

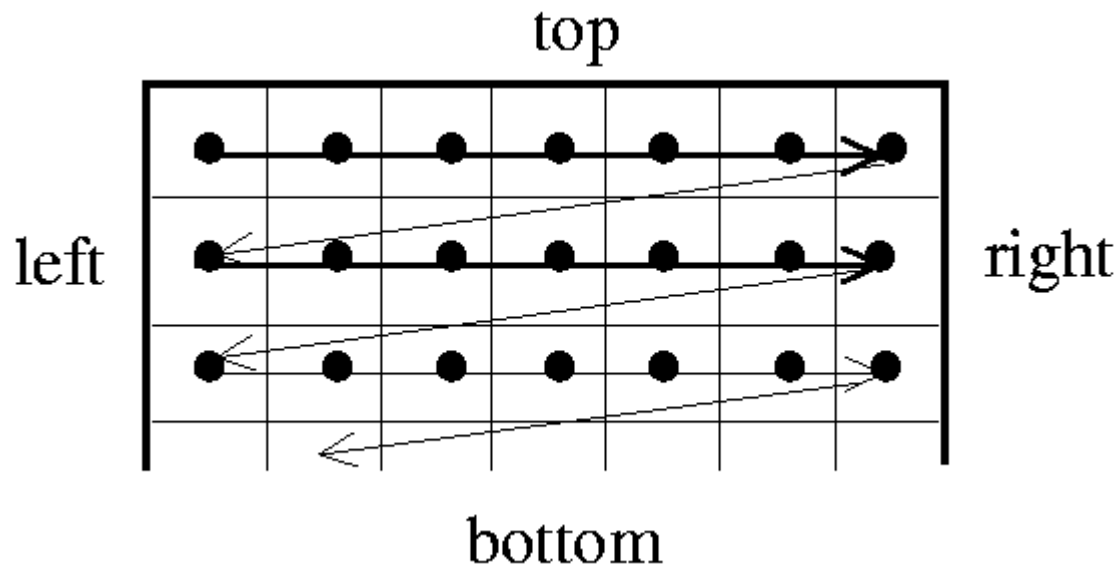
JPEG

- “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.



Encoding Order

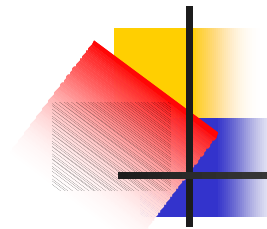
- Each block of 8×8 is treated separately.
- Order of blocks is:





