# Glossary

**AC coefficient**

A **DCT** coefficient that corresponds to nonconstant data in the original image.

**additive color system**

A color model in which colors are built by adding together primary colors. **RGB** is an example.

**affine transform**

Affine transforms include such operations as scaling, rotation, translation, and shearing. What these operations have in common is that they can change the lengths and angles of lines, but not their parallelism. The XIL library contains a function `xil_affine()` that in one operation can scale, rotate, and translate an image.

**atom**

A single XIL function. Because of its deferred-execution scheme, the library is sometimes able to find groups of atoms that it can replace with an optimized routine that does the jobs of all the atoms in the group. These optimized routines are called **molecules**.

**attribute**

A term from object-oriented programming that refers to a characteristic of an object.

**Baseline sequential codec**

A sequential coder-decoder defined by the **JPEG** standard. It is designed to handle images with 8-bit samples and uses Huffman coding for its entropy coding.

**bicubic interpolation**

Bicubic **interpolation** is the most time consuming of the XIL library's interpolation methods, but produces the best results. When this type of interpolation is requested, the library calculates the value of a pixel in the destination image by determining the point in the source image to which that pixel maps and then examining the values of the sixteen pixels closest to that point.

**bilinear interpolation**

When the XIL library uses bilinear (or first-order) **interpolation** to determine the value of an image at noninteger coordinates, it calculates the value by looking at the values of the four pixels surrounding the point of interest and then using a bilinear equation. This type of interpolation yields better results than **nearest neighbor interpolation**, but can itself have an undesirable smoothing effect on an image. To alleviate this problem, you can use **bicubic interpolation**.

**Block Truncation Coding**

The image compression method on which the **Cell encoding** scheme is based. In Cell compression, a 4-by-4 region of pixels from an image is represented by two colors and a 16-bit mask.

**CCIR**

The Committée International de la Radio. This treaty organization, part of the International Telecommunications Union (ITU), is responsible for obtaining international agreement on standards for radio and television transmission and the international exchange of programs.

**CCIR Recommendation 601**

An international standard for digitizing PAL, NTSC, and SECAM analog video.

**CCITT**

International Telegraph and Telephone Consultative Committee, an international association and standards body composed primarily of representatives from national telephone agencies. The CCITT promulgates telephony standards, such as X.25, the **Group 3** facsimile standard, and the **H.261** videoconferencing standard (p*x*64).

**CCITT Group 3**

A standard that specifies how facsimile machines must compress and decompress image data. The compression method relies heavily on **run-length encoding**. Runs of white pixels and runs of black pixels are represented with codes from a **Huffman** table.

**CCITT Group 4**

A standard for compressing document images. This compression technique takes advantage of a characteristic of document images called vertical coherence. This means that transitions from black to white or vice versa generally occur in almost the same place on adjacent scanlines. The Group 4 encoder compresses an image by recording information about the relative locations of these transitions.

**cell**

A 4-by-4 block of pixels. (*See also* **Cell encoding**.)

**Cell encoding**

A video compression algorithm developed by Sun. In Cell encoding, a 4-by-4 region of pixels is represented by two colors and a 16-bit mask that indicates which of the two colors to place at each of the 16 pixel positions. The colors and mask are chosen to preserve the mean and variance of the luminance and the average chrominance for the 4-by-4 block.

Cell decoding takes advantage of the fonting hardware commonly found in bitmapped displays.

**CellB encoding**

A video-compression algorithm derived from the **Cell** algorithm. As with a Cell compressor, the most fundamental task of a CellB compressor is to encode 4-by-4 cells of pixels in four bytes. The first two bytes of each cell code are a 16-bit mask that indicates which of two colors will represent each pixel in the cell. The third and fourth bytes contain indexes into tables of luminance and chrominance values and define the two colors to be used in encoding the cell.

CellB coders and decoders are intended primarily for use in videoconferencing applications.

