























# Intel Image Processing Library

---

*Reference Manual*

Copyright © 1996, 1997, Intel Corporation  
All Rights Reserved  
Issued in U.S.A.  
Order Number 663791B-002

## How to Use This Online Manual


- |   |   |   |  |
|---|---|---|--|
|    | Click to hide or show subtopics when the bookmarks are shown.     |    | Click to go to the previous page.  |
|    | Double-click to jump to a topic when the bookmarks are shown.     |    | Click to go to the next page.  |
|    | Click to display bookmarks.                                       |    | Click to go to the last page.  |
|    | Click to display thumbnails.                                      |    | Click to return back to the previous view. Use this button when you need to go back after using the jump button (see below). |
|    | Click to close bookmark or thumbnail view.                        |    | Click to go forward from the previous view.  |
|   | Click and use on the page to drag the page in vertical direction. |   | Click to set 100% of the page view.  |
|  | Click and drag to the page to magnify the view.                   |  | Click to display the entire page within the window.  |
|  | Click and drag to the page to reduce the view.                    |  | Click to fill the width of the window.   |
|  | Click and drag the selection cursor to the page.                  |  | Click to open a dialog to search for a word or multiple words.   |
|  | Click to go to the first page of the manual.                      |  | Click jump button on manual pages to jump to the related subjects. Use the return back icon above to go back.                |

**Printing an Online File.** Select **Print** from the **File** menu to print an online file. The dialog that opens allows you to print full text, range of pages, or selection.

**Viewing Multiple Online Manuals.** Select **Open** from the **File** menu, and open a .PDF file you need. Select **Cascade** from the **Window** menu to view multiple files.

**Resizing the Bookmark Area.** Drag the double-headed arrow that appears on the area's border as you pass over it.

**Jumping to Topics.** Throughout the text of this manual, you can jump to different topics by clicking on keywords printed in green color, underlined style or on page numbers in a box.

To return to the page from which you jumped, use the  icon in the tool bar. Try this example:

This software is briefly described in the [Overview](#); see page 1-1.

If you click on the phrase printed in green color, underlined style, or on the page number, the Overview opens.

# *Intel Image Processing Library Reference Manual*

---

Order Number: 663791B-002

<b>Revision</b>	<b>Revision History</b>	<b>Date</b>
B-001	Documents release 1.0 Beta of the library.	01/20/97
B-002	Documents release 1.0 Beta 2 of the library. The iplBitonalToGray, iplCreateColorTwist, iplDeleteColorTwist, iplShiftS, iplMultiplyScale, iplMultiplySScale, and iplNot functions have been added. Miscellaneous edits have been made.	04/04/97

Information in this manual is provided in connection with Intel products. Intel assumes no liability whatsoever, including infringement of any patent or copyright, for sale and use of Intel products except as provided in Intel's Terms and Conditions of the license agreement for such products.

Intel Corporation retains the right to make changes to these documents at any time, without notice.

Contact your local sales office to obtain the latest specifications before placing your order.

Since publication of documents referenced in this document, registration of the Pentium and iCOMP trademarks has been issued to Intel Corporation.

Intel, the Intel logo, iCOMP, Pentium, and Indeo are registered trademarks of Intel Corporation.

MMX and i486 are trademarks of Intel Corporation.

\*Other brands and names are the property of their respective owners.

Copyright © 1996, 1997, Intel Corporation. All Rights Reserved.

# Contents

---

## Chapter 1 Overview

About This Software .....	1-1
Hardware and Software Requirements .....	1-1
About This Manual.....	1-2
Manual Organization.....	1-2
Function Descriptions .....	1-4
Audience for This Manual .....	1-4
Online Version .....	1-4
Related Publications .....	1-5
Notational Conventions .....	1-5
Font Conventions.....	1-5
Naming Conventions .....	1-5
Function Name Conventions.....	1-6

## Chapter 2 Image Architecture

Data Architecture .....	2-1
Color Models.....	2-1
Data Types and Palettes .....	2-2
The Sequence and Order of Color Channels .....	2-3
Coordinate Systems .....	2-4
Image Regions.....	2-4
Setting an ROI for Multi-Image Operations .....	2-5
Alpha (Opacity) Channel.....	2-6
Scanline Alignment .....	2-6
Image Dimensions .....	2-7
Execution Architecture .....	2-7
Handling Overflow and Underflow .....	2-7
In-Place and Out-of-Place Operations .....	2-7

Image Tiling .....	2-8
Tile Size .....	2-8
IplCoord Structure .....	2-9
Call-backs .....	2-9
ROI and Tiling .....	2-10
In-Place Operations and Tiling .....	2-10
<b>Chapter 3 Error Handling</b>	
Error-handling Functions .....	3-2
Error .....	3-2
GetErrStatus .....	3-3
SetErrStatus .....	3-3
GetErrMode .....	3-4
SetErrMode .....	3-4
ErrorStr .....	3-5
RedirectError .....	3-6
Error Macros .....	3-7
Status Codes .....	3-8
Error Handling Example .....	3-10
Adding Your Own Error Handler .....	3-13
<b>Chapter 4 Image Creation and Access</b>	
Image Header and Attributes .....	4-3
Tiling Fields in the IplImage Structure .....	4-7
IplTileInfo Structure .....	4-7
Creating Images .....	4-8
CreateImageHeader .....	4-8
AllocatImage .....	4-10
DeallocatImage .....	4-11
Deallocate .....	4-11
Setting Regions of Interest .....	4-12
CreateROI .....	4-13

SetROI .....	4-14
Image Borders and Image Tiling .....	4-15
SetBorderMode .....	4-15
CreateTileInfo .....	4-17
SetTileInfo .....	4-18
DeleteTileInfo .....	4-18
Memory Allocation Functions .....	4-19
Malloc .....	4-19
wMalloc .....	4-20
iMalloc .....	4-21
sMalloc .....	4-21
dMalloc .....	4-22
Free .....	4-23
Image Data Exchange .....	4-24
Set .....	4-24
Copy .....	4-25
Exchange .....	4-26
Convert .....	4-27
Working in the Windows DIB Environment .....	4-28
TranslateDIB .....	4-30
ConvertFromDIB .....	4-32
ConvertToDIB .....	4-33

## **Chapter 5 Image Arithmetic and Logical Operations**

Monadic Arithmetic Operations .....	5-3
AddS .....	5-3
SubtractS .....	5-4
MultiplyS .....	5-4
MultiplySScale .....	5-5
Square .....	5-6

Dyadic Arithmetic Operations .....	5-7
Add .....	5-7
Subtract .....	5-8
Multiply.....	5-9
MultiplyScale.....	5-10
Monadic Logical Operations .....	5-11
LShiftS .....	5-11
RShiftS.....	5-12
Not .....	5-12
AndS .....	5-13
OrS .....	5-14
XorS .....	5-15
Dyadic Logical Operations .....	5-15
And .....	5-16
Or .....	5-16
Xor .....	5-17
Image Compositing Based on Opacity .....	5-18
Using Pre-multiplied Alpha Values .....	5-19
AlphaComposite.....	5-19
AlphaCompositeC .....	5-19
PreMultiplyAlpha .....	5-25

## **Chapter 6 Image Filtering**

Linear Filters .....	6-1
Blur .....	6-2
2D Convolution .....	6-3
CreateConvKernel .....	6-4
GetConvKernel .....	6-5
DeleteConvKernel .....	6-6
Convolve2D .....	6-7
ConvolveSep2D .....	6-8



Non-linear Filters .....	6-9
MedianFilter .....	6-9
MaxFilter .....	6-10
MinFilter .....	6-11
<b>Chapter 7 Linear Image Transforms</b>	
Fast Fourier Transform .....	7-1
Real-Complex Packed (RCPack2D) Format .....	7-1
RealFft2D.....	7-3
CcsFft2D.....	7-5
Discrete Cosine Transform .....	7-6
DCT2D.....	7-7
<b>Chapter 8 Morphological Operations</b>	
Erode .....	8-2
Dilate .....	8-5
Open .....	8-6
Close .....	8-7
<b>Chapter 9 Color Space Conversion</b>	
Reducing the Image Bit Resolution .....	9-2
ReduceBits .....	9-2
Conversion of Bitonal Images to Gray Scale .....	9-4
BitonalToGray.....	9-4
Conversion of Absolute Colors to and from Palette Colors	9-5
Conversion from Color to Gray Scale .....	9-5
ColorToGray .....	9-6
Conversion from Gray Scale to Color (Pseudo-color) .....	9-7
GrayToColor .....	9-7
Conversion of Color Models .....	9-8
RGB2HSV.....	9-9
HSV2RGB.....	9-10

RGB2HLS .....	9-10
HLS2RGB .....	9-11
Using Color-Twist Matrices .....	9-12
CreateColorTwist .....	9-12
SetColorTwist.....	9-13
ApplyColorTwist .....	9-14
DeleteColorTwist.....	9-16
<b>Chapter 10 Histogram and Thresholding Functions</b>	
Thresholding .....	10-2
Threshold.....	10-2
Lookup Table (LUT) and Histogram Operations .....	10-3
The IplLUT Structure.....	10-3
ContrastStretch .....	10-5
ComputeHisto .....	10-6
HistoEqualize .....	10-7
<b>Chapter 11 Linear Geometric Transforms</b>	
Changing Image Size .....	11-1
Zoom .....	11-2
Decimate.....	11-3
Changing Image Orientation .....	11-4
Rotate .....	11-5
Mirror .....	11-6
<b>Appendix A Supported Image Attributes and Operation Modes</b>	A-1
<b>Bibliography</b>	
<b>Glossary</b>	
<b>Index</b>	

## Figures

Figure 2-1 Setting an ROI for Multi-Image Operations .....	2-5
Figure 4-1 RGB Image with a Rectangular ROI and a COI .....	4-5
Figure 8-1 Erosion in a Rectangular ROI .....	8-3

## Tables

Table 2-1 Data Ordering .....	2-3
Table 3-1 ipLError() Status Codes .....	3-8
Table 4-1 Image Creation, Data Exchange and Windows DIB Environment Functions .....	4-1
Table 4-2 IPL Image Header Attributes .....	4-3
Table 5-1 Image Arithmetic and Logical Operations .....	5-1
Table 5-2 Types of Image Compositing Operations .....	5-23
Table 6-1 Image Filtering Functions .....	6-1
Table 7-1 Linear Image Transform Functions .....	7-1
Table 7-2 Arrangement of Output Samples in RCPack2D Format .....	7-3
Table 8-1 Morphological Operation Functions .....	8-1
Table 9-1 Color Space Conversion Functions .....	9-1
Table 9-2 Source and Resultant Image Data Types for Reducing the Bit Resolution .....	9-4
Table 9-3 Source and Resultant Image Data Types for Conversion from Color to Gray Scale .....	9-6
Table 9-4 Source and Resultant Image Data Types for Conversion from Gray Scale to Color .....	9-8
Table 10-1 Point Operation Functions .....	10-1
Table 11-1 Image Geometric Transform Functions .....	11-1

## Examples

Example 3-1 Error Functions .....	3-10
Example 3-2 Output for the Error Function Program .....	3-12
Example 3-3 Output for the Error Function Program .....	3-12
Example 3-4 A Simple Error Handler .....	3-14

Example 4-1 IPLImage Definition .....	4-6
Example 4-2 IpIROI Definition.....	4-12
Example 8-1 Code Used to Produce Erosion in a Rectangular ROI .....	8-4
Example 10-1 IpLUT Definition .....	10-3

# Overview

---

# 1

This manual describes the structure, operation and functions of the Intel Image Processing Library (IPL). This library supports many functions whose performance can be significantly enhanced on the Intel Architecture (IA), particularly the MMX™ technology.

The manual describes the architecture of the IPL data and execution and provides detailed descriptions of the functions included in the Intel Image Processing Library.

This chapter introduces the Intel Image Processing Library and explains the organization of this manual.

## About This Software

The Intel Image Processing Library focuses on taking advantage of the parallelism of the SIMD (single-instruction, multiple-data) instructions that comprise the MMX technology. This technology improves the performance of computationally intensive image processing functions. Thus this library includes a set of functions whose performance significantly improves when used with the Intel Architecture MMX technology. The library does not support the reading and writing of a wide variety of image file formats or the display of images.

## Hardware and Software Requirements

The Intel Image Processing Library runs on personal computers that are based on Intel Architecture processors and running Microsoft\* Windows\*, Windows 95\*, or Windows NT\*. The library integrates into the customer's application or library written in C or C++.

## About This Manual

This manual provides a background of the image and execution architecture of the Intel Image Processing Library as well as detailed descriptions of the IPL functions. The IPL functions are combined in groups by their functionality. Each group of functions is described in a separate chapter (chapters 3 through 11).

## Manual Organization

This manual contains eleven chapters:

- |           |   |
|-----------|---|
| Chapter 1 | “ <a href="#">Overview</a> .” Introduces the Image Processing Library software, provides information on manual organization, and explains notational conventions.   |
| Chapter 2 | “ <a href="#">Image Architecture</a> .” Describes the data formats supported, the execution architecture, and image tiling. The data formats include color models, data types, data order, and coordinate systems. The execution architecture discussion focuses on overflow and underflow issues and in-place and out-of-place operations. |
| Chapter 3 | “ <a href="#">Error Handling</a> .” Provides information on the error-handling functions included with the library. User-defined error handler is also described.   |
| Chapter 4 | “ <a href="#">Image Creation and Access</a> .” Describes the functions used to: create, set, and access image attributes; set image border and tiling; and allocate the memory for different data types. The chapter also describes the functions that facilitate operations in the window environment.                                     |

- Chapter 5      “[Image Arithmetic and Logical Operations.](#)”  
Describes image processing operations that modify pixel values using simple arithmetic or logical operations. These operations include monadic operations (single input image) and dyadic operations (two input images).
- Chapter 6      “[Image Filtering.](#)” Describes linear and non-linear filtering operations that can be applied to images.
- Chapter 7      “[Linear Image Transforms.](#)” Describes the fast Fourier transform (FFT) and Discrete Cosine Transform (DCT) implemented in the IPL.
- Chapter 8      “[Morphological Operations.](#)” Describes the morphological operations supported in the library: simple Erosion, Dilation, Opening and Closing
- Chapter 9      “[Color Space Conversion.](#)” Describes the color space conversions supported in the library; for example, color reduction from high resolution color to low resolution color; conversion from Palette to Absolute color and vice versa; conversion to different color models.
- Chapter 10     “[Histogram and Thresholding Functions.](#)”  
Describes functions that treat an image on a pixel-by-pixel basis: operations that alter the histogram of the image; contrast stretching, histogram computation, histogram equalization and thresholding.
- Chapter 11     “[Image Geometric Transforms.](#)” Describes geometric transforms: Zoom, Decimate, Rotate, and Mirror.

The manual also includes [Appendix](#) that lists supported image attributes and operation modes, [Glossary](#) of terms, [Bibliography](#), and [Index](#).

## Function Descriptions

In Chapters 3 through 11, each function is introduced by name (without the `ipl` prefix) and a brief description of its purpose. This is followed by the function call sequence, more detailed description of the function's purpose, and definitions of its arguments. The following sections are included in each function description:

<i>Arguments</i>	Describes all the function arguments.
<i>Discussion</i>	Defines the function and describes the operation performed by the function. Often, code examples and the equations the function implements are included.
<i>Return Value</i>	If present, describes a value indicating the result of the function execution.
<i>Application Notes</i>	If present, describe any special information which application programmers or other users of the function need to know.
<i>See Also</i>	If present, lists the names of functions which perform related tasks.

## Audience for This Manual

The manual is intended for the developers of image processing applications and image processing libraries. Both parts of the audience are expected to be experienced in using C and to have a working knowledge of the vocabulary and principles of image processing. The developers of image processing software can use the Intel IPL capabilities to improve performance on IA with MMX technology.

## Online Version

This manual is available in an online hypertext format. To obtain a hard copy of the manual, print the online file using the printing capability of Adobe\* Acrobat, the tool used for the online presentation of the document.



## Related Publications

For more information about computer graphics concepts and objects, refer to the books and materials listed in the [Bibliography](#).

## Notational Conventions

In this manual, notational conventions include:

- Fonts used for distinction between the text and the code
- Naming conventions
- Function name conventions

## Font Conventions

The following font conventions are used:

<code>UPPERCASE COURIER</code>	Used in the text for IPL constant identifiers; for example, <code>IPL_DEPTH_1U</code> .
<code>lowercase courier</code>	Mixed with the uppercase in function names as in <code>SetExecutionMode</code> ; also used for key words in code examples; for example, in the function call statement <code>void iplSquare()</code> .
<i>lowercase mixed with UpperCase Courier italic</i>	Variables in arguments and parameters discussion; for example, <i>mode</i> , <i>dstImage</i> .

## Naming Conventions

The following data type conventions are used by the IPL library:

- Constant identifiers are in uppercase; for example, `IPL_SIDE_LEFT`.
- All constant identifiers have the `IPL` prefix.
- All of the function names have the `ipl` prefix. In code examples, you can distinguish the IPL interface functions from the application functions by this prefix.



**NOTE.** In this manual, the `ipl` prefix in function names is always used in the code examples. In the text, this prefix is sometimes omitted.

- All image header structures have the `Ipl` prefix; for example, `IplImage`, `IplROI`.
- Each new part of a function name starts with an uppercase character, without underscore; for example, `iplAlphaComposite`.

## Function Name Conventions

The function names in the IPL library typically begin with the `ipl` prefix and have the following general format:

```
ipl < action > < target > < mod >()
```

where

<i>action</i>	indicates the core functionality; for example, <code>-Set-</code> , <code>-Create-</code> , or <code>-Convert-</code> .
<i>target</i>	indicates the area where image processing is being enacted; for example, <code>-ConvKernel</code> or <code>-FromDIB</code> .  In a number of cases, the target consists of two or more words; for example, <code>-ConvKernel</code> in the function <code>CreateConvKernel</code> . Some function names consist of an <i>action</i> or <i>target</i> only; for example, the functions <code>Multiply</code> or <code>RealFft2D</code> , respectively.
<i>mod</i>	The <i>mod</i> field is optional and indicates a modification to the core functionality of a function. For example, in the name <code>iplAlphaCompositeC()</code> , <code>C</code> indicates that this function is using constant alpha values.

This chapter describes the data architecture and the execution architecture of the Intel Image Processing Library (IPL). The data formats supported by IPL define the image architecture as explained in the sections that follow.

## **Data Architecture**

An IPL image has a header that describes the image as a list of attributes and pointers to the data associated with the image. Library functions use the image header to get the format and characteristics of the image(s) passed to the functions. Based on the information obtained from the header, the functions make appropriate calls to set the data structures. Images can have different organization of data. IPL supports numerous data formats that use different color models, data types, data order, and coordinate systems.

## **Color Models**

The IPL image format supports the following color models:

- Monochrome or gray scale image (one color channel)
- Color image (3 or 4 color channels)
- Multi-spectral image (any number of channels).

Color models are defined by the number of channels and the colors they contain. Examples of three-channel models are RGB, HSV, CMY, and YCC. Examples of four-channel color models are CMYK and RGBA.

Image processing operations can be performed on one or all channels in the image. The operations are performed without specific identification of the colors, unless it is a certain color conversion operation where color identification is required.

The multi-spectral image (MSI) model is used for general purpose images. It is used for any kind of multi-spectral data and any kind of image. For example, the Fourier transform operation writes transform coefficients of color or monochrome images to this model—one channel for each channel in the input. The result can be viewed as an MSI image. An MSI image can contain any number of color channels; they may even correspond to invisible parts of the spectrum. The library functions do not need to identify any specific MSI image channels.

### Data Types and Palettes

The parameter that determines the image data type is the pixel depth in bits. The data could be signed or unsigned. The following data types are supported for various color models (s = signed, u = unsigned):

Gray scale	1, 8s, 8u, 16s, and 16u bits per pixel
Color (three-channel)	8u and 16u bits per channel
Four-channel and MSI	8s, 8u, 16s, 16u, 32s bits per channel

IPL supports only absolute color images in which each pixel is represented by the channel intensities. For example, in an absolute color 24-bit RGB image, three bytes (24 bits) per pixel represent the three channel intensities. LUT (lookup table) images, that is, palette color images are not supported. You must convert palette images to absolute color images for further processing by the IPL functions. The library provides functions to convert DIB palette images to IPL absolute color images.

Color images with 8, 16, or 32 bits per channel simply pack each channel, respectively, into a byte, word, or doubleword. All channels within a given image have the same data type.

Signed data (8s, 16s, or 32s) are used for storing the output of some image processing operations; for example, this is the case for transforms such as FFT. Unless specified otherwise, signed data cannot be used as input to image processing operations.

## The Sequence and Order of Color Channels

Channel sequence corresponds to the order of the color channels in absolute color images. For example, in an RGB image the channels could be stored in the sequence RGB or in the sequence BGR. The sequence of the channels is not important to the image processing functions which do not identify the channels.



**NOTE.** *For the color conversion operations, the information about channel sequence is required and therefore must be provided.*

For images with pixel-oriented data, the channel sequence corresponds to the color data order for each pixel. Data ordering corresponds to the way the color data is arranged: by planes or by pixels. Table 2-1 lists the orderings that are supported for planes and for pixels.

**Table 2-1 Data Ordering**

<b>Data Ordering</b>	<b>Description</b>	<b>RGB Example (channel ordering = RGB)</b>
Pixel-oriented	All channels for each pixel are clustered.	RGBRGRGB (line 1) RGBRGRGB (line 2) RGBRGRGB (line 3)
Plane-oriented	All image data for each channel is contiguous followed by the next channel.	RRRRRRRR (line 1) RRRRRRRR (line 2) R plane RRRRRRRR (line 3)  GGGGGGGG (line 1) GGGGGGGG (line 2) G plane GGGGGGGG (line 3) ...

## Coordinate Systems

Two coordinate systems are supported by the IPL image format.

- The origin of the image is in the top left corner, the x values increase from left to right, and y values increase from top to bottom.
- The origin of the image is in the bottom left corner, the x values increase from left to right, and y values increase from the bottom to the top.

## Image Regions

A very important concept in the IPL library architecture is an image region of interest (ROI). All image processing functions can operate not only on entire images but also on image regions.

Depending on the processing needs, the following image regions can be specified:

- A channel of interest (COI). A COI can be one or all channels of the image. By default, unless the COI is changed by the `SetROI()` function, processing will be carried out on all channels in the image.
- A rectangular region of interest (rectangular ROI). A rectangular ROI is a portion of the image or, possibly, the entire image. By default, unless changed by the `SetROI()` function, the entire image is the rectangular region of interest.

An IPL image can simultaneously have a rectangular ROI and a channel of interest. If this is the case, operations will be performed on the “intersection” of these two ROIs.

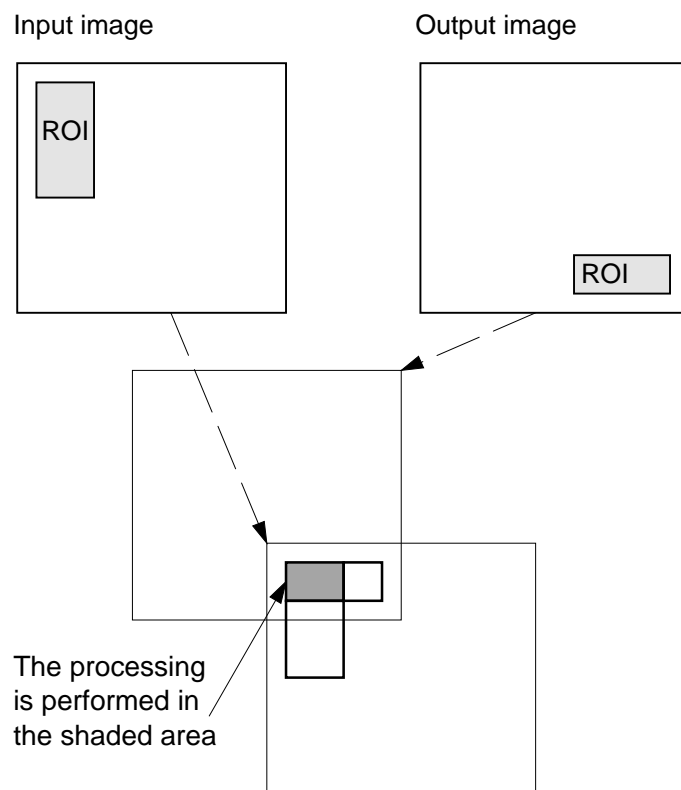
Thus an image region specifies some part of an image or the entire image. Once set, the region information of the IPL image remains the same until changed by the function `SetROI()`.

### Setting an ROI for Multi-Image Operations

Figure 2-1 illustrates image processing operations that take one or more input images and store the results onto an output image.