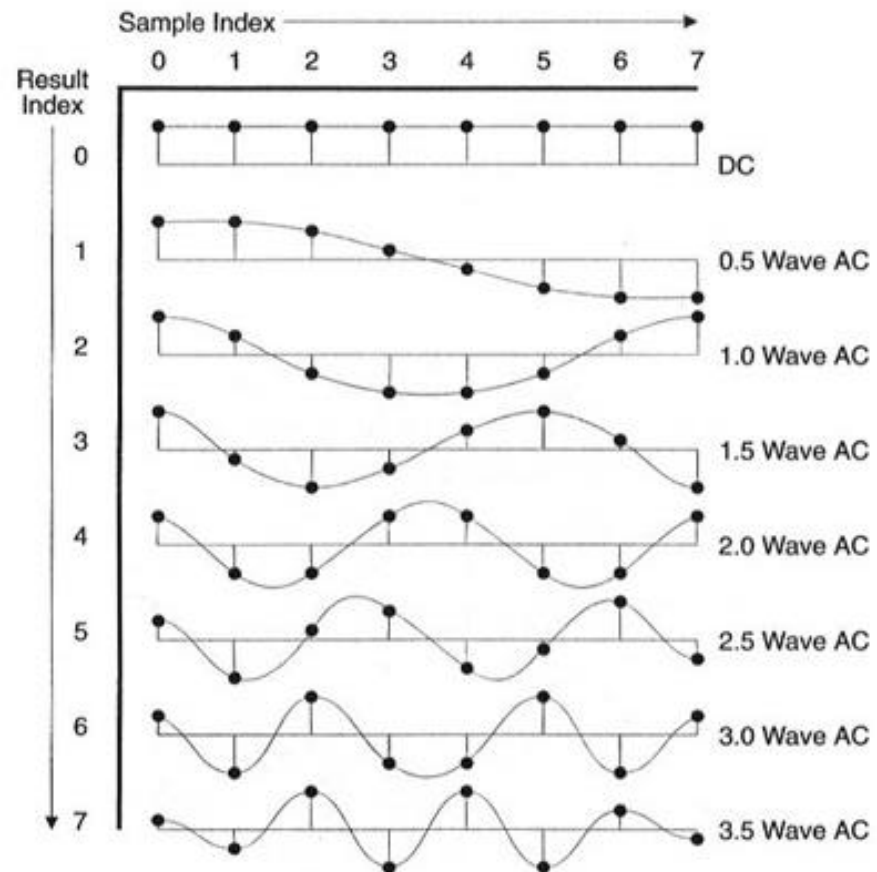
Baseline JPEG

- 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 loss of data.
- Compress the data by a version of Canonical Huffman coding.
 - ▶ Non-Baseline JPEG may use also Arithmetic coding.



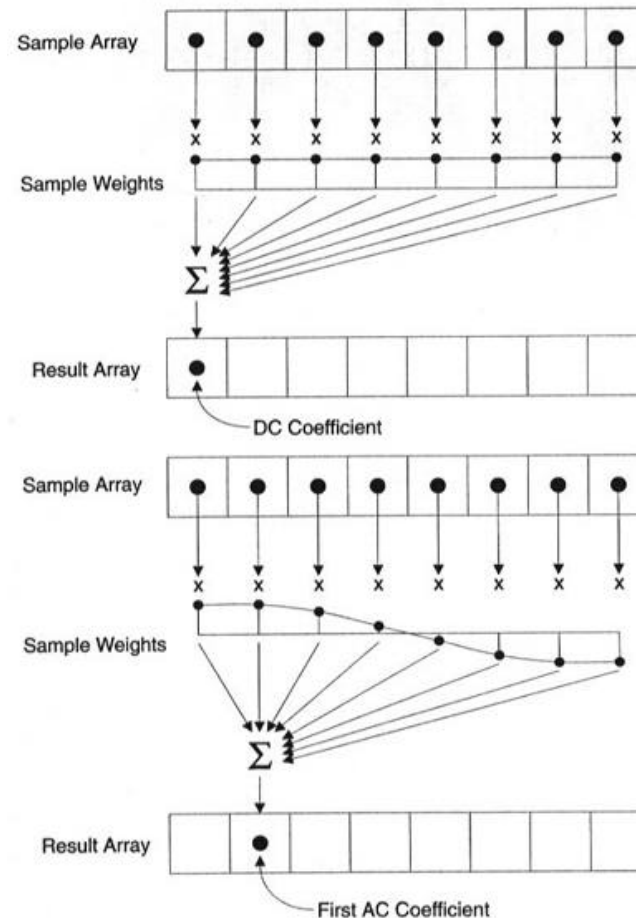
DCT

- 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.



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.





One dimensional DCT

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

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

$$C(u) = 1 \quad \text{If } u \neq 0$$

- 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 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 |



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 columns. 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 formula

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

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

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



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.



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 |

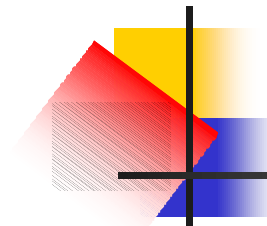
(a) Source image samples

| | | | | | | | |
|-------|-------|-------|------|------|------|------|------|
| 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