**chrominance**

The portion of a composite signal that carries color information. For example, the $C_b$ or $C_r$ component of a $YC_bC_r$ signal represents part of a pixel's chrominance. (*See also* **luminance**.)

**CIF**

Commmon Interchange Format. CIF format images contain $YC_bC_r$ data and are 352 pixels wide and 288 pixels high. Images in this format are one of the two types of images that may be supported by **H.261** codecs.

**CIS**

Compressed image sequence. The XIL library's compressors store (generally related) compressed images in structures called CIS buffers. The images may represent frames in a movie, pages in a document, and so on.

**CMY color model**

In the CMY color model, the subtractive primaries cyan, magenta, and yellow are used to filter their complements (red, green, and blue) from white light. You use this color model when working with devices like color printers that put colored ink on paper. For example, to create the color red on a color printer, you put down a mixture of magenta and yellow ink. The magenta filters out its complement green; the yellow filters out its complement blue; and you see only red.

**CMYK color model**

Similar to the CMY color model, but uses a fourth color: black. Black is used to replace equal amounts of cyan, magenta, and yellow. This color model is used in offset color printing.

**codec**

A coder-decoder.

**colormap**

A color lookup table that stores a set of **RGB** values. Applications index into the colormap to get the values to drive the red, blue, and green guns of an RGB monitor.

**compressed image sequence**

*See* **CIS**.

**compression**

The process of converting data from its original format to a format that requires fewer bits. Compressed data uses less storage than uncompressed data and can be transmitted over a network more quickly. Some compressors, such as the **CCITT Group 3** compressor, are called **lossless** compressors because they compress data in a way such that a decompressor can regenerate the original data exactly. Other compressors are called **lossy** compressors because they compress data in a way that prevents the original data set from being regenerated exactly.

**convolution**

An image-processing operation frequently used to sharpen an image, blur an image, or highlight the edges in an image. The operation calculates the values of pixels in the destination image using the values of a neighborhood of pixels in the source image and the values in a special filter called a **convolution kernel**.

**convolution kernel**

A two-dimensional array of weighted values used in a **convolution** operation. In the XIL library, a kernel is a data structure of data type `XilKernel`.

**DC coefficient**

A **DCT** coefficient that corresponds to the average level of the input image.

**DCT**

Discrete Cosine Transform. Many encoders, including those that conform to the **JPEG**, **MPEG-1**, and **H.261** standards, perform a DCT on an 8-by-8 block of image data as part of the image-compression process. The DCT converts the video data from the spatial domain to the frequency domain. The DCT takes an 8-by-8 matrix, whose values represent brightness levels at particular *x,y* coordinates, and produces an 8-by-8 matrix whose values represent relative amounts of the 64 spatial frequencies that make up the input data's spectrum.

The DCT provides a basis for compression because most of the frequency levels for a block will be zero or close to zero and do not need to be encoded.

**decoder**

A program that takes data that has been encoded, or compressed, by an **encoder** and decompresses it. A decoder can be implemented in hardware, software, or a combination of both. The decompressed data may or may not match the original data set exactly, depending on how the data was encoded. (*See also* **lossless compression** and **lossy compression**.)

**decompression**

The restoring of data that has undergone compression to its original state, or to something close to its original state. How closely the decompressed data matches the original data depends on the compression algorithm used. (*See also* **lossless compression** and **lossy compression**.)

**display image**

The XIL library treats displays as special images, and operations allow display images to serve as destinations. This strategy enables an operation to draw on a screen directly, without requiring an intermediate copy. In some cases, display images can also serve as source images.

**dithering**

The XIL library includes a table-lookup function that enables you to convert a single-band image of any data type to a single-band or multiband image of any data type. Dithering can be thought of as an *inverse* lookup operation; that is, dithering enables you to convert a single-band or multiband image of any data type to a single-band image of any data type. The most common dithering operations convert 3-band, 8-bit images to 1-band, 8-bit images and 1-band, 8-bit images to 1-band, 1-bit images.

