Encoding Order

- Each block of 8x8 is treated separately.
- Order of blocks is:

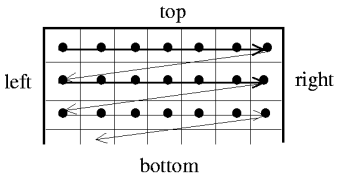

JPEG

- "JPEG" is Joint Photographic Experts Group.
- compresses pictures which don't have sharp changes e.g. landscape pictures.
- May lose some of the data in order to compress better.
- Both color or grayscale images.

DCT

- The DCT is a mathematical operation that transform a set of data, which is sampled at a given sampling rate, to it's frequency components.

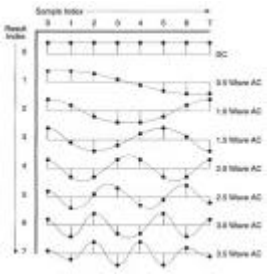

Baseline JPEG

- Transfer to frequency space using Discrete Cosine Transform (DCT).
- Quantization: Divide and round the results according to the required quality.
  - This stage may cause some lose of data.
- Compress the data by a version of Canonical Huffman coding.
  - Non-Baseline JPEG may use also Arithmetic coding.

One dimensional DCT

$$F(u) = \frac{1}{2}\sum_{x=0}^{7} f(x)C(u)\cos[(2x+1)up/16]$$

$C(u) = 1/\sqrt{2}$ If u = 0

$C(u) = 1$ If $u \neq 0$

- If f(x) (the intensity of each pixel) is equal in the whole row, each F(u) which holds u>0, will be zero. F(0) will be the sum of the row's values divided by $2\sqrt{2}$

DCT (cont.)

- The first element in the result array is a simple sum of all the samples in the input array and is referred to as DC coefficient.
- The remaining elements in the result array each indicate the amplitude of a specific frequency component of the input array, and are known as AC coefficients. The frequency content of the sample set at each frequency is calculated by taking a weighted sum of the entire set.

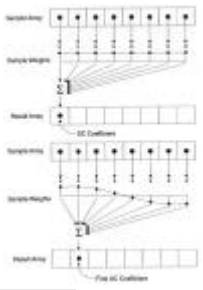

Two dimensional DCT

- one-dimensional DCT is applied separately to each row of eight pixels. The result will be eight rows of frequency coefficients.
- These 64 coefficients are then taken as eight column. The first column will contain all DC coefficients, the second column will contain the first AC coefficient from each row, and so on.
- One-dimensional DCT is applied to each of these columns.

DCT values

- These are the values of weight for one row in a 8x8 matrix (considering f(x) is 1):

Result\Sample Index	0	1	2	3	4	5	6	7
0	+0.707	+0.707	+0.707	+0.707	+0.707	+0.707	+0.707	+0.707
1	+0.981	+0.831	+0.556	+0.195	-0.195	-0.556	-0.831	-0.981
2	+0.924	+0.383	-0.383	-0.924	-0.924	-0.383	+0.383	+0.924
3	+0.831	-0.195	-0.981	-0.556	+0.556	+0.981	+0.195	-0.831
4	+0.707	-0.707	-0.707	+0.707	+0.707	-0.707	-0.707	+0.707
5	+0.556	-0.981	+0.195	+0.831	-0.831	-0.195	+0.981	-0.556
6	+0.383	-0.924	+0.924	-0.383	-0.383	+0.924	-0.924	+0.383
7	+0.195	-0.556	+0.831	-0.981	+0.981	-0.831	+0.556	-0.195

Biased Values

- JPEG allows samples of 8 bits or 12 bits.
- All samples within the same source image must have the same precision.
- The samples are shifted from unsigned integers with range $[0,2^p-1]$ to signed integers with range $[-2^{p-1},2^{p-1}-1]$, by reducing 2^{p-1} from the original values, where p can be either 8 or 12.
- These biased values are sent to the DCT function.

DCT formula

$$F(u,v)=\frac{1}{4}\sum_{y=0}^{7}\sum_{x=0}^{7}f(x,y)C(u)\cos[\frac{(2x+1)up}{16}]C(v)\cos[\frac{(2y+1)vp}{16}]$$

$$C(u), C(v) = \frac{1}{\sqrt{2}} \quad u, v = 0$$

$$= 1 \quad u, v \neq 0$$

- Index 0,0 contains the DC of the DCs. This value is called the DC of the 8x8 matrix.

Example (cont.)

