JPEG Lossless Codec 17

Like the JPEG baseline sequential compressor, the JPEG lossless compressor is designed to compress still continuous-tone images. The difference between the two is that the baseline sequential compressor is a lossy compressor: once an image has been compressed, you cannot recover the original image samples. The JPEG lossless compressor does allow you to recover these samples; however, the lossless compressor produces much lower compression ratios than the lossy compressor. The compression ratio for the lossless compressor is in the neighborhood of 2:1. One other difference between the two compressors is that while the baseline sequential compressor works only with `XIL_BYTE` images, the lossless compressor can work with both `XIL_BYTE` and `XIL_SHORT` images.

The remainder of this chapter is divided into three sections. The first section provides an overview of how the JPEG lossless compressor works. The second explains briefly how to create a JPEG lossless CIS. The third discusses a set of CIS attributes that apply specifically to the JPEG lossless codec (as opposed to the general CIS attributes covered in the section "General CIS Attributes" on page 257).

How the JPEG Lossless Codec Works

The JPEG lossless compressor is not based on the DCT like the baseline sequential encoder. Instead, it uses a predictive method, as shown in Figure 17-1.

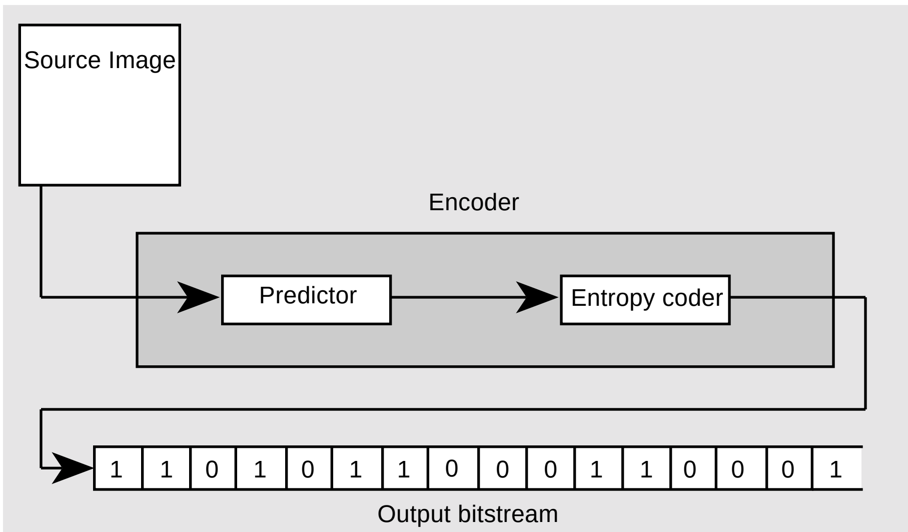

Figure 17-1 JPEG Lossless Compressor

Each sample in the source image is encoded as follows: The predictor makes a guess at the sample's value based on its knowledge of the values of neighboring samples and then subtracts the predicted value of the sample from its actual value. The difference calculated by the predictor is then passed on to the entropy coder, which does a lookup in a Huffman table and writes encoded data to the bitstream.

The prediction and entropy-coding steps in the encoding process are discussed in more detail in the following sections.

Note – The XIL library's JPEG lossless encoder actually lets you perform an operation called a point transformation on your source image before the encoding summarized above begins. The point transformation can be done on a band-by-band basis. What happens is this. If you supply a point-transform value other than 0 for a band of an image, each sample in the band is divided by 2 raised to the power *x*, where *x* is the point-transform value. This operation leads to a higher compression ratio because smaller values will need to be encoded; however, it is not actually part of the lossless-encoding process because it can result in the loss of some data.

Prediction

As mentioned above, the predictor predicts the value of an image sample based on its knowledge of the values of neighboring samples. The neighboring samples it can take into consideration are shown in Figure 17-2 and are labeled A, B, and C. The sample whose value is being predicted is labeled P.

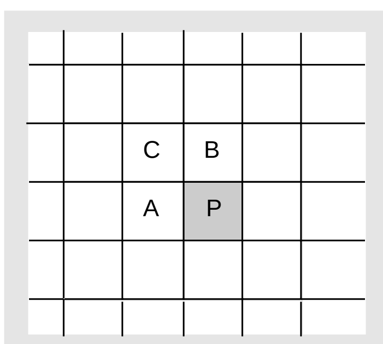

Figure 17-2 Predicting Values in the JPEG Lossless Compressor