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 |

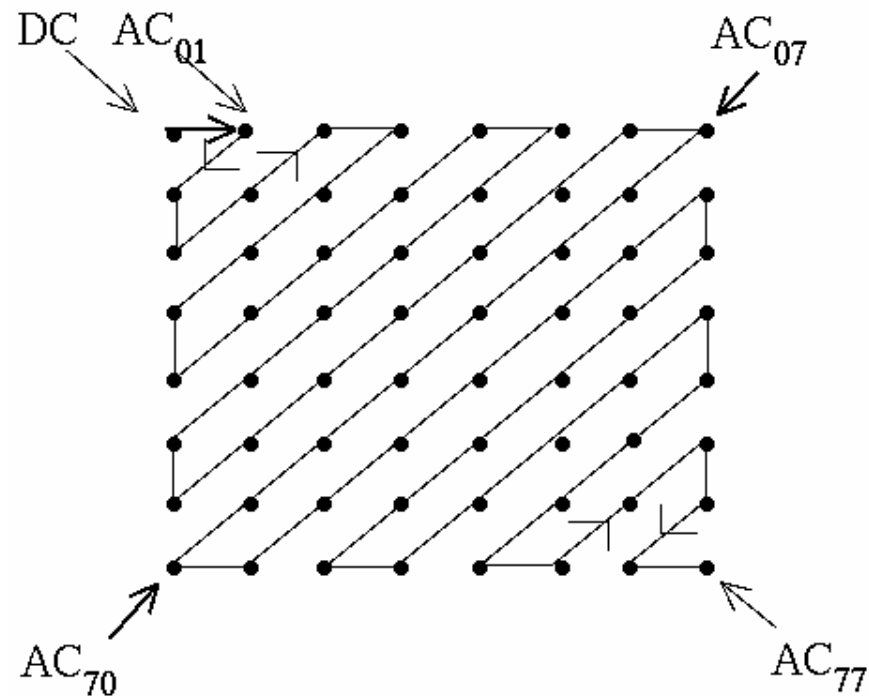
(b) denormalized quantized coefficients

| | | | | | | | |
|-----|-----|-----|-----|-----|-----|-----|-----|
| 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

Zig-Zag Sequence

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



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 |



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 |



AC standard table

Table K.5 – Table for luminance AC coefficients (sheet 1 of 4)

Table K.5 (sheet 2 of 4)

Table K.5 (sheet 3 of 4)

Table K.5 (sheet 4 of 4)

| RunSize | Code length | Code word |
|-----------|-------------|-----------------|
| 0/0 (EOB) | 4 | 1010 |
| 0/1 | 2 | 00 |
| 0/2 | 2 | 01 |
| 0/3 | 3 | 100 |
| 0/4 | 4 | 1011 |
| 0/5 | 5 | 11010 |
| 0/6 | 7 | 1111000 |
| 0/7 | 8 | 11111000 |
| 0/8 | 10 | 1111110110 |
| 0/9 | 16 | 11111111000010 |
| 0/A | 16 | 111111110000011 |
| 1/1 | 4 | 1100 |
| 1/2 | 5 | 11011 |
| 1/3 | 7 | 1111001 |
| 1/4 | 9 | 111110110 |
| 1/5 | 11 | 11111101110 |
| 1/6 | 16 | 111111110000100 |
| 1/7 | 16 | 111111110000101 |
| 1/8 | 16 | 111111110000110 |
| 1/9 | 16 | 111111110000111 |
| 1/A | 16 | 111111110001000 |
| 2/1 | 5 | 11100 |
| 2/2 | 8 | 11111001 |
| 2/3 | 10 | 1111110111 |
| 2/4 | 12 | 111111101100 |
| 2/5 | 16 | 111111110001001 |
| 2/6 | 16 | 111111110001010 |
| 2/7 | 16 | 111111110001011 |
| 2/8 | 16 | 111111110001100 |
| 2/9 | 16 | 111111110001101 |
| 2/A | 16 | 111111110001110 |
| 3/1 | 6 | 111010 |
| 3/2 | 9 | 111110111 |
| 3/3 | 12 | 11111110101 |
| 3/4 | 16 | 111111110001111 |
| 3/5 | 16 | 111111110001000 |
| 3/6 | 16 | 111111110001001 |
| 3/7 | 16 | 111111110001010 |
| 3/8 | 16 | 111111110001011 |
| 3/9 | 16 | 111111110001100 |
| 3/A | 16 | 111111110001101 |