**Figure 2-1** Setting an ROI for Multi-Image Operations

---



---

All images (input and output) in Figure 2-1 have rectangular ROIs that specify either the entire image or specific regions set by the `SetROI()` function. The first step is to align the rectangular ROIs of all the images so that their top left corners coincide. The operation is, then, performed in the

rectangular region where all the images overlap. This scheme gives much flexibility, effectively enabling translation of image data (even for equal-size images) from one region of an input image to another region of an output image.

To successfully perform an image processing operation, one of the following conditions must be met for the channel of interest (COI):

- Each image (input and output) has one channel in the channel COI,
- Each image (input and output) has all channels included in the ROI (COI = 0) and all images (input and output) have the same number of channels (one or more).

If one image (input or output) has one channel in its COI and another image (input or output) has more than one channel included in its COI, an error will occur.

## Alpha (Opacity) Channel

In addition to the color channels, an IPL image can have one alpha channel, also known as an opacity channel, which is mainly used for image compositing operations (see “[Image Compositing Based on Opacity](#)” in Chapter 5).

The alpha channel is treated like any other channel in the IPL format; you are allowed not to identify it when not required. However, alpha channels must be explicitly specified for the functions that require them (see, for example, the [iplAlphaComposite](#) function).

## Scanline Alignment

Image row data (scanline) can be aligned on doubleword (32-bit) or quadword (64-bit) boundaries. Each row is padded with zeros if required. For maximum performance with MMX™ technology, it is important to have the image data aligned on quadword boundaries.



## Image Dimensions

There is no practical limit of the image size. An unsigned long integer is used for the height and width of the image. This allows you to create images of size up to  $2^{31}$  by  $2^{31}$  pixels, which is much beyond the hardware and OS constraints of today's PCs or workstations.

## Execution Architecture

This section describes the execution time issues such as overflow/underflow handling and in-place and out-of-place operations. Cache optimization and deferred execution mode will be discussed in the next release.

### Handling Overflow and Underflow

Overflow and underflow are handled in each image processing function. The default mode of operation is saturation which is a mode that prevents from potential overflow or underflow of the values. In saturation mode, when an overflow of a value is about to happen, this value is clamped to the maximum permissible value (for example, 255 for an unsigned byte). Similarly, when underflow of a value is about to happen, it is clamped to the minimum permissible value, which is always zero for the case of unsigned bytes.

### In-Place and Out-of-Place Operations

All image processing operations in the library can be in-place or out-of-place operations, unless it is explicitly specified that a particular operation must belong to one of these categories only. With an in-place operation, the output image is one of the input images modified (that is, the pointer to the output image is the same as the pointer to the input one). With an out-of-place operation, the output image is a new image, not the same as any of the input images.

## Image Tiling

Tiling is a method of image representation in which, for reasons of efficiency, the image is broken up into smaller images, or tiles. The whole image is reconstructed by arranging the individual tiles in a grid. Usually, the tiles are of a uniform size; special techniques can be used to accommodate an overall image height or width that is not an even multiple of the tile size. Image processing applications frequently use square tiles with sizes that are multiples of two; for example, a 64-by-64 tile is typical.

In most IPL functions, tiled images may be used in the same way as non-tiled images, subject to some restrictions. The effect of functions on tiled images is always the same as that of the same function on an image of the same size and content that is not tiled. The behavior varies as stated below, particularly in the call-back requirement.

This section gives a short overview of image tiling in the IPL. In Chapter 4 you can find more information about tiling, namely, the descriptions of the [TileInfo](#) structure, the [imageID](#) parameter, and the functions [CreateTileInfo](#), [SetTileInfo](#), and [DeleteTileInfo](#).

### Tile Size

In the IPL, all tiles must be of the same size, including those on the edge of an image. The tiles on the edge of an image must contain valid data up to the border of the image; beyond that, the pixels are ignored, and the border mode is used instead.

The size of the image tiles is contained within the [IplTileInfo](#) structure. It is restricted to being an even multiple of 8 in each dimension.

For functions that take more than one source image, either all source images must be tiled with equally-sized tiles or they must all be non-tiled. The source and destination images tiling and tile sizes need not be the same.

## IplCoord Structure

The arguments of every function with both an input and output image include a pointer to the following structure:

```
iplFunction(..., IplCoord*);  
  
typedef struct _IplCoord {  
    int xDst, yDst;  
} IplCoord
```

The `xDst` and `yDst` fields are the offsets (in pixels) of the origin of the destination image from the origin of the source image.

If there are multiple source images, these images are assumed to have their ROI origins aligned with each other.

The `IplCoord` structure is not specific to tiling in that it is supported whether the source/destination images are tiled or not. Its intent is to support for tiling schemes **other** than the IPL tiling scheme by providing a method of identifying the location of the destination with respect to the source. Such a method is important if the argument images are, unbeknownst to the IPL functions, actually tiles in a user-implemented tiling scheme. The `IplCoord` parameter is ignored if `NULL`.

## Call-backs

Since the `IplImage` structure does not contain any image data, functions operating on tiled images must acquire data tile-by-tile. To do this, the IPL uses a system of call-backs, in which the IPL function requests pointers to individual tiles based on need.

The call-back system is implemented (by the library user) as a single function, the prototype and behavior of which are specified below. When called **by the library**, this function provides or releases one tile's worth of data. The function is specified to the library in the `callBack` field of the `IplTileInfo` structure. The prototype is as follows:

```
void (*IplCallBack) (const IplImage* img, int xIndex,  
                    int yIndex, int mode);
```

where

`img` is the header provided for the parent image;  
`xIndex` and `yIndex` are the indices of the requested tile; they refer to the tile number, not pixel number, and count from the origin at (0,0);  
`mode` is one of the following:

<code>IPL_GET_TILE_TO_READ</code>	get a tile for reading; the tile data is returned in <code>img-&gt;tileInfo-&gt;tileData</code> and must not be changed;
<code>IPL_GET_TILE_TO_WRITE</code>	get a tile for writing; the tile data is returned in <code>img-&gt;tileInfo-&gt;tileData</code> and may be changed; changes will be reflected in the image;
<code>IPL_RELEASE_TILE</code>	release tile; commit writes.

Memory pointers provided by a get function will not be used after the corresponding release function has been called.

## ROI and Tiling

The meaning and behavior of ROI for tiled images are identical to those for a non-tiled image.

## In-Place Operations and Tiling

Functions that are called with identical source and destination images (header pointers equal) are handled correctly by the library, even with tiling. If the source and destination image pointers are not equal, no support for source and destination overlap is provided.

Note that the presence of `IplROI` and/or `IplCoord` structures does not affect this restriction.

## Error Handling

---

This chapter describes the error handling facility of the Image Processing Library. The IPL functions report a variety of errors including bad arguments and out-of-memory conditions. When a function detects an error, instead of returning a status code, the function signals an error by calling `iplSetErrStatus()`. This allows the error handling mechanism to work separately from the normal flow of the image processing code. Thus, the image processing code is cleaner and more compact as shown in this example:

```
ColorTwist = iplSetColorTwist(data, scalingValue);  
if(iplGetErrStatus() $<$ 0) // check for errors
```

The error handling system is hidden within the function `iplSetColorTwist()`. As a result, this statement is uncluttered by error handling code and closely resembles a mathematical formula.

Your application should assume that every library function call may result in some error condition. The Image Processing Library performs extensive error checks (for example, `NULL` pointers, out-of-range parameters, corrupted states) for every library function.

Error macros are provided to simplify the coding for error checking and reporting. You can modify the way your application handles errors by calling `iplRedirectError()` with a pointer to your own error handling function. For more information, see “[Adding Your Own Error Handler](#)” later in this chapter. For even more flexibility, you can replace the whole error handling facility with your own code. The source code of the default error handling facility is provided.

The Image Processing Library does not process numerical exceptions (for example, overflow, underflow, and division by zero). The underlying floating point library or processor has the responsibility for catching and

reporting these exceptions. A floating-point library is needed if a processor that handles floating-point is not present. You can attach an exception handler using an underlying floating-point library for your application, if your system supports such a library.

## Error-handling Functions

The following sections describe the error functions in the Image Processing Library.

---

### Error

*Performs basic error handling.*

---

```
void iplError(IPLStatus status, const char *func,  
             const char *context);
```

<i>status</i>	Code that indicates the type of error (see Table 3-1, “ <a href="#">iplError() Status Codes</a> ”).
<i>func</i>	Name of the function where the error occurred.
<i>context</i>	Additional information about the context in which the error occurred. If the value of <i>context</i> is <code>NULL</code> or empty, this string will not appear in the error message.

### Discussion

The `iplError()` function must be called whenever any of the IPL functions encounters an error. The actual error reporting is handled differently, depending on whether the program is running in Windows mode or in console mode. Within each invocation mode, you can set the

error mode flag to alter the behavior of the `iplError()` function. For more information on the defined error modes, see “[SetErrMode](#)” section.

To simplify the coding for error checking and reporting, the error handling system supplied by the IPL Library supports a set of error macros. See “[Error Macros](#)” for a detailed description of the error handling macros.

The `iplError()` function calls the default error reporting function. You can change the default error reporting function by calling `iplRedirectError()`. For more information, see “[RedirectError](#)” (for `iplRedirectError()`).

---

## GetErrStatus

## SetErrStatus

*Gets and sets the error codes that describe the type of error being reported.*

---

```
typedef int IPLStatus;
IPLStatus iplGetErrStatus();
void iplSetErrStatus(IPLStatus status);
```

*status* Code that indicates the type of error (see Table 3-1, “[iplError\(\) Status Codes](#)”).

### Discussion

The `iplGetErrStatus()` and `iplSetErrStatus()` functions get and set the error status codes that describe the type of error being reported. See “[Status Codes](#)” for descriptions of each of the error status codes.

## GetErrMode SetErrMode

*Gets and sets the error modes that describe how an error is processed.*

---

```
#define IPL_ErrModeLeaf    0
#define IPL_ErrModeParent 1
#define IPL_ErrModeSilent 2
int iplGetErrMode();
void iplSetErrMode(int errMode);
```

*errMode* Indicates how errors will be processed. The possible values for *errMode* are `IPL_ErrModeLeaf`, `IPL_ErrModeParent`, or `IPL_ErrModeSilent`.

### Discussion



---

**NOTE.** *This section describes how the default error handler handles errors for applications which run in console mode. If your application has a custom error handler, errors will be processed differently than described below*

---

The `iplSetErrMode()` function sets the error modes that describe how errors are processed. The defined error modes are `IPL_ErrModeLeaf`, `IPL_ErrModeParent`, and `IPL_ErrModeSilent`.

If you specify `IPL_ErrModeLeaf`, errors are processed in the “leaves” of the function call tree. The `iplError()` function (in console mode) prints an error message describing *status*, *func*, and *context*. It then terminates the program.



If you specify `IPL_ErrModeParent`, errors are processed in the “parents” of the function call tree. When `iplError()` is called as the result of detecting an error, an error message will print, but the program will not terminate. Each time a function calls another function, it must check to see if an error has occurred. When an error occurs, the function should call `iplError()` specifying `IPL_StsBackTrace`, and then return. The macro `IPL_ERRCHK()` may be used to perform both the error check and back-trace call. This passes the error “up” the function call tree until eventually some parent function (possibly `main()`) detects the error and terminates the program.

`IPL_ErrModeSilent` is similar to `IPL_ErrModeParent`, except that error messages are not printed.

`IPL_ErrModeLeaf` is the default, and is the simplest method of processing errors. `IPL_ErrModeParent` requires more programming effort, but provides more detailed information about where and why an error occurred. All of the functions in the library support both options (that is, they use `IPL_ERRCHK()` after function calls). If an application uses the `IPL_ErrModeParent` option, it is essential that it check for errors after all library functions that it calls.

The status code of the last detected error is stored into the global variable `IplLastStatus` and can be returned by calling `iplGetErrStatus()`. The value of this variable may be used by the application during the back-trace process to determine what type of error initiated the back trace.

---

## ErrorStr

*Translates an error or status code into a textual description.*

---

```
const char* iplErrorStr(IPLStatus status);
```

*status*

Code that indicates the type of error (see Table 3-1, “[iplError\(\) Status Codes](#)”).

## Discussion

The function `iplErrorStr()` returns a short string describing *status*. Use this function to produce error messages for users. The returned pointer is a pointer to an internal static buffer that may be overwritten on the next call to `iplErrorStr()`.

---

## RedirectError

*Assigns a new error handler to call when an error occurs.*

---

```
IPLErrCallback iplRedirectError(IPLErrCallback func);
```

*func* Pointer to the function that will be called when an error occurs.

## Discussion

The `iplRedirectError()` function assigns a new function to be called when an error occurs in the IPL Library. If *func* is `NULL`, `iplRedirectError()` installs the IPL Library's default error handler.

The return value of `iplRedirectError()` is a pointer to the previously assigned error handling function.

For the definition of the function typedef `IPLErrCallback`, see the include file `iplerror.h`. See [“Adding Your Own Error Handler”](#) for more information on the `iplRedirectError()` function.

## Error Macros

The error macros associated with the `iplError()` function are described below.

```
#define IPL_ERROR(status, func, context) \
    iplError((status),(func),(context));

#define IPL_ERRCHK(func, context)\
    ( (iplGetErrStatus())>=0) ? IPL_StsOk \
    : IPL_ERROR(IPL_StsBackTrace,(func),(context)) )

#define IPL_ASSERT(expr, func, context)\
    ( ( expr) ? IPL_StsOk\
    : IPL_ERROR(IPL_StsInternal,(func),(context)) )

#define IPL_RSTERR()      (iplSetErrStatus(IPL_StsOk))
```

<i>context</i>	Provides additional information about the context in which the error has occurred. If the value of <i>context</i> is <code>NULL</code> or empty, this string does not appear in the error message.
<i>expr</i>	An expression that checks for an error condition and returns <code>FALSE</code> if an error has occurred.
<i>func</i>	Name of the function where the error occurred.
<i>status</i>	Code that indicates the type of error (see Table 3-1, “ <a href="#">iplError() Status Codes</a> .”)

## Discussion

The `IPL_ASSERT()` macro checks for the error condition *expr* and sets the error status `IPL_StsInternal` if the error occurred.

The `IPL_ERRCHK()` macro checks to see if an error has occurred by checking the error status. If an error has occurred, `IPL_ERRCHK()` creates an error back trace message and returns a non-zero value. This macro should normally be used after any call to a function that might have signaled an error.

The `IPL_ERROR()` macro simply calls the `iplError()` function by default. This macro is used by other error macros. By changing `IPL_ERROR()` you can modify the error reporting behavior without changing a single line of source code.

The `IPL_RSTERR()` macro resets the error status to `IPL_StsOk`, thus clearing any error condition. This macro should be used by an application when it decides to ignore an error condition.

## Status Codes

The status codes used by the IPL Library are described in Table 3-1. Status codes are integers, not an enumerated type. This allows an application to extend the set of status codes beyond those used by the library itself. Negative codes indicate errors, while non-negative codes indicate success.

**Table 3-1** `iplError()` Status Codes

Status Code	Value	Description
<code>IPL_StsOk</code>	0	No error. The <code>iplError()</code> function does nothing if called with this status code.
<code>IPL_StsBackTrace</code>	-1	Implements a back-trace of the function calls that lead to an error. If <code>IPL_ERRCHK()</code> detects that a function call resulted in an error, it calls <code>IPL_ERROR()</code> with this status code to provide further context information for the user.
<code>IPL_StsError</code>	-2	An error of unknown origin, or of an origin not correctly described by the other error codes.
<code>IPL_StsInternal</code>	-3	An internal “consistency” error, often the result of a corrupted state structure. These errors are typically the result of a failed assertion.

continued 

**Table 3-1** **iplError() Status Codes** (continued)

Status Code	Value	Description
<code>IPL_StsNoMem</code>	-4	A function attempted to allocate memory using <code>malloc()</code> or a related function and was unsuccessful. The message <code>context</code> indicates the intended use of the memory.
<code>IPL_StsBadArg</code>	-5	One of the arguments passed to the function is invalid. The message <code>context</code> indicates which argument and why.
<code>IPL_StsBadFunc</code>	-6	The function is not supported by the implementation, or the particular operation implied by the given arguments is not supported.
<code>IPL_StsNoConv</code>	-7	An iterative convergence algorithm failed to converge within a reasonable number of iterations.

## Application Notes

The global variable `IplLastStatus` records the status of the last error reported. Its value is initially `IPL_StsOk`. The value of `IplLastStatus` is not explicitly set by the library function detecting an error. Instead, it is set by `iplSetErrStatus()`.

If the application decides to ignore an error, it should reset `IplLastStatus` back to `IPL_StsOk` (see `IPL_RSTERR()` under “[Error Macros](#)”). An application-supplied error-handling function must update `IplLastStatus` correctly; otherwise the Image Processing Library might fail. This is because the macro `IPL_ERRCHK()`, which is used internally to the library, refers to the value of this variable.

## Error Handling Example

The following example describes the default error handling for a console application. In the example program, `test.c`, assume that the function `libFuncB()` represents a library function such as `ipl?TranslatedDIB()`, and the function `libFuncD()` represents a function that is called internally to the library. In this scenario, `main()` and `appFuncA()` represent application code.

The value of the error mode is set to `IPL_ErrModeParent`. The `IPL_ErrModeParent` option produces a more detailed account of the error conditions.

### Example 3-1 Error Functions


---

```
/* application main function */
main() {
    iplSetErrMode(IPL_ErrModeParent);
    appFuncA(5, 45, 1.0);
    if (IPL_ERRCHK("main","compute something")) exit(1);
    return 0;
}

/* application subroutine */
void appFuncA(int order1, int order2, double a) {
    libFuncB(a, order1);
    if (IPL_ERRCHK("appFuncA","compute using order1")) return;
    libFuncB(a, order2);
    if (IPL_ERRCHK("appFuncA","compute using order2")) return;
}

/* do some more work */
```

---

continued 

**Example 3-1 Error Functions (continued)**

---

```
/* library function */
void libFuncB(double a, int order) {
    float *vec;
    if (order > 31) {
        IPL_ERROR(IPL_StsBadArg, "libFuncB",
            "order must be less than or equal to 31");
        return;
    }
    if ((vec = libFuncD(a, order)) == NULL) {
        IPL_ERRCHK("libFuncB", "compute using a");
        return;
    }
    /* code to do some real work goes here */
    free(vec);
} // next: library function called internally
double *libFuncD(double a, int order) {
    double *vec;
    if ((vec=(double*)malloc(order*sizeof(double))) == NULL) {
        IPL_ERROR(IPL_StsNoMem, "libFuncD",
            "allocating a vector of doubles");
        return NULL;
    }
    /* do something with vec */
    return vec;
}
```

---

When the program is run, it produces the output illustrated in Example 3-2.

---

**Example 3-2 Output for the Error Function Program (IPL\_ErrModeParent)**

---

```
IPL Library Error: Invalid argument in function libFuncB: order must
be less than or equal to 31
    called from function appFuncA: compute using order2
    called from function main: compute something
```

---

If the program runs with the `IPL_ErrModeLeaf` option instead of `IPL_ErrModeParent`, only the first line of the above output is produced before the program terminated.

If the program in Example 3-1 runs out of heap memory while using the `IPL_ErrModeParent` option, then the output illustrated in Example 3-3 is produced.

---

**Example 3-3 Output for the Error Function Program (IPL\_ErrModeParent)**

---

```
IPL Library Error: Out of memory in function libFuncD: allocating a
vector of doubles
    called from function libFuncB: compute using a
    called from function appFuncA: compute using order1
    called from function main[]: compute something
```

---

Again, if the program is run with the `IPL_ErrModeLeaf` option instead of `IPL_ErrModeParent`, only the first line of the output is produced.



## Adding Your Own Error Handler

The Image Processing Library allows you to define your own error handler. User-defined error handlers are useful if you want your application to send error messages to a destination other than the standard error output stream. For example, you can choose to send error messages to a dialog box if your application is running under a Windows system or you can choose to send error messages to a special log file.

There are two methods of adding your own error handler. In the first method, you can replace the `iplError()` function or the complete error handling library with your own code. Note that this method can only be used at link time.

In the second method, you can use the `iplRedirectError()` function to replace the error handler at run time. The steps below describe how to create your own error handler and how to use the `iplRedirectError()` function to redirect error reporting.

1. Define a function with the function prototype, `IPLErrCallBack`, as defined by the IPL.
2. Your application should then call the `iplRedirectError()` function to redirect error reporting for your own function. All subsequent calls to `iplError()` will call your own error handler.
3. To redirect the error handling back to the default handler, simply call `iplRedirectError()` with a `NULL` pointer.

Example 3-4 illustrates a user-defined error handler function, `ownError()`, which simply prints an error message constructed from its arguments and exits.

**Example 3-4 A Simple Error Handler**

---

```
IPLStatus ownError(IPLStatus status, const char *func,
    const char *context, const char *file, int line);
{
    fprintf(stderr, "IPL Library error: %s, ", iplErrorStr(status));
    fprintf(stderr, "function %s, ", func ? func : "<unknown>");
    if (line > 0) fprintf(stderr, "line %d, ", line);
    if (file != NULL) fprintf(stderr, "file %s, ", file);
    if (context) fprintf(stderr, "context %s\n", context);
    IplSetErrStatus(status);
    exit(1);
}
main () {
    extern IPLErrCallBack ownError;
    /* Redirect errors to your own error handler */
    iplRedirectError( ownError);
    /* Redirect errors back to the default error handler */
    iplRedirectError(NULL);
}
```

---


# Image Creation and Access

This chapter describes the functions that provide the following functionalities:

- Creating and accessing attributes of images (both tiled and non-tiled)
- Allocating memory for data of required type (see also the functions [CreateConvKernel](#) in Chapter 6 and [CreateColorTwist](#) in Chapter 9)
- Manipulating the image
- Working in the Windows\* DIB (device-independent bitmap) environment.

**Table 4-1 Image Creation, Data Exchange and Windows DIB Environment Functions**

Group	Function Name	Description
Creating Images	<a href="#">iplCreateImageHeader</a>	Creates an image header according to the specified attributes.
	<a href="#">iplAllocateImage</a>	Allocates memory for image data.
	<a href="#">iplDeAllocateImage</a>	Deallocates or frees memory for image data pointed to in the image header.
	<a href="#">iplCreateROI</a>	Creates a region of interest (ROI) header with specified attributes.
	<a href="#">iplDeallocate</a>	Deallocates header attributes or image data or ROI or all of the above.
	<a href="#">iplSetROI</a>	Sets a region of interest for an image.
	<a href="#">iplSetBorderMode</a>	Sets the mode for handling the border pixels.
	<a href="#">iplCreateTileInfo</a>	Creates the <a href="#">IplTileInfo</a> structure.
	<a href="#">iplSetTileInfo</a>	Sets the tiling information.
	<a href="#">iplDeleteTileInfo</a>	Deletes the <a href="#">IplTileInfo</a> structure.

continued 

# 4

**Table 4-1 Image Creation, Data Exchange and Windows DIB Environment Functions (continued)**


Group	Function Name	Description
Memory Allocation	<a href="#"><u>iplMalloc</u></a>	Allocates memory aligned to 8 bytes boundary.
	<a href="#"><u>iplwMalloc</u></a>	Allocates memory aligned to 8 bytes boundary for 16-bit words.
	<a href="#"><u>ipliMalloc</u></a>	Allocates memory aligned to 8 bytes boundary 32-bit double words.
	<a href="#"><u>iplsMalloc</u></a>	Allocates memory aligned to 8 bytes boundary for single float elements.
	<a href="#"><u>ipldMalloc</u></a>	Allocates memory aligned to 8 bytes boundary for double float elements.
	<a href="#"><u>iplFree</u></a>	Frees memory allocated by the <a href="#"><u>ipl?Malloc</u></a> functions.
Data Exchange	<a href="#"><u>iplSet</u></a>	Sets a value for the image pixel data.
	<a href="#"><u>iplCopy</u></a>	Copies image data from one IPL image to another.
	<a href="#"><u>iplExchange</u></a>	Exchanges image data between two IPL images.
	<a href="#"><u>iplConvert</u></a>	Converts an IPL image based on the input and output image requirements.
Windows DIB	<a href="#"><u>iplTranslateDIB</u></a>	Translates a DIB image into an IPL image.
	<a href="#"><u>iplConvertFromDIB</u></a>	Converts a DIB image to an IPL image with specified attributes.
	<a href="#"><u>iplConvertToDIB</u></a>	Converts an IPL image to a DIB image with specified attributes.

## Image Header and Attributes

The IPL library functions operate on a single format for images in memory henceforth called the IPL image format. The IPL image format consists of a header of type `IPLImage` containing the information for all the attributes of the image. The header finally contains a pointer to the image data. (See the attributes description in Chapter 2, section “[Data Architecture.](#)”) The values that these attributes can assume are listed in Table 4-2.