**encoder**

An encoder is a program that encodes data for the purpose of achieving data compression. The encoders included with the XIL library are designed to compress images.

**entropy coding**

Entropy coding is the final step in the compression process in **DCT**-based encoders (such as the **JPEG** baseline sequential encoder). In this step, the encoder compresses quantized **DCT** coefficients using **Huffman** coding. Therefore, values that occur frequently are encoded with fewer bits than are values that occur infrequently.

**error diffusion**

A technique for removing some of the artifacts produced during the dithering process. In its most common form, Floyd-Steinberg error diffusion, this technique involves (1) determining the amount of error produced in dithering a particular pixel and (2) distributing fractions of that error to the pixels to the right of and below the pixel just dithered.

**frame**

A single image taken from a movie.

**full-motion video**

The showing of a series of related digital images at a rate sufficient to give the illusion that objects in the images are moving naturally.

**gamma correction**

In a linear color space, color levels are equally spaced throughout a gamut. The problem with this type of color space is that level-to-level changes at the low end of the gamut seem greater to the eye than equal changes at the high end. In a gamma-corrected color space, color levels are spaced logarithmically so that level-to-level changes seems consistent throughout the range.

**geometric operation**

An image-processing operation that changes the size, shape, or orientation of objects in an image. XIL geometric operations include scaling, rotation, and general affine transforms.

**H.261**

A video compression standard developed by the **CCITT** for use in encoding video to be transferred over an **ISDN** (Integrated Service Digital Network). In this compression method, data is compressed so that the output bit rate is p x 64 Kbits per second, where *p* can range from 1 to 30 depending on the number of ISDN channels used. This standard was developed primarily to support video phone and videoconferencing applications.

**handle**

An identifier that refers to a unique object. In the XIL library, for example, an identifier of type `XilImage` is a handle to an actual image.

**histogram**

A collection of information about how frequently certain gray levels or colors appear in an image. The XIL library contains a data structure of type `XilHistogram` to hold this type of information.

**Huffman encoding**

A method of compressing a given set of data based on the relative frequency of the individual elements: the more often an element occurs, the shorter (in bits) its corresponding code. Huffman encoding is often used to compress text files, with the coding based on letter frequency. Huffman encoding is also used in many **DCT**-based video-compression algorithms.

**image**

A two-dimensional array of **pixels** that represents an object.

**indexed color**

*See* **pseudocolor**.

**in-place operations**

Operations that require both a source and a destination image but that allow them to be the same image. If a function doesn't support in-place operations, the source and destination images must be different images.

**interframe compression**

In image sequences, consecutive frames generally have more similarities than differences. These similarities, or shared elements, are sometimes referred to as being temporally redundant. This redundancy is important because it allows groups of individually compressed frames to be compressed further. That is, if five frames in a group look the same in the upper-left corner, that area needs to be encoded only once; then the encoder can simply note that the same data appears in the next four frames. This type of encoding is called interframe compression.

**interpolation**

A way of calculating a value that falls between other, known values. In image processing, interpolation frequently plays a part in geometric operations such as rotation. After that type of spatial transformation, pixel locations in the output image will correspond to noninteger coordinates in the input image. Therefore, the pixel values in the output must be calculated by looking at the values of the pixels surrounding the point of interest in the input. The XIL library supports several types of interpolation, including **nearest neighbor**, **bilinear**, and **bicubic interpolation**.

**intraframe compression**

The compression that can take place within a single image. Contrast with **interframe compression**.

**ISDN**

Integrated Service Digital Network. A worldwide public telecommunications network designed to handle many types of data, including voice, text, graphics, and video. The **CCITT** designed the **H.261** video encoder to produce compact bitstreams that can be sent over ISDN lines.

