encodes the actual YC_bC_r values in the macroblock. In interframe mode, the compressor

1. Looks at the YC_bC_r values in the macroblock it is encoding
2. Calculates the difference between the predicted values for the macroblock and the actual values in the macroblock. The predicted values are taken from the most recently compressed image, which is stored in a history buffer.
3. Encodes the difference values if they are significant

In general, the H.261 compressor relies very heavily on interframe encoding because this type of encoding leads to greater rates of compression than intraframe compression. However, Recommendation H.261 requires that the encoder intraframe encode each macroblock at least once every 132 frames. This requirement ensures that if you join a videoconference in progress or your videoconference is disrupted by data transmission problems, all macroblocks will be updated properly within a few seconds.

Prediction

As noted in the last section, in interframe-encoding mode, the encoder calculates the difference between YC_bC_r values in the macroblock it is encoding currently and the values in the corresponding macroblock in the preceding picture. Before performing this calculation, the encoder may perform either a motion-compensation operation or a motion-compensation operation followed by a loop-filter operation. Both of these operations are optional.

When the encoder performs motion compensation, it compares the YC_bC_r values in the current macroblock not only with those in the spatially corresponding macroblock in the preceding picture, but also with the values in macroblocks that neighbor the spatially corresponding macroblock in the preceding picture. See Figure 18-3.

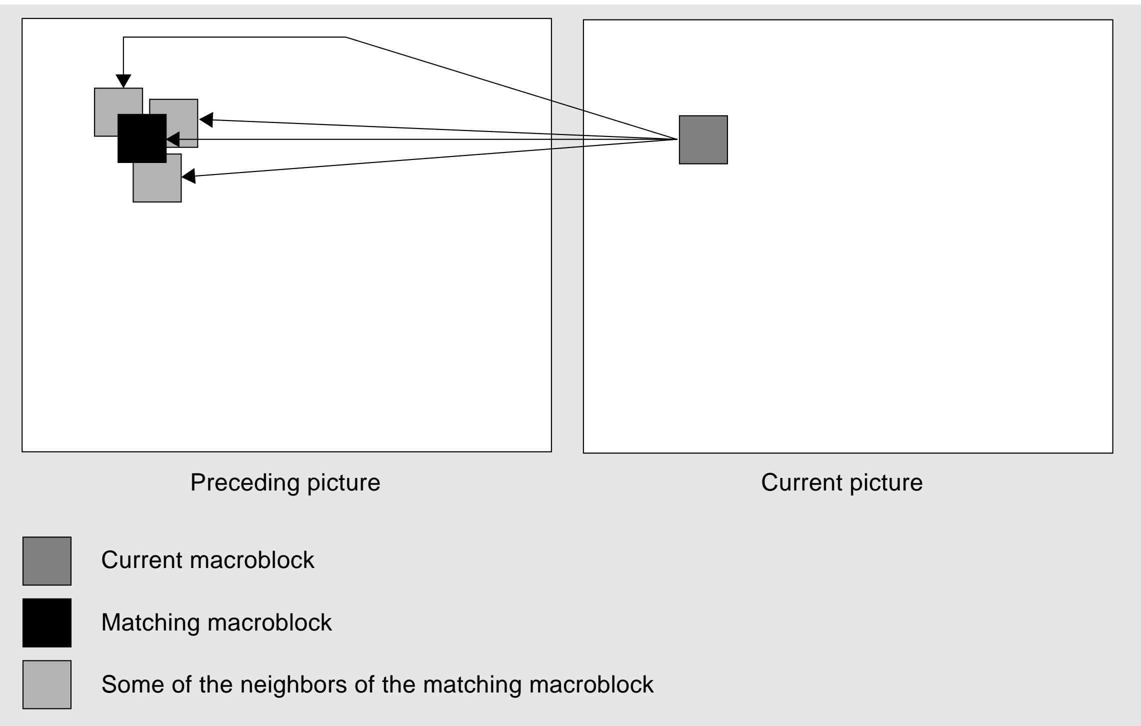

Figure 18-3 Motion Compensation in H.261

The neighboring macroblocks being examined can be offset from the matching macroblock by a maximum of ±15 pixels in both the *x* and *y* directions. The macroblock in the history image that best matches the macroblock in the current image is used in calculating the difference values to encode.

If the macroblock in the history image used to calculate difference values is not the matching macroblock, the encoder must record the number of pixels by which the former is offset from the latter. These *x* and *y* offsets are written to a motion vector, which is later variable-length coded and written to the H.261 bitstream along with any encoded difference values.

The second operation that the encoder may perform before calculating the difference between macroblocks is a filter of the macroblock of interest in the history image. This filtering operation is designed to remove high-frequency information from the macroblock. Generally, this filtering leads to smaller differences between the macroblocks and, thus, to a more compact bitstream.

Encoding YC_bC_r or Difference Values

Whether the encoder is encoding the actual values in a macroblock from the current picture or difference values calculated during the prediction step, it uses the procedure depicted in Figure 18-4.

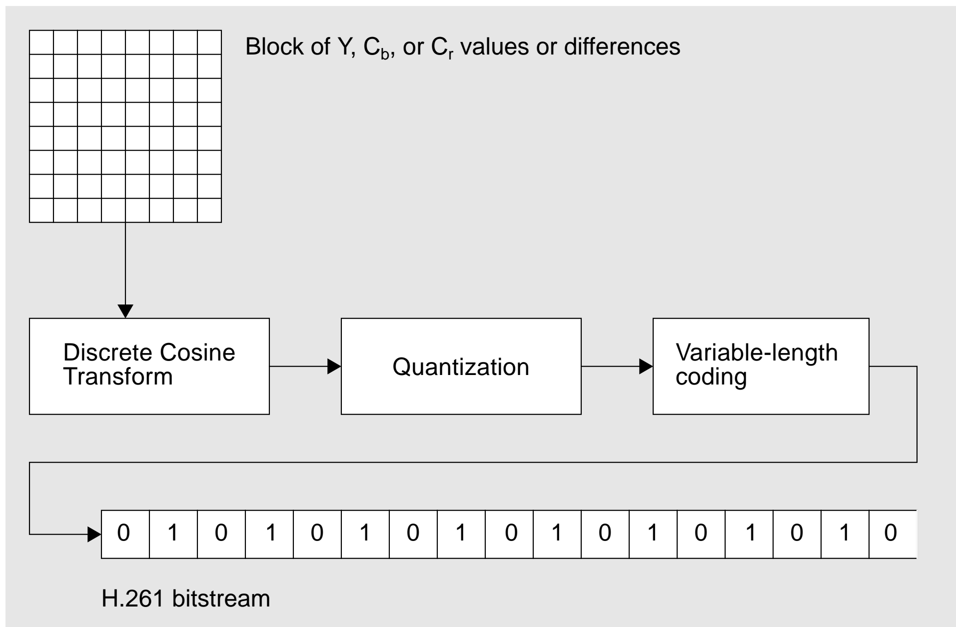

Figure 18-4 Encoding of YC_bC_r or Difference Values in H.261

Each of the six blocks in a macroblock (four blocks of luminance values and two blocks of color values) is encoded separately. The values for each block are

1. Transformed from the spatial to the transform domain using a Discrete Cosine Transform
2. Linearly quantized
3. Encoded with variable length codes or, for less frequently occurring values, with 20-bit codes

This encoding scheme is very similar to the one used in the JPEG still-image compression standard. For a more detailed discussion of the steps involved in this type of encoding, see the section "How the JPEG Baseline Sequential Codec Works" on page 306.

Bit-Rate Control

As mentioned earlier, the H.261 codec is intended primarily for use in videophone and videoconferencing applications. Because these applications need to send data at a constant rate over a network, the encoder must use a constant number of bits to encode, say, a second's worth of video. The encoder can achieve this constant bit rate using any combination of the following techniques.

- Altering the criterion that determines whether a macroblock that is to be interframe encoded actually needs to be encoded. The macroblock needs to be encoded only if its luminance values differ from those in the corresponding macroblock in the preceding picture by a certain amount. By increasing this amount, the encoder decreases the number of macroblocks it must compress.
- Changing the values in the quantizer. To produce a lower bit rate, the encoder can increase the size of the values in the quantizer. This strategy results in quantized coefficients with relatively low values, which can be encoded with relatively short code words.
- Using the loop filter operation described in the section "Prediction" on page 342.

Provisions for Multipoint Conferencing

The CCITT's specification of the H.261 codec includes several features designed to facilitate multipoint conferencing. In a multipoint conference, the receiver may elect to switch between two or more sources of video. These features include:

- A freeze-picture request. This request is an external signal that causes the decoder to stop updating the currently displayed picture. The picture remains frozen until the decoder sees a freeze-picture-release flag in the bitstream or until a timeout period of six seconds or more has elapsed.
- A fast-update request. This request is an external signal that causes the encoder to compress the next picture using intraframe encoding exclusively. The encoder must compress this frame without overflowing its output buffer.
- A freeze-picture release. When an encoder receives a fast-update request, it sets a bit in the header of the next picture it encodes. This bit tells a decoder that has frozen its display to resume displaying pictures in the normal way.

Creating an H.261 CIS

Before you can use the H.261 decompressor to decompress an H.261 bitstream, you must create an H.261 CIS (and write an H.261 bitstream to the CIS). You create this CIS by passing the decompressor name `H261` to the function `xil_cis_create()`. See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, "H261");
```

H.261 Codec Attributes

As discussed in the section “General CIS Attributes” on page 257, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for CISs attached to an H.261 compressor or decompressor. You set compressor-specific attributes using the function `xil_cis_set_attribute()`, and you read compressor-specific attributes using the function `xil_cis_get_attribute()`.

The H.261 attributes can be broadly grouped into those that affect compression and those that affect decompression. The attributes are discussed under these headings below.

Compression Attributes

Setting any of the following attributes affects how an H.261 compressor compresses images.

COMPRESSOR_BITS_PER_IMAGE

This attribute controls the number of bits the encoder uses to encode pictures. Normally, you arrive at this value by taking the number of bits per second you can move over a network and dividing that number by the number of pictures you’re encoding each second. For example, if you have 2 ISDN channels available for transporting data (2 * 65536 == 131,072 bits per second) and you want to encode 30 pictures a second, you would set `BITS_PER_IMAGE` to 4369. This setting would cause the compressor to encode a QCIF picture using about 0.17 bits per pixel.

The code below shows the bits-per-image attribute being set to 4369.

```
XilCis cis;
int bits_per_image = 4369;

xil_cis_set_attribute(cis, “COMPRESSOR_BITS_PER_IMAGE”,
    (void *)bits_per_image);
```

`COMPRESSOR_BITS_PER_IMAGE` can be set to any `int` greater than or equal to 0. Its default value is 5069, the number of bits needed to encode a QCIF picture at 0.2 bits per pixel.

You can read the value of this attribute using code similar to that shown below.

```
XilCis cis;
int bits_per_image;

xil_cis_get_attribute(cis, "COMPRESSOR_BITS_PER_IMAGE",
    (void **)&bits_per_image);
```

COMPRESSOR_DOC_CAMERA

The setting of this attribute determines whether the encoder sets the document-camera bit in the picture header of the pictures it encodes. If the attribute is set to `FALSE`, the encoder does not set this bit, and if it is set to `TRUE`, the encoder does set the bit. By default, the attribute is set to `FALSE`.

The code fragment below show the attribute being set to `TRUE`.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSOR_DOC_CAMERA",
    (void *)TRUE);
```

You can also read the value of this attribute using code similar to the following.

```
XilCis cis;
Xil_boolean doc_camera_status;

xil_cis_get_attribute(cis, "COMPRESSOR_DOC_CAMERA",
    (void **)&doc_camera_status);
```

COMPRESSOR_ENCODE_INTRA

When this attribute is set, a call to `xil_compress()` causes the encoder to compress a whole picture in *intraframe* mode. The encoder must compress this picture without causing its output buffer to overflow.

In general, this attribute should be set to `FALSE`, its default value. An application should set the attribute to `TRUE` only upon receiving a fast-update request and should set the attribute back to `FALSE` after compressing one frame.

The code below shows this attribute being set to `TRUE`.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSOR_ENCODE_INTRA",
    (void *)TRUE);
```

You can read the value of this attribute using code similar to that shown below.

```
XilCis cis;
Xil_boolean encode_intra_status;

xil_cis_get_attribute(cis, "COMPRESSOR_ENCODE_INTRA",
    (void **)&encode_intra_status);
```

COMPRESSOR_FREEZE_RELEASE

As was mentioned in the section "Provisions for Multipoint Conferencing" on page 346, an H.261 decoder can, in response to an external signal, freeze the picture currently being displayed. This freeze remains in effect until a timeout period has elapsed or until the decoder sees a freeze-release bit set in the bitstream.

Setting the encoder's `COMPRESSOR_FREEZE_RELEASE` attribute to `TRUE` causes the encoder to set the freeze-release bit in the header of the subsequent pictures it encodes. If a decoder has been in freeze-display mode, it will exit that mode when it reads the freeze-release bit. By default, `COMPRESSOR_FREEZE_RELEASE` is set to `FALSE`.

The code below shows FREEZE_RELEASE being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSOR_FREEZE_RELEASE",
    (void *)TRUE);
```

You can also read the value of this attribute, as is done in the code below.

```
XilCis cis;
Xil_boolean freeze_release_status;

xil_cis_get_attribute(cis, "COMPRESSOR_FREEZE_RELEASE",
    (void **)&freeze_release_status);
```

COMPRESSOR_IMAGE_SKIP

This attribute does not control the number of pictures the compressor skips between compressed pictures: the application controls that. However, if the application is skipping pictures, the attribute should be set to the number of pictures being skipped. This value figures in a calculation the compressor performs before filling in the 5-bit temporal-reference field in the header of each compressed picture. The compressor sets the value of this field to the value of the temporal-reference field in the last compressed picture plus the value of COMPRESSOR_IMAGE_SKIP plus 1.

The attribute can have a value in the range 0 to 31, and its default value is 0. The code below shows IMAGE_SKIP being set to 1. This would be the appropriate setting if your application were compressing every other frame of video input.

```
XilCis cis;
int images_skipped = 1;

xil_cis_set_attribute(cis, "COMPRESSOR_IMAGE_SKIP",
    (void *)images_skipped);
```

You can also read the value of this attribute using code similar to that shown below.

```
XilCis cis;
int images_skipped;

xil_cis_get_attribute(cis, “COMPRESSOR_IMAGE_SKIP”,
    (void **)&images_skipped);
```

COMPRESSOR_LOOP_FILTER

When an H.261 encoder is interframe encoding a macroblock and is using motion compensation, the encoder has the option of filtering the macroblock in the history image that is being used in the operation. (For more information about interframe encoding in H.261, see the section “Prediction” on page 342.) The setting of the COMPRESSOR_LOOP_FILTER attribute provides a hint to the encoder concerning whether it should filter the macroblock in the history image or not: a setting of TRUE suggests that the encoder perform the filtering if necessary, and a setting of FALSE suggests that it not use the filter.

The default value for COMPRESSOR_LOOP_FILTER is TRUE. This setting gives the encoder the flexibility to produce the most compact bitstream. Use the FALSE setting when you need to minimize compression and decompression time.

The code below shows COMPRESSOR_LOOP_FILTER being set to FALSE.

```
XilCis cis;

xil_cis_set_attribute(cis, “COMPRESSOR_LOOP_FILTER”,
    (void *)FALSE);
```

You can also read the value of this attribute. Use code similar to that shown below.

```
XilCis cis;
Xil_boolean loop_filter_status;

xil_cis_get_attribute(cis, "COMPRESSOR_LOOP_FILTER",
    (void **)&loop_filter_status);
```

COMPRESSOR_MV_SEARCH_RANGE

As discussed in the section "Prediction" on page 342, when an H.261 encoder is interframe encoding a macroblock, it may use a technique called motion compensation in determining which macroblock in the history image to use for the operation. This macroblock in the history image may be offset by up to ±15 pixels in both the *x* and *y* directions from the macroblock in the history image that matches exactly the position of the macroblock that is being encoded. This means that the encoder may have to look at the contents of 961 (31 * 31) macroblocks to find the best match for the macroblock it is encoding. The `COMPRESSOR_MV_SEARCH_RANGE` attribute enables you limit the extent of this search and, consequently, to speed up the encoding process.

Before setting the value of this attribute, you must declare and assign values to the members of an `XilH261MVSearchRange` structure.

```
typedef struct {
    int x;
    int y;
} XilH261MVSearchRange;
```

The value of `x` determines the extent of the search horizontally; that is, if `x` is set to 5, the encoder can only search five pixels to the right or left. The value of `y` determines the extent of the search vertically. Setting both values to 0 means that the encoder should not perform motion compensation at all.

Note – The value of this attribute is actually only a hint to the encoder and could be ignored.

The code below shows COMPRESSOR_MV_SEARCH_RANGE being set so that the encoder will search only 7 pixels horizontally and vertically.

```
XilCis cis;
XilH261MVSearchRange search_range;

search_range.x = 7;
search_range.y = 7;
xil_cis_set_attribute(cis, "COMPRESSOR_MV_SEARCH_RANGE",
    (void *)&search_range);
```

You can also read the value of this attribute. See the code fragment below.

```
XilCis cis;
XilH261MVSearchRange search_range;

xil_cis_get_attribute(cis, "COMPRESSOR_MV_SEARCH_RANGE",
    (void **)&search_range);
```

COMPRESSOR_SPLIT_SCREEN

The setting of this attribute determines whether the encoder sets the split-screen bit in the picture header of the pictures it encodes. If the attribute is set to FALSE, the encoder does not set this bit, and if it is set to TRUE, the encoder does set the bit. By default, the attribute is set to FALSE.

The code fragment below shows the attribute being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "COMPRESSOR_SPLIT_SCREEN",
    (void *)TRUE);
```

You can also read the value of this attribute using code similar to the following.

```
XilCis cis;
Xil_boolean split_screen_status;

xil_cis_get_attribute(cis, "COMPRESSOR_SPLIT_SCREEN",
    (void **)&split_screen_status);
```

Decompression Attributes

Of the decompression attributes, only one, `IGNORE_HISTORY`, affects the behavior of the decoder. The remaining attributes simply enable you to read values from the header of the most recently decompressed picture.

IGNORE_HISTORY

The `IGNORE_HISTORY` attribute affects your ability to seek forward and backward in an H.261 bitstream. These seeks are somewhat problematic in H.261 because the codec relies so heavily on interframe encoding and does not require periodic key frames.

By default, `IGNORE_HISTORY` is set to `FALSE`. In this case, backward seeks are illegal because there is no way to seek backward and still have a valid history frame to use in decoding a picture. Forward seeks are possible, but to ensure that it can decode the frame you seek to properly, the decoder must actually decode all the frames you "skip." This is the only way the decoder can maintain the correct data in its history buffer.

If you set `IGNORE_HISTORY` to `TRUE`, you're telling the decoder that it's OK if, after a seek, it does not have the correct data in its history buffer to decode the next picture. In this case, backward seeks are legal, and forward seeks are faster than they would be otherwise. The drawback to ignoring the history image is that you may have to decode as many as 132 frames after a seek before you get a properly reconstructed picture.

Note – There is one exception to what was said in the preceding paragraph. If `IGNORE_HISTORY` is set to `TRUE` and you seek backward by one frame and then call `xil_decompress()`, you will get a good picture. The decompressor can copy this picture from its history buffer.

The code below shows IGNORE_HISTORY being set to TRUE.

```
XilCis cis;

xil_cis_set_attribute(cis, "IGNORE_HISTORY", (void *)TRUE);
```

You can also read the value of this attribute. To do this, use code similar to that shown below.

```
XilCis cis;
Xil_boolean ignore_history_status;

xil_cis_get_attribute(cis, "IGNORE_HISTORY",
    (void **)&ignore_history_status);
```

DECOMPRESSOR_DOC_CAMERA

Reading this attribute tells you whether the document-camera bit was set in the header of the most recently decompressed picture. If the value of the attribute is TRUE, the bit was set, and if the value is FALSE, it was not set.

The code fragment below illustrates how to read the attribute.

```
XilCis cis;
Xil_boolean doc_camera_status;

xil_cis_get_attribute(cis, "DECOMPRESSOR_DOC_CAMERA",
    (void **)&doc_camera_status);
```

DECOMPRESSOR_FREEZE_RELEASE

This attribute enables you to read the value of the freeze-release bit for the most recently decompressed picture. If the bit was set, the value of the attribute will be TRUE; otherwise, it will be FALSE.

Here's how an application uses this information: If the application has frozen the display in response to a freeze-picture request and DECOMPRESSOR_FREEZE_RELEASE is set to FALSE, the application does not

display the picture it just decoded (unless a timeout period has elapsed). On the other hand, if the attribute is set to TRUE, it does display that picture, and all ensuing pictures until it receives another freeze-picture request.

The code below shows how to read the value of DECOMPRESSOR_FREEZE_RELEASE.

```
XilCis cis;
Xil_boolean stop_freeze;

xil_cis_get_attribute(cis, "DECOMPRESSOR_FREEZE_RELEASE",
    (void **)&stop_freeze);
```

DECOMPRESSOR_SOURCE_FORMAT

This attribute enables you to determine whether the most recently decompressed picture was in CIF (352 by 288 pixels) or QCIF (176 by 144 pixels) format. The value of the attribute will be one of the enumeration constants shown below.

```
typedef enum {
    QCIF, CIF
} XilH261SourceFormat;
```

The code below shows how you might read the value of DECOMPRESSOR_SOURCE_FORMAT.

```
XilCis cis;
XilH261SourceFormat source_format;

xil_cis_get_attribute(cis, "DECOMPRESSOR_SOURCE_FORMAT",
    (void **)&source_format);
```

DECOMPRESSOR_SPLIT_SCREEN

Reading this attribute tells you whether the split-screen bit was set in the header of the most recently decompressed picture. If the value of the attribute is TRUE, the bit was set, and if the value is FALSE, it was not set.

The code fragment below illustrates how to read the attribute.

```
XilCis cis;
Xil_boolean split_screen_status;

xil_cis_get_attribute(cis, "DECOMPRESSOR_SPLIT_SCREEN",
    (void **)&split_screen_status);
```

DECOMPRESSOR_TEMPORAL_REFERENCE

This attribute enables you to retrieve the value of the temporal-reference field in the header of the last picture decoded. Since this is a 5-bit field, the value will be in the range 0 to 31.

By looking at the temporal reference fields of the last two decompressed frames, an application can determine how many pictures the encoder skipped between these pictures. In most instances, the number of skipped pictures equals

temporal reference for picture_x - temporal reference for picture_{x-1} - 1

However, if the temporal reference value for picture_x is less than the temporal reference for picture_{x-1}, the calculation becomes

31 + temporal reference for picture_x - temporal reference for picture_{x-1}

The code below shows how to read the value of `DECOMPRESSOR_TEMPORAL_REFERENCE`.

```
XilCis cis;
int temp_ref;

xil_cis_get_attribute(cis, "DECOMPRESSOR_TEMPORAL_REFERENCE",
    (void **)&temp_ref);
```

18

H.261 Molecules

The XIL library includes a series of *molecules* that accelerate the playback of H.261 bitstreams. These molecules are optimized routines that perform the jobs of two or more functions from the XIL API. You do not call such an optimized routine directly; rather, the library calls a molecule when your program calls a predefined sequence of XIL functions, sometimes with specific arguments.

For example, if your program calls `xil_decompress()` to decode an image stored in an H.261 CIS and then calls `xil_ordered_dither()` to dither the decoded image from a 24- to an 8-bit image, the library may not call these two functions. Instead, it may call a molecule that performs the decompression and the dithering in an optimized way.

For information about the H.261 molecules that are available and information about how to call those molecules, see the section "XIL Molecules" on page 398.

MPEG-1 Codec 19

This chapter discusses the MPEG-1 video codec specified by the Moving Picture Experts Group, an ISO working group. This group has produced a standard that is similar to the H.261 standard developed by the CCITT (see Chapter 18, “H.261 Codec”), but places less emphasis on low bit rates. By accepting a higher bit rate—for example, 1.5 Mbits per second—an MPEG-1 codec is able to recreate very high-quality pictures and to produce a bitstream that is easily editable.

The rate of 1.5 Mbits/s makes the MPEG-1 codec especially viable in applications that read compressed data from CD-ROMs because even older CD-ROM readers can read data at this speed. Thus, putting an MPEG-1 bitstream on a CD-ROM is an effective way to distribute movies, business presentations, and training videos.

The remainder of this chapter is divided into four sections. The first section provides an overview of how an MPEG-1 codec works. The second explains briefly how to create an MPEG-1 CIS. The third discusses CIS attributes that apply specifically to a CIS associated with an XIL MPEG-1 compressor or decompressor (as opposed to the general CIS attributes covered in the section “General CIS Attributes” on page 257). And the fourth section introduces the subject of accelerating the playback of MPEG-1 bitstreams. For further information on this subject, see Chapter 21, “Acceleration in XIL Programs.”

Note – This chapter discusses both MPEG-1 compression and decompression. However, the current release of the XIL library includes only an MPEG-1 decompressor. The compressor interface is defined for third parties who want to implement XIL MPEG-1 compressors.

How an MPEG-1 Codec Works

Since the work done by the Motion Picture Experts Group grew out of the work done by the CCITT in developing the H.261 video codec, this section explains how the MPEG-1 codec works by comparing and contrasting it with the H.261 codec. (If you're unfamiliar with how an H.261 codec works, see the section "How an H.261 Codec Works" on page 338.) The present section also includes a subsection that describes the organizational structure that MPEG-1 imposes on a video sequence.

Similarities Between MPEG-1 and H.261

The key similarities between MPEG-1 and H.261 are listed below:

- Both compressors work with YC_bC_r pictures in which the information in the chroma channels has been subsampled so that there is one C_b and one C_r value for each 2-by-2 block of luma values.
- Like the H.261 compressor, the MPEG-1 compressor can compress a macroblock (a 16-by-16 block of pixels in a picture) by encoding the actual Y, C_b, and C_r values in the macroblock (intraframe encoding) or by encoding the differences between values in the current block and values in the corresponding macroblock in the previous picture (forward prediction). In addition, when using the forward-prediction encoding method, the MPEG-1 compressor, like the H.261 compressor, can employ motion compensation.
- Both compressors encode 8-by-8 blocks of pixel values or difference values using the same method. The compressors first perform a Discrete Cosine Transform (DCT) on the 8-by-8 block of values. This operation transforms the values in the block from the spatial to the transform domain. Second, the compressors quantize the coefficients produced by the DCT. Finally, the compressors use entropy coding to encode the quantized coefficients.

Differences Between MPEG-1 and H.261

Although the Moving Picture Experts Group drew heavily on the work of the CCITT, there are also very significant differences between the MPEG-1 and H.261 compressors. Most of these differences provide for random access to the MPEG-1 bitstream and make the bitstream easily editable.

I Pictures and P Pictures

In H.261, the ideas of intraframe encoding and predictive encoding are applied for the most part at the macroblock level. MPEG-1, on the other hand, includes the notion of intraframe-encoded pictures and predicted pictures.

In an intraframe-encoded picture, or I picture, all macroblocks are intraframe encoded. This, of course, means that the decoder needs no information from a preceding picture to decode an I picture. For this reason, seeks to I frames can be performed very quickly.

In a predicted picture, or P picture, each macroblock can be intraframe encoded or encoded using the forward-prediction method. This type of picture is very similar to an H.261 picture.

Typically, an MPEG-1 bitstream will contain more P pictures than I pictures because the P picture can be encoded using fewer bits. Encoding difference values generally requires fewer bits than encoding pixel values. Also, the encoder does not have to encode macroblocks that are very similar to their counterparts in the preceding I or P picture.

B Pictures

From what we've said so far, you might picture an MPEG-1 bitstream as containing periodic I pictures followed by a number of P pictures. See Figure 19-1.

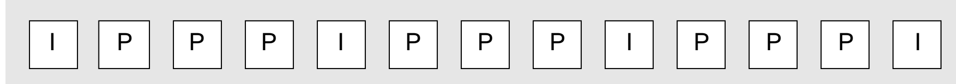

Figure 19-1 An MPEG-1 Bitstream Containing I and P Pictures

This does constitute a legal MPEG-1 bitstream. However, the bitstream may also contain one or more bidirectionally predicted pictures, or B pictures, between any pair of I or P pictures. In a B picture, each macroblock may be

- Intraframe encoded
- Forward predicted from the nearest preceding I or P picture
- Backward predicted from the nearest succeeding I or P picture
- Bidirectionally predicted from the nearest preceding I or P picture and the nearest succeeding I or P picture

We've already discussed intraframe encoding and forward prediction.

Backward prediction is strictly analogous to forward prediction. A backward-predicted macroblock is encoded with respect to the values in the corresponding macroblock in a picture that follows its own picture in the video sequence. This option may seem unintuitive at first thought because it results in a situation where pictures are not transmitted in the order in which they will be displayed. See Figure 19-2.

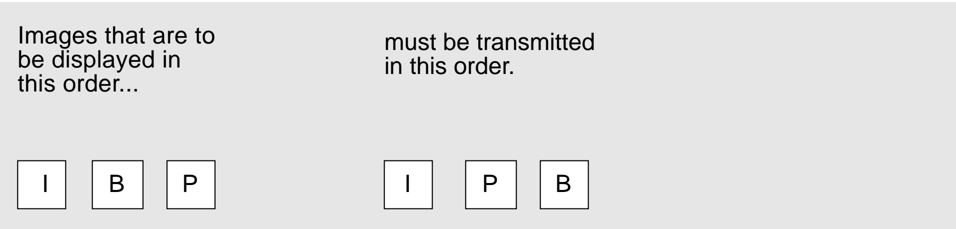

Figure 19-2 MPEG-1 Display Order Versus Decoding Order

However, backward prediction is an option because it can lead to a more compact bitstream in instances where the macroblocks in a B picture match their counterparts in the closest succeeding I or P picture more closely than they match their counterparts in the closest preceding I or P picture.

A bidirectionally predicted macroblock is coded with respect to the corresponding macroblock in both the closest preceding and succeeding I/P pictures. See Figure 19-3.

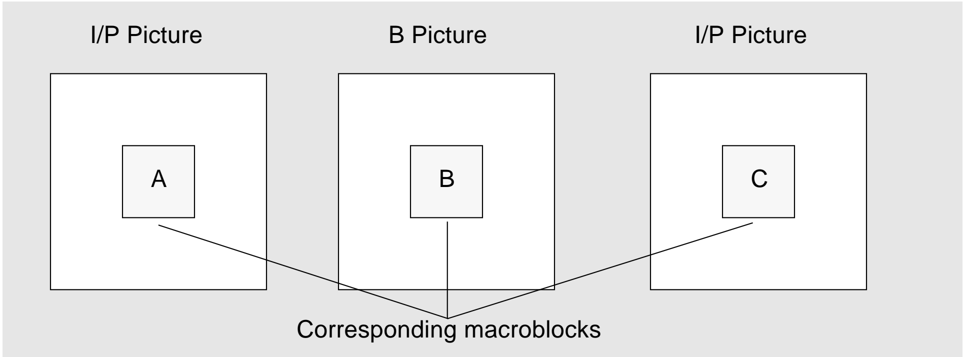

Figure 19-3 Bidirectional Prediction in MPEG-1

In this case, the encoder averages the values in macroblocks A and C and then encodes the difference between values in macroblock B and the averaged macroblock. Bidirectional prediction is effective because the difference mentioned above is often so small that it does not need to be encoded.

For MPEG-1, the behavior of the `xil_cis_get_bits_ptr()` function differs from its usual behavior. For a bitstream with out-of-order frames (that is, a bitstream with B frames), the actual number of frames in the data returned by `xil_cis_get_bits_ptr()` may not equal the value of its `nframes` parameter. The value of `nframes` is the number of frames the CIS read frame has been advanced by the `xil_cis_get_bits_ptr()` call. A seek back by `nframes` by calling `xil_cis_seek()` will restore the read position to the original read frame (before the `xil_cis_get_bits_ptr()` call). This is useful for an application that uses a preview mode, where the compress and write-to-file is followed by a decompress.

Since the number of frames reported may not represent the actual number of frames, if the chunk is subsequently used for an `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()` call, the `nframes` parameter must be set to -1, which indicates an unknown number of frames.

Groups of Pictures

Another concept used in MPEG-1 that does not apply in H.261 is that of a group of pictures. This is a series of consecutive pictures from a video sequence. Generally, the group of pictures provides a unit of the bitstream that can be removed without destroying the integrity of the bitstream and a unit

that can be decoded independently of the rest of the bitstream. As you'll see in a moment, however, the group of pictures does not always have these characteristics.

By definition, a group of pictures (considered in display order) must begin with an I picture or with one or more B pictures followed by an I picture. It must end with an I or a P picture. A closed group of pictures can be decoded without any information from the preceding group of pictures. Thus, a closed group is one that begins with an I picture or one that begins with one or more B pictures whose macroblocks have been encoded using only intraframe encoding and backward prediction. An open group of pictures begins with one or more B pictures, at least one of which contains macroblocks encoded using forward or bidirectional prediction. This type of group can be decoded only if the preceding group of pictures is available. If that group is not available—for example, if an MPEG-1 bitstream editing program has removed the group—a broken-link bit must be set in the header for the open group of pictures.

How MPEG-1 Organizes a Video Sequence

We've already mentioned some of the organizational units used in MPEG-1, such as the group of pictures and the macroblock. This section provides a complete overview of these units. This information should be useful to you as you read about MPEG-1 CIS attributes later in this chapter because many of these attributes are associated with a particular unit.

The largest unit MPEG-1 defines is the video sequence. You might think of this unit as an entire movie or presentation. Each MPEG-1 bitstream includes a sequence header, which defines several attributes for the entire sequence. These include the pixel aspect ratio for the sequence, the picture rate for the sequence in pictures per second, and the bit rate of the data channel over which the compressed sequence will be moved. An additional sequence header is written to the bitstream each time one of the encoder's quantization tables is changed.

Each sequence is divided into a series of groups of pictures. See the section "Groups of Pictures" on page 363 for a definition of this unit. If you're compressing video using an XIL-compliant MPEG-1 encoder, you can control the makeup of each group of pictures using the attribute `COMPRESSOR_PATTERN`.

Each group of pictures consists of pictures, which are individual frames of video. The major attributes associated with each picture are a picture type and a temporal reference. The major picture types are intraframe-encoded pictures, forward-predicted pictures, and bidirectionally predicted pictures. The temporal reference is an integer identifying a picture's place within a group of pictures.

The largest subdivision of a picture is called a slice, which consists of a series of consecutive macroblocks. Slices within a picture may vary in size, but each macroblock in a picture must be part of a slice. This unit is designed primarily to help a decoder recover from a bitstream error. If a decoder detects an error, one way to recover is to skip to the next slice header.

Slices, as mentioned in the preceding paragraph, are built from macroblocks. These are 16-by-16 block of pixels. The macroblock is the level at which motion compensation is performed. Also, the type of encoding used—intraframe, forward prediction, backward prediction, or bidirectional prediction—can change from macroblock to macroblock.

Finally, each macroblock is divided into six 8-by-8 blocks. Four of these blocks contain luma values, one contains C_b values, and one contains C_r values. This is the level at which the DCT is performed.

Creating an MPEG-1 CIS

Before you can use the MPEG-1 decompressor to decompress an MPEG-1 bitstream, you must create an MPEG-1 CIS (and write an MPEG-1 bitstream to the CIS). You create this CIS by passing the decompressor name `Mpeg1` to the function `xil_cis_create()`. See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, "Mpeg1");
```