**Table 4-2 IPL Image Header Attributes**

Description	Variable or Value	Corresponding DIB Attribute
Size of the IPL image header (for internal use)	<code>nSize</code> in bytes	
IPL Image Header Revision ID (internal use)	ID number	
Number of Channels	1 to <code>N</code> (including alpha channel, if any)	1 (Gray) 3 (RGB) 4 (RGBA)
Alpha channel number	0 (if not present) <code>n</code>	4 (RGBA)
Bits per channel		
Gray only	<code>IPL_DEPTH_1U</code> (1-bit)	Supported
All images: color, gray, and multi-spectral	<code>IPL_DEPTH_8U</code> (8-bit unsigned)	Supported (RGB, RGBA)
(The signed data is used only as output for some image output operations.)	<code>IPL_DEPTH_8S</code> (8-bit signed)	Not supported
	<code>IPL_DEPTH_16U</code> (16-bit unsign.)	Not supported
	<code>IPL_DEPTH_16S</code> (16-bit signed)	Not supported
	<code>IPL_DEPTH_32S</code> (32-bit signed)	Not supported
Color model	4 character string: “Gray”, “RGB,” “RGBA”, “CMYK,” etc.	Not supported. Implicitly, RGB color model.


continued 

**Table 4-2 IPL Image Header Attributes (continued)**

<b>Description</b>	<b>Value</b>	<b>Corresponding DIB Attribute</b>
Channel sequence	4-character string: "Gray", "RGB," "RGBA", "CMYK," etc.	Not supported. Always implicitly BGR for RGB images.
Data Ordering	<code>IPL_DATA_ORDER_PIXEL</code> <code>IPL_DATA_ORDER_PLANE</code>	Supported Not supported
Origin	<code>IPL_ORIGIN_TL</code> (top left corner) <code>IPL_ORIGIN_BL</code> (bottom left corner)	Supported Supported
Scanline alignment	<code>IPL_ALIGN_DWORD</code> <code>IPL_ALIGN_QWORD</code>	Supported Not Supported
Image dimensions		
Image Height	m	m
Image Width	n	n
Region of interest (ROI)	Pointer to structure	Not supported
Image size (bytes)	Integer	
Image data pointer	Pointer to data	
Aligned width	Width (row length) in bytes of aligned image	
Border mode of the top, bottom, left, and right sides of the image.	BorderMode [4]	
Border constant on the top, bottom, left, and right side of the image.	BorderConst [4]	
Original Image	Pointer to original image data	
Image ID for tiling	For application use. Ignored by IPL	
Tiling information	Describes tiles for IPL	

Figure 4-1 presents a graphical depiction of an RGB IPL image with a rectangular ROI and a COI.

**Figure 4-1 RGB Image with a Rectangular ROI and a COI**



OSD05559

Example 4-1 presents a C language definition for the `IPLImage` structure.

#### Example 4-1 IPLImage Definition

---

```

typedef struct _IplImage {
    IPL.H
    int      nSize          /* size of iplImage struct */
    int      ID            /* image header version   */
    int      nChannels;
    int      alphaChannel;
    int      depth;        /* pixel depth in bits    */
    char     colorModel[4];
    char     channelSeq[4];
    int      dataOrder;
    int      origin;
    int      align;        /* 4- or 8-byte align     */
    int      height;
    int      width;
    struct   _IplROI *roi;  /* pointer to ROI if any  */
    void     *imageId;     /* use of the application  */
    struct   _IplTileInfo *tileInfo;
                                /* contains information on tiling */
    int      imageSize;    /* useful size in bytes   */
    char     *imageData;   /* pointer to aligned
                                image                       */
    int      widthStep;    /* size of aligned line in
                                bytes                       */
    int      BorderMode[4]; /*
    int      BorderConst[4]; /*
    char     *imageDataOrigin; /* ptr to full, nonaligned
                                image                       */
} IplImage;

```

---



## Tiling Fields in the `IplImage` Structure

[Image tiling](#) in the IPL was briefly described in Section 2.

The following fields serve for tiling purposes in the `IplImage` structure:

```
struct IplImage {
    ...
    void* imageId;
    IplTileInfo *tileInfo;
    ...
}
```

The `imageId` field can be used by the application, and is ignored by the library. The `tileInfo` field contains information on tiling. It is described in the next section.

The library expects either the `tileInfo` pointer or the `imageData` pointer to be `NULL`. If the former is `NULL`, the image is not tiled; if the latter is `NULL`, the image is tiled. It is an error condition if both or neither of the two are `NULL`.

## `IplTileInfo` Structure

This structure provides information for image tiling:

```
typedef struct _IplTileInfo
{
    IplCallback callback;
    void *id;
    char* tileData
    int width, height;
} IplTileInfo;
```

Here `callback` is the call-back function (see “[Call-backs](#)” in Chapter 2); `id` is an additional identification field; `width` and `height` are the tile sizes for the image.

## Creating Images

The following are the ways to create an IPL image:

- Construct an IPL image header by setting the attributes to appropriate values, then call the function `iplAllocateImage()` to allocate memory for the image or set the image data pointer to image data (in a compatible format) that already exists.
- Call `iplCreateImageHeader()` to create an IPL image header, then call the function `iplAllocateImage()` to allocate memory for the image or set the image data pointer to image data (in a compatible format) that already exists.
- Convert a DIB image to an IPL image using the functions `iplTranslateDIB()` or `iplConvertFromDIB()`. See the section [“Working in the Windows DIB Environment.”](#)

---

## CreateImageHeader

*Creates an IPL image header according to the specified attributes.*

---

```
IplImage* iplCreateImageHeader(int nChannels,  
int alphaChannel, int depth, char* colorModel,  
char* channelSeq, int dataOrder, int origin, int align,  
int height, int width, IplROI* roi);
```

<i>nChannels</i>	Number of channels in the image.
<i>alphaChannel</i>	Alpha channel number (0 if no alpha channel in the image).
<i>depth</i>	Bit depth of pixels. Can be one of <code>IPL_DEPTH_1U</code> , <code>IPL_DEPTH_8U</code> , <code>IPL_DEPTH_8S</code> , <code>IPL_DEPTH_16U</code> , <code>IPL_DEPTH_16S</code> , or <code>IPL_DEPTH_32S</code> . See Table 4-2.

<i>colorModel</i>	A four-character string describing the color model: “RGB”, “GRAY”, “MSI” etc.
<i>channelSeq</i>	The sequence of channels in the image; for example, “BGR” for an RGB image.
<i>dataOrder</i>	<code>IPL_DATA_ORDER_PIXEL</code> or <code>IPL_DATA_ORDER_PLANE</code> .
<i>origin</i>	The origin of the image. Can be <code>IPL_ORIGIN_TL</code> or <code>IPL_ORIGIN_BL</code> .
<i>align</i>	Alignment of image data. Can be <code>IPL_ALIGN_DWORD</code> or <code>IPL_ALIGN_QWORD</code> .
<i>height</i>	Height of the image in pixels.
<i>width</i>	Width of the image in pixels.
<i>roi</i>	Pointer to an ROI (region of interest) structure. This argument can be <code>NULL</code> , which implies that a region of interest comprises all channels and the entire image area.

## Discussion

The function `iplCreateImageHeader()` creates an IPL image header according to the specified attributes. The image data pointer is set to `NULL`; no memory for image data is allocated. To allocate memory for image data, call the function `iplAllocateImage()`. The image size attribute (set by the `iplAllocateImage()` function) in the header is set to zero.

## Return Value

The newly constructed IPL image header.

## AllocateImage

*Allocates memory for image data according to the specified header.*

---

```
void iplAllocateImage(IplImage* image, int fillValue)
/* */
```

*image* An IPL image header with a `NULL` image data pointer. The pointer will be set to newly allocated image data memory after calling this function.

*fillValue* The initial value to use for pixel data. Use a value of `0xFFFFFFFF` (hexadecimal) not to initialize the pixel data.

### Discussion

The function `iplAllocateImage()` is used to allocate image data on the basis of a specified image header. The header must be properly constructed before calling this function. Memory is allocated for the image data according to the attributes specified in the image header (see Example 4-1).

The image data pointer will then point to the allocated memory. It is highly preferable, for efficiency considerations, that the scanline alignment attribute (argument *align*) in the image header be set to `IPL_ALIGN_QWORD`. This will force the image data to be aligned on a quadword (64-bit) memory boundary.

This function sets the image size attribute in the header to the number of bytes allocated for the image.

## DeallocateImage

*Deallocates (frees) memory for image data pointed to in the image header.*

---

```
void iplDeallocateImage(IplImage* image)
```

*image*

An IPL image header with a pointer to the allocated image data memory. The image data pointer will be set to `NULL` after this function executes.

### Discussion

The function `iplDeallocateImage()` is used to free image data memory pointed to by the `imageData` member of the image header. The respective pointer to image data or ROI data is set to `NULL` after the memory is freed up.

---

## Deallocate

*Deallocates or frees memory for image header or data or region of interest or all three.*

---

```
void iplDeallocate (IplImage* image, int flag)
```

*image*

An IPL image header with a pointer to allocated image data memory. The image data pointer will be set to `NULL` after this function executes.

*flag*

Flag indicating what memory area to free:

`IPL_IMAGE_HEADER` Free header structure.

<code>IPL_IMAGE_IMAGE</code>	Free image data, set pointer to <code>NULL</code> .
<code>IPL_IMAGE_ROI</code>	Free image ROI, set pointer to <code>NULL</code> .
<code>IPL_IMAGE_ALL</code>	Free header, image data, and ROI.

## Discussion

The function `iplDeallocate()` is used to free or destroy memory allocated for header structure, image data, ROI data, or all three. The respective pointer is set to `NULL` after the memory is freed up.

## Setting Regions of Interest

To set a region of interest, the function `iplSetROI()` uses a ROI structure `IplROI` presented in Example 4-2. The `IplROI` member of the IPL image header must point to this `IplROI` structure to be effective. This can be done by a simple assignment. The application may choose to construct the ROI structure explicitly without the use of the function.

### Example 4-2 IplROI Definition

---

```
typedef struct _IplROI {
    unsigned int coi;
                //Channel to effect in original image
    int xOffset;
    int yOffset;
    int height;
    int width;
} IplROI;
```

---

The members in the above `IplROI` structure define:

<code>coi</code>	The channel of interest number. This parameter indicates which channel in the original image will be affected by processing taking place in the region of interest; <code>coi</code> equal to 0 indicates that all channels will be affected.
<code>xOffset</code> and <code>yOffset</code>	The offset from the origin of the rectangular ROI. (See section “ <a href="#">Image Regions</a> ” in Chapter 2 for the description of image regions.)
<code>height</code> and <code>width</code>	The size of the rectangular ROI.

---

## CreateROI

*Allocates and sets the region of interest (ROI) structure.*

---

```
IplROI* iplCreateROI(int coi, int xOffset, int yOffset,
int height, int width);
```

<code>coi</code>	The channel of interest. It can be set to 0 (for all channels) or to a specific channel number.
<code>xOffset</code> , <code>yOffset</code>	The offsets from the origin of the rectangular region.
<code>height</code> , <code>width</code>	The size of the rectangular region.

## Discussion

The function `iplCreateROI()` allocates a new ROI structure with the specified attributes and returns a pointer to this structure. If the IPL image pointer is `NULL`, then only a rectangular ROI is defined.

## Return Value

A pointer to the newly constructed ROI structure.

---

## SetROI

*Sets the region of interest (ROI) structure.*

```
void iplSetROI(IplROI* roi, int coi, int xOffset, int yOffset, int height, int width);
```

<i>roi</i>	The pointer to the ROI structure to modify in the original image.
<i>coi</i>	The channel of interest in the original image. It can be set to 0 (for all channels) or to a specific channel number.
<i>xOffset, yOffset</i>	The offset from the origin of the rectangular region.
<i>height, width</i>	The size of the rectangular region.

## Discussion

The function `iplSetROI()` sets the channel of interest and the rectangular region of interest in the structure *roi*.

The argument *coi* defines the number of the channel of interest. The arguments *xOffset* and *yOffset* define the offset from the origin of the rectangular ROI. The members *height* and *width* define the size of the rectangular ROI.



## Image Borders and Image Tiling

Many neighborhood operators need intensity values for pixels that lie outside the image, that is, outside the borders of the image. For example, a 3 by 3 filter, when operating on the first row of an image, needs to assume pixel values of the preceding (non-existent) row. A larger filter will require more rows from the border. These border issues therefore exist at the top and bottom, left and right sides, and the four corners of the image. The library provides a function `iplSetBorderMode` that the application can use to set the border mode within the image. This function specifies the behavior for handling border pixels.

For tiled images, the border mode is handled in the same way as for non-tiled images, except that in the outer tiles there might be extra data which is ignored.

---

### SetBorderMode

*Sets the mode for handling the border pixels.*

```
void iplSetBorderMode(IplImage *src, int mode,
                    int border, int constVal)
```

`src` The image where the border mode is to be set.

`mode` The following modes are supported:

- |                                   |  |
|-----------------------------------|--|
| <code>IPL_BORDER_CONSTANT</code>  | The value <code>constVal</code> is used for all pixels.                                      |
| <code>IPL_BORDER_REPLICATE</code> | The last row or column is replicated for the border.   |
| <code>IPL_BORDER_REFLECT</code>   | The last <code>n</code> rows or columns are reflected in reverse order to create the border. |

<code>IPL_BORDER_WRAP</code>	The required border rows or columns are taken from the opposite side of the image.								
<code>border</code>	The side that this function is called for. Can be an OR of one or more of the following four sides of an image: <table> <tr> <td><code>IPL_SIDE_TOP</code></td> <td>Top side.</td> </tr> <tr> <td><code>IPL_SIDE_BOTTOM</code></td> <td>Bottom side.</td> </tr> <tr> <td><code>IPL_SIDE_LEFT</code></td> <td>Left side.</td> </tr> <tr> <td><code>IPL_SIDE_RIGHT</code></td> <td>Right side.</td> </tr> </table> <p>If no mode has been set for a side, the default <code>IPL_BORDER_CONSTANT</code> is assumed with a value of 0 for <code>constVal</code>. The top side is also used to define all border pixels in the top left and right corners. Similarly, the bottom side is used to define the border pixels in the bottom left and right corners.</p>	<code>IPL_SIDE_TOP</code>	Top side.	<code>IPL_SIDE_BOTTOM</code>	Bottom side.	<code>IPL_SIDE_LEFT</code>	Left side.	<code>IPL_SIDE_RIGHT</code>	Right side.
<code>IPL_SIDE_TOP</code>	Top side.								
<code>IPL_SIDE_BOTTOM</code>	Bottom side.								
<code>IPL_SIDE_LEFT</code>	Left side.								
<code>IPL_SIDE_RIGHT</code>	Right side.								
<code>constVal</code>	The value to use for the border when the mode is set to <code>IPL_BORDER_CONSTANT</code> .								

## Discussion

The function `iplSetBorderMode()` is used to set the border handling mode of one or more of the four sides of an image. If the mode is not set for any side, then a constant value of 0 is used for all border pixels on that side. Intensity values for the border pixels are assumed or created based on the mode.

## CreateTileInfo

Creates the `IplTileInfo` structure.

---

```
IplTileInfo* iplCreateTileInfo(IplCallback callback,  
void* id, int width, int height);
```

<code>callback</code>	The call-back function.
<code>id</code>	The image ID (for application use).
<code>width, height</code>	The tile sizes.

### Discussion

The function `iplCreateTileInfo()` allocates a new `IplTileInfo` structure with the specified attributes and returns a pointer to this structure.

### Return Value

The pointer to the created `IplTileInfo` structure.

## SetTileInfo

*Sets the `IplTileInfo` structure fields.*

---

```
void iplSetTileInfo(IplTileInfo* tileInfo, IplCallback  
callback, int width, int height);
```

<code>tileInfo</code>	The pointer to the <code>IplTileInfo</code> structure.
<code>callback</code>	The call-back function.
<code>id</code>	The image ID (for application use).
<code>width, height</code>	The tile sizes.

### Discussion

This function sets attributes for an existing [IplTileInfo](#) structure.

---

## DeleteTileInfo

*Deletes the `IplTileInfo` structure.*

---

```
void iplDeleteTileInfo(IplTileInfo* tileInfo);
```

<code>tileInfo</code>	The pointer to the <code>IplTileInfo</code> structure.
-----------------------	--

### Discussion

This function deletes the [IplTileInfo](#) structure previously created by the [CreateTileInfo](#) function.

## Memory Allocation Functions

Functions of the `ipl?Malloc()` group allocate aligned memory blocks for IPL image data. The size of allocated memory is specified by the `size` parameter. The “?” in `ipl?Malloc()` stands for `w`, `i`, `s`, or `d`; these letters indicate the data type in the function names as follows:

`iplMalloc()`    byte  
`iplwMalloc()`    16-bit word  
`ipliMalloc()`    32-bit double word  
`iplsMalloc()`    4-byte single floating-point element  
`ipldMalloc()`    8-byte double floating-point element



---

**NOTE.** *The only function to free the memory allocated by any of these functions is `iplFree()`.*

---

---

## Malloc

*Allocates memory aligned to 8 bytes boundary.*

```
void* iplMalloc(int size);
```

`size`                      Size (in bytes) of memory block to allocate.

### Discussion

The `iplMalloc()` function allocates memory block aligned to 8 bytes boundary.

## Return Value

The return value of `iplMalloc()` is a pointer to aligned memory block. To free this block, only the function `iplFree()` must be used. If no memory is available in the system, then the `NULL` value is returned.

---

## wMalloc

*Allocates memory aligned to 8 bytes boundary for 16-bit words.*

```
short* iplwMalloc(int size);
```

*size*                      Size in words (16 bits) of memory block to allocate.

## Discussion

The `iplwMalloc()` function allocates memory block aligned to 8 bytes boundary for 16-bit words.

## Return Value

The return value of `iplwMalloc()` is a pointer to aligned memory block. To free this block only the function `iplFree()` must be used. If no memory is available in the system, then the `NULL` value is returned.

## iMalloc

*Allocates memory aligned to 8 bytes boundary for 32-bit double words.*

---

```
int* ipliMalloc(int size);
```

**size** Size in double words (32 bits) of memory block to allocate.

### Discussion

The `ipliMalloc()` function allocates memory block aligned to 8 bytes boundary for 32-bit double words.

### Return Value

The return value of `ipliMalloc()` is a pointer to aligned memory block. To free this block only the function `iplFree()` must be used. If no memory available in the system, then the `NULL` value is returned.

---

## sMalloc

*Allocates memory aligned to 8 bytes boundary for floating-point elements.*

---

```
float * iplsMalloc(int size);
```

**size** Size in float elements (4 bytes) of memory block to allocate.

## Discussion

The `iplsMalloc()` function allocates memory block aligned to 8 bytes boundary for floating-point elements.

## Return Value

The return value of `iplsMalloc()` is a pointer to aligned memory block. To free this block only the function `iplFree()` must be used. If no memory is available in the system, then the `NULL` value is returned.

---

## dMalloc

*Allocates memory aligned to 8 bytes boundary for double floating-point elements.*

---

```
double* ipldMalloc(int size);
```

*size*                      Size in double elements (8 bytes) of memory block to allocate.

## Discussion

The `ipldMalloc()` function allocates memory block aligned to 8 bytes boundary for double floating-point elements.

## Return Value

The return value of `ipldMalloc()` is a pointer to aligned memory block. To free this block only the function `iplFree()` must be used. If no memory is available in the system, then the `NULL` value is returned.



## iplFree

*Frees memory allocated by one of the `ipl?Malloc` functions.*

---

```
void iplMalloc(void * ptr);
```

`ptr`      Pointer to memory block to free.

### Discussion

The `iplFree()` function frees the aligned memory block allocated by one of the functions `iplMalloc()`, `iplwMalloc()`, `ipliMalloc()`, `iplsMalloc()`, or `ipldMalloc()`.



---

**NOTE.** *The function `iplFree()` can't be used to free memory allocated by standard functions like `malloc()` or `calloc()`.*

---

## Image Data Exchange

The functions described in this section provide image manipulation capabilities, such as setting the image pixel data, copying data from one image to another, exchanging the data between the images, and converting one IPL image to another according to the attributes defined in the source and resultant image headers.

---

### Set

*Sets a value for an IPL image pixel data.*

---

```
void iplSet(IplImage* image, int fillValue,  
IplCoord* map);
```

<i>image</i>	An IPL image header with allocated image data.
<i>fillValue</i>	The value to set the pixel data.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplSet()` sets an IPL image pixel data. Before calling this function, the IPL image header must be properly constructed and image data must be allocated. For images with the bit depth lower than the *fillValue*, the *fillValue* is truncated when assigned to pixel. If an ROI is specified, only that ROI is filled.

## Copy

*Copies image data from one IPL image to another.*

---

```
void iplCopy(IplImage* srcImage, IplImage* dstImage,  
            IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplCopy()` copies image data from a source image to a resultant image. Before calling this function, the source and resultant headers must be properly constructed and image data for both images must be allocated. The following constraints apply to the copying:

- The bit depth per channel of the source image should be equal to that of the resultant image.
- The number of channels of interest in the source image should be equal to the number of channels of interest in the resultant image; that is, either the source `coi` = the resultant `coi` = 0 or both `cois` are nonzero.
- The data ordering (by pixel or by plane) of the source image should be the same as that of the resultant image.

The `origin`, `align`, `height`, and `width` field values (see Table 4-2) may differ in source and resultant images. Copying applies to the areas that intersect between the source ROI and the destination ROI.

## Exchange

*Exchanges image data  
between two IPL images.*

---

```
void iplExchange(IplImage* ImageA, IplImage* ImageB,  
                IplCoord* map);
```

<i>ImageA</i>	The first image.
<i>ImageB</i>	The second image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplExchange()` exchanges image data between two images, the first and the second. The image headers must be properly constructed before calling this function, and image data for both images must be allocated. The following constraints apply to the data exchanging:

- The bit depths per channel of both images should be equal.
- The numbers of channels of interest in both images should be equal.
- The data ordering of both images should be the same (either pixel- or plane-oriented) .

The *origin*, *align*, *height*, and *width* field values (see Table 4-2) may differ in the first and the second images. The data are exchanged at the areas of intersection between the ROI of the first image and the ROI of the second image.

## Convert

*Converts source IPL image data to resultant IPL image according to the source and resultant image headers.*

---

```
void iplConvert(IplImage* srcImage, IplImage* dstImage,  
int convertMode, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>convertMode</i>	A flag indicating how to perform the image conversion by reducing the bit depth.  The flag values are:  <b>IPL_BITS_HIGH</b> Reduce by scaling  <b>IPL_BITS_LOW</b> Reduce by using the lower part of the pixel values
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplConvert()` converts image data from the source IPL image to the resultant IPL image according to the attributes defined in the source and resultant image headers. The images that can be converted may have the following different characteristics:

- Bit depth per channel
- Data ordering
- Origins

For the above data description, see [Table 4-2](#).

The following constraints apply to the conversion:

- If the source image has a bit depth per channel equal to 1, the resultant image should also have the bit depth equal to 1.
- The number of channels in the source image should be equal to the number of channels in the resultant image.
- The height and width of the source image should be equal to that of the resultant image.

All ROIs are ignored; `IplCoord` must be `NULL`.

## Working in the Windows DIB Environment

The IPL library provides functions to convert a DIB (device-independent bitmap) image to an IPL image and vice versa. [Table 4-2](#) shows that the DIB image format is a subset of the IPL image format. Not included in this subset are the DIB palette images and DIB 8-bit- and 16-bit-per-pixel absolute color images because they have no equivalent IPL images.

The DIB palette images must be first converted to IPL absolute color images. DIB 8-bit- and 16-bit-per-pixel images have to be unpacked into IPL 8-bit, 16-bit- or 32-bit-per-channel images.

However, any DIB 24-bit absolute color image can be directly converted to an IPL image. You just need to create an IPL image header corresponding to the DIB attributes. The DIB image data can be pointed to by the IPL header or it can be duplicated.

The IPL functions can perform conversion from a DIB image to an IPL image and vice versa with additional useful capabilities:

`iplTranslateDIB()` Performs a simple translation of a DIB image to an IPL image as described above. Also converts a DIB palette image to an IPL absolute color image.

While this is the most efficient way of converting a DIB image, it is not the most efficient format for the IPL functions to manipulate because the DIB image data is doubleword-aligned, and not quadword-aligned.

`iplConvertFromDIB()` Provides more control of the conversion and can convert a DIB image to an IPL image with a prepared IPL image header. The IPL image header must be then set to the desired attributes. The bit depth of the channels in the IPL image header must be equal to or greater than that in the DIB header.