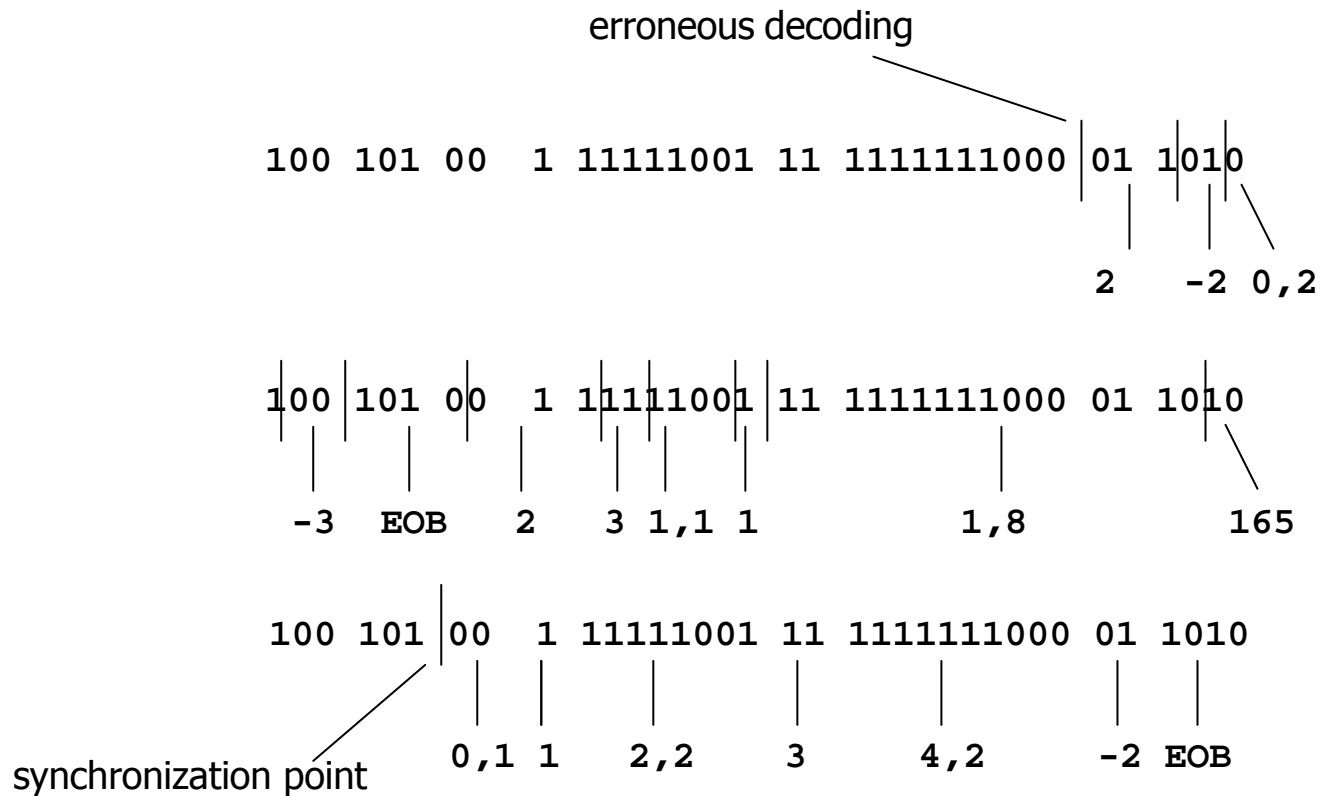
| RunSize | Code length | Code word |
|---------|-------------|------------------|
| 4/1 | 6 | 1111011 |
| 4/2 | 10 | 11111110000 |
| 4/3 | 16 | 1111111100010110 |
| 4/4 | 16 | 1111111100010111 |
| 4/5 | 16 | 1111111100011000 |
| 4/6 | 16 | 1111111100011001 |
| 4/7 | 16 | 1111111100011010 |
| 4/8 | 16 | 1111111100011011 |
| 4/9 | 16 | 1111111100011100 |
| 4/A | 16 | 1111111100011101 |
| 5/1 | 7 | 1111010 |
| 5/2 | 11 | 11111110111 |
| 5/3 | 16 | 1111111100011110 |
| 5/4 | 16 | 1111111100011111 |
| 5/5 | 16 | 1111111110100000 |
| 5/6 | 16 | 1111111110100001 |
| 5/7 | 16 | 1111111110100010 |
| 5/8 | 16 | 1111111110100011 |
| 5/9 | 16 | 1111111110101000 |
| 5/A | 16 | 1111111110101001 |
| 6/1 | 7 | 1111011 |
| 6/2 | 12 | 111111110110 |
| 6/3 | 16 | 1111111110100110 |
| 6/4 | 16 | 1111111110100111 |
| 6/5 | 16 | 1111111110101000 |
| 6/6 | 16 | 1111111110101001 |
| 6/7 | 16 | 1111111110101010 |
| 6/8 | 16 | 1111111110101011 |
| 6/9 | 16 | 1111111110101100 |
| 6/A | 16 | 1111111110101101 |
| 7/1 | 8 | 11111010 |
| 7/2 | 12 | 111111110111 |
| 7/3 | 16 | 1111111110101110 |
| 7/4 | 16 | 1111111110101111 |
| 7/5 | 16 | 1111111110110000 |
| 7/6 | 16 | 1111111110110001 |
| 7/7 | 16 | 1111111110110010 |
| 7/8 | 16 | 1111111110110011 |
| 7/9 | 16 | 1111111110110100 |
| 7/A | 16 | 1111111110110101 |
| 8/1 | 9 | 111111000 |
| 8/2 | 15 | 111111111000000 |

