MPEG-1 Codec 19

This chapter discusses the MPEG-1 video codec specified by the Moving Picture Experts Group, an ISO working group. This group has produced a standard that is similar to the H.261 standard developed by the CCITT (see Chapter 18, “H.261 Codec”), but places less emphasis on low bit rates. By accepting a higher bit rate—for example, 1.5 Mbits per second—an MPEG-1 codec is able to recreate very high-quality pictures and to produce a bitstream that is easily editable.

The rate of 1.5 Mbits/s makes the MPEG-1 codec especially viable in applications that read compressed data from CD-ROMs because even older CD-ROM readers can read data at this speed. Thus, putting an MPEG-1 bitstream on a CD-ROM is an effective way to distribute movies, business presentations, and training videos.

The remainder of this chapter is divided into four sections. The first section provides an overview of how an MPEG-1 codec works. The second explains briefly how to create an MPEG-1 CIS. The third discusses CIS attributes that apply specifically to a CIS associated with an XIL MPEG-1 compressor or decompressor (as opposed to the general CIS attributes covered in the section “General CIS Attributes” on page 257). And the fourth section introduces the subject of accelerating the playback of MPEG-1 bitstreams. For further information on this subject, see Chapter 21, “Acceleration in XIL Programs.”

Note – This chapter discusses both MPEG-1 compression and decompression. However, the current release of the XIL library includes only an MPEG-1 decompressor. The compressor interface is defined for third parties who want to implement XIL MPEG-1 compressors.

How an MPEG-1 Codec Works

Since the work done by the Motion Picture Experts Group grew out of the work done by the CCITT in developing the H.261 video codec, this section explains how the MPEG-1 codec works by comparing and contrasting it with the H.261 codec. (If you're unfamiliar with how an H.261 codec works, see the section "How an H.261 Codec Works" on page 338.) The present section also includes a subsection that describes the organizational structure that MPEG-1 imposes on a video sequence.

Similarities Between MPEG-1 and H.261

The key similarities between MPEG-1 and H.261 are listed below:

- Both compressors work with YC_bC_r pictures in which the information in the chroma channels has been subsampled so that there is one C_b and one C_r value for each 2-by-2 block of luma values.
- Like the H.261 compressor, the MPEG-1 compressor can compress a macroblock (a 16-by-16 block of pixels in a picture) by encoding the actual Y, C_b, and C_r values in the macroblock (intraframe encoding) or by encoding the differences between values in the current block and values in the corresponding macroblock in the previous picture (forward prediction). In addition, when using the forward-prediction encoding method, the MPEG-1 compressor, like the H.261 compressor, can employ motion compensation.
- Both compressors encode 8-by-8 blocks of pixel values or difference values using the same method. The compressors first perform a Discrete Cosine Transform (DCT) on the 8-by-8 block of values. This operation transforms the values in the block from the spatial to the transform domain. Second, the compressors quantize the coefficients produced by the DCT. Finally, the compressors use entropy coding to encode the quantized coefficients.

Differences Between MPEG-1 and H.261

Although the Moving Picture Experts Group drew heavily on the work of the CCITT, there are also very significant differences between the MPEG-1 and H.261 compressors. Most of these differences provide for random access to the MPEG-1 bitstream and make the bitstream easily editable.

I Pictures and P Pictures

In H.261, the ideas of intraframe encoding and predictive encoding are applied for the most part at the macroblock level. MPEG-1, on the other hand, includes the notion of intraframe-encoded pictures and predicted pictures.

In an intraframe-encoded picture, or I picture, all macroblocks are intraframe encoded. This, of course, means that the decoder needs no information from a preceding picture to decode an I picture. For this reason, seeks to I frames can be performed very quickly.

In a predicted picture, or P picture, each macroblock can be intraframe encoded or encoded using the forward-prediction method. This type of picture is very similar to an H.261 picture.

Typically, an MPEG-1 bitstream will contain more P pictures than I pictures because the P picture can be encoded using fewer bits. Encoding difference values generally requires fewer bits than encoding pixel values. Also, the encoder does not have to encode macroblocks that are very similar to their counterparts in the preceding I or P picture.

B Pictures

From what we've said so far, you might picture an MPEG-1 bitstream as containing periodic I pictures followed by a number of P pictures. See Figure 19-1.