`iplConvertToDIB()` Performs conversion in the opposite direction: an IPL image to a DIB image. This function performs dithering if the bit depth of the DIB is less than that of the IPL image. It can also be used to create a DIB palette image from an IPL absolute color image. The function can optionally create a new palette.

## TranslateDIB

*Translates a DIB image into the corresponding IPL image.*

---

```
iplImage* iplTranslatedDIB(BITMAPINFOHEADER* dib,  
    BOOL cloneData)
```

*dib*                      The DIB image.

*cloneData*              An output flag (Boolean): if false, indicates that the image data pointer in the IPL image will point to the DIB image data; if true, indicates that the data was copied.

### Discussion

The function `iplTranslatedDIB()` translates a DIB image into an IPL image. The IPL image attributes corresponding to the DIB image are automatically chosen (see [Table 4-2](#)), so no explicit control of the conversion is provided. A DIB palette image will be converted to an IPL absolute color image with a bit depth of 8 bits per channel, and the image data will be copied, returning `cloneData = true`.

A 24-bit-per-pixel DIB RGB image will be converted to an 8-bit-per-channel RGB IPL image.

A 32-bit-per-pixel DIB RGBA image will be converted to an 8-bit-per-channel RGBA IPL image with an alpha channel.

An 8-bit-per-pixel or 16-bit-per-pixel DIB absolute color RGB image will be converted (by unpacking) into an 8-bit-per-channel RGB IPL image. The image data will be copied, returning `cloneData = true`.

A 1-bit-per-pixel or 8-bit-per-pixel DIB gray scale image with a [standard gray palette](#) will be converted to a 1-bit-per-channel or 8-bit-per-channel IPL gray-scale image, respectively.



A 4-bit-per-pixel DIB gray-scale image with a standard gray palette will be converted into an 8-bit-per-pixel IPL gray-scale image and the image data will be copied, returning `cloneData = true`.

Note that in the cases above where the image data is not copied, it will result in inefficient access of the image by the IPL image processing functions. This is because DIB image data is aligned on doubleword (32-bit) boundaries. Alternatively, when `cloneData` is true, the DIB image data is replicated into newly allocated image data memory and automatically aligned to quadword boundaries which results in a better memory access.

### **Return Value**

A constructed IPL image.

## ConvertFromDIB

*Converts a DIB image to an IPL image with specified attributes.*

---

```
void iplConvertFromDIB(BITMAPINFOHEADER* dib,  
                      IplImage* image)
```

*dib* The input DIB image.

*image* The IPL image header with specified attributes. If the data pointer is `NULL`, image data memory will be allocated and the pointer set to it.

### Discussion

The function `iplConvertFromDIB()` converts a DIB image into an IPL image according to the attributes set in the IPL image header. Explicit control of the conversion is therefore provided. The following constraints apply to the conversion:

- The bit depth per channel of the IPL image should be greater than or equal to that of the DIB image.
- The number of channels (not including the alpha channel) in the IPL image should be greater than or equal to the number of channels in the DIB image (not including the alpha channel if present).
- The dimensions of the IPL image should be greater than or equal to that of the DIB image. When the IPL image is larger than the DIB image, the origins of the IPL and DIB images are made coincident for the purposes of copying.
- When converting a DIB RGBA image, the IPL image should also contain an alpha channel.

---

## ConvertToDIB

*Converts an IPL image to a DIB image with specified attributes.*

---

```
void iplConvertToDIB(iplImage* image, BITMAPINFOHEADER*
                    dib, int dither, int paletteConversion)
```

*image*                    The input IPL image.

*dib*                        The output DIB image.

*dither*                    The dithering algorithm to use if applicable. Dithering will be done if the bit depth in the DIB is less than that of the IPL image. The following algorithms are supported corresponding to these *dither* identifiers:

`IPL_DITHER_STUCKEY`    The Stucki dithering algorithm is used.

`IPL_DITHER_NONE`        No dithering is done. The most significant bits in the IPL image pixel data are retained.

*paletteConversion*    Applicable when the DIB is a palette image. Specifies the palette algorithm to use when converting the IPL absolute color image. The following options are supported:

`IPL_PALCONV_NONE`    The existing palette in the DIB is used.

`IPL_PALCONV_POPULATE`    The popularity palette conversion algorithm is used.

`IPL_PALCONV_MEDCUT`    The median cut algorithm for palette conversion is used.

## Discussion

The function `iplConvertToDIB()` converts an IPL image to a DIB image. The conversion takes place according to the IPL image and DIB image attributes. While IPL images are always in absolute color, DIB images can be in absolute or palette color. When the DIB is a palette image, the absolute color IPL image is converted to a palette image according to the palette conversion option specified. When the bit depth of an absolute color DIB image is less than that of the IPL image, then dithering according to the specified option is performed.

The following constraints and considerations apply when using this function:

- The number of channels (not including the alpha and ROI channels) in the IPL image should be equal to the number of channels in the DIB image.
- The alpha channel in an IPL image will be passed on only when the DIB is an RGBA image.

# Image Arithmetic and Logical Operations

# 5

This chapter describes image processing functions that modify pixel values using simple arithmetic or logical operations. It also includes the library functions that perform image compositing based on opacity (alpha-blending). All these operations can be broken into two categories: monadic operations, which use single input images, and dyadic operations, which use two input images. Table 5-1 lists the functions that perform arithmetic and logical operations.

**Table 5-1** Image Arithmetic and Logical Operations

Group	Function Name	Description
Arithmetic operations	<a href="#"><code>iplAddS</code></a>	Adds a constant to the image pixel values.
	<a href="#"><code>iplSubtractS</code></a>	Subtracts a constant from the pixel values or the values from a constant.
	<a href="#"><code>iplMultiplyS</code></a>	Multiplies pixel values by a constant.
	<a href="#"><code>iplMultiplySScale</code></a>	Multiplies pixel values by a constant and scales the product.
	<a href="#"><code>iplSquare</code></a>	Squares the pixel values of an image.
	<a href="#"><code>iplAdd</code></a>	Adds pixel values of two images.
	<a href="#"><code>iplSubtract</code></a>	Subtracts pixel values of one image from those of another image.
	<a href="#"><code>iplMultiply</code></a>	Multiplies pixel values of two images.
	<a href="#"><code>iplMultiplyScale</code></a>	Multiplies pixel values of two images and scales the product.


Continued 

Table 5-1 Image Arithmetic and Logical Operations (continued)

Group	Function Name	Description
Logical operations	<a href="#"><code>iplAndS</code></a>	Performs a bitwise AND operation on each pixel with a constant.
	<a href="#"><code>iplOrS</code></a>	Performs a bitwise OR operation on each pixel with a constant.
	<a href="#"><code>iplXorS</code></a>	Performs a bitwise XOR operation on each pixel with a constant.
	<a href="#"><code>iplNot</code></a>	Performs a bitwise NOT operation on each pixel
	<a href="#"><code>iplLShiftS</code></a>	Multiplies pixel values by a constant power of 2 by shifting bits to the left.
	<a href="#"><code>iplRShiftS</code></a>	Divides pixel values by a constant power of 2 by shifting bits to the right.
	<a href="#"><code>iplAnd</code></a>	Combines corresponding pixels of two images by a bitwise AND operation.
	<a href="#"><code>iplOr</code></a>	Combines corresponding pixels of two images by a bitwise OR operation.
Alpha-blending	<a href="#"><code>iplXor</code></a>	Combines corresponding pixels of two images by a bitwise XOR operation.
	<a href="#"><code>iplPreMultiplyAlpha</code></a>	Pre-multiplies pixel values of an image by alpha values.
	<a href="#"><code>iplAlphaComposite</code></a>	Composites two images using alpha (opacity) values.
	<a href="#"><code>iplAlphaCompositeC</code></a>	Composites two images using constant alpha (opacity) values.

The functions `iplSquare()`, `iplNot()`, and `iplPreMultiplyAlpha()` as well as all functions with names containing an additional `S` use single input images (perform monadic operations). All other functions in the above table use two input images (perform dyadic operations).

## Monadic Arithmetic Operations

The sections that follow describe the IPL functions that perform monadic arithmetic operations (note that the `iplPreMultiplyAlpha` function is described in the “[Image Compositing Based on Opacity](#)” section of this chapter). All these functions use a single input image to create an output image.

---

### AddS

*Adds a constant to pixel values of the source image.*

---

```
void iplAddS(IplImage* srcImage, IplImage* dstImage, int value, IplCoord* map);
```

<code>srcImage</code>	The source image.
<code>dstImage</code>	The resultant image.
<code>value</code>	The value to be added to the pixel values.
<code>map</code>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplAddS()` changes the image intensity by adding the `value` to pixel values. A positive `value` brightens the image (increases the intensity); a negative `value` darkens the image (decreases the intensity).

## SubtractS

*Subtracts a constant from pixel values, or pixel values from a constant.*

---

```
void iplSubtractS(IplImage* srcImage, IplImage* dstImage,  
int value, BOOL flip, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>value</i>	The value to be subtracted from the pixel values.
<i>flip</i>	A Boolean used to change the order of subtraction.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplSubtractS()` changes the intensity of an image as follows:

If *flip* is false, the *value* is subtracted from the image pixel values.  
If *flip* is true, the image pixel values are subtracted from the *value*.

---

## MultiplyS

*Multiplies pixel values by a constant.*

---

```
void iplMultiplyS(IplImage* srcImage, IplImage* dstImage,  
unsigned int value, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.



*value* A positive value by which to multiply the pixel values.  
*map* The structure specifying offsets for tiling purposes.

### Discussion

The function `iplMultiplyS()` increases the intensity of an image by multiplying each pixel by a positive constant *value*.

---

## MultiplySScale

*Multiplies pixel values  
by a constant and scales  
the products.*

---

```
void iplMultiplySScale(IplImage* srcImage, IplImage*  
dstImage, int value, IplCoord* map);
```

*srcImage* The source image.  
*dstImage* The resultant image.  
*value* A positive value by which to multiply the pixel values.  
*map* The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplMultiplySScale()` multiplies the input image pixel values by *value* and scales the products using the following formula:

$$dst\_pixel = src\_pixel * value / max\_val$$

where *src\_pixel* is a pixel value of the source images, *dst\_pixel* is the resultant pixel value, and *max\_val* is the maximum presentable pixel value. The source and resultant images must have the same pixel depth.

---

## Square

*Squares the pixel values of the image.*

---

```
void iplSquare(IplImage* srcImage, IplImage* dstImage,  
IplCoord* map);
```

*srcImage*      The source image.

*dstImage*      The resultant image.

*map*            The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplSquare()` increases the intensity of an image by squaring each pixel value.

## Dyadic Arithmetic Operations

The sections that follow describe the IPL functions that perform dyadic arithmetic operations. These functions use two input images to create an output image.

---

### Add

*Combines corresponding pixels of two images by addition.*

---

```
void iplAdd(IplImage* srcImageA, IplImage* srcImageB,  
IplImage* dstImage, IplCoord* map);
```

*srcImageA*      The first source image.

*srcImageB*      The second source image.

*dstImage*        The resultant image obtained as  
*dstImage = srcImageA + srcImageB.*

*map*             The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplAdd()` adds corresponding pixels of two input images to produce an output image.

## Subtract

*Combines corresponding pixels of two images by subtraction.*

---

```
void iplSubtract(IplImage* srcImageA, IplImage*  
srcImageB, IplImage* dstImage, BOOL flip, IplCoord* map);
```

*srcImageA*      The first source image.

*srcImageB*      The second source image.

*dstImage*      If *flip* (see below) is false, the resultant image is  
 $dstImage = srcImageA - srcImageB$ ,  
otherwise it is  
 $dstImage = srcImageB - srcImageA$ .

*flip*            A Boolean flag to indicate the order in which the input  
images are subtracted. See *dstImage* above.

*map*            The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplSubtract()` subtracts corresponding pixels of two input images to produce an output image.

## Multiply

*Combines corresponding pixels of two images by multiplication.*

---

```
void iplMultiply(IplImage* srcImageA, IplImage*  
srcImageB, IplImage* dstImage, IplCoord* map);
```

*srcImageA*      The first source image.

*srcImageB*      The second source image.

*dstImage*        The resultant image.

*map*              The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplMultiply()` multiplies corresponding pixels of two input images to produce an output image.

## MultiplyScale

*Multiplies pixel values of two images and scales the products.*

---

```
void iplMultiplyScale(IplImage* srcImageA, IplImage*  
srcImageB, IplImage* dstImage, IplCoord* map);
```

*srcImageA*      The first source image.  
*srcImageB*      The second source image.  
*dstImage*        The resultant image.  
*map*              The structure specifying offsets for tiling purposes;  
see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplMultiplyScale()` multiplies corresponding pixels of two input images and scales the products using the following formula:

$$dst\_pixel = srcA\_pixel * srcB\_pixel / max\_val$$

where *srcA\_pixel* and *srcB\_pixel* are pixel values of the source images, *dst\_pixel* is the resultant pixel value, and *max\_val* is the maximum presentable pixel value. Both source images and the resultant image must have the same pixel depth.

## Monadic Logical Operations

The sections that follow describe the IPL functions that perform monadic logical operations. All these functions use a single input image to create an output image.

---

### LShiftS

*Shifts pixel values' bits to the left.*

---

```
void iplLShiftS(IplImage* srcImage, IplImage* dstImage,  
unsigned int nShift, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nShift</i>	The number of bits by which to shift each pixel value to the left.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplLShiftS()` changes the intensity of the source image by shifting the bits in each pixel value by *nShift* bits to the left. The positions vacated after shifting the bits are filled with zeros.

## RShiftS

*Divides pixel values by a constant power of 2 by shifting bits to the right.*

---

```
void iplRShiftS(IplImage* srcImage, IplImage* dstImage,  
unsigned int nShift, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nShift</i>	The number of bits by which to shift each pixel value to the right.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplRShiftS()` decreases the intensity of the source image by shifting the bits in each pixel value by *nShift* bits. The positions vacated after shifting the bits are filled with zeros.

---

## Not

*Performs a bitwise NOT operation on each pixel.*

---

```
void iplNot(IplImage* srcImage, IplImage* dstImage,  
IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.



*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplNot()` performs a bitwise NOT operation on each pixel value.

---

## AndS

*Performs a bitwise AND operation of each pixel with a constant.*

---

```
void iplAndS(IplImage* srcImage, IplImage* dstImage,  
unsigned int value, IplCoord* map);
```

*srcImage* The source image.

*dstImage* The resultant image.

*value* The bit sequence used to perform the bitwise AND operation on each pixel.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplAndS()` performs a bitwise AND operation between each pixel value and `value`. The least significant bit(s) of the `value` are used.

## OrS

*Performs a bitwise OR operation of each pixel with a constant.*

---

```
void iplOrS(IplImage* srcImage, IplImage* dstImage,  
            unsigned int value, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>value</i>	The bit sequence used to perform the bitwise OR operation on each pixel.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplOrS()` performs a bitwise OR between each pixel value and *value*. The least significant bit(s) of the *value* are used.

## XorS

*Performs a bitwise XOR operation of each pixel with a constant.*

---

```
void iplXorS(IplImage* srcImage, IplImage* dstImage,  
            unsigned int value, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>value</i>	The bit sequence used to perform the bitwise XOR operation on each pixel.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplXorS()` performs a bitwise XOR between each pixel value and *value*. The least significant bit(s) of the *value* are used.

## Dyadic Logical Operations

This section describes the IPL functions that perform dyadic logical operations. These functions use two input images to create an output image.

## And

*Combines corresponding pixels of two images by a bitwise AND operation.*

---

```
void iplAnd(IplImage* srcImageA, IplImage* srcImageB,  
IplImage* dstImage, IplCoord* map);
```

<i>srcImageA</i>	The first source image.
<i>srcImageB</i>	The second source image.
<i>dstImage</i>	The image resulting from the bitwise operation between input images <i>srcImageA</i> and <i>srcImageB</i> .
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplAnd()` performs a bitwise AND operation between the values of corresponding pixels of two input images.

---

## Or

*Combines corresponding pixels of two images by a bitwise OR operation.*

---

```
void iplOr(IplImage* srcImageA, IplImage* srcImageB,  
IplImage* dstImage, IplCoord* map);
```

<i>srcImageA</i>	The first source image.
<i>srcImageB</i>	The second source image.

- dstImage* The image resulting from the bitwise operation between input images *srcImageA* and *srcImageB*.
- map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplOR()` performs a bitwise OR operation between the values of corresponding pixels of two input images.

---

## Xor

*Combines corresponding pixels of two images by a bitwise XOR operation.*

---

```
void iplXor(IplImage* srcImageA, IplImage* srcImageB,  
IplImage* dstImage, IplCoord* map);
```

- srcImageA* The first source image.
- srcImageB* The second source image.
- dstImage* The image resulting from the bitwise operation between input images *srcImageA* and *srcImageB*.
- map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplXor()` performs a bitwise XOR operation between the values of corresponding pixels of two input images.

## Image Compositing Based on Opacity

The IPL library provides functions to composite two images using either the opacity (alpha) channel in the images or a provided alpha value. Alpha values range from 0 (100% translucent, 0% coverage) to full range (0% translucent, 100% coverage). Coverage is the percentage of the pixel's own intensity that is visible.

Using the opacity channel for image compositing provides the capability of overlaying the arbitrarily shaped and transparent images in arbitrary positions. It also reduces aliasing effects along the edges of the combined regions by allowing some of the bottom image's color to show through.

Let us consider the example of RGBA images. Here each pixel is a quadruple  $(r, g, b, \alpha)$  where  $r, g, b,$  and  $\alpha$  are the red, green, blue and alpha channels, respectively. In the formulas that follow, the Greek letter  $\alpha$  with subscripts always denotes the normalized (scaled) alpha value in the range 0 to 1, no matter what the full range of the alpha channel value.

There are many ways of combining images using alpha values. In all compositing operations a resultant pixel  $(r_C, g_C, b_C, \alpha_C)$  in image C is created by overlaying a pixel  $(r_A, g_A, b_A, \alpha_A)$  from the foreground image A over a pixel  $(r_B, g_B, b_B, \alpha_B)$  from the background image B. The resulting pixel values for an OVER operation (A OVER B) are computed as shown below.

$$r_C = \alpha_A * r_A + (1 - \alpha_A) * \alpha_B * r_B$$

$$g_C = \alpha_A * g_A + (1 - \alpha_A) * \alpha_B * g_B$$

$$b_C = \alpha_A * b_A + (1 - \alpha_A) * \alpha_B * b_B$$

The above three expressions can be condensed into one as follows:

$$C = \alpha_A * A + (1 - \alpha_A) * \alpha_B * B$$

In this example, the color of the background image B influences the color of the resultant image through the second term  $(1 - \alpha_A) * \alpha_B * B$ . The resulting alpha value is computed as

$$\alpha_C = \alpha_A + (1 - \alpha_A) * \alpha_B$$

### Using Pre-multiplied Alpha Values

In many cases it is computationally more efficient to store the color channels pre-multiplied by the alpha values. In the RGBA example, the pixel  $(r, g, b, \alpha)$  would actually be stored as  $(r*\alpha, g*\alpha, b*\alpha, \alpha)$ . This storage format reduces the number of multiplications required in the compositing operations. In interactive environments, when an image is composited many times, this capability is especially efficient.

One known disadvantage of the pre-multiplication is that once a pixel is marked as transparent, its color value is gone because the pixel's color channels are multiplied by 0.

The function `iplPreMultiplyAlpha()` implements various alpha compositing operations between two images. One of them is converting the pixel values to pre-multiplied form.

The color channels in images with the alpha channel can be optionally pre-multiplied with the alpha value. This saves a significant amount of computation for some of the alpha compositing operations. For example, in an RGBA color model image, if  $(r, g, b, \alpha)$  are the channel values for a pixel, then upon pre-multiplication they are stored as  $(r*\alpha, g*\alpha, b*\alpha, \alpha)$ .

---

## AlphaComposite

### AlphaCompositeC

*Composite two images using alpha (opacity) values.*

---

```
void iplAlphaComposite(IplImage* srcImageA, IplImage*
srcImageB, IplImage* dstImage, int compositeType,
IplImage* alphaImageA, IplImage* alphaImageB, IplImage*
alphaImageDst, BOOL premulAlpha, BOOL divideMode,
IplCoord* map);
```

# 5

```
void iplAlphaCompositeC(IplImage* srcImageA, IplImage*  
srcImageB, IplImage* dstImage, int compositeType, int aA,  
int aB, BOOL premulAlpha, BOOL divideMode, IplCoord* map);
```

- srcImageA* The foreground input image.
- srcImageB* The background input image.
- dstImage* The resultant output image.
- compositeType* The composition type to perform. See [Table 5-2](#) for the type value and description.
- aA* The constant alpha value to use for the source image *srcImageA*. Should be a positive number.
- aB* The constant alpha value to use for the source image *srcImageB*. Should be a positive number.
- alphaImageA* The image to use as the alpha channel for *srcImageA*. If the image contains an alpha channel, that channel is used. Otherwise channel 1 in the image is used as the alpha channel. If this is not suitable for the application, then the alpha channel number in the IPL header for the image should be set appropriately before calling this function. If the argument *alphaImageA* is *NULL*, then the internal alpha channel of *srcImageA* is used. If *srcImageA* does not contain an alpha channel, an error message is issued.
- alphaImageB* The image to use as the alpha channel for *srcImageB*. If the image already contains an alpha channel, that channel is used. Otherwise channel 1 in the image is used as the alpha channel. If this is not suitable for the application, then the alpha channel number in the IPL header for the image should be set appropriately before calling this function. If the argument *alphaImageB* is *NULL*, then the internal alpha channel of *srcImageB* is used.



If *srcImageB* does not contain an alpha channel, then the value  $(1 - \alpha_A)$  is used for the alpha, where  $\alpha_A$  is a scaled alpha value of *srcImageA* in the range 0 to 1.

*alphaImageDst* The image to use as the alpha channel for *dstImage*. If the image already contains an alpha channel, that channel is used. Otherwise channel 1 in the image is used as the alpha channel. If this is not suitable for the application, then the alpha channel number in the IPL header for the image should be set appropriately before calling this function. This argument can be `NULL`, in which case the resultant alpha values are not saved.

*premulAlpha* A Boolean flag indicating whether or not the input images contain pre-multiplied alpha values. If true, they contain these values.

*divideMode* A Boolean flag set to false by default. When true, the resultant pixel color (see [Table 5-2](#)) is further divided by the resultant alpha value to get the final resultant pixel color.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplAlphaComposite()` performs an image compositing operation by overlaying the foreground image *srcImageA* with the background image *srcImageB* to produce the resultant image *dstImage*.

The function `iplAlphaComposite()` executes under one of the following conditions for the alpha channels:

- If *alphaImageA* and *alphaImageB* are both `NULL`, then the internal alpha channels of the two input images specified by their respective IPL image headers are used. The application has to ensure that these are set to the proper channel number prior to calling this function. If *srcImageB* does not have an alpha channel, then its alpha value is set to  $(1 - \alpha_A)$  where  $\alpha_A$  is the scaled alpha value of image *srcImageA* in the range 0 to 1.
- If both alpha images *alphaImageA* and *alphaImageB* are not `NULL`, then they are used as the alpha values for the two input images. If *alphaImageB* is `NULL`, then its alpha value is set to  $(1 - \alpha_A)$  where  $\alpha_A$  is the scaled alpha value of image *alphaImageA* in the range 0 to 1.

It is an error if none of the above conditions is satisfied.

If *alphaImageDst* is not `NULL`, then the resultant alpha values are written to it. If it is `NULL` and the output image *imageDst* contains an alpha channel (specified by its IPL image header), then it is set to the resulting alpha values. Otherwise (that is, alpha channel number is zero), the output pixel values are multiplied by the resulting alpha values before final storage in the output image occurs.

The function `iplAlphaCompositeC()` is used to specify constant alpha values  $\alpha_A$  and  $\alpha_B$  to be used for the two input images (usually  $\alpha_B$  is set to the value  $1 - \alpha_A$ ). The output pixel values are multiplied by the resultant alpha values before final storage in the output image. The resultant alpha values (also constant) are not saved.

The type of compositing is specified by the argument *compositeType* which can assume the values shown in [Table 5-2](#).

Table 5-2 Types of Image Compositing Operations

Type	Output Pixel (see Note)	Output Pixel (pre-mult. $\alpha$ )	Resultant Alpha	Description
OVER	$\alpha_A * A +$ $(1 - \alpha_A) * \alpha_B * B$	$A + (1 - \alpha_A) * B$	$\alpha_A +$ $(1 - \alpha_A) * \alpha_B$	A occludes B
IN	$\alpha_A * A * \alpha_B$	$A * \alpha_B$	$\alpha_A * \alpha_B$	A within B. A acts as a matte for B. A shows only where B is visible.
OUT	$\alpha_A * A * (1 - \alpha_B)$	$A * (1 - \alpha_B)$	$\alpha_A * (1 - \alpha_B)$	A outside B. NOT-B acts as a matte for A. A shows only where B is not visible.
ATOP	$\alpha_A * A * \alpha_B +$ $(1 - \alpha_A) * \alpha_B * B$	$A * \alpha_B +$ $(1 - \alpha_A) * B$	$\alpha_A * \alpha_B +$ $(1 - \alpha_A) * \alpha_B$	Combination of (A IN B) and (B OUT A). B is both back-ground and matte for A.
XOR	$\alpha_A * A * (1 - \alpha_B) +$ $(1 - \alpha_A) * \alpha_B * B$	$A * (1 - \alpha_B) +$ $(1 - \alpha_A) * B$	$\alpha_A * (1 - \alpha_B) +$ $(1 - \alpha_A) * \alpha_B$	Combination of (A OUT B) and (B OUT A). A and B mutually exclude each other.
PLUS	$\alpha_A * A + \alpha_B * B$	$A + B$	$\alpha_A + \alpha_B$	Blend without precedence



**NOTE.** In Table 5-2, the resultant pixel value is divided by the resultant alpha when `divideMode` is set to true (see the argument descriptions for the `iplAlphaComposite()` function). The Greek letter  $\alpha$  here and below denotes normalized (scaled) alpha values in the range 0 to 1.

For example, for the OVER operation, the output C for each pixel in the inputs A and B is determined as

$$C = \alpha_A * A + (1 - \alpha_A) * \alpha_B * B$$

# 5

The above operation is done for each color channel in A, B, and C. When the images A and B contain pre-multiplied alpha values, C is determined as

$$C = A + (1 - \alpha_A) * B$$

The resultant alpha value  $\alpha_C$  (alpha in the resultant image C) is computed as (both pre-multiplied and not pre-multiplied alpha cases) from  $\alpha_A$  (alpha in the source image A) and  $\alpha_B$  (alpha in the source image B):

$$\alpha_C = \alpha_A + (1 - \alpha_A) * \alpha_B$$

Thus, to perform an OVER operation, use the `IPL_COMPOSITE_OVER` identifier for the argument `compositeType`. For all other types, use `IPL_COMPOSITE_IN`, `IPL_COMPOSITE_OUT`, `IPL_COMPOSITE_ATOP`, `IPL_COMPOSITE_XOR`, and `IPL_COMPOSITE_PLUS`, respectively.

The argument `divideMode` is typically set to false to give adequate results as shown in the above example for an OVER operation and in [Table 5-2](#). When `divideMode` is set to true, the resultant pixel color is divided by the resultant alpha value. This gives an accurate result pixel value, but the division operation is expensive. In terms of the OVER example without pre-multiplication, the final value of the pixel C is computed as

$$C = (\alpha_A * A + (1 - \alpha_A) * \alpha_B * B) / \alpha_C$$

There is no change in the value of  $\alpha_C$ , and it is computed as shown above. When both A and B are 100% transparent (that is,  $\alpha_A$  is zero and  $\alpha_B$  is zero),  $\alpha_C$  is also zero and the result cannot be determined. In many cases, the value of  $\alpha_C$  is 1, so the division has no effect.

## PreMultiplyAlpha

*Pre-multiplies alpha values of an image.*

---