MPEG-1 Codec Attributes

As discussed in the section "General CIS Attributes" on page 257, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for CISs attached to an MPEG-1 compressor or

decompressor. You set codec-specific attributes using the function `xil_cis_set_attribute()`, and you read codec-specific attributes using the function `xil_cis_get_attribute()`.

The MPEG-1 attributes can be broadly grouped into those that affect compression and those that affect decompression. The attributes are discussed under these headings below.

Compression Attributes

Setting any of the following attributes affects how an MPEG-1 compressor compresses images. It is also possible to read the value of each of these attributes.

COMPRESSOR_BITS_PER_SECOND

You use this attribute to tell the encoder how many bits it can use to encode one second's worth of pictures. This value is written to the sequence header for a picture sequence and is assumed to remain constant for the duration of the sequence. You cannot change the value of this attribute after the encoder compresses the first picture in the sequence.

The value you assign to the attribute must fall in the range 1 to 104,856,800. Before inserting this value in the sequence header, the encoder rounds it up to the nearest multiple of 400. The default value of the attribute is 1,152,000.

The code below shows how you might set the value of `COMPRESSOR_BITS_PER_SECOND` to 1,856,000.

```
XilCis cis;
int bits_per_second = 1856000;

xil_cis_set_attribute(cis, "COMPRESSOR_BITS_PER_SECOND",
    (void *)bits_per_second);
```

To read the value of the attribute, you might use code similar to this.

```
XilCis cis;
int bits_per_second;

xil_cis_get_attribute(cis, "COMPRESSOR_BITS_PER_SECOND",
    (void **)&bits_per_second);
```

COMPRESSOR_INSERT_VIDEO_SEQUENCE_END

Normally, the value of this attribute is `FALSE`. If you set it to `TRUE`, any time a subsequent call is made to `xil_cis_flush()`, the encoder writes a video-sequence-end (*eos*) code to the last frame in the CIS. This action is added to the normal actions taken by the flush routine. When set to `FALSE`, this attribute doesn't affect the normal actions of the flush routine.

The code below shows the attribute being set to `TRUE`.

```
XilCis cis;

xil_cis_set_attribute(cis,
    "COMPRESSOR_INSERT_VIDEO_SEQUENCE_END", (void *)TRUE);
```

Note – The library prevents multiple *eos* codes from being written to the same frame. Thus, when the write-frame's number doesn't change and `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` is `TRUE`, multiple calls to `xil_cis_flush()` result in only one *eos* code in the frame; this code is written after the first call to `xil_cis_flush()`.

You can read the value of this attribute using code similar to the following:

```
XilCis cis;
Xil_boolean sequence_end_status;

xil_cis_get_attribute(cis,
    "COMPRESSOR_INSERT_VIDEO_SEQUENCE_END",
    (void **)&sequence_end_status);
```

Table 19-1 shows two command sequences that result in the same CIS pattern; these sequences assume the call to `xil_cis_set_attribute()` sets the `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` attribute to `TRUE`.

Note – On this page and the following pages, the term *CIS pattern* refers to the order that frames appear in the bitstream. Don't confuse this with the `COMPRESSOR_PATTERN` attribute (see page 376), which lets you specify the makeup of a group of pictures. For example, the `COMPRESSOR_PATTERN` IBBPBBP will result in the bitstream CIS pattern IPBBPBB.

Table 19-1 Two Command Sequences: CIS Pattern = IPBB

Command Sequence 1	Command Sequence 2
xil_compress	xil_cis_set_attribute
xil_compress	xil_compress
xil_compress	xil_compress
xil_compress	xil_compress
xil_cis_flush /* normal */	xil_compress
xil_cis_set_attribute	xil_cis_flush /* normal + eos */
xil_cis_flush /* normal + eos */	

For the two sequences shown in Table 19-1, the resulting CIS sequence is:

```
vsh.I0.P3.B1.B2.eos
```

where *vsh* represents the video sequence header information, and *eos* represents the end-of-sequence code. Between the *vsh* and the *eos* code are the frame types and their display id's; among the frame types, I = interframe encoding, P = predicted pictures, and B = bidirectionally predicted pictures.

From the XIL point of view, the *vsh* is bundled with the first frame, *I0*, and the *eos* is bundled with the last frame, *B2*.

There may be multiple *vsh* components associated with one *eos*, since the *vsh* changes as certain CIS attributes change (within the XIL limitations that there are no width/height changes).

In addition, there may be multiple sequences within a bitstream; thus,

```
vsh-----eos.vsh----eos.vsh----eos
```

constitutes three sequences.

If frames are compressed into the CIS after the call to `xil_cis_flush()`, it's the compressor's responsibility to provide the *vsh* per sequence. Before the application changes an attribute that would result in a new sequence header, it must first output an *eos* for the current sequence by calling `xil_cis_flush()` with `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END = TRUE`. For example, before you can change the attribute `COMPRESSOR_BITS_PER_SECOND` (see page 366), you must first set `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` to `TRUE` and call `xil_cis_flush()` to write the *eos* code to the bitstream.

An MPEG-1 sequence isn't valid without the *eos* code; therefore, the last frame in the sequence must be followed by the *eos* code. Since it cannot be predicted when an application will end a sequence, the MPEG-1 codec *reserves* the last frame or subgroup of frames in the CIS so that you can write an *eos* to that frame or subgroup. The reserved frame or subgroup must be released before it can be retrieved with `xil_cis_get_bits_ptr()` or `xil_decompress()`.

Note – The reserved frame or subgroup affects the logic you need for calling `xil_cis_get_bits_ptr()` and `xil_decompress()`. For example, typically you call `xil_decompress()` from within a loop that executes only when `xil_cis_has_frame()` returns `TRUE`. When decompressing an MPEG-1 CIS, you must first release the reserved frames before you can expect a loop control condition that depends on the return value of `xil_cis_has_frame()` to evaluate to `TRUE`.

The following paragraphs tell you how to release the reserved frames. Beginning on page 371, the discussion moves to some example function-call sequences and shows how these sequences would affect both the return value of `xil_cis_has_frame()` and `xil_cis_get_bits_ptr()`.

A frame or subgroup is released when:

- It is followed by an *eos*, thus providing a valid sequence. This is done by setting the `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` attribute to `TRUE` and then calling `xil_cis_flush()`, as described above.
- It is followed by another frame or subgroup. For an all I bitstream or a mixed I-P bitstream, this just means that another frame is added to the CIS. Since there are no out-of-order frames, a subgroup for this type of CIS is just a frame.

The following sequences release a frame:

```
I0.I1          I1 releases I0
I0.P1          P1 releases I0
```

For any bitstream with B frames, releasing a frame or subgroup is more complicated. B frames have a future-predictive frame that appears in the bitstream *previous* to the B frame, and multiple B frames may share the same predictive frame. So you can get the following (in these examples, any of the I's can be replaced by a P):

```
I2.B1
I3.B1.B2
I4.B1.B2.B3
I6.B1.B2.B3.B4.B5
I9.B1.B2.B3.B4.B5.B6.B7.B8
```

A sequence is considered a complete *subgroup* when its predictive frame has the display id N, and the predictive frame is followed by one or more B frames, the last of whose display id is N-1. Thus, all of the above bitstream fragments are complete subgroups.

The following are examples of bitstreams with incomplete subgroups:

```
I3.B1          missing B2
I6.B1.B2.B3    missing B4,B5
```

When there are incomplete subgroups, the compressor must buffer the compressed frames until the buffer contains a complete subgroup, at which point the compressor adds the completed subgroup to the bitstream. In this manner, the compressor can handle interruptions to the compression-frame sequence as necessary. For example, if, halfway through the IPBBB group, the video player's Stop/Output button is pressed, the compressor must be able to take the resulting I0.P4.B1 bitstream and rework it into a legal output form.

Once the subgroup is complete, it cannot be released until it is followed by an *eos*, or it is followed by another valid subgroup. The following sequences show a subgroup being released:

```
I0.I2.B1.eos             I2.B1 released by eos
I0.I2.B1.I4.B3           I2.B1 released by I4.B3
I0.I3.B1.B2.P5.B4        I3.B1.B2 released by P5.B4
I0.I3.B1.B2.P5.B4.eos    P5.B4 released by eos
```

The tables below show how subgroups are affected by a particular sequence of function calls. In each of these tables:

- Column heading **WF** is an abbreviation for Write Frame
- Column heading **has_frame** is an abbreviation for the `xil_cis_has_frame()` function and the column text shows what the return value for that function would be
- Column heading **get_bits_ptr** is an abbreviation for the `xil_cis_get_bits_ptr()` function and the column text shows the frames that would be returned by a call to that function
- Column heading **RF** is an abbreviation for Read Frame and the column text shows what the read frame would be after the call to `xil_cis_get_bits_ptr()`
- A call to `xil_cis_set_attribute()` in the sequence sets the `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` attribute to `TRUE`

In Table 19-2, the CIS pattern is all I frames and the result is that subgroup I2 is released by an appended *eos*. Notice that the first `xil_cis_flush()` call has no affect on the subgroup since the default setting for the `COMPRESSOR_INSERT_VIDEO_SEQUENCE_END` attribute is `FALSE`.

Table 19-2 Releasing a Frame: CIS Pattern = All I Frames

Function	WF	has_frame	get_bits_ptr	RF	CIS
xil_compress	1	FALSE	NULL	0	I0
xil_compress	2	TRUE	I0	1	I0.I1
xil_compress	3	TRUE	I1	2	I0.I1.I2
xil_cis_flush	3	FALSE	NULL	2	I0.I1.I2
xil_cis_set_attribute					
xil_cis_flush	3	TRUE	I2.eos	3	I0.I1.I2.eos

In Table 19-3, the CIS pattern is IPB and the result is that subgroup P2.B1 is released by the appended *eos*.

Table 19-3 Releasing a Frame: CIS Pattern = IPB

Function	WF	has_frame	get_bits_ptr	RF	CIS
xil_compress	1	FALSE	NULL	0	I0
xil_compress	2	FALSE	NULL	0	I0
xil_compress	3	TRUE	I0	1	I0.P2.B1
xil_cis_set_attribute					
xil_cis_flush	3	TRUE	P2.B1.eos	3	I0.P2.B1.eos

Finally, in Table 19-4, the CIS pattern is IPBBPBB and the result is that subgroup P3.B1.B2 is released by the next subgroup, P6.B4.B5.

Table 19-4 Releasing a Frame: CIS Pattern = IPBBPBB

Function	WF	has_frame	get_bits_ptr	RF	CIS
xil_compress	1	FALSE	NULL	0	I0
xil_compress	2	FALSE	NULL	0	I0
xil_compress	3	FALSE	NULL	0	I0
xil_compress	4	TRUE	I0	1	I0.P3.B1.B2
xil_compress	5	FALSE	NULL	1	I0.P3.B1.B2
xil_compress	6	FALSE	NULL	1	I0.P3.B1.B2
xil_compress	7	TRUE	P3.B1.B2	4	I0.P3.B1.B2.P6.B4.B5

COMPRESSOR_INTRA_QUANTIZATION_TABLE

This attribute enables you to specify an 8-by-8 matrix of values to be used in quantizing the DCT coefficients for blocks in intraframe-encoded macroblocks. If you set this attribute, the table will be written to your bitstream's sequence header. If you don't set it, the encoder will use a default implementation-specific table.

If you supply your own table, it must meet these requirements:

- The first value in the table must be an 8, the fixed quantization level of the DC coefficient for the block.
- The remaining values must fall in the range 1 to 255.
- The values must be listed in the zigzag order shown in Figure 19-4.

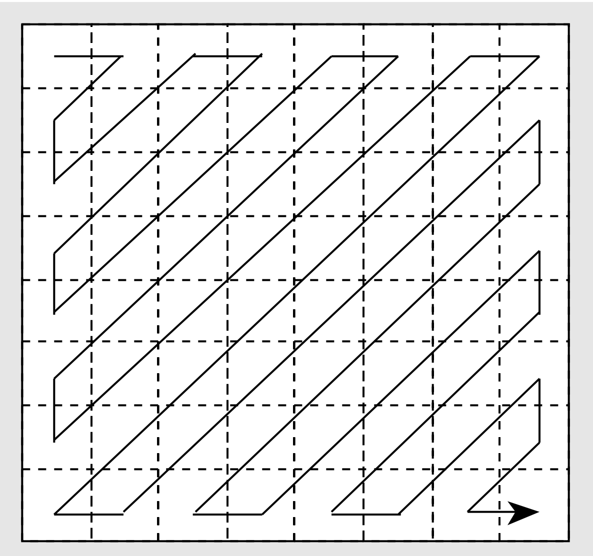

Figure 19-4 Zigzag Ordering of Quantization Table Values

The code below shows how to specify a table. This particular table is the one that is used as an example in the MPEG-1 standard and may well be the encoder's default table.

```
XilCis cis;
Xil_unsigned8 intra_quant_table[64] = {
    8, 16, 19, 22, 26, 27, 29, 34,
    16, 16, 22, 24, 27, 29, 34, 37,
    19, 22, 26, 27, 29, 34, 34, 38,
    22, 22, 26, 27, 29, 34, 37, 40,
    22, 26, 27, 29, 32, 35, 40, 48,
    26, 27, 29, 32, 35, 40, 48, 58,
    26, 27, 29, 34, 38, 46, 56, 69,
    27, 29, 35, 38, 46, 56, 69, 83 };

xil_cis_set_attribute(cis,
    "COMPRESSOR_INTRA_QUANTIZATION_TABLE",
    (void *)intra_quant_table);
```

You can read the value of the attribute using code similar to this.

```
XilCis cis;
Xil_unsigned8 *intra_quant_table;

xil_cis_get_attribute(cis,
    "COMPRESSOR_INTRA_QUANTIZATION_TABLE",
    (void **)&intra_quant_table);
```

If, after this call, `intra_quant_table` is a null pointer, the default table is in use; otherwise, `intra_quant_table` is a pointer to a user-supplied quantization matrix. In the latter case, your application must free the matrix pointed to by `intra_quant_table` when it is no longer needed.

COMPRESSOR_NON_INTRA_QUANTIZATION_TABLE

This attribute enables you to specify an 8-by-8 matrix of values to be used in quantizing the DCT coefficients for blocks in nonintraframe-encoded macroblocks. If you set this attribute, the table will be written to your bitstream's sequence header. If you don't set it, the encoder will use an implementation-specific default table.

If you supply an alternate table, it must meet these requirements:

- The values in the table must fall in the range 1 to 255.
- The values must be listed in the zigzag order shown in Figure 19-4 on page 373.

The code below is an example of how you might supply a custom table. This table, which contains all 16's, is used as an example in the MPEG-1 specification.

```
XilCis cis;
Xil_unsigned8 non_intra_quant_table[64];
int i;


for (i = 0; i < 64; i++)
    non_intra_quant_table[i] = 16;

xil_cis_set_attribute(cis,
    "COMPRESSOR_NON_INTRA_QUANTIZATION_TABLE",
    (void *)non_intra_quant_table);
```

You can read the value of the attribute using code similar to this.

```
XilCis cis;
Xil_unsigned8 *non_intra_quant_table;

xil_cis_get_attribute(cis,
    "COMPRESSOR_NON_INTRA_QUANTIZATION_TABLE",
    (void **)&non_intra_quant_table);
```

COMPRESSOR_PATTERN

This attribute enables you to specify the makeup of each group of pictures. You dictate the number of pictures in each group, the types of the pictures in the group, and the display order of those pictures. Your specification must conform to the following rules:

- A group of pictures may begin with one or more B pictures or an I picture.
- If the group starts with one or more B pictures, those pictures must be followed by an I picture.
- The first I picture in the group may be followed by any number of I and P pictures, and any combination of an I and a P picture may have intervening B pictures.
- The group of pictures must end with an I or a P picture.

Note – A group of pictures can also contain all D pictures. (A D picture contains only low-frequency information and is meant to be used for fast visible searches.) However, the XIL library's MPEG-1 decompressor cannot decode bitstreams containing D pictures.

You specify the makeup of the group by assigning values to a structure of type `XilMpeg1Pattern`:

```
typedef struct {
    char *pattern;
    Xil_unsigned32 repeat_count;
} XilMpeg1Pattern;
```

The string `pattern` determines the types of the pictures (in display order) that will appear at the beginning of the group of pictures. This string must contain a combination of the characters shown in Table 19-5.

Table 19-5 Characters Representing Picture Types

Character	Picture Type
I	Intraframe picture
P	Forward-predicted picture
B	Bidirectionally predicted picture
D	DC-coefficient picture

For example, suppose you want each group of pictures to begin with the thirteen pictures shown in Figure 19-5.

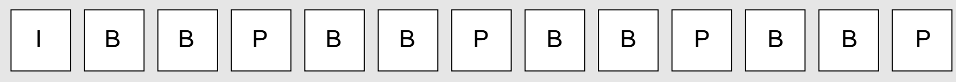

Figure 19-5 Sample Group of Pictures

You would set `pattern` to `"IBBPBBPBBPBBP"`.

The parameter `repeat_count` is a number greater than 0 that determines the total number of pictures in the group of pictures. For example, suppose that you set `repeat_count` to 3 as shown in the following example.

```
XilCis cis;
XilMpeg1Pattern picture_pattern = {"IBBPBBPBBPBBP", 3};

xil_cis_set_attribute(cis, "COMPRESSOR_PATTERN",
    (void *)&picture_pattern);
```

This call requests that the pattern shown in Figure 19-5 be repeated 3 times, so that there will be a total of 39 pictures in the group.

When you set this attribute, the encoder terminates the group of pictures it was working on immediately. Thus, the next call to `xil_compress()` causes the encoder to compress the first picture in a new group of the type you just specified.

If you do not set this attribute, a default implementation-specific pattern will be used. To select the default pattern after having selected a custom pattern, you pass a null pointer to `xil_cis_set_attribute()`.

You can also read the value of `COMPRESSOR_PATTERN`. To do this, you might use code similar to the following.

```
XilCis cis;
XilMpeg1Pattern *picture_pattern;

xil_cis_get_attribute(cis, "COMPRESSOR_PATTERN",
    (void **)&picture_pattern);
```

If, after this call, `picture_pattern` is a null pointer, the encoder is using the default picture pattern. Otherwise, `picture_pattern` points to a structure containing a pattern string and a repeat count. In the latter case, when you have finished using the structure, your application should free the storage that the XIL library allocated to hold the pattern string and then free the storage the library allocated to hold the structure.

COMPRESSOR_PEL_ASPECT_RATIO

The value of this attribute, along with the width and height of a picture in pixels, indicates the shape a picture is meant to have when it is displayed. Basically, the ratio of the width to the height of the displayed picture should be:

width in pixels / (height in pixels * aspect ratio)

If you supply an aspect ratio, this value is written to the sequence header for your bitstream and can be read later by an XIL-compliant MPEG-1 decompressor. The default value is implementation specific.

The valid values for COMPRESSOR_PEL_ASPECT_RATIO are contained in the enumeration shown below.

```
typedef enum {
    NullDefault,
    Ratio_1_0,      /* 1.0    */
    Ratio_0_6735,  /* 0.6735 */
    Ratio_0_7031,  /* 0.7031 */
    Ratio_0_7615,  /* 0.7615 */
    Ratio_0_8055,  /* 0.8055 */
    Ratio_0_8437,  /* 0.8437 */
    Ratio_0_8935,  /* 0.8935 */
    Ratio_0_9157,  /* 0.9157 */
    Ratio_0_9815,  /* 0.9815 */
    Ratio_1_0255,  /* 1.0255 */
    Ratio_1_0695,  /* 1.0695 */
    Ratio_1_0950,  /* 1.0950 */
    Ratio_1_1575,  /* 1.1575 */
    Ratio_1_2015   /* 1.2015 */
} XilMpeg1PelAspectRatio;
```

The most important values in the enumeration are Ratio_1_0, Ratio_0_9157, and Ratio_1_0950. The value Ratio_1_0 indicates that the pictures in the sequence can be shown without distortion on a computer monitor, which has square pixels. Ratio_0_9157 indicates that pictures can be displayed without distortion on CCIR Rec. 601 625-line televisions. And Ratio_1_0950 indicates that pictures can be displayed without distortion on CCIR Rec. 601 525-line televisions.

You can use the value NullDefault to reset the aspect ratio to its default value.

The code below shows COMPRESSOR_PEL_ASPECT_RATIO being set to Ratio_1_0.

```
XilCis cis;
XilMpeg1PelAspectRatio aspect_ratio = Ratio_1_0;

xil_cis_set_attribute(cis, "COMPRESSOR_PEL_ASPECT_RATIO",
    (void *)aspect_ratio);
```

You can read the current value of this attribute using code similar to that shown below.

```
XilCis cis;
XilMpeg1PelAspectRatio aspect_ratio;

xil_cis_get_attribute(cis, "COMPRESSOR_PEL_ASPECT_RATIO",
    (void **)&aspect_ratio);
```

If `aspect_ratio` equals `NullDefault` after this call, the aspect ratio for the sequence is the default aspect ratio defined by the implementors of the encoder.

COMPRESSOR_PICTURE_RATE

You use this attribute to specify the picture rate of the sequence being encoded in pictures per second. This value and the value of `COMPRESSOR_BITS_PER_SECOND` enable the encoder to determine how many bits it can use to encode each picture.

The valid values for `COMPRESSOR_PICTURE_RATE` are the enumeration constants shown below. These constants correspond to the rates of commonly available sources of analog and digital video.

```
typedef enum {
    NullDefault,
    Rate_23_976,   /* 23.976 */
    Rate_24,       /* 24.0   */
    Rate_25,       /* 25.0   */
    Rate_29_97,    /* 29.97  */
    Rate_30,       /* 30.0   */
    Rate_50,       /* 50.0   */
    Rate_59_94,    /* 59.94  */
    Rate_60        /* 60.0   */
} XilMpeg1PictureRate;
```

The default value of this attribute is implementation specific. The code below shows the value being set to `Rate_29_97`.

```
XilCis cis;
XilMpeg1PictureRate picture_rate = Rate_29_97;

xil_cis_set_attribute(cis, "COMPRESSOR_PICTURE_RATE",
    (void *)picture rate);
```

You can read the current value of this attribute using code similar to that shown below.

```
XilCis cis;
XilMpeg1PictureRate picture_rate;

xil_cis_get_attribute(cis, "COMPRESSOR_PICTURE_RATE",
    (void **)&picture_rate);
```

If `picture_rate` equals `NullDefault` after this call, the rate for the sequence is the default rate defined by the implementors of the encoder.

COMPRESSOR_SLICES_PER_PICTURE

The value of this attribute provides a hint to the encoder concerning the number of slices into which it should divide the next picture it compresses; however, an XIL-compliant encoder may ignore this suggested value.

The valid values for this attribute range from 1 to the number of macroblocks in the picture. (The default value of the attribute is implementation specific.) Because slices are designed to aid in error recovery from bitstream errors, the value you choose should be dictated by your assessment of the likelihood that data will be corrupted. If bitstream errors will be rare, one slice per picture may be sufficient. On the other hand, if errors will be common, one slice per row of macroblocks may be needed. Don't request more slices than you need because there is a certain amount of overhead—for instance, a 40-bit slice header—associated with each slice.

The code below shows COMPRESSOR_SLICES_PER_PICTURE being set to 13.

```
XilCis cis;
int slices = 13;

xil_cis_set_attribute(cis, "COMPRESSOR_SLICES_PER_PICTURE",
    (void *)slices);
```

You can also read the value of this attribute. To do this, you might use the code shown below.

```
XilCis cis;
int slices;

xil_cis_get_attribute(cis, "COMPRESSOR_SLICES_PER_PICTURE",
    (void **)&slices);
```

If slices contains 0 after this call, the value of the attribute is set to an implementation-specific default.