**JPEG**

Joint Photographic Experts Group. A joint venture of the **CCITT** and ISO that has developed a standard for compressing grayscale or color *still* images. Actually, the standard defines a number of methods for compressing images. Several of these are **lossy** methods based on the Discrete Cosine Transform (**DCT**), but one method is **lossless** and is based on a **predictive coding** technique.

**key frame**

If the video portion of a movie has been compressed using both **intraframe compression** and **interframe compression**, a decoder cannot decompress the majority of the video frames without referring to preceding—and sometimes succeeding—frames. However, the first frame in the movie, and usually other frames as well, do not have interframe dependencies and can be decoded in isolation. These frames are called key frames. Key frames, besides the one that starts the movie, enable the decoder to decompress other frames without playing the movie from the beginning.

**linear remapping**

*See* **rescaling**.

**local operation**

An image-processing operation in which more than one pixel in a source image is used in calculating the value of a single pixel in the destination image. Generally, the source-image pixels used include the pixel corresponding to the destination pixel being calculated and some set of pixels surrounding that source pixel.

**lookup table**

In XIL applications, lookup tables are used for general image modification. Each entry in an XIL lookup table contains an index—a value that may appear in the source image—and a value or set of values to be written to the destination image. For each pixel in the source, a table-lookup function finds the pixel's value on the index side of the table and then writes the output value or values for that entry to the corresponding pixel in the destination.

**lossless compression**

The compression of data in such a way that the original data can be restored exactly. **Huffman encoding** is an example of a lossless compression technique. Some compressors, such as the **JPEG** baseline sequential compressor, combine lossy and lossless compression algorithms. (*See also* **lossy compression**.)

**lossy compression**

A type of compression that results in the loss of some of the original data. Lossy compression trades the potential loss of some image quality for the opportunity for greater compression. The **JPEG** baseline sequential and **Cell** methods are examples of lossy compression techniques. (*See also* **lossless compression**.)

**luminance**

The portion of a composite signal that carries brightness information. For example, luminance information is contained in the Y component of a $YC_bC_r$ signal. Video compression techniques take advantage of the fact that the human eye is more sensitive to variations in luminance than it is to variations in color (**chrominance**). Therefore, chrominance values can be compressed (with **lossy** techniques) more than luminance values, resulting in greater overall compression.

**macroblock**

Both the **H.261** and the **MPEG-1** specifications define a unit within an image called a macroblock. This unit is a 16-by-16 blocks of pixels. Both H.261 and MPEG-1 encoders can switch from intraframe encoding to interframe encoding on a macroblock basis.

**molecule**

The XIL term for an optimized routine that performs the work of two or more XIL functions (**atoms**). You don't call a molecule directly. Instead, the XIL runtime system executes a molecule whenever your program calls a sequence of XIL functions that the molecule can replace. This type of substitution is possible because of the library's deferred-execution scheme.

**motion compensation**

Both **H.261** and **MPEG-1** encoders can perform interframe compression by encoding the differences between the values in a **macroblock** in the picture being encoded and the values in the corresponding macroblock in the preceding image. Sometimes, however, it's desirable to encode the differences between the values in a macroblock in the current image and the values in a macroblock in the preceding image that is slightly offset from the one that corresponds spatially with the macroblock being encoded. This approach is desirable because it enables the encoder to track the movement of objects from image to image and, thus, to encode smaller differences. If an encoder uses this approach, it must of course record the extent of the offset mentioned above. This offset is recorded in a motion vector.

**movie**

A contiguous series of video frames (and optionally synchronized audio) that are displayed fast enough to provide the illusion of motion. A frame rate of 30 frames/second is a typical target for a smooth-running movie.

**MPEG**

Moving Picture Experts Group. This group has developed standards for compressing moving pictures and audio data and for synchronizing video and audio datastreams. The XIL library includes a decompression module that can decode MPEG-1 video bitstreams.