```
void iplPreMultiplyAlpha (IplImage* image,  
int alphaValue, IplCoord* map );
```

- |                   |   |
|-------------------|---|
| <i>image</i>      | The image for which the alpha pre-multiplication is performed.  |
| <i>alphaValue</i> | The global alpha value to use in the range 0 to 256. If this value is negative (for example, -1), the internal alpha channel of the image is used. It is an error condition if an alpha channel does not exist. |
| <i>map</i>        | The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.  |

### Discussion

The function `iplPreMultiplyAlpha()` converts an image to the pre-multiplied alpha form. If (R, G, B, A) are the red, green, blue, and alpha values of a pixel, then the pixel is stored as (R\* $\alpha$ , G\* $\alpha$ , B\* $\alpha$ , A) after execution of this function. Here  $\alpha$  is the pixel's normalized alpha value in the range 0 to 1.

Optionally, a global alpha value *alphaValue* can be used for the entire image. Then the pixels are stored as (R\* $\alpha$ , G\* $\alpha$ , B\* $\alpha$ , *alphaValue*) if the image has an alpha channel or (R\* $\alpha$ , G\* $\alpha$ , B\* $\alpha$ ) if the image does not have an alpha channel. Here  $\alpha$  is the normalized *alphaValue* in the range 0 to 1.

# Image Filtering

This chapter describes filtering operations that can be applied to images. IPL uses linear and non-linear filters. The linear filters include a subgroup of 2D convolution filters. Table 6-1 lists IPL image filtering functions.

**Table 6-1 Image Filtering Functions**

Group	Function Name	Description
Linear Filters	<a href="#"><u>iplBlur</u></a>	Applies a simple neighborhood averaging filter to blur the image.
2D Convolution Linear Filters	<a href="#"><u>iplCreateConvKernel</u></a>	Creates a convolution kernel.
	<a href="#"><u>iplGetConvKernel</u></a>	Reads the attributes of a convolution kernel.
	<a href="#"><u>iplDeleteConvKernel</u></a>	Deallocates a convolution kernel.
	<a href="#"><u>iplConvolve2D</u></a>	Convolve an image with one or more convolution kernels.
	<a href="#"><u>iplConvolveSep2D</u></a>	Convolve an image with a separable convolution kernel.
Non-linear Filters	<a href="#"><u>iplMedianFilter</u></a>	Applies a median filter to the image.
	<a href="#"><u>iplMaxFilter</u></a>	Applies a maximum filter to the image.
	<a href="#"><u>iplMinFilter</u></a>	Applies a minimum filter to the image.

## Linear Filters

Linear filtering includes simple neighborhood averaging filter to blur the image and 2D convolution operations.

## Blur

*Applies simple neighborhood averaging filter to blur the image.*

---

```
void iplBlur(IplImage* srcImage, IplImage* dstImage,  
int nRows, int nCols, int anchorX, int anchorY, IplCoord*  
map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nRows</i>	Number of rows in the neighborhood to use.
<i>nCols</i>	Number of columns in the neighborhood to use.
<i>anchorX, anchorY</i>	The [x, y] coordinates of the anchor cell in the neighborhood. In this coordinate system, the top left corner would be [0, 0] and the bottom right corner would be [ <i>nRows</i> -1, <i>nCols</i> -1]. For a 3 by 3 neighborhood, the coordinates of the geometric center would be [1, 1]. This specification allows the neighborhood to be skewed with respect to its geometric center.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplBlur()` sets each pixel in the output image as the average of all the input image pixels in the neighborhood of size *nRows* by *nCols* with the anchor cell at that pixel. This has the effect of smoothing or blurring the input image. The linear averaging filter of an image is also called a box filter.

## 2D Convolution

The 2D convolution is a versatile image processing primitive which can be used in a variety of image processing operations; for example, edge detection, blurring, noise removal, and feature detection. It is also known as mask convolution or spatial convolution.



---

**NOTE.** *In some literature sources, the 2D convolution is referred to as box filtering, which is an incorrect use of the term. A box filter is a linear averaging filter (see function `iplBlur` above). Technically, a box filter can be effectively (although less efficiently) implemented by 2D convolution using a kernel with unit or constant values.*

---

For 2D convolution, a rectangular kernel is used. The kernel is a matrix of signed integers which are, actually, the signed bytes of the C-language “signed char” type. The value range of these bytes is  $-128$  to  $127$ . The kernel could be a single row (a row filter) or a single column (a column filter) or composed of many rows and columns. There is a cell in the kernel called the “anchor,” which is usually a geometric center of the kernel, but can be skewed with respect to the geometric center.

For each input pixel, the kernel is placed on the image such that the anchor coincides with the input pixel. The output pixel value is computed as the matrix dot product of the image matrix (the portion of the input image on which the kernel is overlaid) and the kernel matrix. Optionally, the output pixel value may be scaled.

The convolution function can be used in two ways. The first way uses a single kernel for convolution. The second way uses multiple kernels and allows the specification of a method to combine the results of convolution with each kernel. This enables efficient implementation of multiple kernels which eliminates the need of storing the intermediate results when using each kernel. One IPL function, `iplConvolve2D()`, can implement both ways.



In addition, `iplConvolveSep2D()`, a convolution function that uses separable kernels is also provided. The convolution kernel is separable into the x and y components.

The 2D convolution functions, `iplCreateConvKernel()` and `iplGetConvKernel()`, allow you to create and access kernels, upon which convolving the image with one or more convolution kernels or with a separable kernel can be performed.

---

## CreateConvKernel

*Creates a convolution kernel.*

---

```
IplConvKernel* iplCreateConvKernel(int nRows, int nCols,  
int anchorX, int anchorY, char* values, int nShiftR);
```

<i>nRows</i>	The number of rows in the convolution kernel.
<i>nCols</i>	The number of columns in the convolution kernel.
<i>anchorX, anchorY</i>	The [x, y] coordinates of the anchor cell in the kernel. In this coordinate system, the top left corner would be [0, 0] and the bottom right corner would be [ <i>nRows</i> -1, <i>nCols</i> -1]. For a 3 by 3 kernel, the coordinates of the geometric center would be [1, 1]. This specification allows the kernel to be skewed with respect to its geometric center.
<i>values</i>	A pointer to an array of values to be used for the kernel matrix. The values are read in row-major form starting with the top left corner. There should be exactly <i>nRows*nCols</i> entries in this array. For example, the array [1, 2, 3, 4, 5, 6, 7, 8, 9] would represent the kernel matrix

```
1 2 3
4 5 6
7 8 9
```

*nShiftR*

Scale the resulting output pixel by shifting it to the right *nShiftR* times.

## Discussion

The function `iplCreateConvKernel()` can be used to create a convolution kernel of arbitrary size and arbitrary anchor point.

## Return Value

A pointer to the convolution kernel structure `IplConvKernel`.

---

## GetConvKernel

*Reads the attributes of a convolution kernel.*

---

```
void iplGetConvKernel(IplConvKernel* kernel, int* nRows,
int* nCols, int* anchorX, int* anchorY, char** values,
int* nShiftR);
```

*kernel*

The kernel to get the attributes for. The attributes are returned in the remaining arguments.

*nRows*

A pointer to the number of rows in the convolution kernel. Set by the function.

*nCols*

A pointer to the number of columns in the convolution kernel. Set by the function.

---

<i>anchorX, anchorY</i>	Pointers to the [x, y] coordinates of the anchor cell in the kernel. (See <a href="#">iplCreateConvKernel</a> above.) Set by the function.
<i>values</i>	A pointer to an array of values to be used for the kernel matrix. The values are read in row-major form starting with the top left corner. There will be exactly <i>nRows*nCols</i> entries in this array. For example, the array [1, 2, 3, 4, 5, 6, 7, 8, 9] would represent the kernel matrix <pre> 1 2 3 4 5 6 7 8 9</pre>
<i>nShiftR</i>	A pointer to the number of bits to shift (to the right) the resulting output pixel of each convolution. Set by the function.

### Discussion

The function `iplGetConvKernel()` can be used to read the attributes of a convolution kernel.

---

## DeleteConvKernel

*Deletes a convolution kernel.*

---

```
void iplDeleteConvKernel(IplConvKernel* kernel);
```

*kernel*                    The kernel to delete.

### Discussion

The function `iplDeleteConvKernel()` must be used to delete a convolution kernel which was created by the `iplCreateConvKernel()` function.

## Convolve2D

*Convolve an image with one or more convolution kernels.*

---

```
void iplConvolve2D(IplImage* srcImage, IplImage*  
dstImage, IplConvKernel** kernel, int nKernels,  
int combineMethod, IplCoord* map);
```

<i>srcImage</i>	The source image.										
<i>dstImage</i>	The resultant image.										
<i>kernel</i>	A pointer to an array of pointers to convolution kernels. The length of the array is <i>nKernels</i> .										
<i>nKernels</i>	The number of kernels in the array <i>kernel</i> . The value of <i>nKernels</i> can be 1 or more.										
<i>combineMethod</i>	The way in which the results of applying each kernel should be combined. This argument is ignored when a single kernel is used. The following combinations are supported:  <table><tr><td><i>IPL_SUM</i></td><td>Sums the results.</td></tr><tr><td><i>IPL_SUMSQ</i></td><td>Sums the squares of the results.</td></tr><tr><td><i>IPL_SUMSQROOT</i></td><td>Sums the squares of the results and then takes the square root.</td></tr><tr><td><i>IPL_MAX</i></td><td>Takes the maximum of the results.</td></tr><tr><td><i>IPL_MIN</i></td><td>Takes the minimum of the results.</td></tr></table>	<i>IPL_SUM</i>	Sums the results.	<i>IPL_SUMSQ</i>	Sums the squares of the results.	<i>IPL_SUMSQROOT</i>	Sums the squares of the results and then takes the square root.	<i>IPL_MAX</i>	Takes the maximum of the results.	<i>IPL_MIN</i>	Takes the minimum of the results.
<i>IPL_SUM</i>	Sums the results.										
<i>IPL_SUMSQ</i>	Sums the squares of the results.										
<i>IPL_SUMSQROOT</i>	Sums the squares of the results and then takes the square root.										
<i>IPL_MAX</i>	Takes the maximum of the results.										
<i>IPL_MIN</i>	Takes the minimum of the results.										
<i>map</i>	The structure specifying offsets for tiling purposes.										

## Discussion

The function `iplConvolve2D()` is used to convolve an image with a set of convolution kernels. The results of using each kernel are then combined using the `combineMethod` argument.

---

## ConvolveSep2D

*Convolve an image with a separable convolution kernel.*

---

```
void iplConvolveSep2D(IplImage* srcImage, IplImage*  
dstImage, IplConvKernel* xKernel, IplConvKernel* yKernel,  
IplCoord* map);
```

<code>srcImage</code>	The source image.
<code>dstImage</code>	The resultant image.
<code>xKernel</code>	The x or row kernel. Must contain only one row.
<code>yKernel</code>	The y or column kernel. Must contain only one column.
<code>map</code>	The structure specifying offsets for tiling purposes.

## Discussion

The function `iplConvolveSep2D()` is used to convolve the input image `srcImage` with the separable kernel specified by the row kernel `xKernel` and column kernel `yKernel`. The resulting output image is `dstImage`.

## Non-linear Filters

Non-linear filtering involves performing non-linear operations on some neighborhood of the image. Most common are the minimum, maximum and median filters.

---

### MedianFilter

*Apply a median filter to the image.*

---

```
void iplMedianFilter(IplImage* srcImage, IplImage*  
dstImage, int nRows, int nCols, int anchorX,  
int anchorY, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nRows</i>	Number of rows in the neighborhood to use.
<i>nCols</i>	Number of columns in the neighborhood to use.
<i>anchorX, anchorY</i>	The [x, y] coordinates of the anchor cell in the neighborhood. In this coordinate system, the top left corner would be [0, 0] and the bottom right corner would be [ <i>nRows</i> -1, <i>nCols</i> -1]. For a 3 by 3 neighborhood, the coordinates of the geometric center would be [1, 1]. This specification allows the neighborhood to be skewed with respect to its geometric center.
<i>map</i>	The structure specifying offsets for tiling purposes.

## Discussion

The function `iplMedianFilter()` sets each pixel in the output image as the median value of all the input image pixel values in the neighborhood of size `nRows` by `nCols` with the anchor cell at that pixel. This has the effect of removing the noise in the image.

---

## MaxFilter

*Apply a max filter to the image.*

---

```
void iplMaxFilter(IplImage* srcImage, IplImage* dstImage,
int nRows, int nCols, int anchorX, int anchorY, IplCoord*
map);
```

`srcImage` The source image.

`dstImage` The resultant image.

`nRows` Number of rows in the neighborhood to use.

`nCols` Number of columns in the neighborhood to use.

`anchorX, anchorY` The [x, y] coordinates of the anchor cell in the neighborhood. In this coordinate system, the top left corner would be [0, 0] and the bottom right corner would be [`nRows-1`, `nCols-1`]. For a 3 by 3 neighborhood, the coordinates of the geometric center would be [1, 1]. This specification allows the neighborhood to be skewed with respect to its geometric center.

`map` The structure specifying offsets for tiling purposes.

## Discussion

The function `iplMaxFilter()` sets each pixel in the output image as the maximum value of all the input image pixel values in the neighborhood of size `nRows` by `nCols` with the anchor cell at that pixel. This has the effect of increasing the contrast in the image.

---

## MinFilter

*Apply a min filter to the image.*

---

```
void iplMinFilter(IplImage* srcImage, IplImage* dstImage,  
int nRows, int nCols, int anchorX, int anchorY, IplCoord*  
map);
```

<code>srcImage</code>	The source image.
<code>dstImage</code>	The resultant image.
<code>nRows</code>	Number of rows in the neighborhood to use.
<code>nCols</code>	Number of columns in the neighborhood to use.
<code>anchorX, anchorY</code>	The [x, y] coordinates of the anchor cell in the neighborhood. (In this coordinate system, the top left corner would be [0, 0] and the bottom right corner would be [ <code>nRows-1</code> , <code>nCols-1</code> ]. For a 3 by 3 neighborhood the coordinates of the geometric center would be [1, 1]). This specification allows the neighborhood to be skewed with respect to its geometric center.
<code>map</code>	The structure specifying offsets for tiling purposes.



## Discussion

The function `iplMinFilter()` sets each pixel in the output image as the minimum value of all the input image pixel values in the neighborhood of size `nRows` by `nCols` with the anchor cell at that pixel. This has the effect of decreasing the contrast in the image.

# Linear Image Transforms

This chapter describes the linear image transforms implemented in the library: Fast Fourier Transform (FFT) and Discrete Cosine Transform (DCT). Table 7-1 lists IPL functions performing linear image transform operations.

**Table 7-1 Linear Image Transform Functions**

Group	Function Name	Description
Fast Fourier Transform (FFT)	<a href="#">iplRealFft2D</a>	Computes the forward or inverse 2D FFT of an image.
	<a href="#">iplCcsFft2D</a>	Computes the forward or inverse 2D FFT of an image in a complex-conjugate format.
Discrete Cosine Transform (DCT)	<a href="#">iplDCT2D</a>	Computes the forward or inverse 2D DCT of an image.

## Fast Fourier Transform

This section describes the functions that implement the forward and inverse Fast Fourier Transform (FFT) on the 2-dimensional (2D) image data.

### Real-Complex Packed (RCPack2D) Format

The FFT of any real 2D signal, in particular, the FFT of an image is conjugate-symmetric. Therefore, it can be fully specified by storing only half the output data. A special format called `RCPack2D` is provided for this purpose.

The function `iplRealFft2D()` transforms a 2D image and produces the Fourier coefficients in the `RCPack2D` format. To complement this, function `iplCcsFft2D()` is provided that uses its input in `RCPack2D` format, performs the Fourier transform, and produces its output as a real 2D image. The functions `iplRealFft2D()` and `iplCcsFft2D()` together can be used to perform frequency domain filtering of images.

`RCPack2D` format is defined based on the following equations:

$$A_{s,j} = \sum_{l=0}^{L-1} \sum_{k=0}^{K-1} f_{k,l} \exp\left(-\frac{2\pi i j l}{L}\right) \exp\left(-\frac{2\pi i k s}{K}\right)$$

$$f_{k,l} = \frac{1}{LK} \sum_{j=0}^{L-1} \sum_{s=0}^{K-1} A_{s,j} \exp\left(\frac{2\pi i j l}{L}\right) \exp\left(\frac{2\pi i k s}{K}\right)$$

where  $i = \sqrt{-1}$ ,  $f_{k,l} = f(x_k, y_l)$ ,  $x_k = k/K$ ,  $y_l = l/L$ .

Note that the Fourier coefficients have the following relationship:

$$\tilde{A}_{s,j} = A_{K-s, L-j}$$

The symbol “~” denotes complex-conjugate. Hence, the  $L * K$  real values can be used to reconstruct the  $L * K$  complex coefficients  $A_{k,l}$ . Thus it is enough to store only  $L * K$  real values.

Using the DFT (real or complex) function  $a_s(y)$ , we have the following Fourier coefficients:

$$A_{s,j} = \sum_{l=0}^{L-1} a_s(l) \exp\left(-\frac{2\pi i l j}{L}\right) = \sum_{l=0}^{L-1} \sum_{k=0}^{K-1} f_{k,l} \exp\left(-\frac{2\pi i j l}{L}\right) \exp\left(-\frac{2\pi i k s}{K}\right)$$

where

$$\begin{aligned} s &= 0, \dots, K/2 && \text{for even } K \\ s &= 0, \dots, (K-1)/2 && \text{for odd } K \\ j &= 0, \dots, (L-1) \end{aligned}$$

Other Fourier coefficients can be found using complex-conjugate relations. Fourier coefficients  $A_{0,j}$  for real function  $a_0$  (Fourier coefficients  $A_{\kappa/2,j}$  for even  $\kappa$ ) can be stored in the **RCPack** format, other coefficients are standard. The **RCPack** format is a convenient compact representation of a complex conjugate-symmetric sequence. The disadvantage of this format is that it is not the natural format used by the real FFT algorithms (“natural” in the sense that bit-reversed order is natural for radix-2 complex FFTs).

In the **RCPack2D** format, the output samples of the FFT are arranged as shown in Table 7-2 where **Re** corresponds to Real and **Im** corresponds to Imaginary.

**Table 7-2 Arrangement of Output Samples in RCPack2D Format**

$A_{0,0}$	Re $A_{0,1}$ Im $A_{0,1}$	Re $A_{0,2}$ Im $A_{0,2}$	...	Re $A_{0,j}$ Im $A_{0,j}$
Re $A_{1,0}$ Im $A_{1,0}$	Re $A_{1,1}$ Im $A_{1,1}$	Re $A_{1,2}$ Im $A_{1,2}$	...	Re $A_{1,j}$ Im $A_{1,j}$
...	...	...	...	...
Re $A_{\kappa/2,0}$ Im $A_{\kappa/2,0}$	Re $A_{\kappa/2,1}$ Im $A_{\kappa/2,1}$	Re $A_{\kappa/2,2}$ Im $A_{\kappa/2,2}$	...	Re $A_{\kappa/2,j}$ Im $A_{\kappa/2,j}$

## RealFft2D

*Computes the forward or  
inverse 2D FFT of an image.*

```
void iplRealFft2D(IplImage* srcImage, IplImage* dstImage,
                 int flags, IplCoord* map);
```

*srcImage*                      The source image.

---

<i>dstImage</i>	<p>The resultant image in <code>RCPack2D</code> format containing the Fourier coefficients. This image must be a multi-channel image containing the same number of channels as <i>srcImage</i>. The data type for the image must be 8, 16 or 32 bits.</p> <p>This image cannot be the same as the input image <i>srcImage</i> (that is, an in-place operation is not allowed).</p>																		
<i>flags</i>	<p>Specifies how to perform FFT . This is an integer in which every bit can be assigned the following values using logical OR:</p> <table> <tr> <td><code>IPL_FFT_Forw</code></td> <td>Do forward transform</td> </tr> <tr> <td><code>IPL_FFT_Inv</code></td> <td>Do inverse transform</td> </tr> <tr> <td><code>IPL_FFT_NoScale</code></td> <td>Do inverse transform without scaling</td> </tr> <tr> <td><code>IPL_FFT_DoAlpha</code></td> <td>Transform alpha channel (if <i>alphaChannel</i> is not 0)</td> </tr> <tr> <td><code>IPL_FFT_UseInt</code></td> <td>Use only integer core</td> </tr> <tr> <td><code>IPL_FFT_UseFloat</code></td> <td>Use only float core</td> </tr> <tr> <td><code>IPL_FFT_OnlyOffsetROIandCalc</code></td> <td>Take only offset ROI and calc</td> </tr> <tr> <td><code>IPL_FFT_Free</code></td> <td>Only free all working arrays and exit</td> </tr> <tr> <td><code>IPL_FFT_Save</code></td> <td>Save all working arrays on exit</td> </tr> </table>	<code>IPL_FFT_Forw</code>	Do forward transform	<code>IPL_FFT_Inv</code>	Do inverse transform	<code>IPL_FFT_NoScale</code>	Do inverse transform without scaling	<code>IPL_FFT_DoAlpha</code>	Transform alpha channel (if <i>alphaChannel</i> is not 0)	<code>IPL_FFT_UseInt</code>	Use only integer core	<code>IPL_FFT_UseFloat</code>	Use only float core	<code>IPL_FFT_OnlyOffsetROIandCalc</code>	Take only offset ROI and calc	<code>IPL_FFT_Free</code>	Only free all working arrays and exit	<code>IPL_FFT_Save</code>	Save all working arrays on exit
<code>IPL_FFT_Forw</code>	Do forward transform																		
<code>IPL_FFT_Inv</code>	Do inverse transform																		
<code>IPL_FFT_NoScale</code>	Do inverse transform without scaling																		
<code>IPL_FFT_DoAlpha</code>	Transform alpha channel (if <i>alphaChannel</i> is not 0)																		
<code>IPL_FFT_UseInt</code>	Use only integer core																		
<code>IPL_FFT_UseFloat</code>	Use only float core																		
<code>IPL_FFT_OnlyOffsetROIandCalc</code>	Take only offset ROI and calc																		
<code>IPL_FFT_Free</code>	Only free all working arrays and exit																		
<code>IPL_FFT_Save</code>	Save all working arrays on exit																		
<i>map</i>	<p>The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.</p>																		

## Discussion

The function `iplRealFft2D()` performs an FFT on each channel in the specified rectangular ROI of the input image `srcImage` and writes the Fourier coefficients in `RCPack2D` format into the corresponding channel of the output image `dstImage`.

Note that the output data will be clamped (saturated) to the limits 0 and `Max`, where `Max` is determined by the data type of the output image. The 32-bit data type will produce the best results, so at least 16-bit data is recommended.

---

## CcsFft2D

*Computes the forward or inverse 2D FFT of an image in complex-conjugate format.*

---