| RunSize | Code length | Code word |
|---------|-------------|------------------|
| 8/3 | 16 | 1111111110101110 |
| 8/4 | 16 | 1111111110101111 |
| 8/5 | 16 | 1111111110111000 |
| 8/6 | 16 | 1111111110111001 |
| 8/7 | 16 | 1111111110111010 |
| 8/8 | 16 | 1111111110111011 |
| 8/9 | 16 | 1111111110111100 |
| 8/A | 16 | 1111111110111101 |
| 9/1 | 9 | 111111001 |
| 9/2 | 16 | 1111111110111110 |
| 9/3 | 16 | 1111111110111111 |
| 9/4 | 16 | 1111111111000000 |
| 9/5 | 16 | 1111111111000001 |
| 9/6 | 16 | 1111111111000010 |
| 9/7 | 16 | 1111111111000011 |
| 9/8 | 16 | 1111111111000100 |
| 9/9 | 16 | 1111111111000101 |
| 9/A | 16 | 1111111111000110 |
| A/1 | 9 | 111111010 |
| A/2 | 16 | 1111111111000111 |
| A/3 | 16 | 1111111111000100 |
| A/4 | 16 | 1111111111000101 |
| A/5 | 16 | 1111111111000110 |
| A/6 | 16 | 1111111111000111 |
| A/7 | 16 | 1111111111001100 |
| A/8 | 16 | 1111111111001101 |
| A/9 | 16 | 1111111111001110 |
| A/A | 16 | 1111111111001111 |
| B/1 | 10 | 111111001 |
| B/2 | 16 | 1111111110100000 |
| B/3 | 16 | 1111111110100001 |
| B/4 | 16 | 1111111110100010 |
| B/5 | 16 | 1111111110100011 |
| B/6 | 16 | 1111111110101000 |
| B/7 | 16 | 1111111110101001 |
| B/8 | 16 | 1111111110101010 |
| B/9 | 16 | 1111111110101011 |
| B/A | 16 | 1111111110101100 |
| C/1 | 10 | 111111010 |
| C/2 | 16 | 1111111111011001 |
| C/3 | 16 | 1111111111011010 |
| C/4 | 16 | 1111111111011011 |

| RunSize | Code length | Code word |
|-----------|-------------|------------------|
| C/5 | 16 | 1111111110111000 |
| C/6 | 16 | 1111111110111001 |
| C/7 | 16 | 1111111110111010 |
| C/8 | 16 | 1111111110111011 |
| C/9 | 16 | 1111111111000000 |
| C/A | 16 | 1111111111000001 |
| D/1 | 11 | 11111110000 |
| D/2 | 16 | 1111111111000010 |
| D/3 | 16 | 1111111111000011 |
| D/4 | 16 | 1111111111000100 |
| D/5 | 16 | 1111111111000101 |
| D/6 | 16 | 1111111111001100 |
| D/7 | 16 | 1111111111001101 |
| D/8 | 16 | 1111111111010000 |
| D/9 | 16 | 1111111111010001 |
| D/A | 16 | 1111111111010010 |
| E/1 | 16 | 1111111111010011 |
| E/2 | 16 | 1111111111011000 |
| E/3 | 16 | 1111111111011001 |
| E/4 | 16 | 1111111111011100 |
| E/5 | 16 | 1111111111011101 |
| E/6 | 16 | 1111111111010000 |
| E/7 | 16 | 1111111111010001 |
| E/8 | 16 | 1111111111010010 |
| E/9 | 16 | 1111111111010011 |
| E/A | 16 | 1111111111011000 |
| F/0 (ZRL) | 11 | 11111110001 |
| F/1 | 16 | 1111111111011001 |
| F/2 | 16 | 1111111111011010 |
| F/3 | 16 | 1111111111011011 |
| F/4 | 16 | 1111111111011000 |
| F/5 | 16 | 1111111111011001 |
| F/6 | 16 | 1111111111011010 |
| F/7 | 16 | 1111111111011011 |
| F/8 | 16 | 1111111111111000 |
| F/9 | 16 | 1111111111111001 |
| F/A | 16 | 1111111111111010 |