The MPEG-1 video-compression standard is similar to the **H.261** standard developed by the **CCITT**, but places less emphasis on low bit rates. By accepting a higher bit rate—up to 1.5 Mbits per second—an MPEG-1 codec is able to recreate very high-quality pictures and to produce a bitstream that is easily editable.

The rate of 1.5 Mbits/s makes the MPEG-1 codec especially viable in applications that read compressed data from CD-ROMs. For example, putting an MPEG-1 bitstream on a CD-ROM is an effective way to distribute movies, business presentations, and training videos.

**nearest neighbor interpolation**

One of the methods that the XIL library uses to determine the value of an input image at noninteger coordinates. When using this method, the XIL library takes the value it is looking for to be the value of the pixel closest to the point of interest. This type of **interpolation** is the fastest type, but can introduce artifacts in the output image; for example, smooth lines in the input image may show up as jagged lines in the output.

Nearest neighbor interpolation is sometimes called zero-order interpolation.

**origin**

In the XIL library, each image object has an **attribute** called its origin. This origin is a pair of floating-point numbers that represent *x,y* coordinates in the image. (The point 0.0,0.0 is in the upper-left corner.) When an image-processing operation is performed, the origins of the input and output images are aligned, and the rectangle formed by the intersection of these images serves as an implicit region of interest (**ROI**).

**pixel**

Picture element. In a raster grid, a pixel is the smallest unit that can be addressed and given a color or intensity.

**point operation**

An image-processing operation in which the value of a point (or **pixel**) in the destination image depends only on the corresponding point in the source image or images. For example, if you want to add two source images to produce a destination image, the value of the pixel in the upper-left corner of the destination image depends solely on the values of the pixels in the upper-left corner of the two source images.

**predictive encoding**

In predictive encoding, an encoder uses the values of neighboring samples to predict the value of the sample being encoded. The encoder then subtracts this predicted value from the actual value of the sample and encodes the difference. This mode of compression is **lossless**. The XIL library supports a **JPEG** lossless compressor that uses this method.

**pseudocolor**

The distinction between true color and pseudocolor has to do with the design of a monitor's frame buffer. If the frame buffer uses 8 bits per pixel to store color information, the monitor can display 256 colors simultaneously. What you see on such a monitor is called pseudocolor because the colors that can be shown at any one time are a small subset of the colors the eye can distinguish. If the frame buffer uses 24 bits per pixel to store color information, your monitor can display over 16 million colors (true color).

Pseudocolor is sometimes called indexed color because the values stored in the frame buffer on a pseudocolor system are not the **RGB** values needed to drive the red, green, and blue electron guns in a monitor. Rather, they are indexes into a **colormap**, or color lookup table, which stores 256 sets of RGB values.

**QCIF**

Quarter Common Interchange Format. QCIF images contain $YC_bC_r$ data and are 176 pixels wide and 144 pixels high (one-fourth the size of **CIF** images). **H.261** codecs are required to support images in this format.

**quantization**

The technique of scaling down a set of values. In Discrete Cosine Transform (**DCT**)-based encoders, like those that conform to the **JPEG**, **MPEG-1**, and **H.261** standards, quantization is used to ensure that DCT coefficients are represented by the smallest range of numbers needed to produce the desired level of image quality. To bring the coefficients into this range, the quantizer divides each coefficient by the appropriate value from a quantization table and rounds the result to the nearest integer.

**rescaling**

Sometimes called linear remapping, rescaling involves changing the values in an image using a linear equation. Each value in the image is multiplied by a constant; then a second constant is added to the product. Rescaling is useful for mapping the values in an image from one range to another.

**RGB color model**

A color model in which colors are built by mixing the three additive primary colors red, green, and blue. In this model, you construct grays by including equal amounts of each primary: (0,0,0) is black and (1,1,1) is white. The RGB color model is closely associated with color CRT monitors because they use this model to produce their colors.