COMPRESSOR_TIME_CODE

This attribute enables you to associate a time stamp with the first picture in the next group of pictures that is encoded. This time stamp contains the same information as a SMPTE time stamp.

Before actually setting the attribute, you must declare and assign values to the members of a structure of type XilMpeg1TimeCode.

```
typedef struct {
    Xil_boolean drop_frame_flag;
    Xil_unsigned32 hours;
    Xil_unsigned32 minutes;
    Xil_unsigned32 seconds;
    Xil_unsigned32 pictures;
} XilMpeg1TimeCode;
```

An application would obtain the values for most of the members of this structure—hours, minutes, seconds, and pictures—by reading a SMPTE time code in the video being encoded. The member hours can have a value in the range 0 to 23, and minutes and seconds can have values in the range 0 to 59. The member pictures, where a picture is a subdivision of a second, can also have a value in the range 0 to 59.

The other member in the structure, drop_frame_flag, should set to 0 unless the picture rate for the bitstream is 29.97. In this case, the flag can be set to 0 or 1. A setting of 0 indicates that the bitstream should be played back as if the picture rate were 30 pictures per second. A setting of 1 indicates that a play-back application should not count certain pictures so that movie time will not get ahead of wall clock time. Specifically, the application should not count pictures 0 and 1 at the beginning of each minute except minutes 0, 10, 20, 30, 40, and 50.

The code below sets COMPRESSOR_TIME_CODE so that the time stamp 19:07:34:13 is associated with the next group of pictures encoded. The drop-frame flag is set to 0.

```
XilCis cis;
XilMpeg1TimeCode time_code = {0, 19, 7, 34, 13};

xil_cis_set_attribute(cis, "COMPRESSOR_TIME_CODE",
    (void *)&time_code);
```

To read the value of COMPRESSOR_TIME_CODE, you would use code similar to the following.

```
XilCis cis;
XilMpeg1TimeCode *time_code;

xil_cis_get_attribute(cis, "COMPRESSOR_TIME_CODE",
    (void **)&time_code);
```

After this call, time_code will point to a structure of type XilMpeg1TimeCode that contains the current time-code information, or it will be a null pointer, indicating that COMPRESSOR_TIME_CODE has not been set. If the library successfully creates a time-code structure, your application is responsible for freeing the structure's storage.

Decompression Attributes

Only one of the attributes discussed in this section—DECOMPRESSOR_QUALITY—affects how the decoder works. The remaining attributes simply return information from the bitstream being decoded.

DECOMPRESSOR_QUALITY

The value of this attribute provides a hint to the decompressor concerning how it should handle the trade-off between the quality of reconstructed pictures and the speed of decoding those pictures. An XIL-compliant decoder need not act on this hint.

The valid values for this attribute are integers in the range 1 to 100. A value of 100 is a request that the decoder produce the highest quality pictures possible, and a value of 1 is a request that the decompressor decode pictures as fast as possible. By default, DECOMPRESSOR_QUALITY is set to 100.

The code below shows DECOMPRESSOR_QUALITY being set to 90. This setting usually provides satisfactory quality and may result in a significant increase in speed (compared to a setting of 100).

```
XilCis cis;
int quality = 90;

xil_cis_set_attribute(cis, "DECOMPRESSOR_QUALITY",
    (void *)quality);
```

To read the value of this attribute, use code similar to that shown below.

```
XilCis cis;
int quality;

xil_cis_get_attribute(cis, "DECOMPRESSOR_QUALITY",
    (void **)&quality);
```

DECOMPRESSOR_BROKEN_LINK

This is a read-only attribute that indicates whether B pictures at the beginning of the current group of pictures (before any I pictures) can be decoded correctly. Generally, of course, these pictures can be decoded correctly. However, it's possible that during an editing session, the preceding group of pictures—or maybe even just the last I or P picture in that group—was removed. If that type of editing takes place and a B picture at the beginning of the next group of pictures contains macroblocks that were forward predicted or bidirectionally predicted, the decoder will not be able to decode the complete picture. In this case, a broken-link flag should have been set in the group-of-pictures header at the time of the editing.

To read the value of this attribute, you might use code similar to this.

```
XilCis cis;
Xil_boolean broken_link;

xil_cis_get_attribute(cis, "DECOMPRESSOR_BROKEN_LINK",
    (void **)&broken_link);
```

If the broken-link flag for the group of pictures currently being decoded is set, the value of `broken_link` will be `TRUE`. In this case, any B pictures at the beginning of the group cannot be decoded correctly. If the flag is not set, the value will be `FALSE`.

DECOMPRESSOR_CLOSED_GOP

This is a read-only attribute that indicates whether the group of pictures currently being decoded is a closed group or an open group. A closed group is one that begins with an I picture or with one or more B pictures that can be decoded without reference to the last picture in the preceding group. This type of group can be decoded independently. An open group, on the other hand, begins with one or more B pictures, at least one of which contains macroblocks that are forward or bidirectionally predicted. Such a B picture cannot be decoded if the preceding group of pictures is not available (for example, if it has been edited out of the bitstream).

To read the value of this attribute, you might use code similar to this.

```
XilCis cis;
Xil_boolean closed_gop;

xil_cis_get_attribute(cis, "DECOMPRESSOR_CLOSED_GOP",
    (void **)&closed_gop);
```

If the closed-group-of-pictures flag for the group is set, the value of `closed_gop` will be `TRUE`. In this case, any B pictures at the beginning of the group can be decoded independently of the preceding group. If the flag is not set, the value will be `FALSE`.

DECOMPRESSOR_FRAME_TYPE

This is a read-only attribute whose value describes the type of the last picture decoded. This type can be any one of the types included in the following enumeration:

```
typedef enum {
    I,      /* intraframe picture */
    P,      /* forward predicted picture */
    B,      /* bidirectionally predicted picture */
    D       /* DC picture */
} XilMpeg1FrameType;
```

For a description of I, P, and B pictures, see the sections "I Pictures and P Pictures" on page 361 and "B Pictures" on page 361.

The code below illustrates how you might read the value of `DECOMPRESSOR_FRAME_TYPE`.

```
XilCis cis;
XilMpeg1FrameType frame_type;

xil_cis_get_attribute(cis, "DECOMPRESSOR_FRAME_TYPE",
    (void **)&frame_type);
```

DECOMPRESSOR_PEL_ASPECT_RATIO_VALUE

This is a read-only attribute whose value is the pixel aspect ratio stored in the sequence header for the MPEG-1 sequence you're decoding. This value, along with the width and height of the pictures in the sequence, defines the shape of a decoded picture.

The value stored in this attribute will be one of the following constants.

- 0.6735
- 0.7031
- 0.7615
- 0.8055
- 0.8437
- 0.8935
- 0.9157
- 0.9815
- 1.0
- 1.0255
- 1.0695
- 1.0950
- 1.1575
- 1.2015

The key values in this list are 1.0, 0.9157, and 1.0950. A value of 1.0 indicates that the pixels are square and that the pictures in the sequence are meant to be displayed on computer monitors. The value 0.9157 means that decoded pictures can be displayed without distortion on 625-line 50 Hz televisions, and the value 1.0950 indicates that decoded pictures can be displayed without distortion on 525-line 60 Hz televisions.

To read the value of this attribute, use code similar to that shown below.

```
XilCis cis;
float aspect_ratio;

xil_cis_get_attribute(cis,
    "DECOMPRESSOR_PEL_ASPECT_RATIO_VALUE",
    (void **)&aspect_ratio);
```

DECOMPRESSOR_PICTURE_RATE_VALUE

The value of this read-only attribute is the picture-rate value from the sequence header of the MPEG-1 sequence you're decoding. This picture-rate value is the number of pictures that should be decoded and displayed each second.

The value will be one of the constants listed below.

- 23.976
- 24.0
- 25.0
- 29.97
- 30.0
- 50.0
- 59.94
- 60.0

You can use code similar to that shown below to read the value of this attribute.

```
XilCis cis;
float picture_rate;

xil_cis_get_attribute(cis, "DECOMPRESSOR_PICTURE_RATE_VALUE",
    (void **)&picture_rate);
```

DECOMPRESSOR_TEMPORAL_REFERENCE

Because the pictures in a group of pictures may need to be decoded in one order and displayed in another order, each picture contains a temporal-reference value. This value indicates a picture's place in the display order for the group. Picture 0 is meant to be displayed first, picture 1 to be displayed second, and so on. The maximum value the temporal-reference field can accommodate is 1023; therefore, if a group of pictures contains more than 1024 pictures, the counter is reset to 0 for picture 1025 (or more generally, for picture (1024 * *n*) + 1).

The attribute DECOMPRESSOR_TEMPORAL_REFERENCE contains the temporal-reference value for the most recently decompressed picture. To read the value of this attribute, you might use code similar to this.

```
XilCis cis;
int temporal_reference;

xil_cis_get_attribute(cis, "DECOMPRESSOR_TEMPORAL_REFERENCE",
    (void **)&temporal_reference);
```

Note – In a legal MPEG-1 bitstream, the temporal reference values for the pictures in a group must be consecutive values (except in the case where the counter resets). However, it is not uncommon for bitstreams that have been edited to violate this requirement. That is, if the pictures with temporal-reference values between 12 and 24 have been cut from a group, you may find that the picture with a temporal of 11 is followed by one with a temporal reference of 25.

DECOMPRESSOR_TIME_CODE

One field in a group-of-pictures header is a SMPTE time code. This time code is associated with the first picture (in display order) in the group. When you read the attribute DECOMPRESSOR_TIME_CODE, the decoder places information about this time code for the current group of pictures in a structure of type XilMpeg1TimeCode.

```
typedef struct {
    Xil_boolean drop_frame_flag;
    Xil_unsigned32 hours;
    Xil_unsigned32 minutes;
    Xil_unsigned32 seconds;
    Xil_unsigned32 pictures;
} XilMpeg1TimeCode;
```

The member hours will have a value in the range 0 to 23; minutes, seconds, and pictures will have values in the range 0 to 59.

You can read the value of this attribute using code similar to this.

```
XilCis cis;
XilMpeg1TimeCode *time_code;

xil_cis_get_attribute(cis, "DECOMPRESSOR_TIME_CODE",
    (void **)&time_code);
```

Your application is responsible for freeing the storage the library allocates to hold the time-code structure.

MPEG-1 Molecules

The XIL library includes a series of *molecules* that accelerate the playback of MPEG-1 bitstreams. These molecules are optimized routines that perform the jobs of two or more functions from the XIL API. You do not call such an optimized routine directly; rather, the library calls a molecule when your program calls a predefined sequence of XIL functions, sometimes with specific arguments.

For example, if your program calls `xil_decompress()` to decode an image stored in an MPEG-1 CIS and then calls `xil_ordered_dither()` to dither the decoded image from a 24- to an 8-bit image, the library may not call these two functions. Instead, it may call a molecule that performs the decompression and the dithering in an optimized way.

For information about the MPEG-1 molecules that are available and information about how to call those molecules, see the section "XIL Molecules" on page 398.

CCITT Group 3 and Group 4 Codecs 20

The CCITT Group 3 and Group 4 compression methods were developed by the International Telegraph and Telephone Consultative Committee for the compression of document images. Like continuous-tone still images, document images can be stored and transmitted much more inexpensively if they have been compressed. Originally, the CCITT's compression techniques were adopted by the developers of facsimile equipment. Today, however, the techniques are also used heavily by the makers of general document storage and retrieval systems.

The remainder of this chapter is divided into two sections. The first section provides a brief overview of how the Group 3 and Group 4 codecs work. The second discusses a set of CIS attributes that apply specifically to the these codecs (as opposed to the general CIS attributes covered in the section "General CIS Attributes" on page 257).

How CCITT Group 3 and Group 4 Codecs Work

The Group 3 and Group 4 compressors take advantage of a couple of important characteristics of document images. One of these characteristics is that document images tend to consist of small amounts of black (letters or lines) on a white background. Thus, on a given scanline, there are likely to be long stretches of white pixels interrupted by shorter runs of black pixels.

This characteristic has led to the use of run-length encoding in the CCITT compressors. This coding method involves translating information about runs of white and black pixels (within a scanline) into code words stored in a

Huffman table. Run lengths that occur commonly are represented by short code words, and run lengths that occur infrequently are represented by longer code words.

This type of encoding generally produces quite a bit of compression. For instance, a run of 1024 white pixels (128 bytes) might be represented by a code word of 9 bits. This translation results in a compression ratio of about 114:1. However, the compression ratio achievable with run-length encoding is highly dependent on the image being compressed. For example, a noisy image might contain many short runs whose code words contain more bits than the runs themselves.

On standard text, the XIL library's Group 3 compressor achieves a compression ratio of about 5:1, and the library's Group 4 compressor achieves a ratio of about 10:1.

A second characteristic of document images is that the position of a transition from a black pixel to a white pixel (or vice versa) on one scanline is usually not more than a few pixels away from a corresponding transition on the preceding scanline. This characteristic is sometimes called *vertical coherence*. Because document images have this characteristic, once one scanline has been encoded, subsequent lines can be encoded by specifying the position of a black-to-white transition relative to the same transition on the preceding line, or relative to the last transition on the same line.

The XIL library's CCITT Group 3 compressor uses the run-length encoding method described above, and the Group 4 compressor relies almost entirely on the two-dimensional technique.

CCITT Group 3 and Group 4 Decompressor Attributes

Although other compression standards include size information (the image width, height, and number of bands) within the data bitstream, the fax standards do not. Thus, before you can decompress a fax CIS, you must set the decompressor attributes for the width, height, and number of bands. These attributes are discussed separately in the following sections.

WIDTH

If you have put compressed data into your CIS using `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()`, you must set the value of this attribute to the width in pixels of the images to be decompressed. If you do not set it, its value will be 0, and an error will occur when you call `xil_decompress()`.

The legal values for this attribute are integers in the range 0 to 32,767. The code fragment below shows the `WIDTH` attribute being set to 1728.

```
XilCis cis;
int width = 1728;

xil_cis_set_attribute(cis, “WIDTH”, (void *)width);
```

HEIGHT

If you have put compressed data into your CIS using `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()`, you must set the value of this attribute to the height in pixels of the images to be decompressed. If you do not set it, its value will be 0, and an error will occur when you call `xil_decompress()`.

The legal values for this attribute are integers in the range 0 to 32,767. The code fragment below shows the `HEIGHT` attribute being set to 2156.

```
XilCis cis;
int height = 2156;

xil_cis_set_attribute(cis, “HEIGHT”, (void *)height);
```

BANDS

If you have put compressed data into your CIS using `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()`, you must set the value of this attribute to the number of bands in the images to be decompressed. If you do not set it, its value will be 0, and an error will occur when you call `xil_decompress()`.

The legal values for this attribute are integers in the range 0 to 32,767. The code fragment below shows the BANDS attribute being set to 1.

```
XilCis cis;
int bands = 1;

xil_cis_set_attribute(cis, "BANDS", (void *)bands);
```

Acceleration in XIL Programs 21

Obviously, much of the speed of an XIL program is determined by the speed of the individual functions that make up the API. However, XIL applications can also realize big performance improvements when the library is able to replace the execution of a sequence of API-level functions (or *atoms*) with the execution of an optimized function that performs the work of all the atoms in the sequence. This type of optimized function is called a *molecule*.

You do *not* call a molecule directly. Rather, the library itself recognizes sequences of atoms that can be replaced by a molecule and performs the replacement automatically. This type of replacement is made possible by the library's deferred-execution scheme.

What Is Deferred Execution?

Here's how deferred execution works. In general, when an atomic function that affects the state of an image or a compressed image sequence is called, the function is not executed immediately. Instead, information about the operation—such as the function called and the arguments to the function—is stored by the library. This information continues to be stored until the library must produce a particular destination image—for example, because that image is to be displayed.

Let's say that the library has stored the five atomic operations shown in Figure 21-1 and that by performing these operations, it could decompress an image from a Cell CIS and prepare that image for display on a monochrome display.

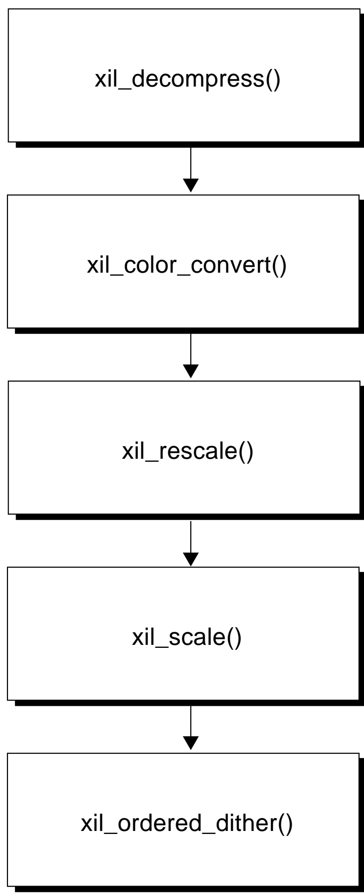

Figure 21-1 Stored Atomic Operations

When it must produce the destination of the ordered dither, instead of simply performing the five atomic functions, the library searches a list of molecules to see if all or part of this sequence can be replaced by a molecule. If the library finds a molecule that can perform this entire sequence of operations, the execution of the program will proceed as shown in Figure 21-2.

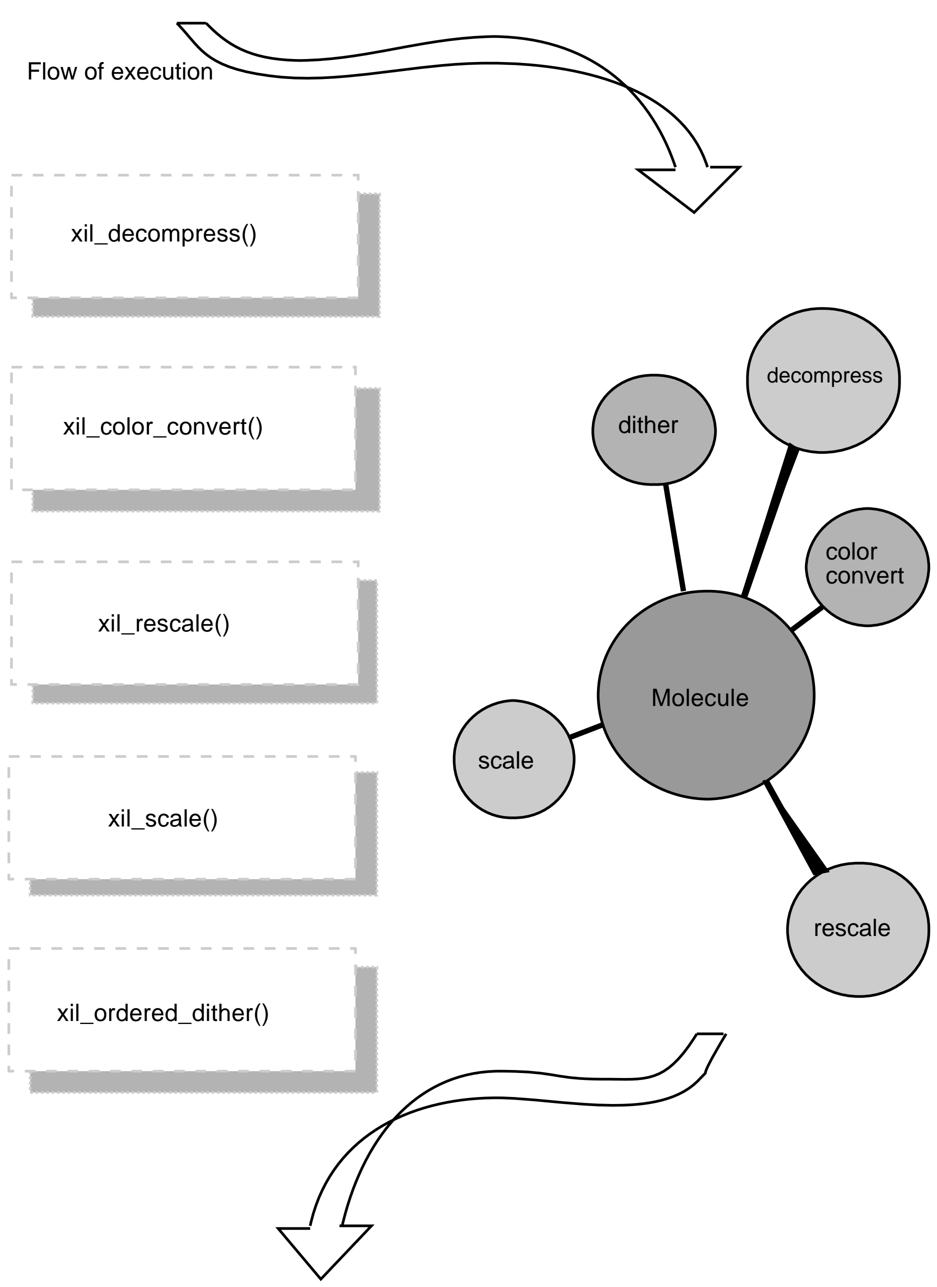

Figure 21-2 Replacing Atomic Functions with a Molecule

The molecule shown is executed and performs the jobs of all five atomic functions. None of the atomic functions is ever executed.

The ability to replace a series of atomic functions with a molecule like this can lead to dramatic increases in performance. Two reasons for this increase are that:

- A molecule may need to page an image into memory only once whereas the equivalent atomic functions would have paged the image into memory several times
- A molecule may not need to create temporary images that would be required by the equivalent atomic functions

XIL Molecules

The molecules available in the XIL library can be grouped into two categories:

- Molecules that involve decompressing images stored in a compressed image sequence. This is the largest category of molecules, and these molecules are available on all platforms that can run XIL applications.
- **SPARC:** Molecules that perform a common XIL operation and then display the results of the operation on a GX frame buffer. These molecules are available only on local GX frame-buffer screens.

This section first looks at some general rules you must follow to execute XIL molecules. (Particular molecules may require that you follow other rules as well.) The section then discusses the molecules that have been implemented in the categories mentioned in the preceding paragraph.

Rules for Executing Molecules

There are a few general rules you must follow to execute XIL molecules. These rules are listed below:

- The images that a molecule works with—source, intermediate, and destination—must have the same width and height. The principal exception to this rule occurs when a molecule performs a scaling operation. In that case, the destination image for the scale operation must have the same dimensions as the scaled source image.

- The images the molecule works with must have the same region of interest. For decompression molecules, this rule dictates that the molecule's destination image and any intermediate images have no region of interest because the images stored in the CIS cannot have a region of interest. In addition, if a molecule includes a scale operation, the images involved cannot have a region of interest. These images will not have a region of interest unless you have explicitly set their region-of-interest attributes.
- The images the molecule works with must have the same origin. For decompression molecules, this rule dictates that the molecule's destination image and any intermediate images have origins of 0.0, 0.0. This is true because the images stored in a CIS cannot have anything other than the default origin.

There are one or two exceptions to these rules among the decompression molecules. These exceptions are covered in the section "Video Decompression Molecules," where these molecules are defined.

Video Decompression Molecules

The library's video decompression molecules decompress an image from a compressed image sequence, process that image in some way, and write it to a destination, usually a display image. This destination may have a depth of 1 bit (a monochrome display), 8 bits (an indexed-color display), or 24 bits (a true-color display). To display video at the best possible speed, it is important that you use these molecules. They provide much better performance than the equivalent set of atomic functions.

Note – The molecules discussed in this section can be used to decompress and display images stored in JPEG baseline sequential, CellB, H.261, and MPEG-1 compressed image sequences. If you're working with a Cell bytestream, you should use the molecules discussed in the section "Cell Molecules" on page 288.

One-Bit Destination

The decompression molecule that displays an image on a 1-bit display works only with the JPEG decompressor and expects either a three-band YC_bC_r image or a one-band Y image as input. The molecule performs the following tasks:

- Decompresses an image from a CIS
- If necessary, converts a 3-band image to a 1-band image by extracting the luminance of the YC_bC_r image
- *Optionally* rescales the values in the Y image
- Dithers the 8-bit image to a 1-bit image
- *Optionally* zooms the dithered image by a factor of 2 in both the *x* and *y* dimensions

A requirement for this molecule is that the images in the CIS have a width and height that are multiples of the bitstream macroblock width and height. For example, single-band images must be divisible into 8-by-8 blocks, 4:2:2 images must be divisible into 16-by-8 blocks, and 4:1:1 images must be divisible into 16-by-16 blocks.

For this molecule to be called in the case where the source image is a YC_bC_r image, your application must include the code shown below.

```
XilColorspace ycc, y;
float scale_values[1], offset_values[1];
XilLookup cmap;
int mults[1] = {-1};
unsigned short dims[1] = 2;
XilKernel distribution;

scale_values[0] = 255.0 / (235.0 - 16.0);
offset_values[0] = -16.0 * scale_values[0];

/* The lookup table cmap must have two entries, and the values of
   those entries on the output side of the table must be 255 and
   0. One way to create this table is to use the call to
   xil_colorcube_create() shown in the example. You could also
   use xil_lookup_create(). */
cmap = xil_colorcube_create(state, XIL_BIT, XIL_BYTE, 1, 0,
    mults, dims);
distribution = xil_kernel_get_by_name(state,
    "floyd-steinberg");
ycc = xil_colorspace_get_by_name(state, "ycc601");
y = xil_colorspace_get_by_name(state, "y601");
xil_set_colorspace(imageYCC_24, ycc);
xil_set_colorspace(imageGRAY_8, y);

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageYCC_24);
xil_color_convert(imageYCC_24, imageGRAY_8);

/* Rescale if the input image has CCIR 601 color space */
if (CCIR601)
    xil_rescale(imageGRAY_8, imageGRAY_8, scale, offset);

xil_error_diffusion(imageGRAY_8, imageMONO_1, cmap,
    distribution);

/* Execute if user has requested zoom */
if (ZOOM)
    xil_scale(imageMONO_1, imageMONO_1_zoom, "nearest", 2.0,
        2.0);
/* MOLECULE ENDS HERE */
```

Note – Arguments shown in boldface must be typed as shown for the molecule to execute correctly.

If the source image is a one-band image, the molecule should look like this.

```
/* MOLECULE STARTS HERE */
xil_decompress(cis, imageGRAY_8);

/* Rescale if the input image has CCIR 601 color space */
if (CCIR601)
    xil_rescale(imageGRAY_8, imageGRAY_8, scale, offset);
xil_error_diffusion(imageGRAY_8, imageMONO_1, cmap,
    distribution);

/* Execute if user has requested zoom */
if (ZOOM)
    xil_scale(imageMONO_1, imageMONO_1_zoom, "nearest", 2.0,
        2.0);
/* MOLECULE ENDS HERE */
```

Eight-Bit Destination

This molecule is optimized where the molecule's destination image is an 8-bit display image associated with a local GX frame buffer. However, the molecule will also accelerate the decompression and subsequent processing of images being written to other 8-bit destinations as well.

The molecule performs the following tasks:

- Decompresses an image from a CIS
- *Optionally* rescales the values in the three bands of the YC_bC_r image
- Dithers the 24-bit YC_bC_r image to an 8-bit pseudocolor image
- *Optionally* zooms the dithered image by a factor of 2 in both the *x* and *y* dimensions

Note – If you are using the CellB decompressor, the zoom operation (if requested) must precede the ordered dither.

There are certain restrictions on the use of this molecule. One is that the molecule can decompress images from a JPEG CIS only if those images are YC_bC_r images and the images are 4:1:1 interleaved in the bitstream. (CellB, H.261, and MPEG-1 CISs always contain YC_bC_r images.) In addition, in most cases, JPEG, H.261, and MPEG-1 CISs must contain images whose width and height are multiples of 16. Exceptions to this rule are noted later.

For this molecule to be executed, your application must include the code shown below.

```
float scale_values[3], offset_values[3];
XilLookup colorcube, colormap;
XilDitherMask dmask;

scale_values[0] = 255.0 / (235.0 - 16.0);
scale_values[1] = 255.0 / (240.0 - 16.0);
scale_values[2] = 255.0 / (240.0 - 16.0);
offset_values[0] = -16.0 * scale_values[0];
offset_values[1] = -16.0 * scale_values[1];
offset_values[2] = -16.0 * scale_values[2];
colorcube = xil_lookup_get_by_name(state, "cc855");
dmask = xil_dithermask_get_by_name(state, "dm443");

/* The lookup table yuv_to_rgb is matched with the colorcube
   cc855. If you use cc855 when dithering your images, you must
   install the 200 sets of RGB values in yuv_to_rgb in your
   application's X colormap before executing the molecule.
   Install the first set of RGB values at colorcell 54. */
colormap = xil_lookup_get_by_name(state, "yuv_to_rgb");

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageYCC_24);

/* Perform this rescale if your image's color space is CCIR Rec.
   601 YCC. The rescale adjusts the range of the values in the
   three bands of the image to 0 to 255. The ordered dither
   operation that takes place later will produce the best results
   if the values are in this range. Note that YCC images produced
   by compressors that use the JFIF format have already been
   rescaled to 0 to 255 and, therefore, should not be rescaled
   here. */
if (CCIR601)
    xil_rescale(imageYCC_24, imageYCC_24, scale_values,
        offset_values);
```

```
if (ZOOM) {
    xil_ordered_dither(imageYCC_24, imageYCC_8, colorcube,
        dmask);
    xil_scale(imageYCC_8, zoom_displayimage, "nearest", 2.0,
        2.0);
}
else
    xil_ordered_dither(imageYCC_24, displayimage, colorcube,
        dmask);
/* MOLECULE ENDS HERE */
```

Exceptions to the Rules

Table 21-1 shows the exceptions to the general rules for executing this molecule. The exceptions have been classified on the basis of the decompressor to which they relate.