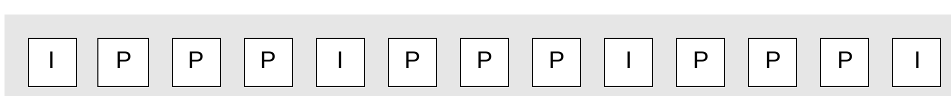

Figure 19-1 An MPEG-1 Bitstream Containing I and P Pictures

This does constitute a legal MPEG-1 bitstream. However, the bitstream may also contain one or more bidirectionally predicted pictures, or B pictures, between any pair of I or P pictures. In a B picture, each macroblock may be

- Intraframe encoded
- Forward predicted from the nearest preceding I or P picture
- Backward predicted from the nearest succeeding I or P picture
- Bidirectionally predicted from the nearest preceding I or P picture and the nearest succeeding I or P picture

We've already discussed intraframe encoding and forward prediction.

Backward prediction is strictly analogous to forward prediction. A backward-predicted macroblock is encoded with respect to the values in the corresponding macroblock in a picture that follows its own picture in the video sequence. This option may seem unintuitive at first thought because it results in a situation where pictures are not transmitted in the order in which they will be displayed. See Figure 19-2.

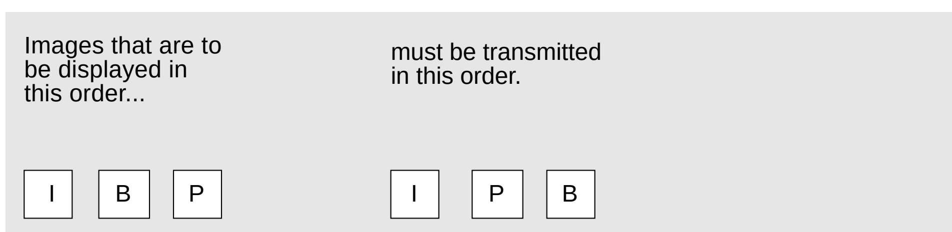

Figure 19-2 MPEG-1 Display Order Versus Decoding Order

However, backward prediction is an option because it can lead to a more compact bitstream in instances where the macroblocks in a B picture match their counterparts in the closest succeeding I or P picture more closely than they match their counterparts in the closest preceding I or P picture.

A bidirectionally predicted macroblock is coded with respect to the corresponding macroblock in both the closest preceding and succeeding I/P pictures. See Figure 19-3.

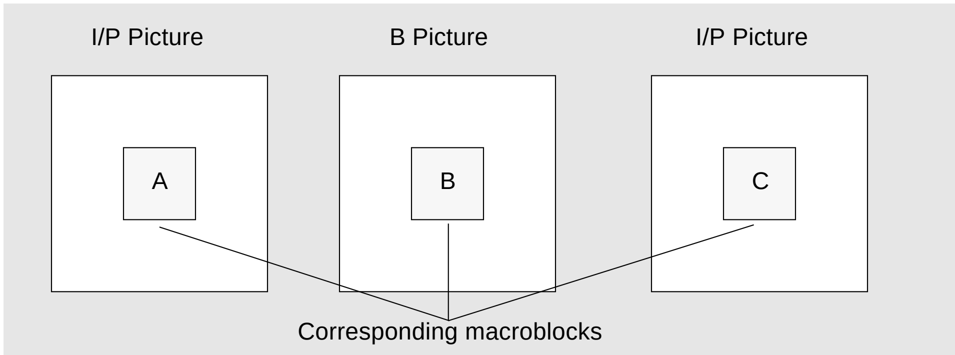

Figure 19-3 Bidirectional Prediction in MPEG-1

In this case, the encoder averages the values in macroblocks A and C and then encodes the difference between values in macroblock B and the averaged macroblock. Bidirectional prediction is effective because the difference mentioned above is often so small that it does not need to be encoded.

For MPEG-1, the behavior of the `xil_cis_get_bits_ptr()` function differs from its usual behavior. For a bitstream with out-of-order frames (that is, a bitstream with B frames), the actual number of frames in the data returned by `xil_cis_get_bits_ptr()` may not equal the value of its `nframes` parameter. The value of `nframes` is the number of frames the CIS read frame has been advanced by the `xil_cis_get_bits_ptr()` call. A seek back by `nframes` by calling `xil_cis_seek()` will restore the read position to the original read frame (before the `xil_cis_get_bits_ptr()` call). This is useful for an application that uses a preview mode, where the compress and write-to-file is followed by a decompress.