**ROI**

A region of interest. In the XIL scheme, an ROI is a 1-bit mask that can be specified as an **attribute** of an image. This mask determines what part of a source image can be read or what part of a destination image can modified during an image-processing operation. The ROI for a given operation is the intersection of the ROIs of all the images involved in the operation.

**run-length encoding**

A compression technique that stores counts of the number of consecutive identical pixels or blocks of pixels in an image.

**sequential encoding**

The encoding of a still image in a single pass through the image data. The sequential encoder first processes a block of pixels in the upper-left corner of the image and then proceeds from left to right and top to bottom until it has processed the entire image.

**skip codes**

Skip codes are employed during the **interframe compression** process. They instruct the decoder to skip over blocks of pixels in the current frame that are identical to, or very similar to, those in the preceding frame.

**slice**

**MPEG-1** terminology for a group of consecutive **macroblocks**. Each macroblock in a picture (or image) must be part of a slice. The slice was designed primarily to help a decoder recover from bitstream errors. If a decoder detects an error, one way to recover is to skip to the next slice header.

**spatial redundancy**

The occurrence of two or more consecutive pixels or blocks of pixels that have the same, or similar, contents. A high level of spatial redundancy in an image can lead to a high rate of compression because only the first pixel or block of pixels in a sequence needs to be encoded fully. A shorthand method can be used to encode succeeding pixels or blocks.

**still-image coding**

The encoding of a single image.

**structuring element**

A two-dimensional array of Boolean values used as a parameter to the XIL dilation and erosion functions.

**subsampling**

A way of mathematically reducing a data set to a subset of its original components. For example, if you have captured a 512-by-512 image, but want to encode it at a resolution of 256-by-256, you can subsample the data before encoding it by discarding every other pixel in both and *x* and *y* directions. On systems that work with $YC_bC_r$ data, it is common to subsample the $C_b$ and $C_r$ **chrominance** information so that each pixel has a unique **luminance** value, but shares chrominance information with one or more nearby pixels.

**subtractive color system**

In a subtractive color system, primary colors are used as filters to subtract their complements from white light. For example, in the **CMY** color model, the primaries cyan, magenta, and yellow subtract red, green, and blue respectively. That is, if light is reflected from a piece of paper coated with cyan ink, no red is reflected, so you see the color cyan.

**temporal redundancy**

*See* **interframe compression**.

**thresholding**

The XIL library's thresholding function sets all the values (in a band) that fall between a low threshold and a high threshold to a value (called a *map* value) that the programmer specifies.

**transcoding**

The conversion of data in one compressed format to another compressed format. For instance, converting a **JPEG**-compressed image to its **Cell**-encoded counterpart is called transcoding. Transcoding an image usually involves decoding it and then reencoding it.

**translation**

A geometric operation in which an image is moved up or down and/or left or right.

**transposition**

A quasi-geometric transformation that involves rotating an image by a multiple of 90 degrees or flipping pixels in an image across a line that passes through the center of the image. The XIL library's transposition function allows for flips across a horizontal or vertical line passing through an image's center and across the main diagonal or antidiagonal.

**XIL library**

The XIL library is a foundation library for imaging and video support. It provides an implementation of imaging functionality that is common to multiple higher-level interfaces, provides imaging capabilities that are not currently available, and provides a way for ISVs to access low-level and hardware functionality.

**$YC_bC_r$ color model**

When data is stored in $YC_bC_r$ format, each pixel is described by a **luminance** value (Y) and two color values. This color model is used in the PAL (European) television format. Also, video analog-to-digital converters often produce data in this format.

Frequently, the color bands of $YC_bC_r$ images are subsampled to produce what are called $YC_bC_r$ 4:2:2 or $YC_bC_r$ 4:1:1 images. In these images, pairs of pixels, or blocks of 4 pixels, share color values; that is, each pixel has its own luminance information, but shares color values with a neighboring pixel or pixels. This subsampling in the color dimension has a very minor impact on image quality.