Note – Among these exceptions, the term *GX window* refers to a local GX frame-buffer screen.

Table 21-1 Exceptions to the General Decompression-Molecule Rules

Decompressor	Exceptions
CellB	If you want the molecule to zoom the image it decompresses, the call to `xil_scale()` that zooms the image must precede the call to `xil_ordered_dither()`.
JPEG	If the destination image is a display image associated with a GX window, the images in the CIS do not have to have a width and height that are multiples of 16. If the destination image is *not* a display image associated with a GX window, the molecule cannot perform the optional zoom. If the destination image is *not* a display image associated with a GX window, the destination cannot be a single-band child of a multiband parent image.
H.261	If the destination image is a display image associated with a GX window, the images in the CIS do not have to have a width and height that are multiples of 16. There is one exception to the general rules that a molecule's destination image must have the same width and height as the images in the CIS and that the destination image must have an origin of 0.0, 0.0. The molecule can execute even though the two rules mentioned above are not met, if the following conditions are satisfied: (1) the destination image is a display image associated with a window on a local GX frame buffer, (2) the coordinates of the destination image's origin are less than or equal to 0.0, and (3) all of the pixels in the display image are written to. This exception allows for the clipping of letterboxed images.
MPEG-1	If the destination image is a display image associated with a GX window, the images in the CIS do not have to have a width and height that are multiples of 16.

Twenty-Four-Bit Destination

The molecule performs the following tasks:

- Decompresses a YC_bC_r image from a CIS
- Converts the color space of the YC_bC_r image to RGB

Obviously, the images in a JPEG CIS must be YC_bC_r images before you can use this molecule. (CellB, H.261, and MPEG-1 CISs always contain YC_bC_r images.) In addition, the images in a JPEG CIS must be 4:1:1 interleaved in the bitstream. One other restriction is that JPEG, H.261, and MPEG-1 CISs must contain images whose width and height are multiples of 16.

For this molecule to be executed, your application must include the code shown below.

```
XilColorspace ycc, rgb;

ycc = xil_colorspace_get_by_name(state, "ycc601");
rgb = xil_colorspace_get_by_name(state, "rgb709");
xil_set_colorspace(imageYCC_24, ycc);
xil_set_colorspace(imageRGB_24, rgb);

/* MOLECULE STARTS HERE */
xil_decompress(cis, imageYCC_24);
xil_color_convert(imageYCC_24, imageRGB_24);
/* MOLECULE ENDS HERE */
```

The molecule can only convert images from `ycc601` to `rgb709` (as shown in the example).

CCITT Group 4 Decompression Molecule

Because it is common to transpose an image immediately after decompressing it from a CCITT compressed image sequence, there is a molecule for accelerating the transposition of images that were compressed with the Group 4 compressor. The molecule takes advantage of the fact that, when decompressing these images, an intermediate representation of the original image is generated, and it is faster to transpose this representation than it is to transpose the bits of the actual image. For images that were compressed with the Group 3 compressor, this intermediate representation is not generated and so the molecule does not execute.

The `xil_transpose()` call must immediately follow the image decompression; otherwise the intermediate representation is lost and the molecule cannot execute. In this case, the transposition is performed on the bit image.

The type of flip designated on the `xil_transpose()` call determines whether the destination size is checked as

```
src_x_size == dest_x_size && src_y_size == dest_y_size
```

or

```
src_x_size == dest_y_size && src_y_size == dest_x_size
```

One of the above expressions must be true for the molecule to be called. Table 21-2 shows the destination image size expected for the source image size (100, 140).

Table 21-2 Type of Flip Designated on the Call to `xil_transpose()`

Type of Flip	Expected Destination Size
XIL_FLIP_X_AXIS	(100, 140)
XIL_FLIP_Y_AXIS	(100, 140)
XIL_FLIP_X_180	(100, 140)
XIL_FLIP_MAIN_DIAGONAL	(140, 100)
XIL_FLIP_ANTIDIAGONAL	(140, 100)
XIL_FLIP_90	(140, 100)
XIL_FLIP_270	(140, 100)

Image-Filtering Molecule

The XIL library provides a molecule for accelerating image filtering when you use a 5-by-5 separable kernel. Assuming you observe the normal rules for executing molecules (see "Rules for Executing Molecules" on page 398), the molecule is executed when you filter either an XIL_BYTE or XIL_SHORT image and make two separate xil_convolve() calls, one call using a 1-by-5 kernel and the other call using a 5-by-1 kernel. It does not matter which kernel is used first; however, both xil_convolve() calls must use the XIL_EDGE_ZERO_FILL edge condition. There is no molecule for filtering XIL_BIT images.

The following code fragment shows an example of how your application code might execute this molecule.

```
XilImage src, dst;
XilKernel filter_1, filter_2;
XilSystemState state;
unsigned int width_1, height_1, keyx_1, keyy_1,
             width_2, height_2, keyx_2, keyy_2;
float data_1[] = {0.1117, 0.2365, 0.3036, 0.2365, 0.1117};
float data_2[] = {0.1117, 0.2365, 0.3036, 0.2365, 0.1117};

width_1 = height_2 = 1;
width_2 = height_1 = 5;
keyx_1 = keyx_2 = 2;
keyy_1 = key2_2 = 0;
filter_1 = xil_kernel_create(state, width_1, height_1, keyx_1,
    keyy_1, data_1);
filter_2 = xil_kernel_create(state, width_2, height_2, keyx_2,
    keyy_2, data_2);

/* MOLECULE STARTS HERE */
xil_convolve(src, dst, filter_1, XIL_EDGE_ZERO_FILL);
xil_convolve(src, dst, filter_2, XIL_EDGE_ZERO_FILL);
/* MOLECULE STARTS HERE */
```

SPARC *Molecules That Result in a Display*

Molecules that combine a common operation and a display are available only on local GX frame-buffer screens. Before you look over the list of molecules in this category, you should understand that operations that read from or write to a device image (including a display image) are treated specially by the XIL library.

Generally, calling an XIL function results in the library's storing one operation. However, operations that read from or write to a device image cause *two* operations to be stored. For example, you might call `xil_scale()` using a scanner as your source image in order to read an image from the scanner and zoom it. The two operations stored will be a *capture* operation, which writes its output to a temporary image, and a scale operation, which writes its output to the destination image named in the call to `xil_scale()`. Likewise, you might call `xil_scale()` using a display as your destination image in order to zoom the image and display it. In this case, the two operations stored will be a scale and a *display* operation.

Several molecules are available that replace two-operation sequences in which the second operation is a display to a GX frame buffer. These molecules are listed below:

- Copy (`xil_copy()`) an 8-bit image to a GX display image.
- Cast (`xil_cast()`) a 1-bit image to an 8-bit image and display the 8-bit image on a GX frame buffer.
- Pass a 1-bit image through a lookup table (`xil_lookup()`) to produce an 8-bit image and display the 8-bit image on a GX frame buffer.
- Pass an 8-bit image through a lookup table to produce a different 8-bit image and display the new image on a GX frame buffer.
- Rescale (`xil_rescale()`) an 8-bit image and display the result on a GX frame buffer.
- Set all the values in an 8-bit image to a constant (`xil_set_value()`) and display the result on a GX frame buffer.
- Translate (`xil_translate()`) an 8-bit image using nearest-neighbor interpolation, and display the result on a GX frame buffer.

21

Troubleshooting Molecules

Once you have coded your application with a replaceable sequence of atomic functions, you may want to make sure a molecule executes when your program runs. And if the molecule is not called, you will want to determine why it's not being called. This section discusses both how to determine whether a molecule is executing and what to check for if the molecule is not being called.

Determining Whether Molecules Are Executing

To determine whether a molecule you want to execute is actually executing, set the environment variable `XIL_DEBUG` as shown below before executing your program.

```
% setenv XIL_DEBUG show_action
```

Once you've done this, the XIL library will print a message to `stderr` each time an operation that affects the state of an XIL image or a compressed image sequence (CIS) is executed. By looking at these messages, you can determine whether the molecules your application is attempting to call are being executed.

For example, try setting `XIL_DEBUG` to `show_action` and then playing a JPEG movie using the example program `xilcis_example`. Instructions for running the example are presented in the section "Running the Movie Player" on page 231. As the program runs, you will see one or more pairs of messages—depending on the length of the movie—displayed on `stderr`.

```
XIL_ACTION[XilDeviceCompJpegMemory]:ordereddither8to8(rescale8(
decompress_Jpeg()))
XIL_ACTION[XilDeviceIOioxlib]:display_ioxlib()
```

The first message indicates that a molecule has been executed. This molecule performs a decompression, a rescale, and an ordered dither. The second message indicates that this molecule was followed by an atomic display operation; this message assumes the display operation uses function calls from `ioxlib`.

If you were now to edit this program so that it did something that prevented the molecule from being executed—like setting the origin of the display image to 1.0, 1.0—you would see the following messages:

```
XIL_ACTION[XilDeviceCompJpegMemory]:ordereddither8to8(rescale8(
decompress_Jpeg()))
XIL_ACTION:FAILED
XIL_ACTION[XilDeviceCompJpegMemory]:decompress_Jpeg()
XIL_ACTION[XilDeviceComputeMemory]:rescale8()
XIL_ACTION[XilDeviceComputeMemory]:ordereddither8to8()
XIL_ACTION[XilDeviceIOioxlib]:display_ioxlib()
```

As you can see, the molecule would be unable to execute, and the decompression, rescale, and ordered dither would have to be executed individually.

Besides the environment variable `XIL_DEBUG`, there are two XIL functions that enable you to control whether the messages discussed above are printed. One of these, `xil_state_set_show_action()` sets a system-state attribute called `SHOW_ACTION`. The setting of this attribute determines whether the library prints the messages, does not print the messages, or prints the messages only if `XIL_DEBUG` is set to `show_action`. The other function, `xil_state_get_show_action()`, reads the value of `SHOW_ACTION`. See the man pages for more information about these functions.

Determining Why a Molecule Is Not Executing

Basically, there are two reasons why a molecule might not execute. Either you didn't call the proper sequence of atomic functions with the proper arguments, or you did something that caused the library to flush some operations from storage before you finished building the molecule. These two situations are discussed in more detail below.

Not Calling a Correct Sequence of Functions

The most likely reason a molecule is not executing is that you haven't followed all the rules for calling the molecule. Remember, you must call the correct functions, in the proper order, sometimes with specific arguments. Any mistake can cause the molecule not be called or to fail.

For example, let's say you are decompressing a JPEG movie and displaying frames on an 8-bit frame buffer, and that you call the code shown below.

```
xil_decompress(cis, imageYCC);
xil_rescale(imageYCC, imageYCC, scale, offset);
xil_ordered_dither(imageYCC, image8, colorcube, dmask);
xil_scale(image8, zoom_displayimage, "bilinear", 2.0, 2.0);
```

If for no other reason, these four functions won't be replaced by a molecule because the call to `xil_scale()` uses an interpolation type of `bilinear`. For this molecule to execute, the interpolation type for the scale must be `nearest` (nearest neighbor). So, if you expect a molecule to execute and it doesn't, check first to make sure that you've met all the guidelines presented in the section "XIL Molecules" on page 398.

A more subtle problem can arise when the calls that make up a molecule are not consecutive calls in your program. For instance, consider the following code fragment taken from the Cell movie playback code discussed in "Playing Cell Movies" on page 241.

```
xil_decompress(cis, image24);

/* Look at the Lookup version number to see if it is time
   to reinstall the colormap. */
if (lu_version != xil_lookup_get_version( xil_cmap )) {
    cell_install_cmap(x_cmap, displayimage, xil_cmap, ilist);
    lu_version = xil_lookup_get_version( xil_cmap );
}
xil_nearest_color(image24, displayimage, xil_cmap);
```

This code is attempting to call a molecule that replaces the sequence

- `xil_decompress()`
- `xil_nearest_color()`

Your first thought may be that this code won't cause the molecule to be executed because the calls to `xil_decompress()` and `xil_nearest_color()` are not consecutive. Between these two calls, the program calls `xil_lookup_get_version()` at least once and may call a user-defined function `cell_install_cmap()`, which doesn't contain any XIL functions. Actually, this code will cause the molecule to be called because only the decompress and nearest-color operations are deferred and become part of the stored chain of operations that the library tries to replace with a molecule.

Calls to non-XIL functions, like `cell_install_cmap()`, of course are never deferred by the library. In addition, there is a class of XIL functions that are never deferred, and calls that read object attributes fall in this class. Therefore, the call to `xil_lookup_get_version()` does not affect the deferred-operation chain.

On the other hand, if you used the following code, the library would not execute a molecule.

```
xil_decompress(cis, image24);
xil_add_const(image24, constants, image24);
xil_nearest_color(image24, displayimage, xil_cmap);
```

The add-constant operation would be deferred (because it changes values in the image) and would prevent the library from finding an `xil_decompress()-xil_nearest_color()` sequence.

Flushing Operations Before a Molecule Is Complete

One other subtle problem to guard against is calling a function that causes the library to execute atomically functions that you intended to be the beginning of a molecule. For example, let's say that you want to execute a molecule that replaces four API-level functions. You call the first three atomic functions, and information about these deferred operations is stored. Then before you call the fourth function, you do something that causes the first three operations to be flushed from storage and executed.

What could you do to cause this to happen? There are a number of things, some obvious and some not so obvious.

At the obvious end of the spectrum, you could

- Turn deferred execution off using the function `xil_state_set_synchronize()`. Once deferred execution is turned off, of course, no operations are deferred.
- Call `xil_sync()` on an intermediate image. This call causes the library to produce that image immediately, which involves flushing all the stored operations needed to produce it.
- Use an intermediate image whose synchronization flag has been set using the function `xil_set_synchronize()`. The operations required to produce this image cannot be deferred.

You can also flush operations that you meant to be part of a molecule by

- Calling `xil_export()` to export an intermediate image. The operations needed to produce the intermediate image must be executed before the image can be exported.
- Calling a function that returns information about the values in an intermediate image—like `xil_extrema()`, `xil_histogram()`, or `xil_get_pixel()`. Such a call causes the operations that will produce the source image to be executed.

Side Effects of Executing Molecules

When a molecule executes, it may not create one or more temporary images that would have been created had the set of atomic functions replaced by the molecule been executed. This section briefly explains the possible side effects of this behavior and how you can deal with them.

As an example, consider the JPEG decompression code below.

```
XilImage imageYCC, displayimage;

xil_decompress(cis, imageYCC);
xil_rescale(imageYCC, imageYCC, scale, offset);
xil_ordered_dither(imageYCC, displayimage, colorcube, dmask);
```

If these functions are executed atomically, decompressed data is written to the 24-bit image `imageYCC`, the data in `imageYCC` is rescaled, and then the data in `imageYCC` is dithered to produce an 8-bit image. However, if a JPEG decompression molecule is executed, `imageYCC` is never actually used. This difference can lead to unexpected results later on.

Say that the decompression molecule executed, and `imageYCC` was never written to. If you later ask the library to destroy the CIS (without having first destroyed `imageYCC`), the library, before it destroys the CIS, will call `xil_decompress()` to decompress data from the CIS and write it to `imageYCC`. The library does this because it knows that it never wrote data to `imageYCC`, that the contents of `imageYCC` depend on the CIS, and that the image may be used again later in the program. If you don't plan to use the image later in your program, the call to `xil_decompress()` is unnecessary and hurts the performance of your program.

Note – The solution to this particular problem would be to destroy the 24-bit image before destroying the CIS or to use the function `xil_toss()` to toss the contents of the 24-bit image before destroying the CIS. When you toss an image, you tell the library that you don't care about the current contents of the image, but that you might write data to the image later.

You can check for unexpected side effects like the one described above by setting `XIL_DEBUG` to `show_action` and monitoring `stderr` to see which functions execute when you run your program.

XIL Functions That Relate to Deferred Execution

Several XIL functions have a direct effect on the way deferred execution works. These are listed in Table 21-3.

Table 21-3 Functions That Affect Deferred Execution

Function Name	What the Function Does
`xil_sync`	Forces the library to compute the value of an image when the operations that will produce the image have been deferred.
`xil_cis_sync`	Forces any outstanding call to `xil_compress()` to complete when it would otherwise have been deferred.
`xil_set_synchronize`	Set an image's synchronization attribute. When an image has this attribute set, operations on it cannot be deferred.
`xil_get_synchronize`	Determines whether an image's synchronization attribute is set.
`xil_state_set_synchronize`	Turns deferred execution on or off.
`xil_state_get_synchronize`	Determines whether deferred execution is being used or not.
`xil_toss`	Tells the library to "toss," or throw away, the contents of an image, but not to destroy the image.

XIL Molecules

This appendix lists in a shorthand form the molecules that are available in the current release of the XIL Imaging Library. As an example of the shorthand used, one of the CellB decompression molecules is listed like this:

decompress_CellB → **colorconvert** → **[scale8nearest]** →

SPARC

[display_ioSUNWgs | display_ioxlib]

x86

[display_ioxlib]

This notation indicates that part of the molecule is platform independent, and part of it is platform dependent. The platform independent portion, only the top line in this example, is performed on all systems. The platform dependent portion can be performed only on the identified system. This example shows that the molecule uses different display functions, depending on whether the application is run on SPARC or on x86 screen hardware. Generally, molecules that don't display an image are entirely platform independent, and those that display an image use platform-dependent display functions.

The generic portion of the molecule performs the following operations:

- Decompresses a compressed image stored in a CellB compressed image sequence (`xil_decompress()`).
- Converts the color space of the decompressed image, in this case from YC_bC_r to RGB (`xil_color_convert()`).

- Optionally, scales the RGB image, which must contain XIL_BYTE data, using nearest-neighbor interpolation (xil_scale()). The square brackets around **scale8nearest** indicate that this function is optional. There is actually one molecule that performs the scale and a corresponding molecule that doesn't.

The platform-dependent portions of the molecule do the following:

- **SPARC:** The molecule optionally displays the image on a 24-bit display. The vertical bar indicates the molecule can perform the job of either of two functions: a display function that uses calls to the Direct Graphics Access library to display an image on local GS frame-buffer screen or a display function that uses Xlib calls to display an image on any 24-bit frame buffer.
- **x86:** The molecule optionally uses Xlib calls to display an image on any 24-bit frame buffer.

Key to the Names Used in the Molecule Definitions

Table A-1 provides a key to the meaning of the function names used to describe the XIL molecules.

Table A-1 Key to Names Used in Molecule Definitions

Name	Meaning
colorconvert	A call to xil_color_convert(), used to convert the color space of an image
convert1to8	A call xil_cast() with an XIL_BIT source image and an XIL_BYTE destination
convolve8	A call to xil_convolve() with a source and destination image of type XIL_BYTE
convolve16	A call to xil_convolve() with a source and destination image of type XIL_SHORT
copy8	A call to xil_copy() with a source and destination image of type XIL_BYTE
decompress_Cell	A call to xil_decompress() that decompresses an image from a Cell compressed image sequence (CIS)
decompress_CellB	A call to xil_decompress() that decompresses an image from a CellB CIS

Table A-1 Key to Names Used in Molecule Definitions

Name	Meaning
decompress_faxG4	A call to `xil_decompress()` that decompresses an image from a faxG4 CIS
decompress_H261	A call to `xil_decompress()` that decompresses an image from an H.261 CIS
decompress_Jpeg	A call to `xil_decompress()` that decompresses an image from a JPEG CIS
decompress_Mpeg1	A call to `xil_decompress()` that decompresses an image from an MPEG-1 CIS
display_ioxlib	A display function that is called implicitly when the destination image for an operation is a display image and no faster display function is available. The display function uses Xlib calls to do its job.
display_ioSUNWgs	A display function that is called implicitly when the destination image for an operation is a display image created from a window on a local GS frame buffer. This is an accelerated display routine. Available only on local GS frame-buffer screens.
display_ioSUNWgx	A display function that is called implicitly when the destination image for an operation is a display image created from a window on a local GX frame buffer. This is an accelerated display routine. Available only on local GX frame-buffer screens.
error_diffusion8_1	A call to `xil_error_diffusion()` with an `XIL_BYTE` source image and an `XIL_BIT` destination
lookup1_8	A call to `xil_lookup()` with an `XIL_BIT` source image and an `XIL_BYTE` destination
lookup8_8	A call to `xil_lookup()` with a source and destination image of type `XIL_BYTE`
nearestcolor8_8	A call to `xil_nearest_color()` with a source and destination image of type `XIL_BYTE`
ordereddither8_1	A call to `xil_ordered_dither()` with an `XIL_BYTE` source image and an `XIL_BIT` destination
ordereddither8_8	A call to `xil_ordered_dither()` with a source and destination image of type `XIL_BYTE`

Table A-1 Key to Names Used in Molecule Definitions

Name	Meaning
rescale8	A call to `xil_rescale()` with a source and destination image of type `XIL_BYTE`
scale1nearest	A call to `xil_scale()` with a source and destination image of type `XIL_BIT` and an interpolation type of `nearest`.
scale8nearest	A call to `xil_scale()` with a source and destination image of type `XIL_BYTE` and an interpolation type of `nearest`.
setvalue8	A call to `xil_set_value()` with a destination image of type `XIL_BYTE`
translate8nearest	A call to `xil_translate()` with a source and destination image of type `XIL_BYTE` and an interpolation type of `nearest`.

Molecule Descriptions

The molecule descriptions in this section summarize what molecules are available in the current release of the library. For complete information on how to execute a molecule, see either Chapter 14, “Cell Codec” (Cell decompression molecules) or Chapter 21, “Acceleration in XIL Programs” (all other molecules).

Cell Decompression

decompress_Cell → colorconvert → rescale8 → [scale8nearest] → ordereddither8_1

---

decompress_Cell → nearestcolor8_8 → [scale8nearest] →

SPARC
[[copy8] → display_ioSUNWgx]

---

decompress_Cell → [scale8nearest] → ordereddither8_8

CellB Decompression

decompress_CellB → [rescale8] → [scale8nearest] → ordereddither8_8 →

SPARC
[display_ioSUNWgx]

---

decompress_CellB → colorconvert → [scale8nearest] →

SPARC
[display_ioSUNWgs | display_ioxlib]

x86
[display_ioxlib]

JPEG Baseline Sequential Decompression

decompress_Jpeg → [colorconvert] → [rescale8] → errordiffusion8_1 → [scale1nearest]

---

decompress_Jpeg → [rescale8] → ordereddither8_8 →

SPARC
[[scale8nearest] → [copy8] → display_ioSUNWgx]

---

decompress_Jpeg → colorconvert

H.261 Decompression

decompress_H261 → [rescale8] → ordereddither8_8 →

SPARC
[[scale8nearest] → [copy8] → display_ioSUNWgx]

---

decompress_H261 → colorconvert

MPEG-1 Decompression

decompress_Mpeg1 → [rescale8] → ordereddither8_8→

SPARC
[[scale8nearest] → [copy8] → display_ioSUNWgx]

---

decompress_Mpeg1 → colorconvert

FaxG4 Decompression

decompress_faxG4 → transpose1

Image Filtering

convolve8 → convolve8

---

convolve16 → convolve16

SPARC *Other*

copy8 → display_ioSUNWgx

---

convert1to8 → display_ioSUNWgx

---

lookup1_8 → display_ioSUNWgx

---

lookup8_8 → display_ioSUNWgx

---

rescale8 → display_ioSUNWgx

---

setvalue8 → display_ioSUNWgx

---

translate8nearest → display_ioSUNWgx

XIL Error Messages B

This appendix contains a listing of the error messages that can be generated by the XIL library. Breaks in the sequential numbering of the errors do not indicate missing errors. These are merely gaps that occurred as the numbers were assigned.

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di -1	Out Of Memory	All functions
di-2	Image source1 bands or type does not match destination	`xil_extrema()`, `xil_fill()`, `xil_add()`, `xil_add_const()`, `xil_and()`, `xil_and_const()`, `xil_threshold()`, `xil_black_generation()`, `xil_convolve()`, `xil_copy()`, `xil_dilate()`, `xil_divide()`, `xil_divide_by_const()`, `xil_divide_into_const()`, `xil_erode()`, `xil_histogram()`, `xil_multiply()`, `xil_multiply_const()`, `xil_not()`, `xil_or()`, `xil_or_const()`, `xil_ordered_dither()`, `xil_paint()`, `xil_rescale()`, `xil_set_value()`, `xil_soft_fill()`, `xil_squeeze_range()`, `xil_subsample_adaptive()`, `xil_subtract()`, `xil_subtract_const()`, `xil_subtract_from_const()`, `xil_transpose()`, `xil_xor()`, `xil_cis_get_output_type()`, `xil_decompress()`, xil_absolute(), xil_edge_detection(), xil_max(), xil_min(), xil_tablewarp(), xil_tablewarp_horizontal(), xil_tablewarp_vertical()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-3	Image source 2 does not match destination	xil_extrema(), xil_fill(), xil_add(), xil_add_const(), xil_and(), xil_and_const(), xil_threshold(), xil_black_generation(), xil_convolve(), xil_copy(), xil_dilate(), xil_divide(), xil_divide_by_const(), xil_divide_into_const(), xil_erode(), xil_histogram(), xil_multiply(), xil_multiply_const(), xil_not(), xil_or(), xil_or_const(), xil_ordered_dither(), xil_paint(), xil_rescale(), xil_set_value(), xil_soft_fill(), xil_squeeze_range(), xil_subsample_adaptive(), xil_subtract(), xil_subtract_const(), xil_subtract_from_const(), xil_transpose(), xil_xor(), xil_cis_get_output_type(), xil_decompress(), xil_absolute(), xil_edge_detection(), xil_max(), xil_min(), xil_tablewarp(), xil_tablewarp_horizontal(), xil_tablewarp_vertical(),

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-4	Image source 3 does not match destination	`xil_extrema()`, `xil_fill()`, `xil_add()`, `xil_add_const()`, `xil_and()`, `xil_and_const()`, `xil_threshold()`, `xil_black_generation()`, `xil_convolve()`, `xil_copy()`, `xil_dilate()`, `xil_divide()`, `xil_divide_by_const()`, `xil_divide_into_const()`, `xil_erode()`, `xil_histogram()`, `xil_multiply()`, `xil_multiply_const()`, `xil_not()`, `xil_or()`, `xil_or_const()`, `xil_ordered_dither()`, `xil_paint()`, `xil_rescale()`, `xil_set_value()`, `xil_soft_fill()`, `xil_squeeze_range()`, `xil_subsample_adaptive()`, `xil_subtract()`, `xil_subtract_const()`, `xil_subtract_from_const()`, `xil_transpose()`, `xil_xor()`, `xil_cis_get_output_type()`, `xil_decompress()`
di-5	Operation not implemented	All functions
di-6	Could not get ROI of an image	Asecondary error caused by an internal error. Could occur with any function.
di-7	Could not intersect ROIs	Asecondary error caused by an internal error. Could occur with any function.
di-8	Could not get region list from ROI	Asecondary error caused by an internal error. Could occur with any function.