```
void iplCcsFft2D(IplImage* srcImage, IplImage* dstImage,
                int flags, IplCoord* map);
```

`srcImage` The source image in `RCPack2D` format.

`dstImage` The resultant image. This image must be a multi-channel image containing the same number of channels as `srcImage`.

This image cannot be the same as the input image `srcImage` (that is, an in-place operation is not allowed).

`flags` Specifies how to perform FFT. This is an integer in which every bit can be assigned the following values using logical OR:

`IPL_FFT_Forw` Do forward transform

`IPL_FFT_Inv` Do inverse transform

<code>IPL_FFT_NoScale</code>	Do inverse transform without scaling
<code>IPL_FFT_DoAlpha</code>	Transform alpha channel (if <code>alphaChannel</code> is not zero)
<code>IPL_FFT_UseInt</code>	Use only integer core
<code>IPL_FFT_UseFloat</code>	Use only float core
<code>IPL_FFT_OnlyOffsetROIandCalc</code>	Take only offset ROI and calc
<code>IPL_FFT_Free</code>	Only free all working arrays and exit
<code>IPL_FFT_Save</code>	Save all working arrays on exit
<code>map</code>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplRealFft2D()` performs an FFT on each channel in the specified rectangle ROI of the input image `srcImage` and writes the output in `RCPack2D` format to the image `dstImage`.

Note that the output data will be clamped (saturated) to the limits 0 and `Max`, where `Max` is determined by the data type of the output image.

## Discrete Cosine Transform

This section describes the functions that implement the forward and inverse Discrete Cosine Transform (DCT) on the 2D image data. The output of the DCT for real input data is real. Therefore, unlike FFT, no special format for the transform output is needed.

---

## DCT2D

*Computes the forward or inverse 2D DCT of an image.*

---

```
void iplDCT2D(IplImage* srcImage, IplImage* dstImage,
int flags, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image containing the DCT coefficients. This image must be a multi-channel image containing the same number of channels as <i>srcImage</i> . The data type for the image must be 8, 16 or 32 bits.  This image cannot be the same as the input image <i>srcImage</i> (that is, an in-place operation is not allowed).
<i>flags</i>	Specifies how to perform FFT. This is an integer whose every bit can be assigned the following values using logical OR: <ul style="list-style-type: none"> <li><b>IPL_DCT_Forward</b> Do forward transform</li> <li><b>IPL_DCT_Inverse</b> Do inverse transform</li> <li><b>IPL_DCT_DoAlpha</b> Transform alpha channel (if <i>alphaChannel</i> is not 0)</li> <li><b>IPL_DCT_Free</b> Only free all working arrays and exit</li> <li><b>IPL_DCT_UseInpBuf</b> Use the input image array for the intermediate calculations. The performance of DCT increases, but the input image is destroyed. This value is a default.</li> </ul>



*map*

The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplDCT2D()` performs a DCT on each channel in the specified rectangular ROI of the input image *srcImage* and writes the DCT coefficients into the corresponding channel of the output image *dstImage*.

Note that the output data will be clamped (saturated) to the limits `Min` and `Max`, where `Min` and `Max` are determined by the data type of the output image. The 32-bit data type will produce the best results, so at least 16-bit data type is recommended.

## *Morphological Operations*

---

The morphological operations of the Image Processing Library are simple erosion and dilation of an image. A specified number of erosions and dilations are performed as part of image opening or closing operations in order to (respectively) eliminate or fill small and thin holes in objects, break objects at thin points or connect nearby objects, and generally smooth the boundaries of objects without significantly changing their area.

Table 8-1 lists the functions that perform these operations.

**Table 8-1 Morphological Operation Functions**

<b>Group</b>	<b>Function Name</b>	<b>Description</b>
Erode, Dilate	<a href="#"><code>iplErode</code></a>	Erodes the image an indicated number of times.
	<a href="#"><code>iplDilate</code></a>	Dilates the image an indicated number of times.
Open, Close	<a href="#"><code>iplOpen</code></a>	Opens the image while smoothing the boundaries of large objects.
	<a href="#"><code>iplClose</code></a>	Closes the image while smoothing the boundaries of large objects.

## Erode

*Erodes the image.*

---

```
void iplErode(IplImage* srcImage, IplImage* dstImage,  
             int nIterations, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nIterations</i>	The number of times to erode the image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplErode()` performs an erosion of the image *nIterations* times. The way the image is eroded depends on whether it is a binary image or not.


- For a binary input image, the output pixel is set to zero if the corresponding input pixel or any of its 8 neighboring pixels is a zero.
- For a gray scale or color image, the output pixel is set to the minimum of the corresponding input pixel and its 8 neighboring pixels.

The effect of erosion is to remove spurious pixels (such as noise) and to thin boundaries of objects on a dark background (whose pixel values are less than those of the objects).

Figure 8-1 shows an example of 8-bit gray scale image before erosion (left) and the same image after erosion of a rectangular ROI (right).

**Figure 8-1** Erosion in a Rectangular ROI: the Source (left) and Result (right)

---



---

The following code (Example 8-1) performs erosion of the image inside the selected rectangular ROI.

**Example 8-1 Code Used to Produce Erosion in a Rectangular ROI**

---

```
/* Create output image header with attributes pointed to
by srcImg */

dstImg = iplCreateImageHeader(
    src->nChannels, src->alphaChannel, src->depth,
    "", "", src->dataOrder, src->align,
    src->height, src->width, 0);

    /* Allocate output image */
iplAllocateImage(dstImage, NOFILL);

    /* Copy source image into output image */
memcpy(      dstImg->imageData, srcImg->imageData,
            src->imageSize);

    /* Set ROI attributes */
roi.coi = 0;          /* channels: all */
                    /* region: */
roi.xOffset = 10;    /* position 10,155 */
roi.yOffset = 155;
roi.width = 192;     /* size: 192x86 */
roi.height = 86;

    /* Set ROI into srcImg and dstImg */
srcImg->roi = &roi;
dstImg->roi = &roi;

    /* Erosion */
IplErode(srcImg, dstImg, 2);
```



**NOTE.** All source image attributes are defined in the image header pointed to by *srcImage*.

---

## Dilate

*Dilates the image.*

---

```
void iplDilate(IplImage* srcImage, IplImage* dstImage,  
int nIterations, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nIterations</i>	The number of times to dilate the image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplDilate()` performs a dilation of the image *nIterations* times. The way the image is dilated depends on whether the image is a binary image or not.

- For a binary input image, the output pixel is set to 1 if the corresponding input pixel is 1 or any of 8 neighboring input pixels is 1.
- For a gray scale or color image, the output pixel is set to the maximum of the corresponding input pixel and its 8 neighboring pixels.

The effect of dilation is to fill up holes and to thicken boundaries of objects on a dark background (whose pixel values are less than those of the objects).

## Open

*Opens the image by performing erosions followed by dilations.*

---

```
void iplOpen(IplImage* srcImage, IplImage* dstImage,  
            int nIterations, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nIterations</i>	The number of times to erode and dilate the image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplOpen()` performs *nIterations* of erosion followed by *nIterations* of dilation performed by `iplErode()` and `iplDilate()`, respectively.

The process of opening has the effect of eliminating small and thin objects, breaking objects at thin points, and generally smoothing the boundaries of larger objects without significantly changing their area.

### See Also

[Erode](#)

[Dilate](#)

## Close

*Closes the image by performing dilations followed by erosions.*

---

```
void iplClose(IplImage* srcImage, IplImage* dstImage,  
int nIterations, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>nIterations</i>	The number of times to dilate and erode the image.
<i>map</i>	The structure specifying offsets for tiling purposes.

### Discussion

The function `iplClose()` performs *nIterations* of dilation followed by *nIterations* of erosion performed by `iplDilate()` and `iplErode()`, respectively.

The process of closing has the effect of filling small and thin holes in objects, connecting nearby objects, and generally smoothing the boundaries of objects without significantly changing their area.

### See Also

[Erode](#)

[Dilate](#)



## Color Space Conversion


This chapter describes the Image Processing Library functions that perform color space conversion. The following color space conversions are supported in the library:

- Reduction from high bit resolution color to low bit resolution color
- Conversion of absolute color images to and from palette color images
- Color model conversion
- Conversion from color to gray scale and vice versa

Table 9-1 lists color space conversion functions. For information on the absolute-to-palette and palette-to-absolute color conversion, see “[Working in the Windows DIB Environment](#)” in Chapter 4.

**Table 9-1 Color Space Conversion Functions**

Conversion Type	Function Name	Description
Reducing Bit Resolution	<a href="#">iplReduceBits</a>	Reduces the number of bits per channel in an image.
Bitonal to gray scale	<a href="#">iplBitonalToGray</a>	Converts bitonal images to 8- and 16-bit gray scale images.
Color to gray scale and vice versa	<a href="#">iplColorToGray</a> <a href="#">iplGrayToColor</a>	Convert color images to and from gray scale images.
Color Models Conversion	<a href="#">iplRGB2HSV</a> , <a href="#">iplHSV2RGB</a>  <a href="#">iplRGB2HLS</a> , <a href="#">iplHLS2RGB</a>	Convert RGB images to and from HSV color model.  Convert RGB images to and from HLS color model.

continued 

**Table 9-1 Color Space Conversion Functions** (continued)

Conversion Type	Function Name	Description
Color Models Conversion (continued)	<a href="#"><code>iplApplyColorTwist</code></a>	Applies a color-twist matrix to an image.
	<a href="#"><code>iplCreateColorTwist</code></a>	Allocates memory for color-twist matrix data structure.
	<a href="#"><code>iplDeleteColorTwist</code></a>	Deletes the color-twist matrix data structure.
	<a href="#"><code>iplSetColorTwist</code></a>	Sets a color-twist matrix data structure.

## Reducing the Image Bit Resolution

This section describes functions that reduce the bit resolution of absolute color and gray scale images.

### ReduceBits

*Reduces the bits per channel in an image.*

```
void iplReduceBits(IplImage* srcImage, IplImage*
dstImage, int ditherType, int jitterType, int levels,
IplCoord* map);
```

*srcImage*

The source image of a higher bit resolution. Refer to the discussion below for a list of valid source and resultant image combinations.

---

<i>dstImage</i>	The resultant image of a lower bit resolution. Refer to the discussion below for a list of valid source and resultant image combinations.
<i>jitterType</i>	The number specifying the noise added; should be in the range 0 to 8.
<i>ditherType</i>	The type of dithering to be used. The following types are allowed: <ul style="list-style-type: none"> <li><code>IPL_DITHER_NONE</code> No dithering is done</li> <li><code>IPL_DITHER_FS</code> The Floyd-Steinberg dithering algorithm is used.</li> <li><code>IPL_DITHER_JJH</code> The Jarvice-Judice-Hinke dithering algorithm is used.</li> <li><code>IPL_DITHER_STUCKEY</code> The Stucki dithering algorithm is used</li> <li><code>IPL_DITHER_BAYER</code> The Bayer dithering algorithm is used</li> </ul>
<i>levels</i>	Number of levels for dithering; should be a power of 2.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplReduceBits()` reduces a higher bit resolution of the absolute color or gray scale source image *srcImage* to a lower resolution of the resultant absolute color or gray scale image *dstImage*. All combinations of jittering and dithering values are valid. If *jitterType* is greater than 0, some random noise is added to all pixels before the reduction, which eliminates the problem of visible color stepping; see [Bragg]. The resultant image can be used as input to a color quantization method for further reduction in the number of colors; see [Thomas] and [Schumacher].

Table 9-2 lists the valid combinations of the source and resultant image bit data types for reducing the bit resolution.

**Table 9-2 Source and Resultant Image Data Types for Reducing the Bit Resolution**

Source Image	Resultant Image
32 bit per channel	1 (for gray image), 8 or 16 bit per channel
16 bit per channel	8 or 1 (for gray image) bit per channel
8 bit per channel	1 bit per channel (for gray image)

Bit reducing uses the equation  $dst = src * (((1 \ll n) - 1) / ((1 \ll m) - 1))$ , where  $m$  is the bit depth of the source and  $n$  is the bit depth of the destination. To reduce a gray scale image to a bitonal (1-bit) image, see the discussion under the thresholding function [iplThreshold](#) in Chapter 10.

## Conversion from Bitonal to Gray Scale Images

This section describes the function that performs the conversion of bitonal images to gray scale.

### BitonalToGray

*Converts a bitonal image to gray scale.*

```
void iplBitonalToGray(IplImage* srcImage, IplImage*
dstImage, int ZeroScale, int OneScale, IplCoord* map);
```

*srcImage*                      The bitonal source image.

<i>dstImage</i>	The resultant gray scale image. (See the discussion below.)
<i>zeroScale</i>	The value that zero pixels of the source image should have in the resultant image.
<i>oneScale</i>	The value given to a resultant pixel if the corresponding input pixel is 1.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplBitonalToGray()` converts the input 1-bit bitonal image *srcImage* to an 8s, 8u, 16s or 16u gray scale image *dstImage*.

If an input pixel is 0, the corresponding output pixel is set to *zeroScale*.

If an input pixel is 1, the corresponding output pixel is set to *oneScale*.

## Conversion of Absolute Colors to and from Palette Colors

Since the IPL image format supports only absolute color images, this functionality is provided only within the context of converting an IPL absolute color image to and from a palette color DIB image. See the section “[Working in the Windows DIB Environment](#)” in Chapter 4.

## Conversion from Color to Gray Scale

This section describes the function that performs the conversion of absolute color images to gray scale.

## ColorToGray

*Converts a color image to gray scale.*

```
void iplColorToGray(IplImage* srcImage,
IplImage* dstImage, IplCoord* map);
```

*srcImage* The source image. See the discussion below for a list of valid source and resultant image combinations.

*dstImage* The resultant image. See the discussion below for a list of valid source and resultant image combinations.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplColorToGray()` converts a color source image *srcImage* to a gray scale resultant image *dstImage*.

Table 9-3 lists the valid combinations of source and resultant image bit data types for conversion from color to gray scale.

**Table 9-3 Source and Resultant Image Data Types for Conversion from Color to Gray Scale**

Source Image (data type)	Result image (data type)
32 bit per channel	Gray scale (1, 8, or 16 bit)
16 bit per channel	Gray scale (1, 8, or 16 bit)
8 bit per channel	Gray scale (1, 8, or 16 bit)

The weights to compute true luminance from linear red, green and blue are these:

$$Y = 0.212671 * R + 0.715160 * G + 0.072169 * B.$$

## Conversion from Gray Scale to Color (Pseudo-color)

This section describes the conversion of gray scale image to pseudo color.

---

### GrayToColor

*Converts a gray scale to color image.*

---

```
void iplGrayToColor (IplImage* srcImage,  
IplImage* dstImage, float FractR, float FractG, float  
FractB, IplCoord* map);
```

*srcImage* The source image. See the discussion below for a list of valid source and resultant image combinations.

*dstImage* The resultant image. See the discussion below for a list of valid source and resultant image combinations.

*FractR, FractG, FractB* The red, green and blue intensities for image reconstruction. See the discussion below for a list of valid *FractR*, *FractG*, and *FractB* values.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

### Discussion

The function `iplGrayToColor()` converts a gray scale source image *srcImage* to a resultant pseudo-color image *dstImage*. Table 9-4 lists the valid combinations of source and resultant image bit data types for conversion from gray scale to color.

**Table 9-4 Source and Resultant Image Data Types for Conversion from Gray Scale to Color**

Source Image (data type)	Result image (data type)
Gray scale 1 bit	8 bit per channel
Gray scale 8 bit	8 bit per channel
Gray scale 16 bit	16 bit per channel
Gray scale 32 bit	32 bit per channel

The equation for chrominance in RGB from luminance  $Y$  is:

$$\begin{aligned}
 R &= FractR * Y; & 0 \leq FractR \leq 1 \\
 G &= FractG * Y; & 0 \leq FractG \leq 1 \\
 B &= FractB * Y; & 0 \leq FractB \leq 1.
 \end{aligned}$$

If  $FractR==0 \ \&\& \ FractG==0 \ \&\& \ FractB==0$ , then the default values are used in above equation so that:

$$R = 0.212671 * Y, \quad G = 0.715160 * Y, \quad B = 0.072169 * Y.$$

## Conversion of Color Models

To convert one color model to another, a color twist matrix can be used. See the “[Color Twist Matrices](#)” section, which describes this method and presents examples of various color model conversions.

Described in this section are conversions of color models when the color twist matrix cannot be used.

Note that conversion of the RGB to the CMY models can be performed by a simple subtraction. The function `iplSubtractS` can be used to accomplish this conversion for two 8-bit per channel images. For example, with maximum pixel value of 255, the `iplSubtractS()` function is used as follows:

```
iplSubtractS(rgbImage, cmyImage, 256, TRUE)
```



This call converts the RGB image *rgbImage* to the CMY image *cmvImage* by setting each channel in the CMY image as follows:

```
C = 255 - R
M = 255 - G
Y = 255 - B
```

The conversion from CMY to RGB is similar: just switch the RGB and CMY images.

---

## RGB2HSV

*Converts from the RGB color model to the HSV color model.*

---

```
void iplRGB2HSV(IplImage* rgbImage, IplImage* hsvImage,
IplCoord* map);
```

<i>rgbImage</i>	The source RGB image.
<i>hsvImage</i>	The resultant HSV.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplRGB2HSV()` converts the RGB image *rgbImage* to the HSV image *hsvImage*. The function checks that the input image is an RGB image. The channel sequence and color model of the output image are set to HSV.

---

## HSV2RGB

*Converts from the HSV color model to the RGB color model.*

---

```
void iplHSV2RGB(IplImage* hsvImage, IplImage* rgbImage,  
IplCoord* map);
```

<i>hsvImage</i>	The source HSV image.
<i>rgbImage</i>	The resultant RGB.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplHSV2RGB()` converts the HSV image *hsvImage* to the RGB image *rgbImage*. The function checks that the input image is an HSV image. The channel sequence and color model of the output image are set to RGB.

---

## RGB2HLS

*Converts from the RGB color model to the HLS color model.*

---

```
void iplRGB2HLS(IplImage* rgbImage, IplImage* hlsImage,  
IplCoord* map);
```

<i>rgbImage</i>	The source RGB image.
<i>hlsImage</i>	The resultant HLS.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplRGB2HLS()` converts the RGB image *rgbImage* to the HLS image *hlsImage*. The function checks that the input image is an RGB image. The function sets the channel sequence and color model of the output image to HLS.

---

## HLS2RGB

*Converts from the HLS color model to the RGB color model.*

---

```
void iplHLS2RGB(IplImage* hlsImage, IplImage* rgbImage, IplCoord* map);
```

*hlsImage* The source HLS image.

*rgbImage* The resultant RGB.

*map* The structure specifying offsets for tiling purposes; see [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplHLS2RGB()` converts the HLS image *hlsImage* to the RGB image *rgbImage*; see [Rogers85]. The function checks that the input image is an HLS image. The channel sequence and color model of the output image are set to RGB.

## Using Color-Twist Matrices

One of the methods of color model conversion is using a color-twist matrix. The color-twist matrix is a generalized 4 by 4 matrix  $[t_{i,j}]$  that converts the three channels (a, b, c) into (d, e, f) according to the following matrix multiplication by a color-twist matrix (the superscript  $T$  is used to indicate the transpose of the matrix).

$$[d, e, f, 1]^T = \begin{bmatrix} t_{11} & t_{12} & t_{13} & t_{14} \\ t_{21} & t_{22} & t_{23} & t_{24} \\ t_{31} & t_{32} & t_{33} & t_{34} \\ 0 & 0 & 0 & t_{44} \end{bmatrix} * [a, b, c, 1]^T$$

To apply a color-twist matrix to an IPL image, use the function `iplApplyColorTwist()`. But first call the `iplCreateColorTwist()` and `iplSetColorTwist()` functions to create the data structure `IplColorTwist`. This data structure contains the color-twist matrix and allows you to store the data internally in a form that is efficient for computation.

The function descriptions that follow provide examples of using the color-twist matrices for color model conversion.

---

### CreateColorTwist

*Creates a color-twist matrix data structure.*

```
IplColorTwist* iplCreateColorTwist(int data[16],
int scalingValue);
```

*data*

An array containing the sixteen values that constitute the color-twist matrix. The values are in row-wise order. Color-twist values that are in the range -1 to 1 should be scaled up to be in the

range  $-2^{31}$  to  $2^{31}$ . (Simply multiply the floating point number in the  $-1$  to  $1$  range by  $2^{31}$ .)

*scalingValue*

The scaling value: an exponent of a power of 2 that was used to convert to integer values; for example, if  $2^{31}$  was used to multiply the values, the *scalingValue* is 31. This value is used for normalization.

## Discussion

The function `iplCreateColorTwist()` allocates memory for the data structure `IplColorTwist` and creates the color-twist matrix that can subsequently be used by the function `iplApplyColorTwist()`.

## Return Value

A pointer to the `IplColorTwist` data structure containing the color-twist matrix in the form suitable for efficient computation by the function `iplApplyColorTwist()`.

---

## SetColorTwist

*Sets a color-twist matrix data structure.*

---

```
void iplSetColorTwist(IplColorTwist* cTwist, int
data[16],
int scalingValue);
```

*data*

An array containing the sixteen values that constitute the color-twist matrix. The values are in row-wise order. Color-twist values that are in the range  $-1$  to  $1$  should be scaled up to be in the

range  $-2^{31}$  to  $2^{31}$ . (Simply multiply the floating point number in the  $-1$  to  $1$  range by  $2^{31}$ .)

*scalingValue*

The scaling value: an exponent of a power of 2 that was used to convert to integer values; for example, if  $2^{31}$  was used to multiply the values, the *scalingValue* is 31. This value is used for normalization.

### Discussion

The function `iplSetColorTwist()` is used to set the values of the color-twist matrix in the data structure `IplColorTwist` that can subsequently be used by the function `iplApplyColorTwist()`.

### Return Value

A pointer to the `IplColorTwist` data structure containing the color-twist matrix in the form suitable for efficient computation by the function `iplApplyColorTwist()`.

---

## ApplyColorTwist

*Applies a color-twist matrix to an image.*

---

```
void iplApplyColorTwist(IplImage* srcImage,
                       IplImage* dstImage, IplColorTwist* cTwist, int offset,
                       IplCoord* map);
```

*srcImage*

The source image.

*dstImage*

The resultant image.

<i>cTwist</i>	The color-twist matrix data structure that was prepared by a call to the function <code>iplSetColorTwist()</code> .
<i>offset</i>	An offset value that will be added to each pixel channel after multiplication by the color-twist matrix.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplApplyColorTwist()` applies the color-twist matrix to each of the first three color channels in the input image to obtain the resulting data for the three channels.

For example, the matrix below can be used to convert normalized `PhotoYCC` to normalized `PhotoRGB` (both with an opacity channel) when the channels are in the order YCC and RGB, respectively:

$$\begin{matrix} 2^{31} & 0 & 2^{31} & 0 \\ 2^{31} & X & Y & 0 \\ 2^{31} & 2^{31} & 0 & 0 \\ 0 & 0 & 0 & 2^{31} \end{matrix}$$

where  $X = -416611827$  (that is,  $-0.194 * 2^{31}$ ) and  $Y = -1093069176$  (that is,  $-0.509 * 2^{31}$ ).

Color-twist matrices may also be used to perform many other color conversions and operations such as

- Lightening an image
- Color saturation
- Color balance
- R, G, and B color adjustments
- Contrast Adjustment

---

## DeleteColorTwist

*Frees memory used for  
a color-twist matrix.*

---

```
void iplDeleteColorTwist(IplColorTwist* cTwist);
```

*cTwist*

The color-twist matrix data structure that was prepared by a call to the function `iplCreateColorTwist()`.