Since the number of frames reported may not represent the actual number of frames, if the chunk is subsequently used for an `xil_cis_put_bits()` or `xil_cis_put_bits_ptr()` call, the `nframes` parameter must be set to -1, which indicates an unknown number of frames.

Groups of Pictures

Another concept used in MPEG-1 that does not apply in H.261 is that of a group of pictures. This is a series of consecutive pictures from a video sequence. Generally, the group of pictures provides a unit of the bitstream that can be removed without destroying the integrity of the bitstream and a unit that can be decoded independently of the rest of the bitstream. As you'll see in a moment, however, the group of pictures does not always have these characteristics.

By definition, a group of pictures (considered in display order) must begin with an I picture or with one or more B pictures followed by an I picture. It must end with an I or a P picture. A closed group of pictures can be decoded without any information from the preceding group of pictures. Thus, a closed group is one that begins with an I picture or one that begins with one or more B pictures whose macroblocks have been encoded using only intraframe encoding and backward prediction. An open group of pictures begins with one or more B pictures, at least one of which contains macroblocks encoded using forward or bidirectional prediction. This type of group can be decoded only if the preceding group of pictures is available. If that group is not available—for example, if an MPEG-1 bitstream editing program has removed the group—a broken-link bit must be set in the header for the open group of pictures.

How MPEG-1 Organizes a Video Sequence

We've already mentioned some of the organizational units used in MPEG-1, such as the group of pictures and the macroblock. This section provides a complete overview of these units. This information should be useful to you as you read about MPEG-1 CIS attributes later in this chapter because many of these attributes are associated with a particular unit.

The largest unit MPEG-1 defines is the video sequence. You might think of this unit as an entire movie or presentation. Each MPEG-1 bitstream includes a sequence header, which defines several attributes for the entire sequence. These include the pixel aspect ratio for the sequence, the picture rate for the sequence in pictures per second, and the bit rate of the data channel over which the compressed sequence will be moved. An additional sequence header is written to the bitstream each time one of the encoder's quantization tables is changed.

Each sequence is divided into a series of groups of pictures. See the section "Groups of Pictures" on page 363 for a definition of this unit. If you're compressing video using an XIL-compliant MPEG-1 encoder, you can control the makeup of each group of pictures using the attribute `COMPRESSOR_PATTERN`.

Each group of pictures consists of pictures, which are individual frames of video. The major attributes associated with each picture are a picture type and a temporal reference. The major picture types are intraframe-encoded pictures, forward-predicted pictures, and bidirectionally predicted pictures. The temporal reference is an integer identifying a picture's place within a group of pictures.

The largest subdivision of a picture is called a slice, which consists of a series of consecutive macroblocks. Slices within a picture may vary in size, but each macroblock in a picture must be part of a slice. This unit is designed primarily to help a decoder recover from a bitstream error. If a decoder detects an error, one way to recover is to skip to the next slice header.

Slices, as mentioned in the preceding paragraph, are built from macroblocks. These are 16-by-16 block of pixels. The macroblock is the level at which motion compensation is performed. Also, the type of encoding used—intraframe, forward prediction, backward prediction, or bidirectional prediction—can change from macroblock to macroblock.

Finally, each macroblock is divided into six 8-by-8 blocks. Four of these blocks contain luma values, one contains C_b values, and one contains C_r values. This is the level at which the DCT is performed.

Creating an MPEG-1 CIS

Before you can use the MPEG-1 decompressor to decompress an MPEG-1 bitstream, you must create an MPEG-1 CIS (and write an MPEG-1 bitstream to the CIS). You create this CIS by passing the decompressor name `Mpeg1` to the function `xil_cis_create()`. See the code fragment below.

```
XilCis cis;
XilSystemState state;

cis = xil_cis_create(state, “Mpeg1”);
```

MPEG-1 Codec Attributes

As discussed in the section “General CIS Attributes” on page 257, there is a class of attributes that can be set for any CIS. There is also a set of attributes that are valid only for CISs attached to an MPEG-1 compressor or