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-9	Could not get src1 as memory	xil_add(), xil_add_const(), xil_affine(), xil_and(), xil_and_const(), xil_black_generation(), xil_blend(), xil_cast(), xil_convolve(), xil_copy(), xil_copy_pattern(), xil_dilate(), xil_divide(), xil_divide_by_const(), xil_divide_into_const(), xil_erode(), xil_error_diffusion(), xil_extrema(), xil_fill(), xil_rotate(), xil_scale(), xil_translate(), xil_transpose(), xil_histogram(), xil_lookup(), xil_multiply(), xil_multiply_const() xil_nearest_color(), , xil_not(), xil_or(), xil_or_const(), xil_ordered_dither(), xil_paint(), xil_rescale(), xil_soft_fill(), xil_subsample_adaptive(), xil_subsample_binary_to_gray(), xil_subtract(), xil_subtract_const(), xil_subtract_from_const(), xil_threshold(), xil_xor(), xil_edge_detection(), xil_tablewarp(), xil_tablewarp_horizontal(), xil_tablewarp_vertical()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-10	Couldn't get dst1 as memory	xil_add(), xil_add_const(), xil_affine(), xil_and(), xil_and_const(), xil_black_generation(), xil_blend(), xil_cast(), xil_choose_colormap(), xil_color_convert(), xil_convolve(), xil_copy(), xil_copy_pattern(), xil_dilate(), xil_divide(), xil_divide_by_const(), xil_divide_into_const(), xil_erode(), xil_error_diffusion(), xil_extrema(), xil_fill(), xil_rotate(), xil_scale(), xil_translate(), xil_transpose(), xil_lookup(), xil_multiply(), xil_multiply_const(), xil_nearest_color(), xil_not(), xil_or(), xil_or_const(), xil_ordered_dither(), xil_paint(), xil_rescale(), xil_soft_fill(), xil_subsample_adaptive(), xil_subsample_binary_to_gray(), xil_subtract(), xil_subtract_const(), xil_subtract_from_const(), xil_threshold(), xil_xor(), xil_absolute(), xil_copy_with_planemask(), xil_max(), xil_min(), xil_tablewarp(), xil_tablewarp_horizontal(), xil_tablewarp_vertical()
di-11	Could not get alpha as memory	xil_blend()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-12	Could not get combined ROI	xil_add(), xil_add_const(),xil_and(), xil_and_const(), xil_black_generation(), xil_blend(), xil_cast(), xil_color_convert(), xil_copy(), xil_dilate(), xil_divide(), xil_divide_by_const(), xil_divide_into_const(), xil_erode(), xil_error_diffusion(), xil_lookup(), xil_multiply(), xil_multiply_const(), xil_nearest_color(), xil_not(), xil_or(), xil_or_const(), xil_ordered_dither(), xil_paint(), xil_rescale(), xil_subtract(), xil_subtract_const(), xil_subtract_from_const(), xil_threshold(), xil_xor(), xil_absolute(), xil_copy_with_planemask(), xil_max(), xil_min()
di-13	Dependent count overflow	xil_compress(), xil_decompress()
di-14	Could not copy ROI	A secondary error caused by memory resource limitations. Could occur with any function.
di-15	Could not get src2 as memory	A secondary error caused by an internal error. Could occur with any function.
di-16	ROI translation failed	xil_copy_pattern(). Secondary error probably due to resource problems.
di-17	Attempted to get an unknown attribute	xil_cis_get_attribute()
di-18	Attempted to get a set-only attribute	xil_cis_get_attribute()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-19	Attempted to set an unknown attribute	`xil_cis_set_attribute()`
di-20	Attempted to set a get-only attribute	`xil_cis_set_attribute()`
di-21	Cannot create op	A secondary error probably due to resource problems. Could occur with any function.
di-22	Image type mismatch	Any function
di-23	Unable to perform the specified function	Should not happen unless an error occurred during the operation
di-24	Could not get the ROI list	`xil_transpose()`
di-55	XilCis: JPEG attribute error: cannot have interleaved bands with different selector(s) and/or pt_transform values	Reported by the JPEG lossless compressor
di-56	XilCis: JPEG attribute error: DECOMPRESSION_ QUALITY value must be >= 1 and <= 100	`xil_cis_set_attribute()`
di-57	XilCis: JPEG attribute error: ENCODE_INTERLEAVED value must be either TRUE or FALSE	`xil_cis_set_attribute()`
di-58	XilCis: JPEG attribute error: ENCODE_VIDEO value must be either TRUE or FALSE	`xil_cis_set_attribute()`
di-59	XilCis: JPEG attribute error: COMPRESSION_QUALITY value must be >=1 and <=100	`xil_cis_set_attribute()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-60	XilCis: JPEG attribute error: TEMPORAL_FILTER value must be either TRUE or FALSE	`xil_cis_set_attribute()`
di-61	XilCis: JPEG attribute error: OPTIMIZE_HUFFMAN_ TABLES value must be either TRUE or FALSE	`xil_cis_set_attribute()`
di-62	XilCis: JPEG attribute error: COMPRESSED_DATA_ FORMAT value must be either INTERCHANGE or ABBREVIATED_FORMAT	`xil_cis_set_attribute()`
di-63	XilCis: Cell attribute error: Must set DECOMPRESSOR_ COLORMAP before RDWR_INDICES	`xil_cis_set_attribute()` (for Cell)
di-64	XilCis: Cell attribute error: RDWR_INDICES: Invalid NULL index list	`xil_cis_set_attribute()` (for Cell)
di-65	XilCis: Cell attribute error: RDWR_INDICES: Invalid index list count	`xil_cis_set_attribute()` (for Cell)
di-66	XilCis: ROIs not allowed for compress image	Reported by JPEG lossless compressor
di-67	XilCis: JPEG bytestream error: Invalid code value index	`xil_decompress()`. This error can also be generated by `xil_cis_has_frame()`, `xil_cis_number_of_frames()`, and `xil_cis_seek()` if user inserts data into the CIS with `xil_cis_put_bits_[ptr]()` with an unknown or partial number of frames.

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-68	XilCis: JPEG bytestream error: Invalid symbol table ID	`xil_decompress()`. This error can also be generated by `xil_cis_has_frame()`, `xil_cis_number_of_frames()`, and `xil_cis_seek()` if user inserts data into the CIS with `xil_cis_put_bits_[ptr]()` with an unknown or partial number of frames.
di-69	XilCis: JPEG bytestream error: Invalid Huffman code length	`xil_decompress()`. This error can also be generated by `xil_cis_has_frame()`, `xil_cis_number_of_frames()`, and `xil_cis_seek()` if user inserts data into the CIS with `xil_cis_put_bits_[ptr]()` with an unknown or partial number of frames.
di-70	XilCis: JPEG bitstream error: Invalid value for prediction selector	`xil_cis_set_attribute()`
di-71	XilCis: JPEG bitstream error: Invalid value for point transform	`xil_cis_set_attribute()`
di-72	XilCis: JPEG bitstream error: Invalid SOF marker for this decompressor	`xil_decompress()`
di-73	XilCis: JPEG bitstream error: Invalid selector in SOF segment	`xil_decompress()`
di-74	XilCis: JPEG bitstream error: Unsupported image precision in SOF segment	`xil_decompress()`
di-75	XilCis: JPEG bitstream error: Invalid image height in SOF segment	`xil_decompress()`
di-76	XilCis: JPEG bitstream error: Invalid image width in SOF segment	`xil_decompress()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-77	XilCis: JPEG bitstream error: Invalid qtable marker 0	`xil_compress()`
di-78	XilCis: JPEG bitstream error: Table not used by any bands	`xil_compress()`
di-79	XilCis: JPEG bitstream error: Invalid quantizer precision identifier	`xil_cis_create()`, `xil_compress()`, `xil_decompress()`
di-80	XilCis: JPEG bitstream error: Number of quantizer tables too large	`xil_cis_create()`
di-81	XilCis: JPEG bitstream error: Table identifier not in use	`xil_decompress()`
di-82	XilCis: JPEG bitstream error: Invalid table identifier	`xil_decompress()`
di-83	XilCis: JPEG bitstream error: Invalid ac table identifier	`xil_decompress()`
di-84	XilCis: JPEG bitstream error: dc table identifier not in use	`xil_compress()`
di-85	XilCis: JPEG bitstream error: Invalid dc table identifier	`xil_compress()`
di-86	XilCis: JPEG bitstream error: Invalid component identifier	`xil_decompress()`
di-87	XilCis: JPEG bitstream error: attempted use of non loaded ac table	`xil_decompress()`
di-88	XilCis: JPEG bitstream error: attempted use of non loaded dc table	`xil_compress()`
di-89	XilCis: JPEG bitstream error: attempted use of non loaded qtable	`xil_compress()`
di-90	XilCis: JPEG bitstream error: invalid htable identifier	`xil_decompress()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-91	XilCis: JPEG bitstream error: invalid qtable identifier	`xil_cis_create()`, `xil_compress()`
di-92	XilCis: JPEG bitstream error: unknown htable type	`xil_cis_set_attribute()`
di-93	XilCis: JPEG bitstream error: invalid band number	`xil_cis_get_attribute()`
di-94	XilCis: JPEG bitstream error: table index too large	`xil_decompress()`
di-95	XilCis: internal error	`xil_cis_create()`, `xil_cis_get_output_type()`, `xil_cis_get_attribute()`, `xil_compress()`, `xil_decompress()`, `xil_cis_has_data()`, `xil_nearest_color()`, `xil_cis_destroy()`
di-96	XilCis: Wrong number of ops in molecule	`xil_nearest_color()`
di-97	XilCis: Cell bitstream error	`xil_nearest_color()`
di-98	XilCis: JPEG bitstream error	`xil_decompress()`
di-100	XilCis: No data to decompress	`xil_cis_get_output_type()`
di-101	XilCis: Could not create CIS destination image	`xil_decompress()`. Secondary error that occurs when the compressor is unable to allocate a temporary image to compress into.
di-102	XilCis: JPEG cannot encode image as 4:1:1 yuv, using ENCODE_INTERLEAVED attribute instead	`xil_compress()`
di-103	XilCis: JPEG cannot interleave image, it has 5 or more bands	`xil_compress()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-104	XilCis: Wrote more than maximum frame size into buffer space	`xil_compress()`. Internal error. This should never happen.
di-105	XilCis: Seek to frame number not in CIS	`xil_cis_seek()`
di-106	XilCis: Desired frame number no longer in CIS	`xil_cis_seek()`
di-107	XilCis: CIS is empty, cannot seek to frame number	`xil_cis_seek()`
di-108	XilCis: No previous desired frame type to seek backward to	`xil_cis_seek()`. Internal error.
di-109	XilCis: Unable to complete seek by skipping (burning) frames	`xil_nearest_color()`
di-110	XilCis: Complete frame does not exist. No data to decompress	`xil_decompress()`
di-111	XilCis: Image width and height must be a multiple of 4	`xil_cis_get_attribute()`, `xil_decompress()`, `xil_cis_get_output_type()`, `xil_cis_has_data()`
di-112	XilCis: Could not create XilLookup for colormap	`xil_decompress()` (for Cell)
di-113	XilCis: Fax: source not a bit image	`xil_compress()`
di-114	XilCis: Incomplete frame in buffer, cannot copy until complete	`xil_compress()`, `xil_cis_put_bits()`
di-115	Xilcis: start_frame adjusted past read_frame. CIS data lost.	`xil_cis_seek()`, `xil_compress()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-116	XilCis: unsupported or illegal precision for compressed image	`xil_decompress()`. Internal error.
di-117	XilCis: Invalid NULL CIS specified	`xil_cis*()` (ny `xil_cis*` routine)
di-118	XilCis: Partial frame end adjustment error on frame position	`xil_compress()`, `xil_cis_put_bits()`
di-119	XilCis: Partial frame end adjustment error on frame list	`xil_compress()`, `xil_cis_put_bits()`
di-120	XilCis: Cannot create copy op	`xil_compress()`, `xil_decompress()`
di-121	XilCis: Invalid compression type or specification	`xil_compress()`, `xil_decompress()`
di-122	XilCis: Can't create compression op	`xil_decompress()`
di-123	XilCis: Decompress destination image's type or nbands does not match CIS	`xil_decompress()`
di-124	XilCis: Illegal relative_to value in xil_cis_seek()	`xil_cis_seek()`
di-125	XilCis: Compression device is unavailable	`xil_cis_create()`
di-126	XilCis: Couldn't create compression device	`xil_cis_create()`
di-127	Lookup would be too large for data type	`xil_lookup_create()`, `xil_colorcube_create()`
di-128	Colorcube multipliers do not match sizes	`xil_colorcube_create()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-129	Object would be too large for 32-bit addressability	xil_create(), xil_create_from_type(), xil_kernel_create(), xil_dithermask_create(), xil_sel_create(), xil_lookup_create(), xil_colorcube_create(), xil_histogram_create()
di-130	Attempted to read or write outside the range of a lookup	xil_lookup_get_values(), xil_lookup_set_values()
di-131	Invalid NULL lookup	xil_ordered_dither(), xil_error_diffusion(), xil_nearest_color(), xil_lookup_convert(), xil_choose_colormap(), xil_lookup(), xil_lookup_destroy(), xil_lookup_get_input_datatype(), xil_lookup_get_output_datatype(), xil_lookup_get_output_nbands(), xil_lookup_get_num_entries(), xil_lookup_get_version(), xil_lookup_get_offset(), xil_lookup_set_offset(), xil_lookup_get_colorcube(), xil_lookup_get_colorcube_info(), xil_lookup_get_values(), xil_lookup_set_values(), xil_lookup_create_copy()
di-132	Look-up type mismatch	xil_lookup_convert()
di-133	Lookup may not have input type of XIL_FLOAT	xil_lookup_create(), xil_colorcube_create()
di-134	Incorrect data type for image	xil_create()
di-135	Attempted to set attribute on a non-device image	xil_set_device_attribute()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-136	Attempted to get attribute of a non-device image	`xil_get_device_attribute()`
di-137	Requested pixel outside range of image	`xil_set_pixel()`, `xil_get_pixel()`
di-138	Could not get image as memory	`xil_get_memory_storage()`
di-139	Could not propagate image storage	`xil_get_memory_storage()`, `xil_set_memory_storage()`. This is probably the result of another error, such as running out of memory.
di-140	Could not get image storage	Internal system error
di-141	Could not create storage device	`xil_get_pixel()`, `xil_set_pixel()`, `xil_get_memory_storage()`
di-142	Cannot get memory storage on non-exported image	`xil_get_memory_storage()`
di-143	Could not access storage device	Internal system error
di-144	Could not copy colormap	`xil_ordered_dither()`
di-145	Could not create ROI	A secondary error caused by an internal error. Could occur with any function.
di-146	Invalid parameters passed to function	`xil_add()`, `xil_create()`, `xil_create_child()`, `xil_create_copy()`, `xil_set_origin()`, `xil_histogram_create()`, `xil_kernel_create()`, `xil_lookup_create()`, `xil_sel_create()`
di-147	Could not create image	Secondary error in `xil_create()`, `xil_create_child()`, `xil_create_copy()`
di-148	Child count exceeded on parent – could not create child image	`xil_create_child()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-149	Could not create input/output device	Internal system error
di-150	Could not access input/output device	Internal system error
di-151	Image already exported	`xil_export()`
di-152	Image is of a type that cannot be exported	`xil_export()`
di-153	Cannot import image that was not previously exported	`xil_import()`
di-154	Image must be exported to set memory storage	`xil_set_memory_storage()`
di-155	Cannot set memory storage on a child image	`xil_set_memory_storage()`
di-156	Image band mismatch	`xil_cast()`, `xil_copy()`, `xil_affine()`, `xil_blend()`, `xil_copy_pattern()`, `xil_rotate()`, `xil_scale()`, `xil_subsample_binary_to_gray()`, `xil_translate()`, `xil_absolute()`, `xil_edge_detection()`, `xil_max()`, `xil_min()`
di-157	Undefined image data type encountered	Internal error
di-158	Could not load `.so` file	`xil_open()`, `xil_create_from_device()`
di-159	Could not extract symbols from `.so` file	`xil_open()`
di-160	Could not expand $XILHOME path	`xil_open()`
di-161	Could not open $XILHOME/lib/ xil.compute file	`xil_open()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-162	XIL Xlib I/O Driver: could not get graphics context	xil_create_from_window()
di-163	Could not copy kernel	xil_convolve(), xil_kernel_create(), xil_error_diffusion()
di-164	Input and output must be in CMYK color space	xil_black_generation()
di-165	Alpha channel image must have only 1 band	xil_blend()
di-166	Invalid size for lookup	xil_choose_colormap(), xil_lookup_create(), xil_error_diffusion(), xil_nearest_color()
di-167	Image must have 3 bands	xil_choose_colormap()
di-169	SEL must be specified	xil_dilate(), xil_erode()
di-170	Could not copy SEL	xil_dilate(), xil_erode(), xil_sel_create_copy()
di-171	Coefficient(s) would cause a divide by zero	xil_divide()
di-172	Supplied coefficient(s) contained NaN	xil_divide(), xil_multiply(), xil_subtract()
di-173	Number of bands in src image does not match look-up table	xil_error_diffusion(), xil_nearest_color(), xil_ordered_dither()
di-174	Data type of dest image does not match look-up table input type	xil_error_diffusion(), xil_nearest_color()
di-175	Data type of src image does not match look-up table output type	xil_error_diffusion(), xil_nearest_color(), xil_ordered_dither()
di-176	Output image must be single banded	xil_error_diffusion(), xil_nearest_color(), xil_ordered_dither()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-177	Could not copy lookup	`xil_error_diffusion()`, `xil_nearest_color()`, `xil_lookup_create_copy()`
di-179	Data type of src image does not match look-up table input type	`xil_lookup()`
di-180	Data type of dest image does not match look-up table output type	`xil_lookup()`
di-181	Input image must be single banded	`xil_lookup()`
di-182	Brush must be specified	`xil_paint()`
di-183	Could not copy brush	`xil_paint()`
di-184	XilCis: Error AdjustStart	Internal error
di-185	Attempt to insert at illegal position in linked list	Internal error
di-186	Attempt to delete at illegal position in linked list	Internal error
di-187	NULL ROI passed to intersect routine	`xil_roi_intersect()`
di-188	Could not create imagetype	`xil_cis_get_output_type()`, `xil_cis_get_input_type()`, `xil_create()`, `xil_create_child()`, `xil_create_copy()`, `xil_create_from_device()`, `xil_create_from_type()`, `xil_create_from_window()`
di-189	Could not create kernel	`xil_kernel_create()`
di-190	Could not create dithermask	`xil_dithermask_create()`
di-191	Could not create SEL	`xil_sel_create()`
di-192	Could not create lookup	`xil_lookup_create()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-193	Could not branch to op	Internal error
di-194	No capture operation defined for this device	`xil_copy()`
di-195	No display operation defined for this device	`xil_copy()`
di-197	Internal error in getting op parameters	Internal error
di-198	Internal error in setting op parameters	Internal error
di-199	Internal error searching for optimized execution path	Internal error
di-202	Source color space not specified	`xil_color_convert()`
di-203	Destination color space not specified	`xil_color_convert()`
di-204	Unsupported color conversion	`xil_color_convert()`
di-205	Could not get colorcube information	Internal error
di-206	Lookup must be a colorcube	`xil_ordered_dither()`
di-207	Invalid NULL image specified	All functions
di-208	Tried to remove an error handler which was not installed	`xil_remove_error_handler()`
di-209	Cannot get system state from image	Internal error
di-210	Cannot create lookup for returning data	`xil_squeeze_range()`
di-211	Cannot squeeze range of multiband image	`xil_squeeze_range()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-212	XIL I/O Driver: could not open specified frame buffer	`xil_create_from_device()`
di-213	XIL I/O Driver: could not perform ioctl on frame buffer	`xil_create_from_device()`
di-214	XIL I/O Driver: frame buffer not of expected type	`xil_create_from_device()`
di-215	XIL I/O Driver: could not mmap device	`xil_create_from_device()`
di-216	Could not get FLOAT image as memory	Internal error
di-217	XIL Memory Driver: Could not get named storage type	A secondary error that can be seen after any call that creates or copies images or CISs
di-218	XIL Memory Driver: Could not create named storage device	A secondary error that can be seen after any call that creates or copies images or CISs
di-219	XIL I/O Driver: could not connect to dga	`xil_create_from_device()`
di-220	XIL I/O Driver: could not get window geometry	`xil_create_from_device()`
di-221	Invalid NULL kernel specified	`xil_convolve()`, `xil_error_diffusion()`, `xil_kernel_get_width()`, `xil_kernel_get_height()`, `xil_kernel_get_key_x()`, `xil_kernel_get_key_y()`, `xil_kernel_create_copy()`
di-222	Invalid NULL dither mask specified	`xil_dithermask_get_width()`, `xil_dithermask_get_height()`, `xil_dithermask_get_nbands()`, `xil_dithermask_create_copy()`, `xil_error_diffusion()`, `xil_ordered_dither()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-223	Tried to write to an invalid CIS	`xil_compress()`
di-224	Tried to read from an invalid CIS	`xil_decompress()`
di-230	XilCis: Cell attribute failed because too few bytes per frame group	`xil_cis_set_attribute()` (`BITS_PER_SECOND`)
di-231	XilCis: Cell key frame interval is too large for bit rate control	`xil_cis_set_attribute()` (`KEYFRAME_INTERVAL`)
di-232	XilCis: Cell key frame interval exceeds maximum	`xil_cis_set_attribute()` (`KEYFRAME_INTERVAL`)
di-233	Must specify histogram	`xil_histogram()`
di-234	Number of bands in image and histogram do not match	`xil_histogram()`
di-235	Zero bins in histogram not allowed	`xil_create_histogram()`
di-236	Cannot create histogram	`xil_create_histogram()`
di-237	Could not initialize compression type	`xil_cis_create()`
di-238	Could not initialize input/output type	`xil_create_from_device()`
di-239	Could not set image attribute	`xil_set_attribute()`. A secondary error.
di-240	Error opening or parsing $XILHOME/lib/xil.modules	`xil_open()`
di-241	Error loading Xlib display driver	`xil_create_from_window()`
di-242	Couldn't create display image	`xil_create_from_window()`. A secondary error.

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-243	Number of bands cannot be 0	`xil_lookup_create()`
di-244	Cannot create system state	`xil_open()`. A secondary error.
di-245	Invalid dither mask dimensions	`xil_dithermask_create()`
di-246	Cannot create global state	`xil_open()`. A secondary error.
di-247	Cannot initialize localization information	`xil_open()`
di-248	Unable to load any compute device handlers	`xil_open()`. Probably a secondary error due to improper configuration.
di-249	XIL Xlib I/O Driver: could not get window attributes	`xil_create_from_window()`
di-250	XIL Xlib I/O Driver: could not create ximage	`xil_create_from_window()`
di-251	Number of bands in src image does not match dither mask	`xil_ordered_dither()`
di-252	Invalid colorcube dimensions	`xil_colorcube_create()`
di-253	XilCis: Cell compress image width is not an even multiple of 4	`xil_compress()`
di-254	XilCis: Cell compress image height is not an even multiple of 4	`xil_compress()`
di-255	Could not copy image	`xil_create_copy()`
di-256	Could not install error handler	`xil_install_error_handler()`
di_257	Invalid NULL ROI specified	Any of the `xil_roi*()` functions
di-258	Seed pixel not within fill region defined by fill color	`xil_soft_fill()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-259	Invalid NULL data pointer	xil_dithermask_create(), xil_histogram_create(), xil_kernel_create(), xil_lookup_get_values(), xil_lookup_set_values(), xil_sel_create(), xil_device_create(), xil_interpolation_table_create()
di-260	Invalid NULL system state	xil_create(), xil_create_from_type(), xil_create_from_device(), xil_create_from_window(), xil_cis_create(), xil_roi_create(), xil_kernel_create(), xil_dithermask_create(), xil_sel_create(), xil_lookup_create(), xil_colorcube_create(), xil_histogram_create(), xil_close(), xil_state_get_synchronize(), xil_state_set_synchronize(), xil_state_get_show_action(), xil_state_set_show_action(), xil_install_error_handler(), xil_remove_error_handler(), xil_image_get_by_name(), xil_lookup_get_by_name(), xil_imageytpe_get_by_name(), xil_cis_get_by_name(), xil_dithermask_get_by_name(), xil_kernel_get_by_name(), xil_sel_get_by_name(), xil_roi_get_by_name(), xil_histogram_get_by_name(), xil_colorspace_get_by_name()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-261	XilCis: Attempted to set Cell encoding type to unknown value	`xil_cis_set_attribute()` (“`CELL_ENCODING_TYPE`”)
di-262	Invalid NULL SEL specified	`xil_sel_get_width()`, `xil_sel_get_height()`, `xil_sel_get_key_x()`, `xil_sel_get_key_y`, `xil_sel_cretae_copy()`
di-263	Invalid kernel value (NaN)	`xil_kernel_create()`
di-264	Invalid kernel value (infinity or -infinity)	`xil_kernel_create()`
di-265	Invalid NULL histogram specified	`xil_histogram_get_values()`, `xil_histogram_get_limits()`, `xil_histogram_get_info()`, `xil_histogram_get_nbins()`, `xil_histogram_get_nbands()`, `xil_histogram_destroy()`
di-266	Invalid image dimensions specified	`xil_create()`, `xil_create_from_type()`
di-267	Could not copy dither mask	`xil_ordered_dither()`, `xil_dithermask_create_copy()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-268	Invalid NULL object pointer	xil_lookup_get_version(), xil_object_get_type(), xil_cis_get_name(), xil_cis_set_name(), xil_dithermask_get_name(), xil_dithermask_set_name(), xil_get_name(), xil_set_name(), xil_histogram_get_name(), xil_histogram_set_name(), xil_imagetype_get_name(), xil_imagetype_set_name(), xil_kernel_get_name(), xil_kernel_set_name(), xil_lookup_get_name(), xil_lookup_set_name(), xil_roi_get_name(), xil_roi_set_name(), xil_sel_get_name(), xil_sel_set_name()
di-269	Could not set image ROI	Internal error
di-271	Invalid dither mask value (NaN)	xil_dithermask_create()
di-272	Invalid dither mask value (greater than 1.0 or less than 0.0)	xil_dithermask_create()
di-273	Invalid SEL value (must be 0 or 1)	xil_sel_create()
di-274	Could not create internal symbol table	xil_cis_create()
di-275	Could not create internal Cell compressor object	xil_cis_create()
di-276	Internal error op number not found	Internal error
di-277	Could not create internal CisBufferManager object	xil_cis_create()

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-278	Could not create internal base XilDeviceCompression object	`xil_cis_create()`
di-279	Could not set object name	`xil_cis_set_name()`, `xil_dithermask_set_name()`, `xil_set_name()`, `xil_histogram_set_name()`, `xil_imagetype_set_name()`, `xil_kernel_set_name()`, `xil_lookup_set_name()`, `xil_roi_set_name()`, `xil_sel_set_name()`. A secondary error.
di-280	Could not create JpegLL compressor object	`xil_cis_create()`
di-281	Could not create JpegLL decompressor object	`xil_cis_create()`
di-282	Could not create XilCis	`xil_cis_create()`
di-283	XilCis: Cell user data size is too large	`xil_cis_set_attribute()` (“`CELL_USER_DATA`”)
di-284	XilCis: bitstream error: Invalid SOF marker	`xil_decompress()`, `xil_cis_number_of_frames()`, `xil_cis_get_attribute()`, `xil_cis_set_attribute()`, `xil_cis_get_input_type()`, `xil_cis_get_output_type()`, `xil_cis_attempt_recovery()`, `xil_cis_get_bits_ptr()`
di-285	XilCis: Invalid bitstream parameters	`xil_decompress()`, `xil_cis_number_of_frames()`, `xil_cis_get_attribute()`, `xil_cis_set_attribute()`, `xil_cis_get_input_type()`, `xil_cis_get_output_type()`, `xil_cis_attempt_recovery()`, `xil_cis_get_bits_ptr()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-286	Could not create full set of standard objects	`xil_open()`; a secondary error
di-287	Insert of operation failed	`xil_error_diffusion()`, `xil_nearest_color()`, `xil_ordered_dither()`, `xil_squeeze_range()`. A secondary error.
di-288	XilCis: image for compress has ROIs or non-zero origin	`xil_compress()`
di-289	Invalid data in kernel for use with error diffusion	`xil_error_diffusion()`
di-290	Partial frame start adjustment error on frame position	`xil_decompress()` (with partial frame in the CIS), `xil_cis_number_of_frames()`, `xil_cis_get_attribute()`, `xil_cis_set_attribute()`, `xil_cis_get_input_type()`, `xil_cis_get_output_type()`, `xil_cis_attempt_recovery()`, `xil_cis_get_bits_ptr()`
di-291	Partial frame flag not set on buffer with partial frame	`xil_decompress()`, `xil_cis_number_of_frames()`, `xil_cis_get_attribute()`, `xil_cis_set_attribute()`, `xil_cis_get_input_type()`, `xil_cis_get_output_type()`, `xil_cis_attempt_recovery()`, `xil_cis_get_bits_ptr()`
di-292	Partial frame start adjustment error--mismatch with buffer start	`xil_decompress()`, `xil_cis_number_of_frames()`, `xil_cis_get_attribute()`, `xil_cis_set_attribute()`, `xil_cis_get_input_type()`, `xil_cis_get_output_type()`, `xil_cis_attempt_recovery()`, `xil_cis_get_bits_ptr()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-293	Current buffer manager read frame unknown	Internal error
di-294	XilCis: could not derive output type from data in CIS	`xil_decompress()`
di-295	Colorspace and image's number of bands mismatch	`xil_set_colorspace()`
di-296	Invalid NULL colorspace specified	`xil_color_convert()`. An internal error.
di-297	Could not copy colorspace	`xil_color_convert()`. An internal error.
di-298	Could not create colorspace	`xil_open()`. An internal error, probably due to running out of memory.
di-299	Number of bands in destination image does not match lookup table	`xil_lookup()`
di-300	Could not create internal fax compressor	`xil_compress()`
di-301	Invalid NULL X region specified	`xil_roi_add_region()`
di-302	Invalid colorcube multipliers	`xil_colorcube_create()`
di-303	XilCis: Invalid Cell key frame interval – negative	`xil_cis_set_attribute()` `(KEYFRAME_INTERVAL)`
di-304	XilCis: Illegal seek backward in a non-random access CIS	`xil_cis_seek()`
di-305	Bad value for Xil_boolean type	`xil_cis_set_attribute()` `(TEMPORAL_FILTERING,` `COLORMAP_ADAPTATION)`
di-306	XilCis: Cell colormaps must contain at least 2 colors	`xil_cis_set_attribute()` `(COMPRESSOR_MAX_CMAP_SIZE,` `COMPRESSOR_COLORMAP)`
di-307	Cell colormaps cannot contain more than 256 colors	`xil_cis_set_attribute()` `(COMPRESSOR_MAX_CMAP_SIZE)`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-308	XilCis: Cell COMPRESSOR_MAX_ CMAP_SIZE attribute cannot be changed after first xil_compress() call	`xil_cis_set_attribute()` `(COMPRESSOR_MAX_CMAP_SIZE)`
di-310	XilCis: Quantizer value out of range	`xil_cis_set_attribute()` `(QUANTIZATION_TABLE)`
di-311	XilCis: Failure decompressing Cell header	`xil_decompress()`
do-312	XilCis: Illegal H.261 bitstream	`xil_decompress()`
di-313	Compressor for px64 device compression not yet implemented	`xil_compress()`
di-314	XilCis: unable to sync up with current frame PSC in H.261 bitstream	`xil_decompress()`
di-315	XIL I/O Driver: uname system call failed	`xil_create_from_window()`
di-316	XIL Xlib I/O Driver: could not get ximage	`xil_get_pixel()`
di-317	XilCis: Error in seeking. No frames to skip/burn	`xil_cis_seek()`
di-318	XilCis: attribute error: value must be either TRUE or FALSE	`xil_cis_set_attribute()`
di-319	XilCis: no user data established for already scanned Mpeg frame	`xil_cis_seek()`
di-320	XIL I/O Driver: could not get frame buffer info	`xil_create_from_window()`
di-321	Cannot resize a child image	Internal error

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-322	Cannot resize non-device images	Internal error
di-323	Compressor for MPEG-1 device compression not yet implemented	`xil_compress()`
di-324	XilCis: attribute error: value out of range	`xil_cis_set_attribute()`
di-325	XilCis: D (DC intra-coded) frame type not supported	`xil_decompress()`
di-326	Because of Bug ID # 1139760, di-326 prints the error message for di-325. The correct error message for di-326 is: Matrix size does not match with src or dst number of bands	`xil_band_combine()`
di-327	Tried to use an invalid image in an operation	Any operations with an image that becomes invalid. An image becomes invalid when an operation that is supposed to write to it fails.
di-328	Internal error in ROI scan conversion	Internal error
di-329	XilCis: Mpeg1 bitstream error, incomplete frame	Any routine that involves a seek operation within a CIS
di-330	XIL Xlib I/O Driver: could not capture image – window not mapped	`xil_copy()`
di-331	Null lookup list or zero number of lookups to combine	`xil_lookup_create_combined()`
di-332	Input type of lookups are not of same type	`xil_lookup_create_combined()`
di-333	Output type of lookups are not of same type	`xil_lookup_create_combined()`

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-334	Not all lookups have a single band	`xil_lookup_create_combined()`
di-335	Requested lookup table doesn't exist -- num_bands too large	`xil_lookup_create_combined()`
di-336	Wrong lookup type -- not a combined lookup object	`xil_lookup_create_combined()`
di-337	XilCis: wrote more than max_frame_size bytes into frame	xil_compress()
di-338	XilCis: moveEndStart error, read frame must be at end of buffer	xil_cis_get_bits_ptr()
di-339	XilCis: moveEndStart error, read frame must have following frame	xil_cis_get_bits_ptr()
di-340	XilCis: removeStartFrame error, read frame must be start of buffer	xil_cis_get_bits_ptr()
di-341	XilCis: mpeg1 seek, prev_nonbframe_id is frame which already has valid display_id	xil_cis_seek(), xil_decompress() for MPEG-1 CIS
di-342	XilCis: doneBufferSpace called with negative number of bytes, illegal	xil_compress()
di-343	Unsupported edge detection method	`xil_edge_detection()`
di-344	Invalid NULL attribute specified	xil_attribute_set_value(), xil_device_create()
di-347	ROI rectangle width/height cannot be <= 0	Any function that passes an image as an argument

Table B-1 XIL Library Error Messages

Error Number	Error Message	Function That Generates Error
di-348	NULL warp table specified	`xil_tablewarp()`, `xil_tablewarp_horizontal()`, `xil_tablewarp_vertical()`
di-349	Warp table datatype mismatch	`xil_tablewarp()`, `xil_tablewarp_horizontal()`, `xil_tablewarp_vertical()`
di-350	Warp table nbands mismatch	`xil_tablewarp()`, `xil_tablewarp_horizontal()`, `xil_tablewarp_vertical()`
di-351	Could not create Interpolation Table	xil_affine(), xil_interpolation_table_create(), xil_rotate(), xil_scale(), xil_tablewarp(), xil_tablewarp_horizontal(), xil_tablewarp_vertical(), xil_translate()
di-352	Invalid NULL Interpolation Table specified	`xil_interpolation_table_create()`
di-354	X & Y skip values must have value greater than 0	xil_histogram(), xil_histogram_create()
di-356	Can not open specified data file(s)	Could occur with any function whose source image is a Photo CD device image.
di-358	width and height must be defined before parsing XilCis	xil_decompress() of a CellB CIS
di-359	Cannot read from a write-only device	Could occur with any function that tries to read from a write-only device
di-360	Cannot write to a read-only device	Could occur with any function that tries to write to a read-only device

XIL-XGL Interoperability

The current release of the XIL imaging library includes two functions that enable you to process an image using both XIL and XGL 3.*x* calls in the same program.