15	0	-1	0	0	0	0	0
-2	-1	0	0	0	0	0	0
-1	-1	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

(a) normalized quantized coefficients

240	0	-10	0	0	0	0	0
-24	-12	0	0	0	0	0	0
-14	-13	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

(b) denormalized quantized coefficients

144	146	149	152	154	156	156	156
148	150	152	154	156	156	156	156
155	156	157	158	158	157	156	155
160	161	161	162	161	159	157	155
163	163	164	163	162	160	158	156
163	164	164	164	162	160	158	157
160	161	162	162	162	161	159	158
158	159	161	161	162	161	159	158

(c) reconstructed image samples

Example

139	144	149	153	155	155	155	155
144	151	153	156	159	156	156	156
150	155	160	163	158	156	156	156
159	161	162	160	160	159	159	159
159	160	161	162	162	155	155	155
161	161	161	161	160	157	157	157
162	162	161	163	162	157	157	157
162	162	161	161	163	158	158	158

(a) Source image samples

235.6	-1.0	-12.1	-5.2	2.1	-1.7	-2.7	1.3
-22.6	-17.5	-6.2	-3.2	-2.9	-0.1	0.4	-1.2
-10.9	-9.3	-1.6	1.5	0.2	-0.9	-0.6	-0.1
-7.1	-1.9	0.2	1.5	0.9	-0.1	0.0	0.3
-0.6	-0.8	1.5	1.6	-0.1	-0.7	0.6	1.3
-1.8	-0.2	1.6	-0.3	-0.8	1.5	-1.0	-1.0
-1.3	-0.4	-0.3	-1.5	-0.5	1.7	1.1	-0.8
-2.6	1.6	-3.8	-1.8	1.9	1.2	-0.6	-0.4

(b) forward DCT coefficients

16	11	10	16	24	40	51	61
12	12	14	19	26	58	60	55
14	13	16	24	40	57	69	56
14	17	22	29	51	87	80	62
18	22	37	56	68	109	103	77
24	35	55	64	81	104	113	92
49	64	78	87	103	121	120	101
72	92	95	98	112	100	103	99

(c) quantization table

An example for compression

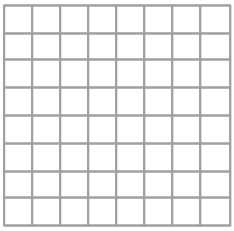

- This matrix is after DCT and after Quantization.
- Different Huffman codes For DC and AC values.
- Suppose last block's DC value was 15.

- JPEG switches zeros and ones in negative numbers

```
100 101 00  1 11111001 11 1111111000 01 1010
 |   |  |  |  |        |  |          |  |  |
 3   5  0,1 1  2,2      3  4,2        -2 EOB
```


Zig-Zag Sequence

- The entropy encoder looks on the coefficients in this order:

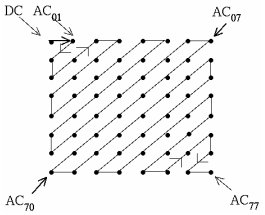

AC standard table

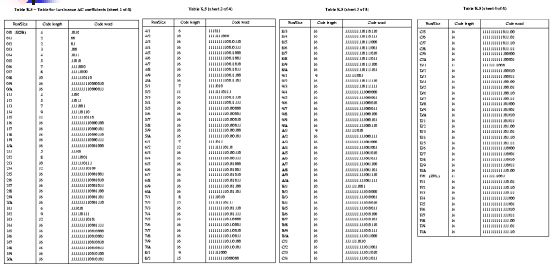

JPEG's Huffman standard tables

- A special Huffman tree can be built for each image.
- These tables (in this and next slide) are the default ones:

Table K.3 – Table for luminance DC coefficient differences

Category	Code length	Code word
0	2	00
1	3	010
2	3	011
3	3	100
4	3	101
5	3	110
6	4	1110
7	5	11110
8	6	111110
9	7	1111110
10	8	11111110
11	9	111111110

An example

- Least significant bit of byte 10000 in this picture changed from 0 to 1.

Original picture

Edited picture

Error in Baseline JPEG

- 3 blocks as above.
  - Synchronization when an EOB decoded correctly.