### Discussion

The function `iplDeleteColorTwist()` frees memory used for the color-twist matrix structure referred to by *cTwist*.



# *Histogram and Thresholding Functions*

# 10

This chapter describes functions that operate on an image on a pixel-by-pixel basis, in particular, the operations that alter the histogram of the image. In addition, the use of color-twist matrices for color model conversions is described. Table 10-1 lists histogram and thresholding functions in the IPL.

**Table 10-1 Histogram and Thresholding Functions**

<b>Group</b>	<b>Function Name</b>	<b>Description</b>
Thresholding	<a href="#"><u>iplThreshold</u></a>	Performs a simple thresholding of an image.
Lookup Table and Histogram	<a href="#"><u>iplContrastStretch</u></a>	Stretches the contrast of an image using intensity transformation.
	<a href="#"><u>iplComputeHisto</u></a>	Computes the intensity histogram of an image.
	<a href="#"><u>iplHistoEqualize</u></a>	Enhances an image by flattening its intensity histogram.

## Thresholding

The thresholding operation changes pixel values depending on whether they are less than, equal to, or greater than the specified *threshold*. If an input pixel value is less than the *threshold*, the corresponding output pixel is set to the minimum presentable value. Otherwise, it is set to the maximum presentable value.

---

### Threshold

*Performs a simple thresholding of an image.*

---

```
void iplThreshold(IplImage* srcImage, IplImage* dstImage,  
                 int threshold, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>threshold</i>	The threshold value to use for each pixel. The pixel value in the output is set to the maximum presentable value if it is greater than or equal to the threshold value (for each channel). Otherwise the pixel value in the output is set to the minimum presentable value.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplThreshold()` thresholds the source image *srcImage* using the value *threshold* to create the resultant image *dstImage*. The pixel value in the output is set to the maximum presentable value (for

example, 255 for an 8-bit-per-channel image) if it is greater than or equal to the threshold value. Otherwise it is set to the minimum presentable value (for example, 0 for an 8-bit-per-channel image). This is done for each channel in the input image.

To convert an image to bi-tonal, do the following:

- use `iplColorToGray()` to convert a color image to gray scale
- use `iplThreshold()` to threshold the image to max and zero values
- use `iplConvert()` with `IPL_BITS_LOW` flag set to convert to 1U depth.

## Lookup Table (LUT) and Histogram Operations

A LUT can be used to specify an intensity transformation. Given an input intensity, LUT can be used to look up an output intensity. Usually a LUT is provided for each channel in the image, although sometimes the same LUT can be shared by many channels.

### The IplLUT Structure

Example 10-1 presents a C language definition for the `IplLUT` structure to set a LUT.

#### Example 10-1 IplLUT Definition

---

```
typedef struct _IplLUT {
    int      num;    /* number of keys or values */
    int*     key;
    int*     value;
    int*     factor;
    int      interpolateType;
} IplLUT;
```

---

The *key* array has the length *num*; the *value* and *factor* are arrays of the same length *num*-1. The *interpolateType* can be either `IPL_LUT_LOOKUP` or `IPL_LUT_INTER`.

Consider the following example of *num* = 4:

<i>key</i>	<i>value</i>	<i>factor</i>
k1	v1	f1
k2	v2	f2
k3	v3	f3
k4		

If *interpolateType* is `LOOKUP`, then any input intensity *D* in the range  $k1 \leq D < k2$  will result in the value *v1*, in the range  $k2 \leq D < k3$  will result in the value *v2* and so on. If *interpolateType* is `INTER`, then an intensity *D* in the range  $k1 \leq D < k2$  will result in the linearly interpolated value

$$v1 + [(v2 - v1)/(k2 - k1)] * (D - k1)$$

The *value*  $(v2-v1)/(k2-k1)$  is pre-computed and stored in the array *factor* in the `IplLUT` data structure.

The data structure described above can be used to specify a piece-wise linear transformation that is ideal for the purpose of contrast stretching.

The histogram is a data structure that shows how the intensities in the image are distributed. The same data structure `IplLUT` is used for a histogram except that *interpolateType* is always `IPL_LUT_LOOKUP` and *factor* is a `NULL` pointer for a histogram. However, unlike the LUT, the *value* array represents counts of pixels falling in the specified ranges in the *key* array.

The sections that follow describe the functions that use the above data structure.

## ConstrastStretch

*Stretches the contrast of an image using an intensity transformation.*

---

```
void iplContrastStretch(IplImage* srcImage,  
IplImage* dstImage, IplLUT** lut, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>lut</i>	An array of pointers to LUTs, one pointer for each channel. Each lookup table should have the <i>key</i> , <i>value</i> and <i>factor</i> arrays fully initialized (see “ <a href="#">The IplLUT Structure</a> ”). One or more channels may share the same LUT. Specifies an intensity transformation.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplContrastStretch()` stretches the contrast in a color source image *srcImage* by applying intensity transformations specified by LUTs in *lut* to produce an output image *dstImage*. Fully specified LUTs should be provided to this function.

## ComputeHisto

*Computes the intensity histogram of an image.*

---

```
void iplComputeHisto(IplImage* srcImage, IplLUT** lut,  
IplCoord* map);
```

<i>srcImage</i>	The source image for which the histogram will be computed.
<i>lut</i>	An array of pointers to LUTs, one pointer for each channel. Each lookup table should have the <i>key</i> array fully initialized. The <i>value</i> array will be filled by this function. (For the <i>key</i> and <i>value</i> arrays, see “ <a href="#">The IplLUT Structure</a> ” above.) The same LUT can be shared by one or more channels.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplComputeHisto()` computes the intensity histogram of an image. The histograms (one per channel in the image) are stored in the array *lut* containing all the LUTs. The *key* array in each LUT should be initialized before calling this function. The *value* array containing the histogram information will be filled in by this function. (For the *key* and *value* arrays, see “[The IplLUT Structure](#)” above.)

## HistoEqualize

*Enhances an image by flattening its intensity histogram.*

---

```
void iplHistoEqualize(IplImage* srcImage,  
IplImage* dstImage, IplLUT** lut, IplCoord* map);
```

<i>srcImage</i>	The source image for which the histogram will be computed.
<i>dstImage</i>	The resultant image after equalizing.
<i>lut</i>	The histogram of the image is represented as an array of pointers to LUTs, one pointer for each channel. Each lookup table should have the <i>key</i> and <i>value</i> arrays fully initialized. (For the <i>key</i> and <i>value</i> arrays, see “ <a href="#">The IplLUT Structure</a> ” above.) These LUTs will contain flattened histograms after this function is executed. In other words, the call of <code>iplHistoEqualize()</code> is destructive with respect to the LUTs.
<i>map</i>	The structure specifying offsets for tiling purposes; see <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplHistoEqualize()` enhances the source image *srcImage* by flattening its histogram represented by *lut* and places the enhanced image in the output image *dstImage*. After execution, *lut* points to the flattened histogram of the output image.

## Image Geometric Transforms

This chapter describes the IPL functions that perform geometric transforms to resize the image or change its orientation. The geometric transforms are performed by resampling (“Zoom,” “Decimate,” and “Rotate”) or flipping the axis of the image (“Mirror”). Table 11-1 lists image geometric transform functions.

**Table 11-1 Image Geometric Transform Functions**

Group	Function Name	Description
Resizing	<a href="#">iplZoom</a>	Zooms or expands an image.
	<a href="#">iplDecimate</a>	Decimates or shrinks an image.
Changing Orientation	<a href="#">iplMirror</a>	Mirrors an image about a horizontal or vertical axis.
	<a href="#">iplRotate</a>	Rotates an image.

### Changing Image Size

The functions that expand or shrink an image perform image resampling by using various kinds of interpolation: nearest neighbor, linear, or cubic convolution.



## Zoom

*Zooms or expands an image.*

```
void iplZoom(IplImage* srcImage, IplImage* dstImage,
int xDst, int xSrc, int yDst, int ySrc, int interpolate,
IplCoord* map);
```

*srcImage*                   The source image.

*dstImage*                   The resultant image.

*xSrc, ySrc*                 X and Y dimensions of the source image.

*xDst, yDst*                 X and Y dimensions of the destination image.

These four integers must be positive, meeting the conditions of  $xDst \geq xSrc$  and  $yDst \geq ySrc$  to specify the fractions  $xDst/xSrc$  and  $yDst/ySrc$ . These fractions indicate the value to magnify the image in the X and Y directions. For example,

$xDst = 2$ ,  $xSrc = 1$ ,  $yDst = 2$ ,  $ySrc = 1$  doubles the image size in each dimension to give an image 4 times larger in area.

*interpolate*               The type of interpolation to perform for resampling. The following are currently supported:

**IPL\_INTER\_NN**            Nearest neighbor interpolation.

**IPL\_INTER\_LINEAR**      Linear interpolation.

**IPL\_INTER\_CUBIC**       Cubic convolution interpolation.

*map* The structure specifying offsets for tiling purposes. See [IplCoord Structure](#) in Chapter 2.

## Discussion

The function `iplZoom()` zooms or expands the source image *srcImage* by *xDst/xSrc* in the X direction and *yDst/ySrc* in the Y direction. The interpolation specified by *interpolate* is used during resampling the input image.

---

## Decimate

*Decimates or shrinks an image.*

---

```
void iplDecimate(IplImage* srcImage, IplImage* dstImage,
int xDst, int xSrc, int yDst, int ySrc, int interpolate,
IplCoord* map);
```

*srcImage* The source image.

*dstImage* The resultant image.

*xSrc, ySrc* X and Y dimensions of the source image.

*xDst, yDst* X and Y dimensions of the destination image.

These four integers should be positive, meeting the conditions  $xDst \leq xSrc$  and  $yDst \leq ySrc$  to specify the fractions  $xDst/xSrc$  and  $yDst/ySrc$ . These fractions indicate the value to shrink the image in the X and Y directions. For example,  $xDst = 1$ ,  $xSrc = 2$ ,  $yDst = 1$ ,  $ySrc = 2$  halves the image size in each dimension to give an image 1/4 times smaller in area.

<i>interpolate</i>	The type of interpolation to perform for resampling. The following are currently supported: <ul style="list-style-type: none"><li><code>IPL_INTER_NN</code> Nearest neighbor interpolation.</li><li><code>IPL_INTER_LINEAR</code> Linear interpolation.</li><li><code>IPL_INTER_CUBIC</code> Cubic convolution interpolation.</li></ul>
<i>map</i>	The structure specifying offsets for tiling purposes. See <a href="#">IplCoord Structure</a> in Chapter 2.

## Discussion

The function `iplDecimate()` decimates or shrinks the source image *srcImage* by *xDst/xSrc* in the X direction and *yDst/ySrc* in the Y direction. The interpolation specified by *interpolate* is used during resampling the input image.

## Changing Image Orientation

The functions described in this section change image orientation by rotating or mirroring the source image. Rotation involves image sampling by using various kinds of interpolation: nearest neighbor, linear, or cubic convolution. Mirroring is performed by flipping the image axis in horizontal or vertical direction.

## Rotate

*Rotates an image.*

---

```
void iplRotate(IplImage* srcImage, IplImage* dstImage,  
int angle, int centerX, int centerY, int interpolate,  
IplCoord* map);
```

<i>srcImage</i>	The source image.						
<i>dstImage</i>	The resultant image.						
<i>angle</i>	The angle in hundredths of degree to rotate the image (for example, 6000 for 60 degrees). The image is rotated about the center specified as <i>centerX</i> and <i>centerY</i> coordinates.						
<i>centerX, centerY</i>	The coordinates of the rotation center.						
<i>interpolate</i>	The type of interpolation to perform for resampling. The following are currently supported:  <table><tr><td><code>IPL_INTER_NN</code></td><td>Nearest neighbor interpolation.</td></tr><tr><td><code>IPL_INTER_LINEAR</code></td><td>Linear interpolation.</td></tr><tr><td><code>IPL_INTER_CUBIC</code></td><td>Cubic convolution interpolation.</td></tr></table>	<code>IPL_INTER_NN</code>	Nearest neighbor interpolation.	<code>IPL_INTER_LINEAR</code>	Linear interpolation.	<code>IPL_INTER_CUBIC</code>	Cubic convolution interpolation.
<code>IPL_INTER_NN</code>	Nearest neighbor interpolation.						
<code>IPL_INTER_LINEAR</code>	Linear interpolation.						
<code>IPL_INTER_CUBIC</code>	Cubic convolution interpolation.						
<i>map</i>	The structure specifying offsets for tiling purposes. See <a href="#">IplCoord Structure</a> in Chapter 2.						

### Discussion

The function `iplRotate()` rotates the source image *srcImage* by *angle* degrees around the origin defined by the coordinates *centerX* and *centerY*. The interpolation specified by *interpolate* is used during resampling the input image.

## Mirror

*Mirrors an image about a horizontal or vertical axis.*

---

```
void iplMirror(IplImage* srcImage, IplImage* dstImage,  
int flipAxis, IplCoord* map);
```

<i>srcImage</i>	The source image.
<i>dstImage</i>	The resultant image.
<i>flipAxis</i>	Specifies the axis to mirror the image. Use 0 for the horizontal axis, 1 for a vertical axis and -1 for both horizontal and vertical axes.
<i>map</i>	The structure specifying offsets for tiling purposes. See <a href="#">IplCoord Structure</a> in Chapter 2.

### Discussion

The function `iplMirror()` mirrors or flips the source image *srcImage* about a horizontal or vertical axis or both.

# Supported Image Attributes and Operation Modes



This appendix contains tables that list the supported image attributes and operation modes for all IPL functions that have input and/or output images. The *ipl* prefixes in the function names are omitted.

**Table A-1 Image Attributes and Modes of Data Exchange Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<u>Set</u>	all <sup>†</sup>	always in-place				x		x	x
<u>Copy</u>	all	x	x	x	x	x		x	x
<u>Exchange</u>	all	x	x	x	x	x		x	x
<u>Convert</u>	all								x

<sup>†</sup> all = 1u, 8s, 8u, 16s, 16u, and 32s bits per channel

**Table A-2 Windows\* DIB Conversion Functions**

Function	Depths		Input & output images have the same			Remarks
	input	output	order	origin	# of channels	
<u>ConvertFromDIB</u>	all <sup>‡</sup>	1u,8u,16u	x	x		Rectangular ROI,
<u>ConvertToDIB</u>	1u,8u,16u	all <sup>‡</sup>	x	x	x	border mode and
<u>TranslateDIB</u>	1bpp	1u	x		x	tiling are not
	other <sup>‡</sup>	8u	x		x	supported

<sup>‡</sup> all = 1, 4, 8, 16, 24, and 32 bits per pixel (DIB images).  
other = 4, 8, 16, 24, and 32 bits per pixel (DIB images).

# A

For `iplConvertFromDIB`, the number of channels, bit depth per channel and the dimensions of the IPL image should be greater than or equal to those of the DIB image. When converting a DIB RGBA image, the IPL image should also contain an alpha channel.

**Table A-3 Image Attributes and Modes of Arithmetic and Logical Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<a href="#">AddS</a>	all <sup>†</sup>	x	x	x	x	x		x	x
<a href="#">SubtractS</a>	all	x	x	x	x	x		x	x
<a href="#">MultiplyS</a>	all	x	x	x	x	x		x	x
<a href="#">MultiplySScale</a>	8u,16u	x	x	x	x	x			x
<a href="#">Square</a>	all	x	x	x	x	x		x	x
<a href="#">Add</a>	all	x	x	x	x	x		x	x
<a href="#">Subtract</a>	all	x	x	x	x	x		x	x
<a href="#">Multiply</a>	all	x	x	x	x	x		x	x
<a href="#">MultiplyScale</a>	8u,16u	x	x	x	x	x			x
<a href="#">LShiftS</a>	all	x	x	x	x	x		x	x
<a href="#">RShiftS</a>	all	x	x	x	x	x		x	x
<a href="#">Not</a>	all	x	x	x	x	x		x	x
<a href="#">AndS</a>	all	x	x	x	x	x		x	x
<a href="#">OrS</a>	all	x	x	x	x	x		x	x
<a href="#">XorS</a>	all	x	x	x	x	x		x	x
<a href="#">And</a>	all	x	x	x	x	x		x	x
<a href="#">Or</a>	all	x	x	x	x	x		x	x
<a href="#">Xor</a>	all	x	x	x	x	x		x	x

<sup>†</sup> all = 1u, 8s, 8u, 16s, 16u, and 32s bits per channel

**Table A-4 Image Attributes and Modes of Alpha-Blending Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<a href="#">PreMultiplyAlpha</a>	8u,16u	x	x	x	x	x			x
<a href="#">AlphaComposite</a>	8u,16u	x	x	x	x	x			x
<a href="#">AlphaCompositeC</a>	8u,16u	x	x	x	x	x			x

**Table A-5 Image Attributes and Modes of Filtering Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<a href="#">Blur</a>	>1bpc	x	x	x	x	x	x	x	x
<a href="#">Convolve2D</a>	all <sup>†</sup>	x	x	x	x	x	x	x	x
<a href="#">ConvolveSep2D</a>	all	x	x	x	x	x	x		x
<a href="#">MaxFilter</a>	all	x	x	x	x	x	x		x
<a href="#">MinFilter</a>	all	x	x	x	x	x	x		x
<a href="#">MedianFilter</a>	all	x	x	x	x	x	x		x

<sup>†</sup> all = 1u, 8s, 8u, 16s, 16u, and 32s bits per channel

**Table A-6 Image Attributes and Modes of Fourier and DCT Functions**

Function	Depths		Input & output images have the same			Rect. ROI	Border Mode	In-place	Tiling
	input	output	order	origin	COI				
<a href="#">DCT2D</a>	>1bpc	16s,32s	x	x		x			
<a href="#">RealFft2D</a>	8u,16u	16s,32s	x	x	x	x			
<a href="#">CcsFft2D</a>	16s,32s	8u,16u	x		x	x			



**Table A-7 Image Attributes and Modes of Morphological Operations**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
s u p p o r t e d ( x )									
<a href="#">Erode</a>	1u,8u,16u	x	x	x	x	x	x		x
<a href="#">Dilate</a>	1u,8u,16u	x	x	x	x	x	x		x
<a href="#">Open</a>	1u,8u,16u	x	x	x	x	x	x		x
<a href="#">Close</a>	1u,8u,16u	x	x	x	x	x	x		x

**Table A-8 Image Attributes and Modes of Color Space Conversion Functions**

Function	Depths		Input & output images have the same				Rect. ROI	Bord. Mode	In-place	Tiling
	input	output	depth	order	origin	COI				
s u p p o r t e d ( x )										
<a href="#">ReduceBits</a>	32s	8u,16u		x	x	x				
	16u	8u		x	x	x				
<a href="#">BitonalToGray</a>	1u	>1bpc								x
<a href="#">RGB2HSV</a>	1u,16u,32s		x	x	x	x				x
<a href="#">HSV2RGB</a>	1u,16u,32s		x	x	x	x				x
<a href="#">RGB2HLS</a>	1u,16u,32s		x	x	x	x				x
<a href="#">HLS2RGB</a>	1u, 16u,32s		x	x	x	x				x
<a href="#">ApplyColorTwist</a>	16u,32s		x	x	x	x	x			x

**Table A-9 Image Attributes and Modes of Histogram and Thresholding Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<a href="#">Threshold</a>	8u,8s,16u,16s,32s	x	x	x	x	x	x	x	
<a href="#">ComputeHisto</a>	1u,8u,16u	no output image				x	x	x	
<a href="#">HistoEqualize</a>	8u,16u	x	x	x	x	x	x		
<a href="#">ContrastStretch</a>	8u,16u	x	x	x	x	x	x		

**Table A-10 Image Attributes and Modes of Geometric Transform Functions**

Function	Depths	Input and output images must have the same				Rect. ROI	Border Mode	In-place	Tiling
		depth	order	origin	COI				
<a href="#">Mirror</a>	1u,8u,16u	x	x	x	x	x	x		
<a href="#">Rotate</a>	1u,8u,16u	x	x	x	x	x	x		
<a href="#">Zoom</a>	1u,8u,16u	x	x	x	x	x	x		
<a href="#">Decimate</a>	1u,8u,16u	x	x	x	x	x	x		

# Bibliography

---

This bibliography provides a list of publications that might be useful to the Image Processing Library users. This list is not complete; it serves only as a starting point. The books [Rogers85], [Rogers90], and [Foley90] are good resources of information on image processing and computer graphics, with mathematical formulas and code examples.

The Image Processing Library is part of Intel Performance Libraries Suite. The manuals [RPL96] and [SPL96] describe Intel Recognition Primitives Library and Intel Signal Processing Library, which are other parts of the Performance Libraries Suite.

- [Bragg] Dennis Bragg. A simple color reduction filter, *Graphic Gems III*: 20–22.
- [Foley90] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics — Principles and Practice*, Second Edition. Addison Wesley, 1990.
- [Rogers85] David Rogers. *Procedural Elements for Computer Graphics*. McGraw-Hill, 1985.
- [Rogers90] David Rogers and J.Alan Adams. *Mathematical Elements for Computer Graphics*. McGraw-Hill, 1990.
- [RPL96] *Intel Recognition Primitives Library Reference Manual*. Intel Corp. Order number 637785, Rev.4, 1996.
- [SPL96] *Intel Signal Processing Library Reference Manual*. Intel Corp. Order number 630508, Rev.6, 1996.
- [Schumacher] Dale A. Schumacher. A comparison of digital halftoning techniques, *Graphic Gems III*: 57–71.
- [Thomas] Spencer W. Thomas and Rod G. Bogart. Color dithering, *Graphic Gems II*: 72–77.

# Glossary

---

absolute colors	Colors specified by each pixel's coordinates in a color space. IPL functions use images with absolute colors. <i>See</i> palette colors.
alpha channel	A color channel, also known as the opacity channel, that can be used in color models; for example, the RGBA model.
arithmetic operation	An operation that adds, subtracts, multiplies, shifts, or squares the image pixel values.
CCS	<i>See</i> complex conjugate-symmetric.
channel of interest	The color channel on which an IPL operation acts (or processing occurs). Channel of interest (COI) can be considered as a separate case of region of interest (ROI).
CMY	Cyan-magenta-yellow. A three-channel color model that uses cyan, magenta, and yellow color channels.
CMYK	Cyan-magenta-yellow-black. A four-channel color model that uses cyan, magenta, yellow, and black color channels.
COI	<i>See</i> channel of interest.
color-twist matrix	A matrix used to multiply the pixel coordinates in one color space for determining the coordinates in another color space.
complex conjugate-symmetric	A kind of symmetry that arises in the Fourier transform of real signals. A complex conjugate-symmetric signal has the property that $x(-n) = x(n)^*$ , where * denotes conjugation.

conjugate	The conjugate of a complex number $a+bj$ is $a-bj$ .
conjugate-symmetric	See complex conjugate-symmetric.
DCT	Acronym for the discrete cosine transform. See “ <a href="#">Discrete Cosine Transform</a> ” in Chapter 7.
decimation	An IPL geometric transform operation that shrinks the source image.
DIB	Device-independent bitmap, an image format used by IPL in Windows* environment.
dilation	An IPL morphological operation whose effect is to fill up holes and thicken object boundaries.
dyadic operation	An operation that has two input images. It can have other input parameters as well.
erosion	An IPL morphological operation that results in less noise and thinner object boundaries.
FFT	Acronym for the fast Fourier transform. See “ <a href="#">Fast Fourier Transform</a> ” in Chapter 7.
four-channel model	A color model that uses four color channels; for example, the RGBA color model.
geometric transform functions	IPL functions that perform geometric transformations of images: zoom, decimation, rotation, and mirror functions.
gray scale image	An image characterized by a single intensity channel so that each intensity value corresponds to a certain shade of gray.
HLS	Hue-lightness-saturation. A three-channel color model that uses hue, lightness, and saturation channels. The HLS and HSV models differ in the way of scaling the image luminance. See HSV.

HSV	Hue-saturation-value. A three-channel color model that uses hue, saturation, and value channels. HSV is often used as a synonym for the HSB (hue-saturation-brightness) and HSI (hue-saturation-intensity) models. <i>See</i> HLS.
hue	A color channel in several color models that measures the “angular” distance (in degrees) from red to the particular color: 60 corresponds to yellow, 120 to green, 180 to cyan, 240 to blue, and 300 to magenta. Hue is undefined for shades of gray.
in-place operation	An operation whose output image is one of the input images. <i>See</i> out-of-place operation.
linear filtering	In IPL, either neighborhood averaging (blur) or 2D convolution operations.
linear image transforms	In IPL, either the fast Fourier transform (FFT) or the discrete cosine transform (DCT).
luminance	A measure of image intensity, as perceived by a “standard observer”. Since human eyes are more sensitive to green and less to red or blue, different colors of equal physical intensity make different contribution to luminance. <i>See</i> <a href="#">ColorToGray</a> in Chapter 9.
LUT	Acronym for lookup table (palette).
MMX™ technology	A major enhancement to the Intel Architecture aimed at better performance in multimedia and communications applications. The technology uses four new data types, eight 64-bit MMX registers, and 57 new instructions implementing the SIMD (single instruction, multiple data) technique.
monadic operation	An operation that has a single input image. It can have other input parameters as well.
morphological operation	In IPL, simple erosion or dilation of an image.