Note – This interface for XIL-XGL interoperability may change or be replaced by another mechanism in a future release.

One of these functions, `xil_to_xgl()`, enables you to convert an XIL memory image to an XGL memory raster *or* to convert an XIL display image to an XGL window raster. The prototype for this function is shown below.

```
Xgl_ras xil_to_xgl(XilImage src, Xgl_sys_state xgl_state);
```

The parameter `src` is a handle to an XIL memory image or display image, and `xgl_state` is a handle to an XGL System State object. If the function is successful, the return value will be a handle to an XGL memory raster or window raster. If the functions fails, the return value will be `NULL`.

The other function, `xgl_to_xil()`, enables you to convert an XGL memory raster to an XIL memory image or to convert an XGL window raster to an XIL display image.

```
XilImage xgl_to_xil(Xgl_ras src, XilSystemState xil_state);
```

The parameter `src` is a handle to an XGL memory raster or window raster, and `xil_state` is an XIL system state. The function's return value is a handle to an XIL memory image or display image. If the function fails, the return value will be `NULL`.

Table C-1 below lists the specific conversions that these functions make possible. Within each row, the XIL image listed in the left column can be converted to the XGL object in the right column and vice versa.

Table C-1 XIL-XGL Interoperability

XIL Image	XGL Object
A one-band `XIL_BYTE` memory image	A memory raster with the attribute `XGL_COLOR_INDEX`
A four-band `XIL_BYTE` memory image	A memory raster with the attribute `XGL_COLOR_RGB`
A one-band `XIL_BYTE` display image	A window raster with the attribute `XGL_COLOR_INDEX`
A three-band `XIL_BYTE` display image	A window raster with the attribute `XGL_COLOR_RGB`

Note – The conversions shown in the first row of Table C-1 are possible only if the width of the XIL image is an even number.

Once you have both an XIL and an XGL handle to an object, you can process it using either XIL functions or XGL functions. The only restriction on combining these calls is that there cannot be any outstanding deferred operations on the object at the time you switch from making XIL calls to XGL calls, or vice versa. This means that if you have been processing an image using XIL functions, you must call `xil_sync()` to flush all outstanding operations on the image before operating on the image using XGL functions. Similarly, if you have been processing a raster using XGL functions, you must call `xgl_context_flush()` before performing any XIL operations on the raster.

One last note. The XIL-XGL conversion functions are not literally part of the XIL library. Instead, they reside in a separate library called `libxil_to_xgl.a`. This means that programs that use the conversion functions must link with

- `-lxil`
- `-lxgl`
- `-lxil_to_xgl`

Cell and CellB Bytestream Definitions

D

This appendix provides an overview of the Cell technology, then focuses on the descriptions of the codes that are used in a Cell bytestream. Although we include some rationale for the existence of the various codes, we do not provide specific implementation information about Cell encoders or decoders.

Introduction to Cell

The Cell image compression technology, which was developed by Sun Microsystems, provides high quality, low bit-rate image compression at low computational cost. Applications where Cell compression can be used include videoconferencing, media distributions on CD-ROM, and multimedia mail applications. The bytestream that is produced by a Cell encoder is variable length and is made up of instructional codes and information about the compressed video image.

There are two versions of the Cell image compression technology: Cell and CellB. Cell compression, which is designed for use with movies, is computationally asymmetric; it takes longer to compress video data than it does to decompress the data. To provide high-quality images, Cell supports Adaptive Colormap Selection, which enables an encoder to change colormaps dynamically. The bytestream of images that are compressed with Cell can be decompressed in software on any SPARCstation™ workstation. Depending on compression ratios, 640x480-resolution movies can be played back at 30 frames/second (fps). Multiple smaller movies can be displayed simultaneously at this rate.

CellB, which is derived from Cell, is designed for use in videoconferencing applications. To reduce computational overhead and meet the timing demands of videoconferencing, CellB compression is more computationally symmetric than Cell. CellB uses a fixed colormap and is designed to use vector quantization techniques. The bytestream of a CellB image is simpler (and more compact) than that of a Cell image, which reduces requirements on network bandwidth.

Encoding Images for Cell

Cell works with 3-band RGB images, with no subsampling, and requires that the width and height of the images be divisible by four. CellB takes industry-standard 3-band YC_bC_r video as input.

A cell encoder breaks the video into cells. A cell is 16 pixels, arranged in a 4-by-4 group (see Figure D-1). Cells are encoded into the bytestream in scanline order, from left to right and from top to bottom.

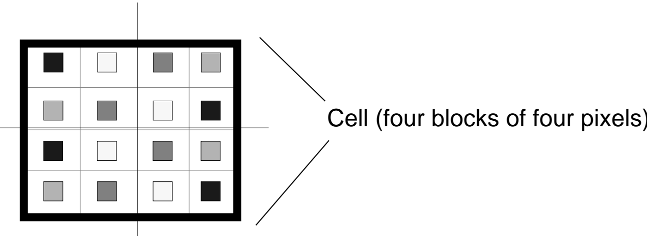

Figure D-1 Cell

The basic encoding scheme used in both versions of Cell is based on an image coding method called Block Truncation Coding (BTC). The 16 pixels in a cell are represented by a 16-bit mask and two colors. The values in the mask specify which color to place at each of the pixel positions. The mask and colors can be chosen to maintain certain statistics of the cell or to reduce contouring in a manner similar to ordered dither.

The primary advantage of BTC is that its decoding process is similar to the operation of character fonting in a color frame buffer. The character display process for a frame buffer takes as input a foreground color, a background color, and a mask that specifies whether to use the foreground or background

color at each pixel. Because this function is so important to the window system, it is often implemented as a display primitive in graphics accelerators. The Cell compression technique leverages these existing primitives to provide full-motion video decoding without special hardware or modifications to the window system.

The basic component of the Cell and CellB bytestreams is the four-byte cell code. The first two bytes of the cell code are a bitmask (Figure D-2). Each bit in the mask represents a pixel in a 16-bit cell. The bitmask is normalized so that the most significant bit is 0. The figure shows the relationship of the bits in the mask to the location of the pixels in a cell. A value of 0 in a mask bit means that the pixel is rendered in the background color (color 0). A value of 1 means that the pixel is rendered in the foreground color (color 1).

The last two bytes of a cell code establish a pixel's color. Cell and CellB differ in the way that pixel colors are derived. In Cell, the values in the color 0 and color 1 bytes are indexes into a colormap. In CellB, the values are indexes into Y/Y and C_b/C_r vector quantization tables.

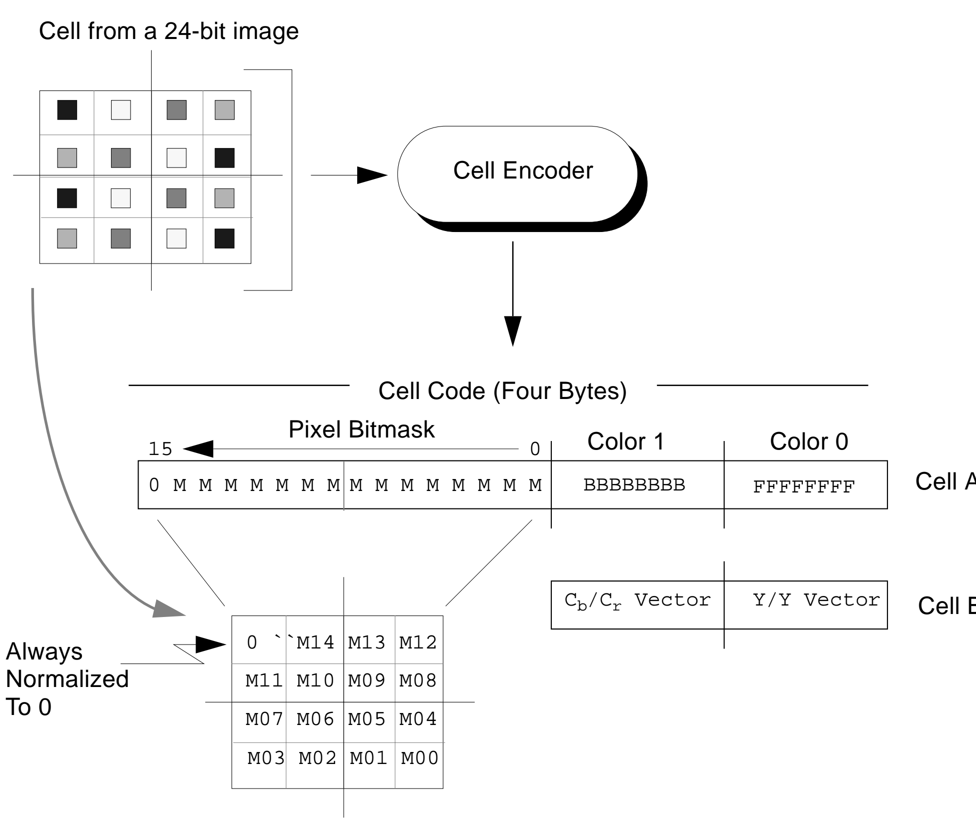

Figure D-2 Encoding a Cell

Cell Bytestream Description

The Cell bytestream can contain the following codes:

- Key-frame headers
- Key parameters
- Cell codes
- Run length codes
- Escape codes

At a minimum, the Cell bytestream must contain a key-frame header, a key frame, and cell codes. More compression is provided by using run length codes for cells. The escape codes enable you to skip frames, load colormaps, change masks and colors, and include non-video (user) data in the bytestream.

The codes are described in the following sections.

Key-Frame Header and Key Parameters

A Cell bytestream contains periodic "key" frames, which are free of interframe escape codes. Every key frame is preceded by a key frame header, which is described below. Immediately after the key frame header, the colormap is encoded. If the colormap has not changed since the previous frame, it is repeated anyway. The first frame of a Cell movie is always a key frame.

Key-Frame Header

The key frame header begins with a series of 8 bytes of all 1's (0xff). This code is followed by a number of parameters, which are introduced by 1-byte codes. The end of a key frame header is marked by an all-zero byte.

```
Start of Key Frame Header
11111111  11111111  11111111  11111111
11111111  11111111  11111111  11111111
        .
        .
        .
   Parameters
        .
        .
        .
End of Key Frame Header
00000000
```

Image Dimensions Key-Frame Code

The following code precedes the 16-bit unsigned width and height of an image.

Code	Width	Height
00000001	WWWWWWWW WWWWWWWW	HHHHHHHH HHHHHHHH

Frame Rate Key-Frame Code

The following code precedes a 32-bit unsigned number that sets the frame rate in microseconds per frame.

Code	Frame Rate
00000010	RRRRRRRR RRRRRRRR RRRRRRRR RRRRRRRR

Size of Colormap Key-Frame Code

The following code precedes a byte that specifies the maximum size (Count+1) of a colormap in the bytestream.

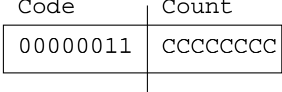

Reserved Key-Frame Codes

The key frame header codes in the range 00000000 to 11111111 (inclusive) are reserved.

Example Key-Frame Header

As an example, a key frame header for a 320-by-240 movie with a frame rate of 30 fps and a maximum colormap size of 240 would look like this (in hex):

```
ff ff ff ff ff ff ff ff        Start of key frame header
01 01 40 00 f0                 Dimensions 320 (0x140) by 240 (0xf0)
02 00 00 82 35                 33,333 (0x8235) usecs per frame
03 ef                          Max colormap size is 239 (0xef) + 1
00                             End of key frame header
c0 ...                         Escape code for colormap
```

Cell Code

The four-byte code that describes a 4-by-4 Cell cell is shown below. The values in the Color1 and Color0 bytes are indexes into the current colormap. If a pixel's bit is set, the color that is indexed in the Color1 byte is used. If not set (0), the pixel is rendered with the color that is indexed by the value of the Color0 byte.

4x4 Bitmask	Color1	Color0
0MMMMMMM MMMMMMMM	FFFFFFFF	BBBBBBBB

The relationship of the bitmask bits to the pixels in a cell is shown below.

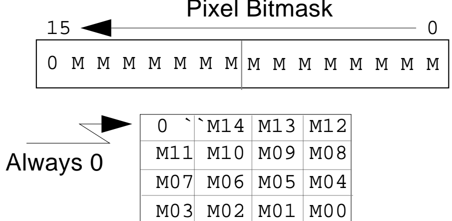

Run Length Code

Runs of cells that use the same color can be described by a run length code, which is shown below. The code causes the next Count+1 cells to be rendered in the color that is specified in the Color0 byte.

Runlength Code	Color0	Count
00000000 00000000	BBBBBBBB	CCCCCCCC

Escape Codes

A Cell encoder can use the following escape codes to perform operations such as loading colormaps, skipping frames, and including user data in the compressed bytestream. All escape codes start with a 1 in the code's most significant bit.

Skip Cells

The following code causes the next N+1 cells to be skipped (to a maximum of 64 cells).

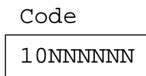

Load a New Colormap

The following code loads a new colormap. Count+1 byte-triples (red byte, green byte, and blue byte) follow the Count byte.

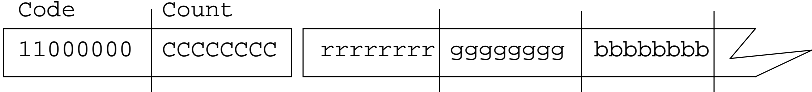

Use Same Mask with New Colors

The following code uses the same mask as the previous cell, but with new colors.

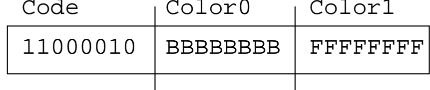

Use Same Mask with Color 0

The following code uses the same mask as the previous cell, but with a new color as indexed by the value of the Color0 byte.

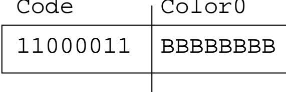

Use Same Mask with Color 1

The following code uses the same mask as the previous cell, but with a new color as indexed by the value of the Color1 byte.

Code	Color1
11000100	FFFFFFFF

Use Same Colors with New Mask

The following code uses the same colors as the previous cell, but with a new bitmask.

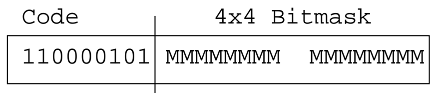

Include User Data

The following code enables you to include your own data in the bytestream. The next N+1 bytes in the bytestream are user data.

Code	Byte Count			
11000110	NNNNNNNN NNNNNNNN NNNNNNNN		uuuuuuuu uuuuuuuu	

Skip an Entire Frame

The following code causes the next frame to be skipped.

Code
11000111

First Byte of Key-Frame Header

The following code is the first byte of a key-frame header, which was described earlier.

Code
11111111

Reserved Codes

The codes in the range 11001000 to 11111110 (inclusive) are reserved.

Summary of Cell Codes

Table D-1 lists the Cell codes.

Table D-1 Cell Bytestream Codes

Code	Description
`0MMMMMMM MMMMMMMM FFFFFFFF BBBBBBBB`	Cell code: M=bit mask, F=color 1 index, B=color 0 index
`00000000 00000000 BBBBBBBB CCCCCCCC`	Run length code: B=color 0 index, cell count = C+1
`10NNNNNN`	Interframe skip: cell count N+1
`11000000 CCCCCCCC`	Load new colormap: count of entries (RGB triples) that follow = C+1

Table D-1 Cell Bytestream Codes

Code	Description
`11000010 BBBBBBBB FFFFFFFF`	Same mask, new colors: B=color 0 index, F=color 1 index
`11000011 BBBBBBBB`	Same mask, new color 0: B=color 0 index
`11000100 FFFFFFFF`	Same mask, new color 1: F=color 1 index
`11000101 0MMMMMMM MMMMMMMM`	Same colors, new mask: M=bit mask
`11000110 NNNNNNNN NNNNNNNN NNNNNNNN`	User data: count of user bytes that follow = C+1
`11001000-11111111`	Reserved
`11111111 11111111 11111111 11111111 11111111 11111111 11111111 11111111`	Start of key frame header
`00000001 WWWWWWWW WWWWWWWW HHHHHHHH HHHHHHHH`	Key frame, image dimensions: W=width, H=height
`00000010 RRRRRRRR RRRRRRRR RRRRRRRR RRRRRRRR`	Key frame, frame rate in usec: R=frame rate
`00000011 CCCCCCCC`	Key frame, maximum colormap size = C+1
`00000000`	End of key frame header

CellB Bytestream Description

CellB compression is designed for use in videoconferencing applications. Rather than indexing into a colormap to determine pixel colors (as in Cell), CellB is designed to be used with vector quantization and dequantization techniques in the YC_bC_r color space.

The CellB bytestream contains no information about image size and frame rates. It's the responsibility of the videoconferencing application to provide this information.

The CellB bytestream consists of:

- Cell codes
- Skip codes
- Quantization-table specification codes

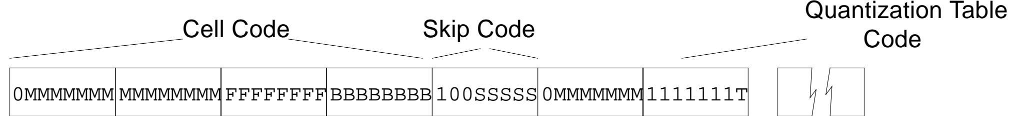

Cell Code

The four-byte code that describes a 4-by-4 CellB cell is shown below. There are two possible luminance (Y) levels for each cell but only one pair of chrominance (C_b and C_r) values.

The value in the C_b/C_r byte represents the chrominance component of the cell. The value in the Y/Y byte represents two luminance values (Y_0 and Y_1) that can represent the cell's luminance.

4x4 Bitmask		U/V Code	Y/Y Code
0MMMMMMM	MMMMMMMM	UVUVUVUV	YYYYYYYY

If a pixel's bit is set in the 4-by-4 bitmask, the color that is rendered for the pixel is $<Y_1,C_b,C_r>$. If a pixel's bit is not set, the pixel is rendered by the color represented by $<Y_0,C_b,C_r>$.

The relationship of the bitmask bits to a cell's pixels is shown below.

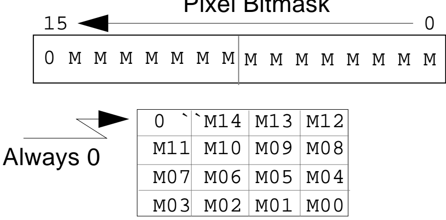

C_b/C_r Quantization Table

The C_b/C_r field of the CellB bytecode represents the chrominance component of the cell. This C_b/C_r code is an index into a table of vectors that represent two independent components of chrominance. Figure D-3 on page 476 shows the default chrominance table. The section “Cb/Cr Table Values” on page 484 contains a list of the values used in the table.

The distribution of values in this default table is based on the observation that, in videoconferencing applications, a cell’s C_b/C_r vectors are clustered around the origin (0,0). Therefore, the C_b/C_r codeword is circularly symmetric, with higher densities near the origin.

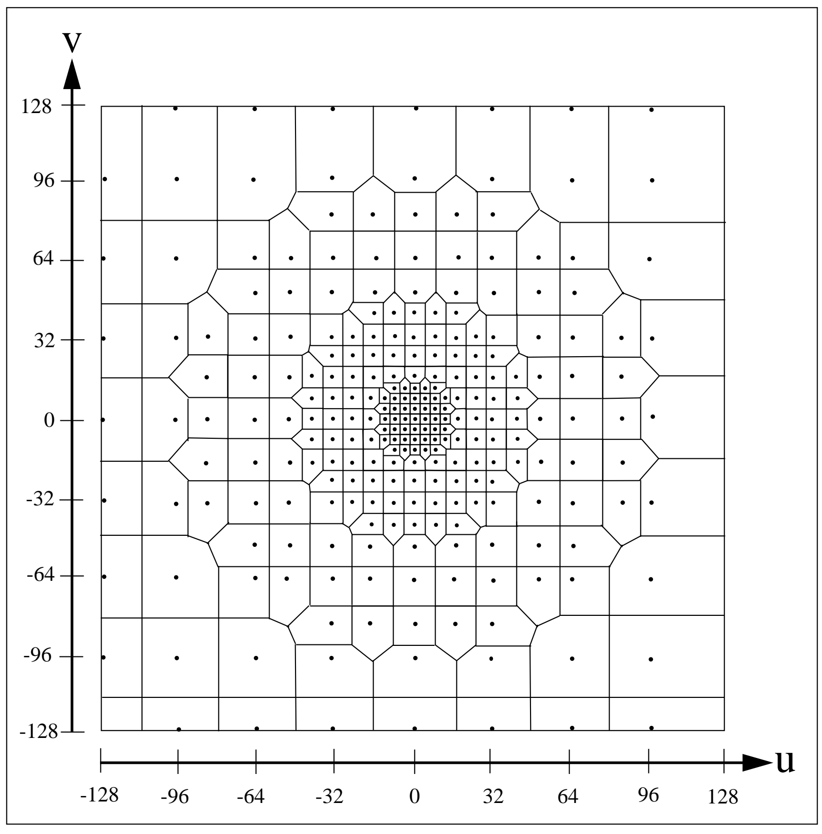

Figure D-3 Default CellB Chrominance Quantization Table

Y/Y Quantization Table

The Y/Y field of the CellB bytecode represents two luminance values of a cell (Y_0 and Y_1). This Y/Y code is an index into a table of two-component luminance vectors. Figure D-4 on page 478 shows the default luminance table. The section "Y/Y Table Values" on page 480 provides the list of values for the table.

The distribution of values in the default luminance table is statistically optimized. The quantizer takes advantage of the high correlation of luminance values within local regions of a cell. This results in a set of representative vectors that are most densely populated around the diagonal, where y1 equals y2.

An observer's sensitivity to contrast is also taken into account, resulting in a distribution of points that is farther apart in regions where the contrast between two values is low.

In this example, the value for the Y/Y code is selected by approximating the mean luminance for the cell. Then, the pixels within the cell are separated into groups that are above and below the mean luminance. The mean luminance of these two groups is used to index a two-dimensional vector quantizer, which returns a byte value for the Y/Y code.

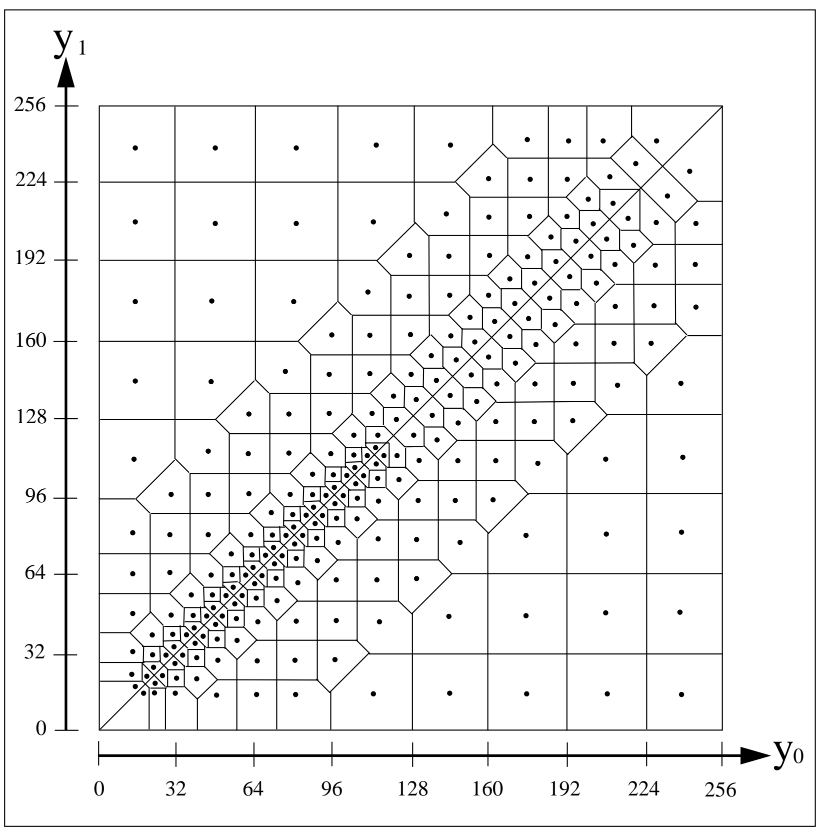

Figure D-4 Default CellB Luminance Quantization Table

Skip Code

The single-byte CellB skip code is shown below.

```
100SSSSS
```

The skip code tells the decoder to skip the next S+1 cells in the frame that is being decoded, which, for CellB, supports a simple form of interframe encoding. There are five skip bits, so that a maximum of 32 cells can be skipped with a single-byte skip code.

New Y/Y Table

The single-byte “new Y/Y table” code is shown below.

```
11111110
```

This code tells the decoder that the next 512 bytes are a new Y/Y quantization table. The bytes are arranged as:

```
Y1_000    Y2_000
Y1_001    Y2_001
.         .
.         .
.         .
Y1_255    Y2_255
```

Note – This code is not implemented in the current CellB compressor and decompressor.

New U/V Table

The single-byte “new U/V table” code is shown below.

```
11111111
```

This code tells the decoder that the next 512 bytes are a new U/V quantization table. The bytes are arranged as:

U_000 V_000
U_001 V_001
. .
. .
. .
U_255 V_255

Note – This code is not implemented in the current CellB compressor and decompressor.

Default CellB Quantization Tables

Y/Y Table Values

Table D-2 lists the default values for the CellB Y/Y quantization table.