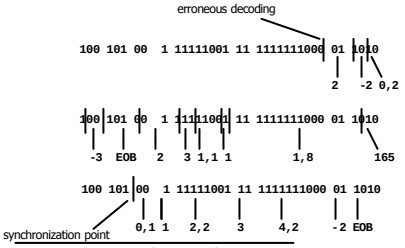

A Color Image - Example

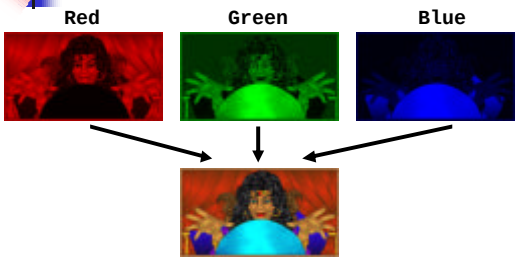

A Color Image

- There are some methods used in encoding color images.
- The most simple one is one data unit for Red, one for Green and one for Blue.
- The RGB components are interleaved together within the compressed data.
- Each component's data unit, can be a block of 8x8, but can be larger.
- YUV is also permitted. Y is the luminance component, while U and V are color difference components.

Grayscale vs. Color

- In fact, for comparable visual quality, a grayscale image needs perhaps 25% less space than a color image. Certainly, not the 66% less that you might naively expect.
- You can afford to lose a lot more information in the chrominance components than you can in the luminance component: the human eye is not as sensitive to high-frequency chrominance information as it is to high-frequency luminance.
- The luminance component is left at full resolution, while the chrominance components are often reduced 2:1 horizontally and either 2:1 or 1:1 (no change) vertically.

Huffman's tree

- The default Huffman table for the chrominance components of an image:

Table K.4 – Table for chrominance DC coefficient differences

Category	Code length	Code word
0	2	00
1	2	01
2	2	10
3	3	110
4	4	1110
5	5	11110
6	6	111110
7	7	1111110
8	8	11111110
9	9	111111110
10	10	1111111110
11	11	11111111110

- There is also a different tree for the chrominance AC components

Progressive Mode

- The progressive mode is intended to support real-time transmission of images.
- With each scan, the decoder can produce a higher-quality rendition of the image. Thus a low-quality preview can be sent very quickly, then refined as time allows.
- The total space needed is roughly the same as for a baseline JPEG image of the same final quality.
- A buffer is needed in the decoder.

Error in a color picture

- Least significant bit of byte 10000 in this picture changed from 0 to 1.
- Since the RGB components are interleaved together within the compressed data, components can be switched.
- In this picture the chrominance component was reduced 2:1 horizontally and 2:1 vertically.

More progressive modes

- The quantized DCT coefficients can be shown as a box:

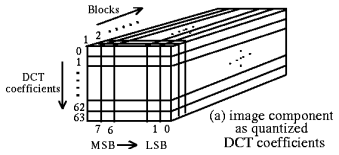

Progressive Mode

- The image is encoded in multiple scans rather than in a single scan.

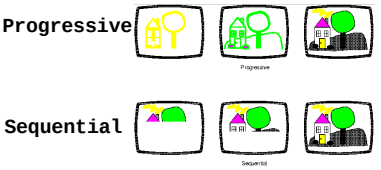

Spectral Selection

- First, the DCs are sent. Then the ACs are sent according to their order.

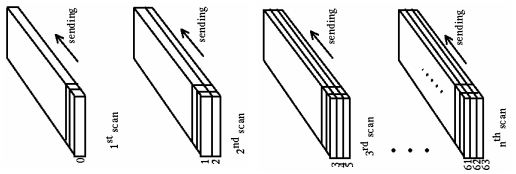

Sequential send

- When a non-progressive mode is used, the sending order is:

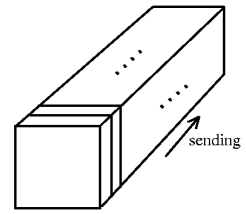

(b) Sequential encoding

Successive Approximation

- First, the MSB is sent. Then, the other lower bits are sent.

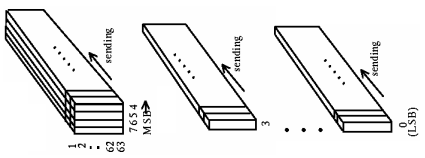

Spectral Selection (cont.)

- When there are just DCs, each 8x8 block is filled with equal pixels.

JPEG quality

- 1% 3203 bytes

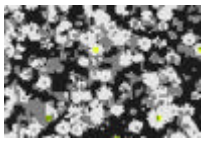

- 5% 11832 bytes

- 20% 32155 bytes

- 100% 284179 bytes

Successive Approximation

- Since ACs are usually low, most of the MSB are zeros. Hence the picture is filled with 8x8 blocks with equal pixels.