MSI	Acronym for multi-spectral image. An MSI can use any number of channels and colors.
non-linear filtering	In IPL, minimum, maximum, or median filtering operation.
opacity channel	See alpha channel.
out-of-place operation	An operation whose output is an image other than the input image(s). See in-place operation.
palette colors	Colors specified by a palette, or lookup table. IPL uses palette colors only in operations of image conversion to and from absolute colors. See absolute colors.
PhotoYCC*	A Kodak* proprietary color encoding and image compression scheme. See YCC.
pixel depth	The number of bits determining a single pixel in the image.
pixel-oriented ordering	Storing the image information in such an order that the values of all color channels for each pixel are clustered; for example, RGBRGB... . See " <a href="#">Channel Sequence</a> " in Chapter 2.
plane-oriented ordering	Storing the image information so that all data of one color channel follow all data of another channel, thus forming a separate "plane" for each channel; for example, RRRRRGGGGG...
point operation	An operation performed on a pixel-by-pixel basis. IPL point operations include applying a color-twist matrix, computing and altering the image histogram, contrast stretching, histogram equalization and thresholding.
region of interest	An image region on which an IPL operation acts (or processing occurs).
RGB	Red-green-blue. A three-channel color model that uses red, green, and blue color channels.
RGBA	Red-green-blue-alpha. A four-channel color model that uses red, green, blue, and alpha (or

	opacity) channels.
ROI	<i>See</i> region of interest.
saturation	A quantity used for measuring the purity of colors. The maximum saturation corresponds to the highest degree of color purity; the minimum (zero) saturation corresponds to shades of gray.
scanline	All image data for one row of pixels.
standard gray palette	A complete palette of a DIB image whose red, green, and blue values are equal for each entry and monotonically increasing from entry to entry.
three-channel model	A color model that uses three color channels; for example, the CMY color model.
YCC	A three-channel color model that uses one luminance channel (Y) and two chroma channels (usually denoted by $C_R$ and $C_B$ ). The term is sometimes used as a synonym for the entire PhotoYCC encoding scheme. <i>See</i> PhotoYCC.
zoom	A geometric transform function that magnifies the source image.



# Index

---

## A

a function that helps you

- add a constant to pixel values, 5-3

- add pixel values of two images, 5-7

- allocate a quadword-aligned memory block, 4-19

- allocate image data, 4-10

- allocate memory for 16-bit words, 4-20

- allocate memory for 32-bit double words, 4-21

- allocate memory for double floating-point elements, 4-22

- allocate memory for floating-point elements, 4-21

- apply a color twist matrix, 9-15

- assign a new error-handling function, 3-6

- average neighboring pixels, 6-2

- change the image orientation, 11-4

- change the image size, 11-1

- composite images using the alpha channel, 5-18

- compute bitwise AND of pixel values and a constant, 5-13

- compute bitwise AND of pixel values of two images, 5-16

- compute bitwise NOT of pixel values, 5-12

- compute bitwise OR of pixel values and a constant, 5-14

- compute bitwise OR of pixel values of two images, 5-16

- compute bitwise XOR of pixel values and a constant, 5-15

- compute bitwise XOR of pixel values of two images, 5-17

- compute CCS fast Fourier transform, 7-5

- compute discrete cosine transform, 7-7

- compute real fast Fourier transform, 7-3

- compute the image histogram, 10-6

- convert a bitonal image to gray scale, 9-4

- convert a color image to gray scale, 9-6

- convert a gray scale image to color, 9-7

- convert an HLS image to RGB, 9-11

- convert an HSV image to RGB, 9-10

- convert an RGB image to HLS, 9-10

- convert an RGB image to HSV, 9-9

- convert images from DIB (changing attributes), 4-32

- convert images from DIB (preserving attributes), 4-30

- convert images to DIB, 4-33

- convert images to different data formats, 4-27

- convolve an image with 2D kernel, 6-7

- convolve an image with a separable kernel, 6-8

- copy image data, 4-25
- create 2D convolution kernel, 6-4
- create a color twist matrix, 9-12
- create a region of interest (ROI), 4-13
- create image header, 4-8
- create the IplTileInfo structure, 4-17
- decimate the image, 11-3
- delete 2D convolution kernel, 6-6
- delete a color twist matrix, 9-16
- delete the IplTileInfo structure, 4-18
- divide pixel values by  $2^N$ , 5-12
- equalize the image histogram, 10-7
- exchange data of two images, 4-26
- fill image's pixels with a value, 4-24
- free memory allocated by Malloc functions, 4-23
- free the image data memory, 4-11
- free the image header memory, 4-11
- get error-handling mode, 3-4
- get the error status code, 3-3
- handle an error, 3-2
- magnify the image, 11-2
- mirror the image, 11-6
- multiply pixel values by a constant, 5-4
- multiply pixel values by a constant and scale the products, 5-5
- multiply pixel values of two images, 5-9
- multiply pixel values of two images and scale the products, 5-10
- pre-multiply pixel values by alpha values, 5-25
- produce error messages for users, 3-5
- read convolution kernel's attributes, 6-5
- reduce the image bit resolution, 9-2
- report an error, 3-2
- rotate the image, 11-5
- set a color twist matrix, 9-13
- set a region of interest (ROI), 4-14
- set error-handling mode, 3-4
- set pixels to the maximum value of the neighbors, 6-10
- set pixels to the median value of the neighbors, 6-9
- set pixels to the minimum value of the neighbors, 6-11
- set the error status code, 3-3
- set the image border mode, 4-15
- set the IplTileInfo structure fields, 4-18
- shift pixel bits to the left, 5-11
- shift pixel bits to the right, 5-12
- shrink the image, 11-3
- square pixel values, 5-6
- stretch the image contrast, 10-5
- subtract pixel values from a constant, 5-4
- subtract pixel values of two images, 5-8
- threshold the source image, 10-2
- zoom the image, 11-2
- about this manual, 1-2
- about this software, 1-1
- absolute color image, 2-2
- add a constant to pixel values, 5-3
- Add function, 5-7
- add pixels of two images, 5-7
- AddS function, 5-3
- align the rectangular ROIs, 2-5
- aligning
  - image data, 2-6
  - rectangular ROIs, 2-5

- scanlines, 2-6
  - allocate a quadword-aligned memory block, 4-19
  - allocate memory for 16-bit words, 4-20
  - allocate memory for 32-bit double words, 4-21
  - allocate memory for double floating-point elements, 4-22
  - allocate memory for floating-point elements, 4-21
  - AllocateImage function, 4-10
  - alpha channel, 2-6
  - alpha pre-multiplication, 5-25
  - alpha-blending, 5-18
    - alpha pre-multiplication, 5-25
    - AlphaComposite function, 5-19
    - AlphaCompositeC function, 5-19
    - ATOP operation, 5-23
    - IN operation, 5-23
    - OUT operation, 5-23
    - OVER operation, 5-23
    - PLUS operation, 5-23
    - PreMultiplyAlpha function, 5-25
    - XOR operation, 5-23
  - AlphaComposite function, 5-19
  - AlphaCompositeC function, 5-19
  - And function, 5-16
  - AndS function, 5-13
  - apply a color twist matrix, 9-15
  - ApplyColorTwist function, 9-14
  - arithmetic operations, 5-1
    - Add, 5-7
    - AddS, 5-3
    - AlphaComposite, 5-19
    - AlphaCompositeC, 5-19
    - Multiply, 5-9
    - MultiplyS, 5-4
    - MultiplyScale, 5-10
    - MultiplySScale, 5-5
    - PreMultiplyAlpha, 5-25
    - Square, 5-6
    - Subtract, 5-8
    - SubtractS, 5-4
  - ATOP compositing operation, 5-23
  - attribute values, 4-3
  - attributes, 4-3
  - audience for this manual, 1-4
  - averaging the neighboring pixels, 6-2
- B**
- bit depths supported, A-1
  - BitonalToGray function, 9-4
  - bitwise AND
    - with a constant, 5-13
    - with another image, 5-16
  - bitwise NOT, 5-12
  - bitwise OR
    - with a constant, 5-14
    - with another image, 5-16
  - bitwise XOR
    - with a constant, 5-15
    - with another image, 5-17
  - Blur function, 6-2
  - brightening the image, 5-3
- C**
- call-backs, 2-9

- CcsFft2D function, 7-5
  - changing the image orientation, 11-4
  - changing the image size, 11-1
  - channel sequence, 2-3
  - channel(s) of interest, 2-4
  - Close function, 8-7
  - COI. *See* channel(s) of interest
  - color data order, 2-3
  - color models, 2-1
    - gray scale, 2-1
    - multi-spectral image, 2-2
    - three or four channels, 2-1
  - color space conversion functions
    - ApplyColorTwist, 9-14
    - BitonalToGray, 9-4
    - ColorToGray, 9-6
    - CreateColorTwist, 9-12
    - DeleteColorTwist, 9-16
    - GrayToColor, 9-7
    - HLS2RGB, 9-11
    - HSV2RGB, 9-10
    - ReduceBits, 9-2
    - RGB2HLS, 9-10
    - RGB2HSV, 9-9
    - SetColorTwist, 9-13
  - ColorToGray function, 9-6
  - compute the image histogram, 10-6
  - ComputeHisto function, 10-6
  - ContrastStretch function, 10-5
  - conventions
    - font, 1-5
    - function names, 1-6
    - naming, 1-5
  - convert a bitonal image to gray scale, 9-4
  - convert a color image to gray scale, 9-6
  - convert a gray scale image to color, 9-7
  - convert an HLS image to RGB, 9-11
  - convert an HSV image to RGB, 9-10
  - convert an RGB image to HLS, 9-10
  - convert an RGB image to HSV, 9-9
  - Convert function, 4-27
  - convert images from DIB (changing attributes), 4-32
  - convert images from DIB (preserving attributes), 4-30
  - convert images to DIB, 4-33
  - convert images to different data formats, 4-27
  - ConvertFromDIB function, 4-32
  - ConvertToDIB function, 4-33
  - convolution, 6-3
  - Convolve2D function, 6-7
  - ConvolveSep2D function, 6-8
  - coordinate systems, 2-4
  - Copy function, 4-24
  - create a color twist matrix, 9-12
  - CreateColorTwist function, 9-12
  - CreateConvKernel function, 6-4
  - CreateImageHeader function, 4-8
  - CreateROI function, 4-13
  - CreateTileInfo function, 4-17
  - creating images, 4-1
- D**
- darkening the image, 5-3
  - data architecture, 2-1
  - data exchange, 4-2

data exchange functions, 4-24  
    Convert, 4-27  
    Copy, 4-24  
    Exchange, 4-25  
    Set, 4-24  
data ordering, 2-3  
data types, 2-2  
DCT. *See* discrete cosine transform  
DCT2D function, 7-7  
Deallocate function, 4-11  
DeallocateImage function, 4-11  
Decimate function, 11-3  
decimating the image, 11-3  
delete a color twist matrix, 9-16  
DeleteColorTwist function, 9-16  
DeleteConvKernel function, 6-6  
DeleteTileInfo function, 4-18  
device-independent bitmap, 4-2  
DIB. *See* device-independent bitmap  
DIB palette image, 2-2  
Dilate function, 8-5  
discrete cosine transform, 7-6  
divide pixel values by  $2^N$ , 5-12  
dMalloc function, 4-22  
dyadic operations, 5-1

## E

equalize the image histogram, 10-7  
erode an image, 8-2  
Erode function, 8-2  
ErrModeLeaf error mode, 3-4  
ErrModeParent error mode, 3-5

ErrModeSilent error mode, 3-5  
error checks, 3-1  
Error function, 3-2  
error handling, 3-1  
    example, 3-10  
    status codes, 3-8  
    user-defined error handler, 3-13  
error handling macros, 3-7  
error processing modes  
    IPL\_ErrModeLeaf, 3-4  
    IPL\_ErrModeParent, 3-5  
    IPL\_ErrModeSilent, 3-5  
error-handling functions, 3-2  
    Error, 3-2  
    ErrorStr, 3-5  
    GetErrMode, 3-4  
    GetErrStatus, 3-3  
    RedirectError, 3-6  
    SetErrMode, 3-4  
    SetErrStatus, 3-3  
ErrorStr function, 3-5  
Exchange function, 4-25  
execution architecture, 2-7  
    in-place and out-of-place operations, 2-7  
    overflow and underflow, 2-7  
    saturation mode, 2-7

## F

fast Fourier and discrete cosine transforms  
    CcsFft2D, 7-5  
    DCT2D, 7-7  
    RealFft2D, 7-3

fast Fourier transform, 7-1  
FFT. *See* fast Fourier transform  
fill image's pixels with a value, 4-24  
filtering functions, 6-1

- Blur, 6-2
- Convolve2D, 6-7
- ConvolveSep2D, 6-8
- CreateConvKernel, 6-4
- DeleteConvKernel, 6-6
- GetConvKernel, 6-5
- MaxFilter, 6-10
- MedianFilter, 6-9
- MinFilter, 6-11

font conventions, 1-5  
Free function, 4-23  
free memory allocated by Malloc functions, 4-23  
function descriptions, 1-4  
function name conventions, 1-6

## G

geometric transform functions

- Decimate, 11-3
- Mirror, 11-6
- Rotate, 11-5
- Zoom, 11-2

get the error processing mode, 3-4  
get the error status code, 3-3  
GetConvKernel function, 6-5  
GetErrMode function, 3-4  
GetErrStatus function, 3-3  
global variables

- IplLastStatus, 3-5

gray scale, 2-1  
GrayToColor function, 9-7

## H

handling overflow and underflow, 2-7  
hardware and software requirements, 1-1  
HistoEqualize function, 10-7  
histogram and thresholding functions

- ComputeHisto, 10-6
- ContrastStretch, 10-5
- HistoEqualize, 10-7

histogram and thresholding functions, 10-1

- Threshold, 10-2

histogram of an image, 10-6  
histogram operations, 10-3  
HLS2RGB function, 9-11  
HSV2RGB function, 9-10

## I

image attributes, A-1  
image compositing, 5-18

- alpha pre-multiplication, 5-25
- AlphaComposite function, 5-19
- AlphaCompositeC function, 5-19
- ATOP operation, 5-23
- IN operation, 5-23
- OUT operation, 5-23
- OVER operation, 5-18, 5-23
- PLUS operation, 5-23
- PreMultiplyAlpha function, 5-25
- XOR operation, 5-23

Image Creation function, 4-8

- image creation functions, 4-1
  - AllocateImage, 4-10
  - CreateImageHeader, 4-8
  - CreateROI, 4-13
  - CreateTileInfo, 4-17
  - Deallocate, 4-11
  - DeallocateImage, 4-11
  - DeleteTileInfo, 4-18
  - SetBorderMode, 4-15
  - SetROI, 4-14
  - SetTileInfo, 4-18
- image dimensions, 2-7
- image filtering functions, 6-1
- image histogram, 10-6
- image processing operation
  - conditions, 2-6
- image row data, 2-6
- image size, 2-7
- image tiling, 2-8, 4-7
  - call-backs, 2-9
  - IplCoord structure, 2-9
  - IplTileInfo structure, 4-7
- iMalloc function, 4-21
- IN compositing operation, 5-23
- in-place operation, 2-7
- invisible colors, 2-2
- IPL functionality
  - 2D convolution, 6-3
  - alpha-blending, 5-1
  - arithmetic operations, 5-1
  - color space conversion, 9-1
  - data exchange, 4-1
  - DIB environment functions, 4-1, 4-28
  - discrete cosine transform, 7-6
  - error handling, 3-1
  - fast Fourier transform, 7-1
  - filtering functions, 6-1
  - geometric transform functions, 11-1
  - histogram and thresholding functions, 10-1
  - image compositing, 5-18
  - image creation, 4-1
  - image tiling, 2-8, 4-7
  - logical operations, 5-1
  - memory allocation, 4-19
  - morphological operations, 8-1
  - supported image attributes and modes, A-1
- IPL image
  - attribute values, 4-3
  - borders and tiling, 4-15
  - channel sequence, 2-3
  - color models, 2-1
  - coordinate systems, 2-4
  - data architecture, 2-1
  - data ordering, 2-3
  - data types, 2-2
  - format, 4-3
  - header, 4-3
  - header attributes, 4-3
  - regions of interest, 2-4
  - size, 2-7
  - tile size, 2-8
  - tiling, 2-8, 4-7
- IPL\_ErrModeLeaf, 3-4
- IPL\_ErrModeParent, 3-5
- IPL\_ErrModeSilent, 3-5
- iplAdd, 5-7

iplAddS, 5-3  
iplAllocateImage, 4-10  
iplAlphaComposite, 5-19  
iplAlphaCompositeC, 5-19  
iplAnd, 5-16  
iplAndS, 5-13  
iplApplyColorTwist, 9-14  
iplBitonalToGray, 9-4  
iplBlur, 6-2  
iplCcsFft2D, 7-5  
iplClose, 8-7  
iplColorToGray, 9-6  
iplComputeHisto, 10-6  
iplContrastStretch, 10-5  
iplConvert, 4-27  
iplConvertFromDIB, 4-32  
iplConvertToDIB, 4-33  
iplConvolve2D, 6-7  
iplConvolveSep2D, 6-8  
IplCoord structure, 2-9  
iplCopy, 4-24  
iplCreateColorTwist, 9-12  
iplCreateConvKernel, 6-4  
iplCreateImageHeader, 4-8  
iplCreateROI, 4-13  
iplCreateTileInfo, 4-17  
iplDCT2D, 7-7  
iplDeallocate, 4-11  
iplDeallocateImage, 4-11  
iplDecimate, 11-3  
iplDeleteColorTwist, 9-16  
iplDeleteConvKernel, 6-6  
iplDeleteTileInfo, 4-18  
iplDilate, 8-5  
ipldMalloc, 4-22  
iplErode, 8-2  
iplError, 3-2  
iplErrorStr, 3-5  
iplExchange, 4-25  
iplFree, 4-23  
iplGetConvKernel, 6-5  
iplGetErrMode, 3-4  
iplGetErrStatus, 3-3  
iplGrayToColor, 9-7  
iplHistoEqualize, 10-7  
iplHLS2RGB, 9-11  
iplHSV2RGB, 9-10  
IplImage structure, 4-6  
ipliMalloc, 4-21  
IplLastStatus global variable, 3-5  
iplLShiftS, 5-11  
iplMalloc, 4-19  
iplMaxFilter, 6-10  
iplMedianFilter, 6-9  
iplMinFilter, 6-11  
iplMirror, 11-6  
iplMultiply, 5-9  
iplMultiplyS, 5-4  
iplMultiplyScale, 5-10  
iplMultiplySScale, 5-5  
iplNot, 5-12  
iplOpen, 8-6  
iplOr, 5-16  
iplOrS, 5-14  
iplPreMultiplyAlpha, 5-25  
iplRealFft2D, 7-3



iplRedirectError, 3-6  
iplReduceBits, 9-2  
iplRGB2HLS, 9-10  
iplRGB2HSV, 9-9  
iplRotate, 11-5  
iplRShiftS, 5-12  
iplSet, 4-24  
iplSetBorderMode, 4-15  
iplSetColorTwist, 9-13  
iplSetErrMode, 3-4  
iplSetErrStatus, 3-3  
iplSetROI, 4-14  
iplSetTileInfo, 4-18  
iplsMalloc, 4-21  
iplSquare, 5-6  
iplSubtract, 5-8  
iplSubtractS, 5-4  
iplThreshold, 10-2  
IplTileInfo structure, 4-7  
iplTranslateDIB, 4-30  
iplwMalloc, 4-20  
iplXor, 5-17  
iplXorS, 5-15  
iplZoom, 11-2

## L

linear filters, 6-1  
logical operations, 5-1  
    And, 5-16  
    AndS, 5-13  
    LShiftS, 5-11  
    Not, 5-12

Or, 5-16  
OrS, 5-14  
RShiftS, 5-12  
Xor, 5-17  
XorS, 5-15

lookup table. *See* palette color image  
lookup table operations, 10-3  
LShiftS function, 5-11

## M

magnifying the image, 11-2  
Malloc function, 4-19  
manual organization, 1-2  
MaxFilter function, 6-10  
maximum permissible value, 2-7  
MedianFilter function, 6-9  
memory allocation functions, 4-2, 4-19  
    dMalloc, 4-22  
    Free, 4-23  
    iMalloc, 4-21  
    Malloc, 4-19  
    sMalloc, 4-21  
    wMalloc, 4-20  
MinFilter function, 6-11  
minimum permissible value, 2-7  
Mirror function, 11-6  
mirroring the image, 11-6  
monadic operations, 5-1  
morphological operations  
    Close, 8-7  
    Dilate, 8-5  
    Erode, 8-2

- Open, 8-6
- MSI. *See* multi-spectral image
- multi-image operations, 2-5
- multiply and scale pixel values
  - by a constant, 5-5
  - in two input images, 5-10
- Multiply function, 5-9
- multiply pixel values
  - by a constant, 5-4
  - by a negative power of 2, 5-12
  - in two input images, 5-9
  - square pixel values, 5-6
- MultiplyS function, 5-4
- MultiplyScale function, 5-10
- MultiplySScale function, 5-5
- multi-spectral image, 2-2
  
- N**
- naming conventions, 1-5
- Not function, 5-12
- notational conventions, 1-5
- numerical exceptions, 3-1
  
- O**
- online version, 1-4
- opacity, 5-18
- opacity channel. *See* alpha channel
- Open function, 8-6
- opening and smoothing the image, 8-6
- operation modes of IPL functions, A-1
- Or function, 5-16
- OrS function, 5-14
  
- OUT compositing operation, 5-23
- out-of-place operation, 2-7
- output samples in RCPack2D format, 7-3
- OVER compositing operation, 5-18, 5-23
  
- P**
- palette color image, 2-2
- parallelism, 1-1
- pixel depth, 2-2
- PLUS compositing operation, 5-23
- PreMultiplyAlpha function, 5-25
- producing error messages for users, 3-6
  
- R**
- RCPack2D format, 7-1
- real-complex packed format, 7-1
- RealFft2D function, 7-3
- rectangular region of interest, 2-4
- RedirectError function, 3-6
- reduce the image bit resolution, 9-2
- ReduceBits function, 9-2
- region of interest, 2-4, 4-12
  - channel, 2-4
  - rectangular, 2-4
- report an error, 3-2
- return value, 1-4
- RGB2HLS function, 9-10
- RGB2HSV function, 9-9
- ROI. *See* region of interest
- Rotate function, 11-5
- rotating the image, 11-5
- RShiftS function, 5-12

## S

- saturation mode, 2-7
- scanline. *See* image row data
- scanline alignment, 2-6
- Set function, 4-24
- set the error processing mode, 3-4
- set the error status code, 3-3
- SetBorderMode function, 4-15
- SetColorTwist function, 9-13
- SetErrMode function, 3-4
- SetErrStatus function, 3-3
- SetROI function, 4-14
- SetTileInfo function, 4-18
- shift pixel bits, 5-11, 5-12
- shrinking the image, 11-3
- signed data, 2-2
- SIMD instructions, 1-1
- sMalloc function, 4-21
- smoothing and closing the image, 8-7
- specify a color twist matrix, 9-13
- Square function, 5-6
- square pixel values, 5-6
- status codes, 3-8
- stretching the image contrast, 10-5
- Subtract function, 5-8
- subtract pixel values
  - from a constant, 5-4
  - two input images, 5-8
- SubtractS function, 5-4
- supported image attributes and modes, A-1

## T

- Threshold function, 10-2
- thresholding the source image, 10-2
- tiling, 2-8, 4-7
  - call-backs, 2-9
  - CreateTileInfo function, 4-17
  - DeleteTileInfo function, 4-18
  - IplCoord Structure, 2-9
  - IplTileInfo structure, 4-7
  - SetTileInfo function, 4-18
- TranslateDIB function, 4-30
- two-dimensional convolution, 6-3

## U-Z

- user-defined error handler, 3-13
- Windows DIB, 4-2
- Windows DIB functions, 4-2, 4-28
  - ConvertFromDIB, 4-32
  - ConvertToDIB, 4-33
  - TranslateDIB, 4-30
- wMalloc function, 4-20
- XOR compositing operation, 5-23
- Xor function, 5-17
- XorS function, 5-15
- Zoom function, 11-2
- zooming the image, 11-2