Table D-2 Default Y/Y Table *(1 of 4)*

Index	Y1	Y2	Index	Y1	Y2	Index	Y1	Y2
0	16	20	87	120	136	174	112	64
1	16	24	88	128	136	175	128	64
2	16	32	89	128	144	176	72	68
3	16	48	90	128	160	177	76	72
4	16	64	91	128	176	178	80	72
5	16	80	92	128	192	179	88	72

Table D-2 Default Y/Y Table *(2 of 4)*

Index	Y1	Y2	Index	Y1	Y2	Index	Y1	Y2
6	16	112	93	136	144	180	80	76
7	16	144	94	136	152	181	84	80
8	16	176	95	144	152	182	88	80
9	16	208	96	144	160	183	96	80
10	16	240	97	144	176	184	112	80
11	20	24	98	144	192	185	128	80
12	24	28	99	144	208	186	144	80
13	24	32	100	144	240	187	176	80
14	24	40	101	152	160	188	208	80
15	28	32	102	152	168	189	240	80
16	32	36	103	160	168	190	88	84
17	32	40	104	160	176	191	92	88
18	32	48	105	160	192	192	96	88
19	32	64	106	160	208	193	104	88
20	32	80	107	160	224	194	96	92
21	32	96	108	168	176	195	100	96
22	36	40	109	168	184	196	104	96
23	40	44	110	176	184	197	112	96
24	40	48	111	176	192	198	128	96
25	40	56	112	176	208	199	144	96
26	44	48	113	176	224	200	160	96
27	48	52	114	176	240	201	104	100
28	48	56	115	184	192	202	108	104
29	48	64	116	184	200	203	112	104
30	48	80	117	192	200	204	120	104
31	48	96	118	192	208	205	112	108
32	48	112	119	192	224	206	116	112

Table D-2 Default Y/Y Table *(3 of 4)*

Index	Y1	Y2	Index	Y1	Y2	Index	Y1	Y2
33	48	144	120	192	240	207	120	112
34	48	176	121	200	208	208	128	112
35	48	208	122	200	216	209	144	112
36	48	240	123	208	216	210	160	112
37	52	56	124	208	224	211	176	112
38	56	60	125	208	240	212	208	112
39	56	64	126	216	232	213	240	112
40	56	72	127	224	240	214	128	120
41	60	64	128	20	16	215	136	120
42	64	68	129	24	16	216	136	128
43	64	72	130	32	16	217	144	128
44	64	80	131	48	16	218	160	128
45	64	96	132	64	16	219	176	128
46	64	112	133	80	16	220	192	128
47	64	128	134	112	16	221	144	136
48	68	72	135	144	16	222	152	136
49	72	76	136	176	16	223	152	144
50	72	80	137	208	16	224	160	144
51	72	88	138	240	16	225	176	144
52	76	80	139	24	20	226	192	144
53	80	84	140	28	24	227	208	144
54	80	88	141	32	24	228	240	144
55	80	96	142	40	24	229	160	152
56	80	112	143	32	28	230	168	152
57	80	128	144	36	32	231	168	160
58	80	144	145	40	32	232	176	160
59	80	176	146	48	32	233	192	160

Table D-2 Default Y/Y Table *(4 of 4)*

Index	Y1	Y2	Index	Y1	Y2	Index	Y1	Y2
60	80	208	147	64	32	234	208	160
61	80	240	148	80	32	235	224	160
62	84	88	149	96	32	236	176	168
63	88	92	150	40	36	237	184	168
64	88	96	151	44	40	238	184	176
65	88	104	152	48	40	239	192	176
66	92	96	153	56	40	240	208	176
67	96	100	154	48	44	241	224	176
68	96	104	155	52	48	242	240	176
69	96	112	156	56	48	243	192	184
70	96	128	157	64	48	244	200	184
71	96	144	158	80	48	245	200	192
72	96	160	159	96	48	246	208	192
73	100	104	160	112	48	247	224	192
74	104	108	161	144	48	248	240	192
75	104	112	162	176	48	249	208	200
76	104	120	163	208	48	250	216	200
77	108	112	164	240	48	251	216	208
78	112	116	165	56	52	252	224	208
79	112	120	166	60	56	253	240	208
80	112	128	167	64	56	254	232	216
81	112	144	168	72	56	255	240	224
82	112	160	169	64	60			
83	112	176	170	68	64			
84	112	208	171	72	64			
85	112	240	172	80	64			
86	120	128	173	96	44			

C_b/C_r Table Values

Table D-3 lists the default values for the CellB C_b/C_r quantization table.

Table D-3 Default C_b/C_r Table *(1 of 4)*

Index	U	V	Index	U	V	Index	U	V
0	16	16	87	128	184	174	160	120
1	16	48	88	128	192	175	160	128
2	16	80	89	128	208	176	160	136
3	16	112	90	128	224	177	160	144
4	16	144	91	132	136	178	160	152
5	16	176	92	132	140	179	160	160
6	16	208	93	132	144	180	160	168
7	16	240	94	132	148	181	160	176
8	48	16	95	132	152	182	160	184
9	48	48	96	136	104	183	160	192
10	48	80	97	136	112	184	160	208
11	48	112	98	136	120	185	160	224
12	48	144	99	136	128	186	168	112
13	48	176	100	136	132	187	168	120
14	48	208	101	136	136	188	168	128
15	48	240	102	136	140	189	168	136
16	64	112	103	136	144	190	168	144
17	64	128	104	136	148	191	168	152
18	64	144	105	136	152	192	168	160
19	64	160	106	136	156	193	168	168
20	64	176	107	136	160	194	168	176
21	80	16	108	136	168	195	176	16
22	80	48	109	136	176	196	176	48
23	80	80	110	136	184	197	176	64
24	80	96	111	140	132	198	176	80

Table D-3 Default C_b/C_r Table *(2 of 4)*

Index	U	V	Index	U	V	Index	U	V
25	80	112	112	140	136	199	176	96
26	80	128	113	140	140	200	176	112
27	80	144	114	140	144	201	176	120
28	80	160	115	140	148	202	176	128
29	80	176	116	140	152	203	176	136
30	80	192	117	140	156	204	176	144
31	80	208	118	144	16	205	176	152
32	80	240	119	144	48	206	176	160
33	96	80	120	144	64	207	176	168
34	96	96	121	144	80	208	176	176
35	96	112	122	144	96	209	176	192
36	96	128	123	144	104	210	176	208
37	96	144	124	144	112	211	176	224
38	96	160	125	144	120	212	176	240
39	96	176	126	144	128	213	184	128
40	96	192	127	144	132	214	184	136
41	96	208	128	144	136	215	184	144
42	104	128	129	144	140	216	184	152
43	104	136	130	144	144	217	184	160
44	104	144	131	144	148	218	192	80
45	104	152	132	144	152	219	192	96
46	104	160	133	144	156	220	192	112
47	112	16	134	144	160	221	192	128
48	112	48	135	144	168	222	192	144
49	112	64	136	144	176	223	192	160
50	112	80	137	144	184	224	192	176
51	112	96	138	144	192	225	192	192

Table D-3 Default C_b/C_r Table *(3 of 4)*

Index	U	V	Index	U	V	Index	U	V
52	112	112	139	144	208	226	192	208
53	112	120	140	144	224	227	208	16
54	112	128	141	144	240	228	208	48
55	112	136	142	148	132	229	208	80
56	112	144	143	148	136	230	208	96
57	112	152	144	148	140	231	208	112
58	112	160	145	148	144	232	208	128
59	112	168	146	148	148	233	208	144
60	112	176	147	148	152	234	208	160
61	112	192	148	148	156	235	208	176
62	112	208	149	152	104	236	208	192
63	112	224	150	152	112	237	208	208
64	112	240	151	152	120	238	208	240
65	120	112	152	152	128	239	224	112
66	120	120	153	152	132	240	224	128
67	120	128	154	152	136	241	224	144
68	120	136	155	152	140	242	224	160
69	120	144	156	152	144	243	224	176
70	120	152	157	152	148	244	240	16
71	120	160	158	152	152	245	240	48
72	120	168	159	152	156	246	240	80
73	120	176	160	152	160	247	240	112
74	128	64	161	152	168	248	240	144
75	128	80	162	152	176	249	240	176
76	128	96	163	152	184	250	240	208
77	128	104	164	156	136	251	240	240
78	128	112	165	156	140	252	0	0

Table D-3 Default C_b/C_r Table *(4 of 4)*

Index	U	V	Index	U	V	Index	U	V
79	128	120	166	156	144	253	0	0
80	128	128	167	156	148	254	0	0
81	128	136	168	156	152	255	0	0
82	128	144	169	160	64			
83	128	152	170	160	80			
84	128	160	171	160	96			
85	128	168	172	160	104			
86	128	176	173	160	112			

Bibliography

Andrews, H. C. *Computer Techniques in Image Processing.* New York: Academic Press, 1970.

Baxes, G. A. *Digital Image Processing: A Practical Primer.* Englewood Cliffs, N.J.: Prentice-Hall, 1984.

A nonmathematical introduction to image processing.

Castleman, K. R. *Digital Image Processing.* Englewood Cliffs, N.J.: Prentice-Hall, 1979.

Foley, J. D., et al. *Computer Graphics: Principles and Practice.* 2nd ed. Reading, Mass.: Addison-Wesley, 1990.

Discusses dithering, fills, geometric transforms, and color models.

Gonzalez, R. C., and P. Wintz. *Digital Image Processing.* 2nd ed. Reading, Mass.: Addison-Wesley, 1987.

Green, W.B. *Digital Image Processing: A Systems Approach.* New York: Van Nostrand Reinhold, 1983.

Groff, V. *The Power of Color in Design for Desktop Publishing.* Portland: MIS Press, 1990.

Hall, E. L. *Computer Image Processing and Recognition*. New York: Academic Press, 1979.

International Organization for Standardization and International Electrotechnical Commission. *Information Technology—Coding of Moving Pictures and Associated Audio for Digital Storage Media up to About 1.5 Mbits/s*. 1992.

The MPEG-1 standard.

International Organization for Standardization and International Electrotechnical Commission. *Information Technology—Digital Compression and Coding of Continuous-Tone Still Images*. 1991.

The JPEG standard.

International Telegraph and Telephone Consultative Committee (CCITT). "Recommendation H.261 - Video Codec for Audiovisual Services at p x 64 Kbits/s." *Study Group XV - Report R 37*. 1990.

The H.261 standard.

Jain, A. *Fundamentals of Digital Image Processing*. Englewood Cliffs, N.J.: Prentice-Hall, 1989.

Magnenat-Thalmann, N., and D. Thalmann. *Image Synthesis: Theory and Practice*. Tokyo: Springer-Verlag, 1987.

Nye, A. *Xlib Programming Manual*. O'Reilly & Associates, Inc., 1988.

OpenWindows Version 3.0.1 Programmer's Guide. Mountain View, Ca.: SunSoft, 1992.

Discusses colormaps in Sun's implementation of X11.

Pennebaker, W. B., and J. L. Mitchell. *JPEG Still Image Data Compression*. Van Nostrand Reinhold, 1991.

A detailed description of the JPEG still-image compression standard.

PostScript Language Reference Manual. 2nd ed. Reading, Mass.: Addison-Wesley, 1990.

Discusses color models and filtering operations.

Pratt, W. K. *Digital Image Processing*. 2nd ed. New York: Wiley, 1991.

Raster Graphics Handbook. 2nd ed. New York: Van Nostrand Reinhold, 1985.

Contains a good appendix on color models.

Rogers, D. F. *Procedural Elements for Computer Graphics*. New York: McGraw-Hill, 1985.

Provides information about dithering and fills.

Rosenfield, A., and A. C. Kak. *Digital Picture Processing*. 2nd ed. 2 vols. New York: Academic Press, 1982.

Scheifler, R. W., and J. Gettys. *X Window System*. Bedford, Mass.: Digital Press, 1988.

Covers color management in the X Window System.

Thorell, L. G., and W. J. Smith. *Using Computer Color Effectively: An Illustrated Reference*. Englewood Cliffs, N.J.: Prentice Hall, 1990.

Ulichney, R. *Digital Halftoning*. Cambridge, Mass.: MIT Press, 1987.

Contains good discussions of halftoning and dithering.

Wallace, G. K. “The JPEG Still Picture Compression Standard.” *Communications of the ACM*. April 1991, p. 31.

A good introduction to JPEG compression.

Wolberg, G. *Digital Image Warping*. Los Alamitos, Ca.: IEEE Computer Society Press, 1990.

Glossary

AC coefficient

A **DCT** coefficient that corresponds to nonconstant data in the original image.

additive color system

A color model in which colors are built by adding together primary colors. **RGB** is an example.

affine transform

Affine transforms include such operations as scaling, rotation, translation, and shearing. What these operations have in common is that they can change the lengths and angles of lines, but not their parallelism. The XIL library contains a function `xil_affine()` that in one operation can scale, rotate, and translate an image.

atom

A single XIL function. Because of its deferred-execution scheme, the library is sometimes able to find groups of atoms that it can replace with an optimized routine that does the jobs of all the atoms in the group. These optimized routines are called **molecules**.

attribute

A term from object-oriented programming that refers to a characteristic of an object.

Baseline sequential codec

A sequential coder-decoder defined by the **JPEG** standard. It is designed to handle images with 8-bit samples and uses Huffman coding for its entropy coding.

bicubic interpolation

Bicubic **interpolation** is the most time consuming of the XIL library's interpolation methods, but produces the best results. When this type of interpolation is requested, the library calculates the value of a pixel in the destination image by determining the point in the source image to which that pixel maps and then examining the values of the sixteen pixels closest to that point.

bilinear interpolation

When the XIL library uses bilinear (or first-order) **interpolation** to determine the value of an image at noninteger coordinates, it calculates the value by looking at the values of the four pixels surrounding the point of interest and then using a bilinear equation. This type of interpolation yields better results than **nearest neighbor interpolation**, but can itself have an undesirable smoothing effect on an image. To alleviate this problem, you can use **bicubic interpolation**.

Block Truncation Coding

The image compression method on which the **Cell encoding** scheme is based. In Cell compression, a 4-by-4 region of pixels from an image is represented by two colors and a 16-bit mask.

CCIR

The Committée International de la Radio. This treaty organization, part of the International Telecommunications Union (ITU), is responsible for obtaining international agreement on standards for radio and television transmission and the international exchange of programs.

CCIR Recommendation 601

An international standard for digitizing PAL, NTSC, and SECAM analog video.

CCITT

International Telegraph and Telephone Consultative Committee, an international association and standards body composed primarily of representatives from national telephone agencies. The CCITT promulgates telephony standards, such as X.25, the **Group 3** facsimile standard, and the **H.261** videoconferencing standard (p*x*64).

CCITT Group 3

A standard that specifies how facsimile machines must compress and decompress image data. The compression method relies heavily on **run-length encoding**. Runs of white pixels and runs of black pixels are represented with codes from a **Huffman** table.

CCITT Group 4

A standard for compressing document images. This compression technique takes advantage of a characteristic of document images called vertical coherence. This means that transitions from black to white or vice versa generally occur in almost the same place on adjacent scanlines. The Group 4 encoder compresses an image by recording information about the relative locations of these transitions.

cell

A 4-by-4 block of pixels. (*See also* **Cell encoding**.)

Cell encoding

A video compression algorithm developed by Sun. In Cell encoding, a 4-by-4 region of pixels is represented by two colors and a 16-bit mask that indicates which of the two colors to place at each of the 16 pixel positions. The colors and mask are chosen to preserve the mean and variance of the luminance and the average chrominance for the 4-by-4 block.

Cell decoding takes advantage of the fonting hardware commonly found in bitmapped displays.

CellB encoding

A video-compression algorithm derived from the **Cell** algorithm. As with a Cell compressor, the most fundamental task of a CellB compressor is to encode 4-by-4 cells of pixels in four bytes. The first two bytes of each cell code are a 16-bit mask that indicates which of two colors will represent each pixel in the cell. The third and fourth bytes contain indexes into tables of luminance and chrominance values and define the two colors to be used in encoding the cell.

CellB coders and decoders are intended primarily for use in videoconferencing applications.

chrominance

The portion of a composite signal that carries color information. For example, the C_b or C_r component of a YC_bC_r signal represents part of a pixel's chrominance. (*See also* **luminance**.)

CIF

Commmon Interchange Format. CIF format images contain YC_bC_r data and are 352 pixels wide and 288 pixels high. Images in this format are one of the two types of images that may be supported by **H.261** codecs.

CIS

Compressed image sequence. The XIL library's compressors store (generally related) compressed images in structures called CIS buffers. The images may represent frames in a movie, pages in a document, and so on.

CMY color model

In the CMY color model, the subtractive primaries cyan, magenta, and yellow are used to filter their complements (red, green, and blue) from white light. You use this color model when working with devices like color printers that put colored ink on paper. For example, to create the color red on a color printer, you put down a mixture of magenta and yellow ink. The magenta filters out its complement green; the yellow filters out its complement blue; and you see only red.

CMYK color model

Similar to the CMY color model, but uses a fourth color: black. Black is used to replace equal amounts of cyan, magenta, and yellow. This color model is used in offset color printing.

codec

A coder-decoder.

colormap

A color lookup table that stores a set of **RGB** values. Applications index into the colormap to get the values to drive the red, blue, and green guns of an RGB monitor.

compressed image sequence

See **CIS**.

compression

The process of converting data from its original format to a format that requires fewer bits. Compressed data uses less storage than uncompressed data and can be transmitted over a network more quickly. Some compressors, such as the **CCITT Group 3** compressor, are called **lossless** compressors because they compress data in a way such that a decompressor can regenerate the original data exactly. Other compressors are called **lossy** compressors because they compress data in a way that prevents the original data set from being regenerated exactly.

convolution

An image-processing operation frequently used to sharpen an image, blur an image, or highlight the edges in an image. The operation calculates the values of pixels in the destination image using the values of a neighborhood of pixels in the source image and the values in a special filter called a **convolution kernel**.

convolution kernel

A two-dimensional array of weighted values used in a **convolution** operation. In the XIL library, a kernel is a data structure of data type `XilKernel`.

DC coefficient

A **DCT** coefficient that corresponds to the average level of the input image.

DCT

Discrete Cosine Transform. Many encoders, including those that conform to the **JPEG**, **MPEG-1**, and **H.261** standards, perform a DCT on an 8-by-8 block of image data as part of the image-compression process. The DCT converts the video data from the spatial domain to the frequency domain. The DCT takes an 8-by-8 matrix, whose values represent brightness levels at particular *x,y* coordinates, and produces an 8-by-8 matrix whose values represent relative amounts of the 64 spatial frequencies that make up the input data's spectrum.

The DCT provides a basis for compression because most of the frequency levels for a block will be zero or close to zero and do not need to be encoded.

decoder

A program that takes data that has been encoded, or compressed, by an **encoder** and decompresses it. A decoder can be implemented in hardware, software, or a combination of both. The decompressed data may or may not match the original data set exactly, depending on how the data was encoded. (*See also* **lossless compression** and **lossy compression**.)

decompression

The restoring of data that has undergone compression to its original state, or to something close to its original state. How closely the decompressed data matches the original data depends on the compression algorithm used. (*See also* **lossless compression** and **lossy compression**.)

display image

The XIL library treats displays as special images, and operations allow display images to serve as destinations. This strategy enables an operation to draw on a screen directly, without requiring an intermediate copy. In some cases, display images can also serve as source images.

dithering

The XIL library includes a table-lookup function that enables you to convert a single-band image of any data type to a single-band or multiband image of any data type. Dithering can be thought of as an *inverse* lookup operation; that is, dithering enables you to convert a single-band or multiband image of any data type to a single-band image of any data type. The most common dithering operations convert 3-band, 8-bit images to 1-band, 8-bit images and 1-band, 8-bit images to 1-band, 1-bit images.

encoder

An encoder is a program that encodes data for the purpose of achieving data compression. The encoders included with the XIL library are designed to compress images.

entropy coding

Entropy coding is the final step in the compression process in **DCT**-based encoders (such as the **JPEG** baseline sequential encoder). In this step, the encoder compresses quantized **DCT** coefficients using **Huffman** coding. Therefore, values that occur frequently are encoded with fewer bits than are values that occur infrequently.

error diffusion

A technique for removing some of the artifacts produced during the dithering process. In its most common form, Floyd-Steinberg error diffusion, this technique involves (1) determining the amount of error produced in dithering a particular pixel and (2) distributing fractions of that error to the pixels to the right of and below the pixel just dithered.

frame

A single image taken from a movie.

full-motion video

The showing of a series of related digital images at a rate sufficient to give the illusion that objects in the images are moving naturally.

gamma correction

In a linear color space, color levels are equally spaced throughout a gamut. The problem with this type of color space is that level-to-level changes at the low end of the gamut seem greater to the eye than equal changes at the high end. In a gamma-corrected color space, color levels are spaced logarithmically so that level-to-level changes seems consistent throughout the range.

geometric operation

An image-processing operation that changes the size, shape, or orientation of objects in an image. XIL geometric operations include scaling, rotation, and general affine transforms.

H.261

A video compression standard developed by the **CCITT** for use in encoding video to be transferred over an **ISDN** (Integrated Service Digital Network). In this compression method, data is compressed so that the output bit rate is p x 64 Kbits per second, where *p* can range from 1 to 30 depending on the number of ISDN channels used. This standard was developed primarily to support video phone and videoconferencing applications.

handle

An identifier that refers to a unique object. In the XIL library, for example, an identifier of type `XilImage` is a handle to an actual image.

histogram

A collection of information about how frequently certain gray levels or colors appear in an image. The XIL library contains a data structure of type `XilHistogram` to hold this type of information.

Huffman encoding

A method of compressing a given set of data based on the relative frequency of the individual elements: the more often an element occurs, the shorter (in bits) its corresponding code. Huffman encoding is often used to compress text files, with the coding based on letter frequency. Huffman encoding is also used in many **DCT**-based video-compression algorithms.

image

A two-dimensional array of **pixels** that represents an object.

indexed color

See **pseudocolor**.

in-place operations

Operations that require both a source and a destination image but that allow them to be the same image. If a function doesn't support in-place operations, the source and destination images must be different images.

interframe compression

In image sequences, consecutive frames generally have more similarities than differences. These similarities, or shared elements, are sometimes referred to as being temporally redundant. This redundancy is important because it allows groups of individually compressed frames to be compressed further. That is, if five frames in a group look the same in the upper-left corner, that area needs to be encoded only once; then the encoder can simply note that the same data appears in the next four frames. This type of encoding is called interframe compression.

interpolation

A way of calculating a value that falls between other, known values. In image processing, interpolation frequently plays a part in geometric operations such as rotation. After that type of spatial transformation, pixel locations in the output image will correspond to noninteger coordinates in the input image. Therefore, the pixel values in the output must be calculated by looking at the values of the pixels surrounding the point of interest in the input. The XIL library supports several types of interpolation, including **nearest neighbor**, **bilinear**, and **bicubic interpolation**.

intraframe compression

The compression that can take place within a single image. Contrast with **interframe compression**.

ISDN

Integrated Service Digital Network. A worldwide public telecommunications network designed to handle many types of data, including voice, text, graphics, and video. The **CCITT** designed the **H.261** video encoder to produce compact bitstreams that can be sent over ISDN lines.

JPEG

Joint Photographic Experts Group. A joint venture of the **CCITT** and ISO that has developed a standard for compressing grayscale or color *still* images. Actually, the standard defines a number of methods for compressing images. Several of these are **lossy** methods based on the Discrete Cosine Transform (**DCT**), but one method is **lossless** and is based on a **predictive coding** technique.

key frame

If the video portion of a movie has been compressed using both **intraframe compression** and **interframe compression**, a decoder cannot decompress the majority of the video frames without referring to preceding—and sometimes

succeeding—frames. However, the first frame in the movie, and usually other frames as well, do not have interframe dependencies and can be decoded in isolation. These frames are called key frames. Key frames, besides the one that starts the movie, enable the decoder to decompress other frames without playing the movie from the beginning.

linear remapping

See **rescaling**.

local operation

An image-processing operation in which more than one pixel in a source image is used in calculating the value of a single pixel in the destination image. Generally, the source-image pixels used include the pixel corresponding to the destination pixel being calculated and some set of pixels surrounding that source pixel.

lookup table

In XIL applications, lookup tables are used for general image modification. Each entry in an XIL lookup table contains an index—a value that may appear in the source image—and a value or set of values to be written to the destination image. For each pixel in the source, a table-lookup function finds the pixel's value on the index side of the table and then writes the output value or values for that entry to the corresponding pixel in the destination.

lossless compression

The compression of data in such a way that the original data can be restored exactly. **Huffman encoding** is an example of a lossless compression technique. Some compressors, such as the **JPEG** baseline sequential compressor, combine lossy and lossless compression algorithms. (*See also* **lossy compression**.)

lossy compression

A type of compression that results in the loss of some of the original data. Lossy compression trades the potential loss of some image quality for the opportunity for greater compression. The **JPEG** baseline sequential and **Cell** methods are examples of lossy compression techniques. (*See also* **lossless compression**.)

luminance

The portion of a composite signal that carries brightness information. For example, luminance information is contained in the Y component of a YC_bC_r signal. Video compression techniques take advantage of the fact that the human eye is more sensitive to variations in luminance than it is to variations

in color (**chrominance**). Therefore, chrominance values can be compressed (with **lossy** techniques) more than luminance values, resulting in greater overall compression.

macroblock

Both the **H.261** and the **MPEG-1** specifications define a unit within an image called a macroblock. This unit is a 16-by-16 blocks of pixels. Both H.261 and MPEG-1 encoders can switch from intraframe encoding to interframe encoding on a macroblock basis.

molecule

The XIL term for an optimized routine that performs the work of two or more XIL functions (**atoms**). You don't call a molecule directly. Instead, the XIL runtime system executes a molecule whenever your program calls a sequence of XIL functions that the molecule can replace. This type of substitution is possible because of the library's deferred-execution scheme.

motion compensation

Both **H.261** and **MPEG-1** encoders can perform interframe compression by encoding the differences between the values in a **macroblock** in the picture being encoded and the values in the corresponding macroblock in the preceding image. Sometimes, however, it's desirable to encode the differences between the values in a macroblock in the current image and the values in a macroblock in the preceding image that is slightly offset from the one that corresponds spatially with the macroblock being encoded. This approach is desirable because it enables the encoder to track the movement of objects from image to image and, thus, to encode smaller differences. If an encoder uses this approach, it must of course record the extent of the offset mentioned above. This offset is recorded in a motion vector.

movie

A contiguous series of video frames (and optionally synchronized audio) that are displayed fast enough to provide the illusion of motion. A frame rate of 30 frames/second is a typical target for a smooth-running movie.

MPEG

Moving Picture Experts Group. This group has developed standards for compressing moving pictures and audio data and for synchronizing video and audio datastreams. The XIL library includes a decompression module that can decode MPEG-1 video bitstreams.

The MPEG-1 video-compression standard is similar to the **H.261** standard developed by the **CCITT**, but places less emphasis on low bit rates. By accepting a higher bit rate—up to 1.5 Mbits per second—an MPEG-1 codec is able to recreate very high-quality pictures and to produce a bitstream that is easily editable.

The rate of 1.5 Mbits/s makes the MPEG-1 codec especially viable in applications that read compressed data from CD-ROMs. For example, putting an MPEG-1 bitstream on a CD-ROM is an effective way to distribute movies, business presentations, and training videos.

nearest neighbor interpolation

One of the methods that the XIL library uses to determine the value of an input image at noninteger coordinates. When using this method, the XIL library takes the value it is looking for to be the value of the pixel closest to the point of interest. This type of **interpolation** is the fastest type, but can introduce artifacts in the output image; for example, smooth lines in the input image may show up as jagged lines in the output.

Nearest neighbor interpolation is sometimes called zero-order interpolation.

origin

In the XIL library, each image object has an **attribute** called its origin. This origin is a pair of floating-point numbers that represent *x,y* coordinates in the image. (The point 0.0,0.0 is in the upper-left corner.) When an image-processing operation is performed, the origins of the input and output images are aligned, and the rectangle formed by the intersection of these images serves as an implicit region of interest (**ROI**).

pixel

Picture element. In a raster grid, a pixel is the smallest unit that can be addressed and given a color or intensity.

point operation

An image-processing operation in which the value of a point (or **pixel**) in the destination image depends only on the corresponding point in the source image or images. For example, if you want to add two source images to produce a destination image, the value of the pixel in the upper-left corner of the destination image depends solely on the values of the pixels in the upper-left corner of the two source images.

predictive encoding

In predictive encoding, an encoder uses the values of neighboring samples to predict the value of the sample being encoded. The encoder then subtracts this predicted value from the actual value of the sample and encodes the difference. This mode of compression is **lossless**. The XIL library supports a **JPEG** lossless compressor that uses this method.

pseudocolor

The distinction between true color and pseudocolor has to do with the design of a monitor's frame buffer. If the frame buffer uses 8 bits per pixel to store color information, the monitor can display 256 colors simultaneously. What you see on such a monitor is called pseudocolor because the colors that can be shown at any one time are a small subset of the colors the eye can distinguish. If the frame buffer uses 24 bits per pixel to store color information, your monitor can display over 16 million colors (true color).

Pseudocolor is sometimes called indexed color because the values stored in the frame buffer on a pseudocolor system are not the **RGB** values needed to drive the red, green, and blue electron guns in a monitor. Rather, they are indexes into a **colormap**, or color lookup table, which stores 256 sets of RGB values.

QCIF

Quarter Common Interchange Format. QCIF images contain YC_bC_r data and are 176 pixels wide and 144 pixels high (one-fourth the size of **CIF** images). **H.261** codecs are required to support images in this format.

quantization

The technique of scaling down a set of values. In Discrete Cosine Transform (**DCT**)-based encoders, like those that conform to the **JPEG**, **MPEG-1**, and **H.261** standards, quantization is used to ensure that DCT coefficients are represented by the smallest range of numbers needed to produce the desired level of image quality. To bring the coefficients into this range, the quantizer divides each coefficient by the appropriate value from a quantization table and rounds the result to the nearest integer.

rescaling

Sometimes called linear remapping, rescaling involves changing the values in an image using a linear equation. Each value in the image is multiplied by a constant; then a second constant is added to the product. Rescaling is useful for mapping the values in an image from one range to another.

RGB color model

A color model in which colors are built by mixing the three additive primary colors red, green, and blue. In this model, you construct grays by including equal amounts of each primary: (0,0,0) is black and (1,1,1) is white. The RGB color model is closely associated with color CRT monitors because they use this model to produce their colors.

ROI

A region of interest. In the XIL scheme, an ROI is a 1-bit mask that can be specified as an **attribute** of an image. This mask determines what part of a source image can be read or what part of a destination image can modified during an image-processing operation. The ROI for a given operation is the intersection of the ROIs of all the images involved in the operation.

run-length encoding

A compression technique that stores counts of the number of consecutive identical pixels or blocks of pixels in an image.

sequential encoding

The encoding of a still image in a single pass through the image data. The sequential encoder first processes a block of pixels in the upper-left corner of the image and then proceeds from left to right and top to bottom until it has processed the entire image.

skip codes

Skip codes are employed during the **interframe compression** process. They instruct the decoder to skip over blocks of pixels in the current frame that are identical to, or very similar to, those in the preceding frame.

slice

MPEG-1 terminology for a group of consecutive **macroblocks**. Each macroblock in a picture (or image) must be part of a slice. The slice was designed primarily to help a decoder recover from bitstream errors. If a decoder detects an error, one way to recover is to skip to the next slice header.

spatial redundancy

The occurrence of two or more consecutive pixels or blocks of pixels that have the same, or similar, contents. A high level of spatial redundancy in an image can lead to a high rate of compression because only the first pixel or block of pixels in a sequence needs to be encoded fully. A shorthand method can be used to encode succeeding pixels or blocks.

still-image coding

The encoding of a single image.

structuring element

A two-dimensional array of Boolean values used as a parameter to the XIL dilation and erosion functions.

subsampling

A way of mathematically reducing a data set to a subset of its original components. For example, if you have captured a 512-by-512 image, but want to encode it at a resolution of 256-by-256, you can subsample the data before encoding it by discarding every other pixel in both and *x* and *y* directions. On systems that work with YC_bC_r data, it is common to subsample the C_b and C_r **chrominance** information so that each pixel has a unique **luminance** value, but shares chrominance information with one or more nearby pixels.

subtractive color system

In a subtractive color system, primary colors are used as filters to subtract their complements from white light. For example, in the **CMY** color model, the primaries cyan, magenta, and yellow subtract red, green, and blue respectively. That is, if light is reflected from a piece of paper coated with cyan ink, no red is reflected, so you see the color cyan.

temporal redundancy

See **interframe compression**.

thresholding

The XIL library's thresholding function sets all the values (in a band) that fall between a low threshold and a high threshold to a value (called a *map* value) that the programmer specifies.

transcoding

The conversion of data in one compressed format to another compressed format. For instance, converting a **JPEG**-compressed image to its **Cell**-encoded counterpart is called transcoding. Transcoding an image usually involves decoding it and then reencoding it.

translation

A geometric operation in which an image is moved up or down and/or left or right.

transposition

A quasi-geometric transformation that involves rotating an image by a multiple of 90 degrees or flipping pixels in an image across a line that passes through the center of the image. The XIL library's transposition function allows for flips across a horizontal or vertical line passing through an image's center and across the main diagonal or antidiagonal.

XIL library

The XIL library is a foundation library for imaging and video support. It provides an implementation of imaging functionality that is common to multiple higher-level interfaces, provides imaging capabilities that are not currently available, and provides a way for ISVs to access low-level and hardware functionality.

YC_bC_r color model

When data is stored in YC_bC_r format, each pixel is described by a **luminance** value (Y) and two color values. This color model is used in the PAL (European) television format. Also, video analog-to-digital converters often produce data in this format.

Frequently, the color bands of YC_bC_r images are subsampled to produce what are called YC_bC_r 4:2:2 or YC_bC_r 4:1:1 images. In these images, pairs of pixels, or blocks of 4 pixels, share color values; that is, each pixel has its own luminance information, but shares color values with a neighboring pixel or pixels. This subsampling in the color dimension has a very minor impact on image quality.

Index

A

absolute value of images, 143
acceleration in XIL programs, 395
adding images, 141
affine transforms, general, 169
alpha masks, 211
AND of images, 144
arithmetic operations
  involving an image and a constant, 144
  involving bit images, 147
  involving two images, 141
attributes, *see* device attributes

B

backing_store attribute, 66
BAND_HUFFMAN_TABLE attribute, 311, 331
BAND_QUANTIZER attribute, 312
BANDS attribute, 393
bibliography, 489
bicubic interpolation, 151
bilinear interpolation, 151
bins, histogram, 179
BITS_PER_SECOND attribute, 276
bitwise logical operations
  involving an image and a constant, 144
  involving bit images, 147
  involving two images, 144
black generation, 125
blending images, 211
blurring images, 190
boundary fills, 183
BYTES_PER_FRAME attribute, 313

C

casting images from one data type to another, 103
CCITT Group 3 and Group 4 codecs
  applications, 391
  decompression attributes of
    BANDS, 393
    HEIGHT, 393
    WIDTH, 393
  how they work, 391
  transposition molecule, 407
Cell codec
  applications, 271
  compression attributes of
    BITS_PER_SECOND, 276
    COLORMAP_ADAPTION, 277
    COMPRESSOR_COLORMAP, 278

COMPRESSOR_FRAME_
RATE, 278
COMPRESSOR_MAX_CMAP_
SIZE, 279
COMPRESSOR_USER_DATA, 280
ENCODING_TYPE, 281
KEYFRAME_INTERVAL, 282
TEMPORAL_FILTERING, 283
creating a CIS associated with a, 275
decompression attributes of
DECOMPRESSOR_
COLORMAP, 244, 284
DECOMPRESSOR_FRAME_
RATE, 285
DECOMPRESSOR_MAX_CMAP_
SIZE, 243, 286
DECOMPRESSOR_USER_
DATA, 286
RDWR_INDICES, 244, 287
how it works, 272
playback molecules, 288
CellB codec
applications, 297
compression rate, 300
creating a CIS associated with a, 301
decompression attributes
HEIGHT, 301
IGNORE_HISTORY, 302
WIDTH, 301
how it works, 298
playback molecules, 303, 399
restriction on size of images to be
compressed, 298
child images, creating, 53
CIF images, *see* Common Interchange
Format images
clearing images, 216
closing the library, 28
color spaces, 52
conversion, 52, 122
identifiers, 123
PhotoYCC, 76
colorcubes
creating, 110
XIL-supplied, 113, 238
COLORMAP_ADAPTION attribute, 277
Common Interchange Format (CIF)
images, 338, 356
compiling code, conditional, 31
compressed image sequences
attributes of
compression type, 258
compressor, 258
error recovery flag, 265
input image type, 259
keep frames, 262
maximum frames, 262
name, 266
output image type, 260
random access flag, 260
read frame, 260
start frame, 260
write frame, 260
checking for unread data, 228, 240,
256
creating, 223, 233, 250
decompressing images from, 240,
246, 255
defined, 249
destroying, 251
determining the dimensions of
images stored in, 235
error recovery, 267
flushing compressed data from, 230,
252
getting a pointer to the compressed
data in, 228, 255
putting compressed data into, 233,
251
resetting, 250
seeking an image in, 257, 260
COMPRESSED_DATA_FORMAT
attribute, 313, 332
compressing fax images, 391
compressing images, 227, 251
COMPRESSION_QUALITY attribute, 314
COMPRESSOR_BITS_PER_IMAGE
attribute, 347

COMPRESSOR_BITS_PER_SECOND attribute, 366
COMPRESSOR_COLORMAP attribute, 278
COMPRESSOR_DOC_CAMERA attribute, 348
COMPRESSOR_ENCODE_INTRA attribute, 349
COMPRESSOR_FRAME_RATE attribute, 278
COMPRESSOR_FREEZE_RELEASE attribute, 349
COMPRESSOR_IMAGE_SKIP attribute, 350
COMPRESSOR_INSERT_VIDEO_SEQUENCE_END attribute, 367
COMPRESSOR_INTRA_QUANTIZATION_TABLE attribute, 373
COMPRESSOR_LOOP_FILTER attribute, 351
COMPRESSOR_MAX_CMAP_SIZE attribute, 279
COMPRESSOR_MV_SEARCH_RANGE attribute, 352
COMPRESSOR_NON_INTRA_QUANTIZATION_TABLE attribute, 375
COMPRESSOR_PATTERN attribute, 376
COMPRESSOR_PEL_ASPECT_RATIO attribute, 378
COMPRESSOR_PICTURE_RATE attribute, 380
COMPRESSOR_SLICES_PER_PICTURE attribute, 381
COMPRESSOR_SPLIT_SCREEN attribute, 353
COMPRESSOR_TIME_CODE attribute, 382
COMPRESSOR_USER_DATA attribute, 280
compressors/decompressors
- reading attributes of, 243, 276
- setting attributes of, 244, 276
- types of images supported, 252
- *See also* individual compressors

conditionally compiling code, 31
ConfigureNotify event, 65
converting
- a 16-bit image to an 8-bit image, 91
- a 24-bit image to a 1-bit image, 93
- a grayscale image to a 1-bit image, 94
- a single-band image to a multiband image, 88
- a true-color image to a pseudocolor image, 92

convolution, 187
convolution filters, *See* kernels
copying
- images to displays, 97
- patterns to images, 217
- plane mask control, 99

creating
- child images, 53
- colorcubes, 110
- compressed image sequences, 223, 233, 250
- device images, 71, 78
- device objects, 68
- display images, 65
- dither masks, 118
- histograms, 179
- images, 16, 58
- kernels, 188
- lookup tables, 89, 108
- Photo CD image, 78
- plane mask, 100
- regions of interest, 49

D

data types of images, 35
decompressing fax images, 391
decompressing images, 240, 246, 255
DECOMPRESSION_QUALITY attribute, 324
DECOMPRESSOR_BROKEN_LINK attribute, 385
DECOMPRESSOR_CLOSED_GOP attribute, 385

DECOMPRESSOR_COLORMAP attribute, 244, 284
DECOMPRESSOR_DOC_CAMERA attribute, 355
DECOMPRESSOR_FRAME_RATE attribute, 285
DECOMPRESSOR_FRAME_TYPE attribute, 386
DECOMPRESSOR_FREEZE_RELEASE attribute, 355
DECOMPRESSOR_MAX_CMAP_SIZE attribute, 243, 286
DECOMPRESSOR_PEL_ASPECT_RATIO_VALUE attribute, 387
DECOMPRESSOR_PICTURE_RATE_VALUE attribute, 388
DECOMPRESSOR_QUALITY attribute, 384
DECOMPRESSOR_SOURCE_FORMAT attribute, 356
DECOMPRESSOR_SPLIT_SCREEN attribute, 356
DECOMPRESSOR_TEMPORAL_REFERENCE attribute, 357, 388
DECOMPRESSOR_TIME_CODE attribute, 389
DECOMPRESSOR_USER_DATA attribute, 286
deferred execution, 395
detecting edges, 187, 192
device attributes
  initializing, 68
  Photo CD, 79
  reading, 72
  setting, 72
device images, 43, 71
  Photo CD, 77
device object
  creating, 68
  destroying, 73
devices
  initializing attributes of, 68
  partial list of, 84
  reading images from, 67
  writing images to, 67
dilating images, 194
Discrete Cosine Transforms, 307
display images, 22, 24, 44, 65, 236
displays
  monochrome, 94
  reading images from, 66
  writing images to, 64
dither masks
  creating, 118
  used by xil_ordered_dither(), 120
  XIL-supplied, 119, 238
dithering operations
  defined, 105
  error diffusion, 114
  nearest color, 108, 246
  ordered dither, 118
  when to use the different dithering functions, 121
dividing images, 141

E

edge detection, 187, 192
ENCODE_411_INTERLEAVED attribute, 316
ENCODE_INTERLEAVED attribute, 316, 332
ENCODING_TYPE attribute, 281
entropy coding, 309, 330
eroding images, 194
error categories, 130
error diffusion, 114
error handling
  installing error handlers, 136
  linking error handlers, 136
  using the default error handler, 128
  writing a custom error handler, 128
error IDs, 425
error messages, 129, 425
events
  ConfigureNotify, 65

Expose, 66
example programs
making a movie, 222
playing a Cell movie, 241
playing a JPEG movie, 230
simple display program, 11
more complete display program, 87
exclusive OR of images, 144
exporting images, 17, 36, 59
Expose event, 66
extrema, finding an image's, 178

F

faxG3, faxG4, *see* CCITT Group 3 and Group 4 Codecs
files
reading images from, 13, 57
writing compressed data to, 227
writing images to, 62
filling regions
boundary fills, 183
soft fills, 185
filtering images, 187
Floyd-Steinberg error-distribution kernel, 115
foundation libraries, 5
frame groups, 276

G

general interpolation, 151
creating kernels, 154
destroying kernels, 159
edge conditions, 154
kernel data, 157
kernel size, 155
key element, 155
key values, 153
setting on system-state object, 157
subsamples, 155
geometric operations, 149
glossary, 493

H

H.261 codec
applications, 337
bit-rate control, 345
compression attributes of
COMPRESSOR_BITS_PER_IMAGE, 347
COMPRESSOR_DOC_CAMERA, 348
COMPRESSOR_ENCODE_INTRA, 349
COMPRESSOR_FREEZE_RELEASE, 349
COMPRESSOR_IMAGE_SKIP, 350
COMPRESSOR_LOOP_FILTER, 351
COMPRESSOR_MV_SEARCH_RANGE, 352
COMPRESSOR_SPLIT_SCREEN, 353
compressor not supplied with XIL library, 338
creating a CIS associated with an H.261 decompressor, 346, 347
decompression attributes of
DECOMPRESSOR_DOC_CAMERA, 355
DECOMPRESSOR_FREEZE_RELEASE, 355
DECOMPRESSOR_SOURCE_FORMAT, 356
DECOMPRESSOR_SPLIT_SCREEN, 356
DECOMPRESSOR_TEMPORAL_REFERENCE, 357
IGNORE_HISTORY, 354
how it works, 338
loop filtering, 343, 351
motion compensation, 342, 352
multipoint conferencing, 346
playback molecules, 358, 399
header file, xil.h, 12
HEIGHT attribute, 301, 393

histograms
creating, 179
destroying, 182
reading data from, 181
Huffman tables, 309, 317, 330, 333, 392
HUFFMAN_TABLE attribute, 317, 333

I

IGNORE_HISTORY attribute, 302, 325, 354
image types, 54, 235
images
attributes of
color space, 52
data type, 35
height, 34
image type, 54
name, 56
number of bands, 34
origin, 46
parent, 53
readable flag, 55
region of interest, 48
width, 34
writable flag, 55
creating, 16, 58
data types supported, 16
determining the dimensions of, 63
exporting, 17, 36, 59
getting information about storage in memory, 17, 59
importing, 22, 36, 62
memory formats of, 37
naming, 56
overlaying, 99
Photo CD, 75
reading from devices other than displays, 67
reading from displays, 66
reading from files, 13, 57
types of
device images, 43
display images, 44
memory images, 43
warping, 172
writing to devices other than displays, 67
writing to displays, 64
writing to files, 62
importing images, 22, 36, 62
initializing the library, 13
in-place operations, 499
input/output, 57
interpolation
bicubic, 151
bilinear, 151
general, 151
kernels, 153
nearest neighbor, 150

J

JPEG baseline sequential codec
applications, 305
compression attributes of
BAND_HUFFMAN_TABLE, 311
BAND_QUANTIZER, 312
COMPRESSED_DATA_FORMAT, 313
COMPRESSION_QUALITY, 314
ENCODE_411_INTERLEAVED, 316
ENCODE_INTERLEAVED, 316
HUFFMAN_TABLE, 317
OPTIMIZE_HUFFMAN_TABLES, 322
QUANTIZATION_TABLE, 322
TEMPORAL_FILTERING, 323
creating a CIS associated with a, 310
decompression attributes of
BYTES_PER_FRAME, 313
DECOMPRESSION_QUALITY, 324
IGNORE_HISTORY, 325
how it works, 306
playback molecules, 326, 399
JPEG lossless codec
attributes of
BAND_HUFFMAN_TABLE, 331

COMPRESSED_DATA_FORMAT, 332
ENCODE_INTERLEAVED, 332
HUFFMAN_TABLE, 333
LOSSLESS_BAND_PT_TRANSFORM, 336
LOSSLESS_BAND_SELECTOR, 334
how it works, 328

K

keep frames, 262
kernels
convolution, 187
creating, 115, 188
destroying, 191
error-distribution, 115
Floyd-Steinberg, 115
interpolation, 153
key values of, 153
keys values of, 187
used for painting, 213
key frames (Cell), 282
KEYFRAME_INTERVAL attribute, 282
Kodak's Photo CD format, *see* Photo CD images

L

libraries to link with, 29, 461
linear remapping of images, 102
linking XIL programs, 29, 461
logical operations
involving an image and a constant, 144
involving bit images, 147
involving two images, 144
lookup tables
creating, 89, 108, 200
destroying, 205
determining the number of entries in, 239
passing images through, 89, 199
reading values from, 239
setting the value of the first index, 239
version numbers, 244
XIL-supplied, 238
LOSSLESS_BAND_PT_TRANSFORM attribute, 336
LOSSLESS_BAND_SELECTOR attribute, 334

M

macroblocks, 338, 365
maximum value
pixel by pixel, 143
memory format of images, 37
memory images, 43
minimum and maximum values, finding an image's, 178
minimum value
pixel by pixel, 143
molecules
CCITT Group 4 transposition, 407
Cell playback, 288
CellB playback, 303, 399
determining whether they're executing, 410
H.261 playback, 358, 399
involving a copy to a GX display, 409
JPEG baseline sequential playback, 326, 399
MPEG-1 playback, 390, 399
side effects of, 414
monochrome displays, 94
MPEG-1 codec
applications, 359
broken links, 364, 385
compression attributes of
COMPRESSOR_BITS_PER_SECOND, 366
COMPRESSOR_INSERT_VIDEO_SEQUENCE_END, 367
COMPRESSOR_INTRA_QUANTIZATION_TABLE, 373

COMPRESSOR_NON_INTRA_QUANTIZATION_TABLE, 375
COMPRESSOR_PATTERN, 376
COMPRESSOR_PEL_ASPECT_RATIO, 378
COMPRESSOR_PICTURE_RATE, 380
COMPRESSOR_SLICES_PER_PICTURE, 381
COMPRESSOR_TIME_CODE, 382
compressor not supplied with XIL library, 360
creating a CIS associated with an MPEG-1 decompressor, 365
decompression attributes of
DECOMPRESSOR_BROKEN_LINK, 385
DECOMPRESSOR_CLOSED_GOP, 385
DECOMPRESSOR_FRAME_TYPE, 386
DECOMPRESSOR_PEL_ASPECT_RATIO_VALUE, 387
DECOMPRESSOR_PICTURE_RATE_VALUE, 388
DECOMPRESSOR_QUALITY, 384
DECOMPRESSOR_TEMPORAL_REFERENCE, 388
DECOMPRESSOR_TIME_CODE, 389
groups of pictures, 363, 385
how it works, 360
playback molecules, 390, 399
releasing reserved frames, 369
reserved frames, 369
sequences, 364
slices, 365, 381
subgroups, 370
multiplying images, 141
multithread programs, 7

N

nearest-neighbor interpolation, 150
NOT of an image, 144
notational conventions, xxix

O

opening the library, 13
operations, in-place, 499
OPTIMIZE_HUFFMAN_TABLES attribute, 322
OR of images, 144
ordered dither, 118
origin of an image, 46
overlaying images, 99

P

painting images, 213
Photo CD images, 75
capturing from disk, 82
creating, 78
FILEPATH attribute, 79
how stored, 76
MAX_RESOLUTION attribute, 81
RESOLUTION attribute, 80
resolutions, 77
ROTATION attribute, 82
setting device attributes, 79
PhotoYCC color space, 76
pixels
reading the values of, 215
setting the values of, 215
plane mask control, 99
predictive encoding, 328
primary errors, 131
px64, *See*H.261 codec

Q

QCIF images, *see* Quarter Common Interchange Format images
quantization, 308
quantization tables, 322
QUANTIZATION_TABLE attribute, 322

Quarter Common Interchange Format (QCIF) images, 338, 356

R

RDWR_INDICES attribute, 244, 287
read frame, 260
reading
  images from devices other than displays, 67
  images from displays, 66
  images from files, 13, 57
regions of interest
  associating with images, 50
  building, 50
  creating, 49
  defined, 48
  destroying, 49
  naming, 51
  performing geometric operations on, 50
related books, xxix
relational operations
  involving two images, 143
rescaling images, 102
resizing a window, 65
rotating images, 167, 175

S

scaling images, 161
secondary errors, 131
seeking images in compressed image sequences, 257
  CellB, 302
  H.261, 354
  JPEG baseline sequential, 325
sharpening images, 189
shearing images, 171
smoothing images, 190
soft fills, 185
start frame, 260
structuring elements
  creating, 196
  defined, 194
  destroying, 198
  key values of, 194
  naming, 197
subimages, creating, 53
subsampling images, 161
subtracting images, 141
system state, 13

T

TEMPORAL_FILTERING attribute, 283, 323
thread library, 7
thresholding images, 183
translating images, 160
transposing images, 174
trapping, 194
typographic changes, xxix

U

undercolor removal, 125

V

version control, 31
version number
  major, 31
  minor, 31

W

warping images, 172
WIDTH attribute, 301, 393
window
  exposing, 66
  resizing, 65
write frame, 260
writing
  compressed data to files, 227
  images to devices other than displays, 67
  images to displays, 64

images to files, 62

X

X colormaps, 66, 238, 242
XGL functions, using in XIL programs, 459
xgl_to_xil(), 459
XIL data types
 Xil_signed16, 18
 Xil_unsigned8, 18
 XilVersionNumber, 244
XIL enumerations
 XilCellEncodingType, 281
 XilDataType, 16, 58
 XilEdgeCondition, 190
 XilEdgeDetection, 192
 XilErrorCategory, 130
 XilFlipType, 175
 XilJpegHTableType, 312
 XilJpegLLBandSelectorType, 335
 XilMpeg1FrameType, 386
 XilMpeg1PelAspectRatio, 379
 XilMpeg1PictureRate, 380
 XilObjectType, 134
 XilPhotoCDResolution, 81
 XilPhotoCDRotate, 82
XIL objects
 XilCis, 223, 250
 XilColorspace, 123
 XilDevice, 68
 XilDitherMask, 119
 XilError, 127
 XilHistogram, 179
 XilImage, 16
 XilImageType, 235
 XilInterpolationTable, 154
 XilKernel, 188
 XilLookup, 89, 200, 202
 XilSel, 196
 XilSystemState, 13
XIL structures
 XilCellUserData, 280
 XilH261MVSearchRange, 352
 XilIndexList, 242, 287
 XilJpegBandHTable, 311, 331
 XilJpegBandQTable, 312
 XilJpegHTable, 317, 333
 XilJpegHTableValue, 318, 333
 XilJpegLLBandPtTransform, 336
 XilJpegLLBandSelector, 334
 XilJpegQTable, 322
 XilMpeg1Pattern, 376
 XilMpeg1TimeCode, 382, 389
xil.h, 12
xil_absolute(), 143
xil_add(), 144
xil_add_const(), 145
xil_affine(), 169
xil_and_const(), 146
xil_black_generation(), 125
xil_blend(), 211
xil_call_next_error_handler(), 139
xil_cast(), 103
xil_choose_colormap(), 109
xil_cis_attempt_recovery(), 267
xil_cis_create(), 223, 233, 250
xil_cis_destroy(), 251
xil_cis_flush(), 230, 252, 367
xil_cis_get_attribute(), 243
xil_cis_get_autorecover(), 266
xil_cis_get_bits_ptr(), 229, 255, 363, 369
xil_cis_get_input_type(), 252
xil_cis_get_output_type(), 235
xil_cis_get_read_invalid(), 268
xil_cis_has_frame(), 228, 240, 256, 369
xil_cis_put_bits(), 253
xil_cis_put_bits_ptr(), 233, 253
xil_cis_reset(), 250
xil_cis_seek(), 257
xil_cis_set_autorecover(), 265
xil_close(), 28
xil_color_convert(), 93, 124

xil_colorcube_create(), 110
xil_compress(), 251
xil_convolve(), 190
xil_copy(), 98
xil_copy_pattern(), 217
xil_copy_with_planemask(), 99
xil_create(), 16, 58
xil_create_child(), 53
xil_create_from_device(), 71, 78
xil_create_from_window(), 24, 65, 236
XIL_DEBUG environment variable, 410
xil_decompress(), 240, 246, 255, 369
xil_default_error_handler(), 139
xil_device_create(), 69
xil_device_destroy(), 73
xil_device_set_value(), 70
xil_dilate(), 194
xil_dithermask_create(), 118
xil_dithermask_get_by_name(), 119, 238
xil_divide(), 141
xil_edge_detection(), 192
xil_erode(), 194
xil_error_diffusion(), 114
xil_error_get_category(), 130
xil_error_get_category_string(), 131
xil_error_get_id(), 129
xil_error_get_location(), 132
xil_error_get_object(), 132
xil_error_get_primary(), 131
xil_error_get_string(), 129
xil_export(), 59
xil_extrema(), 91, 178
xil_fill(), 183
xil_get_device_attribute(), 72, 80, 81, 82
xil_get_info(), 63
xil_get_memory_storage(), 17, 59
xil_get_pixel(), 216
xil_get_readable(), 55
xil_get_writeable(), 56
xil_histogram(), 180
xil_histogram_create(), 180
xil_histogram_destroy(), 182
xil_histogram_get_values(), 181
xil_imagetype_get_info(), 235
xil_import(), 22, 62
xil_install_error_handler(), 136
xil_interpolation_table_create(), 154
xil_kernel_create(), 115
xil_kernel_destroy(), 191
xil_kernel_get_by_name(), 115
xil_lookup(), 204
xil_lookup_create(), 109, 200, 244
xil_lookup_create_combined(), 202
xil_lookup_destroy(), 205
xil_lookup_get_by_name(), 238
xil_lookup_get_num_entries(), 239
xil_lookup_get_values(), 239
xil_lookup_set_offset(), 239
xil_max(), 143
xil_min(), 143
xil_multiply(), 141
xil_nearest_color(), 108, 246
xil_object_get_error_string(), 133
xil_object_get_type(), 133
xil_open(), 13
xil_ordered_dither(), 118
xil_paint(), 215
xil_remove_error_handler(), 138
xil_rescale(), 102
xil_rotate(), 167
xil_scale(), 162
xil_sel_create(), 196
xil_sel_destroy(), 198
xil_sel_get_by_name(), 197

xil_sel_get_name(), 197
xil_sel_set_name(), 197
xil_set_colorspace(), 123
xil_set_device_attribute(), 72, 79
xil_set_interpolation_tables(), 157
xil_set_pixel(), 216
xil_set_value(), 216
Xil_signed16 data type, 18
xil_soft_fill(), 185
xil_subsample_adaptive(), 164
xil_subsample_binary_to_gray(), 165
xil_subtract(), 141
xil_tablewarp(), 173
xil_tablewarp_horizontal(), 172
xil_tablewarp_vertical(), 172
xil_threshold(), 183
xil_to_xgl(), 459
xil_translate(), 160
xil_transpose(), 174, 407
Xil_unsigned8 data type, 18
XilCellEncodingType enumeration, 281
XilCellUserData structure, 280
XilCis data structure, 223, 250
XilColorspace data structure, 123
XilDataType enumeration, 16, 58
XilDevice data structure, 68
XilDitherMask data structure, 119
XilEdgeCondition enumeration, 190
XilError data structure, 127
XilErrorCategory enumeration, 130
XilFlipType enumeration, 175
XilH261MVSearchRange structure, 352
XilHistogram data structure, 179
XilImage data structure, 16
XilImageType data structure, 235
XilIndexList structure, 242, 287
XilJpegBandHTable structure, 311, 331
XilJpegBandQTable structure, 312
XilJpegHTable structure, 317, 333
XilJpegHTableType enumeration, 312
XilJpegHTableValue structure, 318, 333
XilJpegLLBandPtTransform structure, 336
XilJpegLLBandSelector structure, 334
XilJpegLLBandSelectorType enumeration, 335
XilJpegQTable structure, 322
XilKernel data structure, 188
XilLookup data structure, 89, 200, 202
XilMemoryStorage union, 18
XilMpeg1FrameType enumeration, 386
XilMpeg1Pattern structure, 376
XilMpeg1PelAspectRatio enumeration, 379
XilMpeg1PictureRate enumeration, 380
XilMpeg1TimeCode structure, 382, 389
XilObjectType enumeration, 134
XilSel data structure, 196
XilSystemState data structure, 13
XilVersionNumber data type, 244
XOR of images, 144
XResizeWindow(), 65

Z

zooming images, 161