Error in Baseline JPEG

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





An example

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

Original picture



Edited picture





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 8×8 , but can be larger.
- YUV is also permitted. Y is the luminance component, while U and V are color difference components.

A Color Image - Example

Red



Green



Blue





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



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.

Error in a color picture

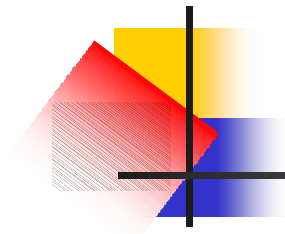
- Least significant bit of byte 10000 in this picture changed from 0 to 1.
- Since the RGB components are interleaved together within the compressed data, components can be switched.
- In this picture the chrominance component was reduced 2:1 horizontally and 2:1 vertically.





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.



Progressive Mode

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

Progressive



Progressive



Sequential

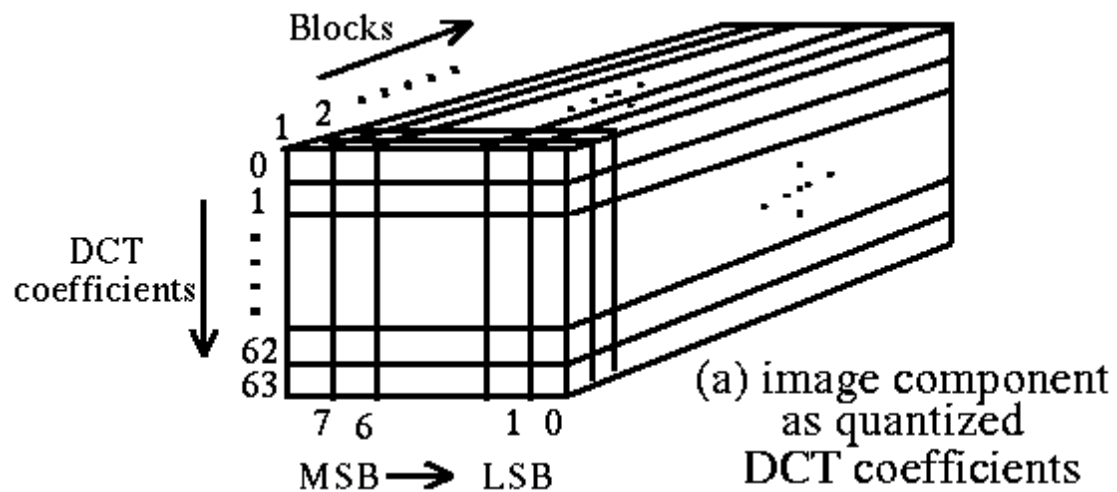


Sequential



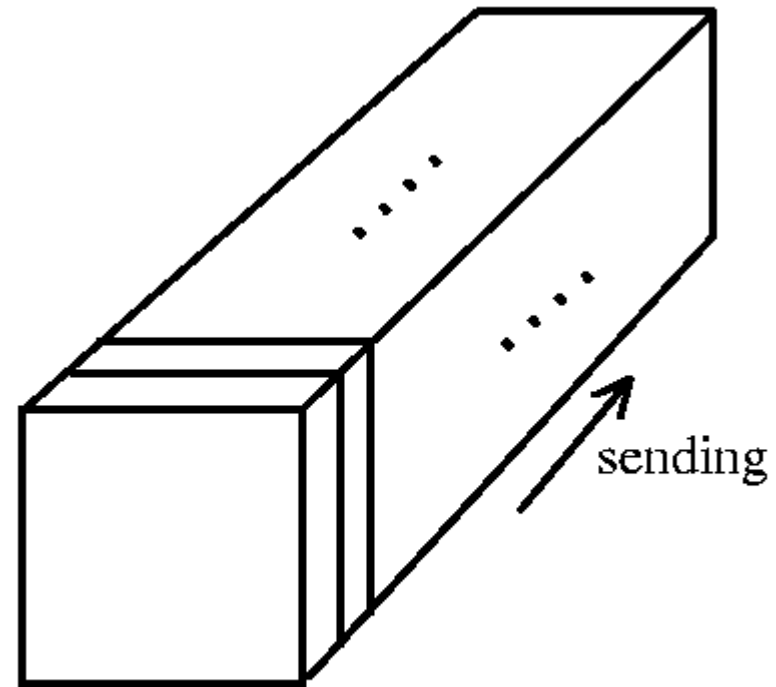
More progressive modes

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

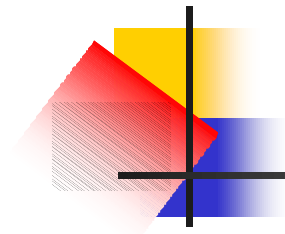


Sequential send

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

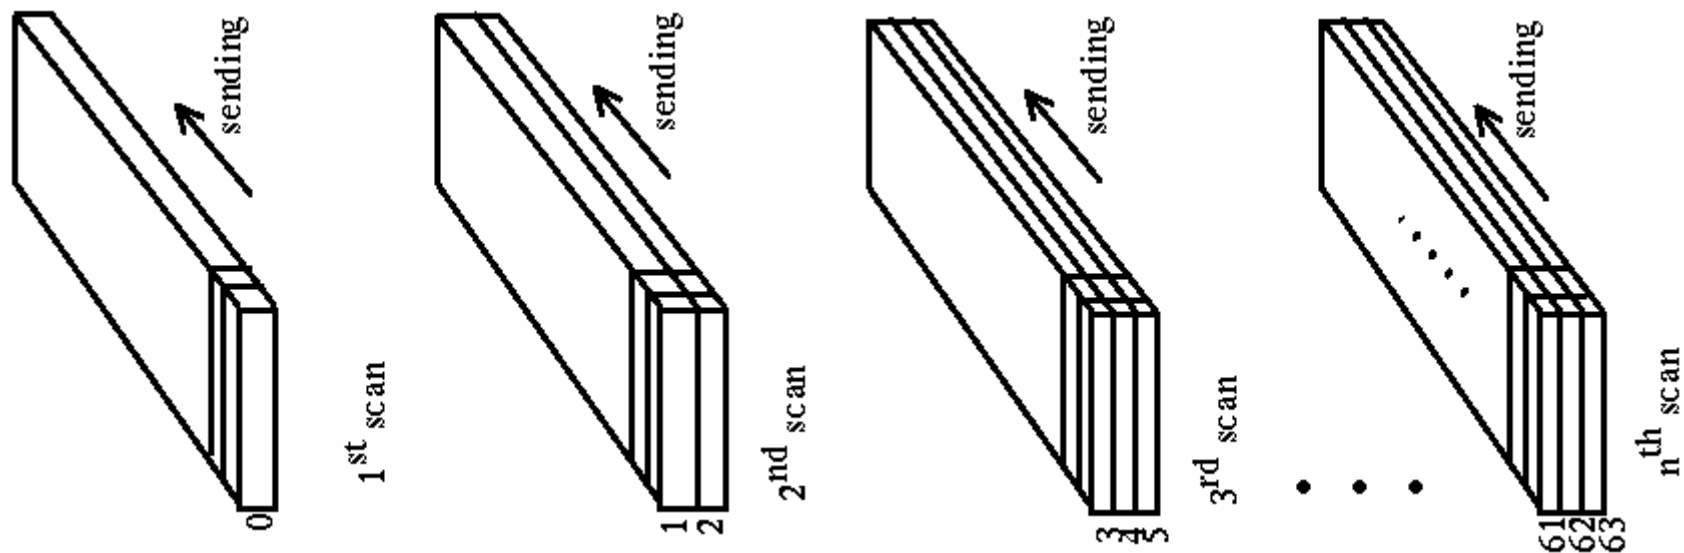


(b) Sequential encoding



Spectral Selection

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



Spectral Selection (cont.)

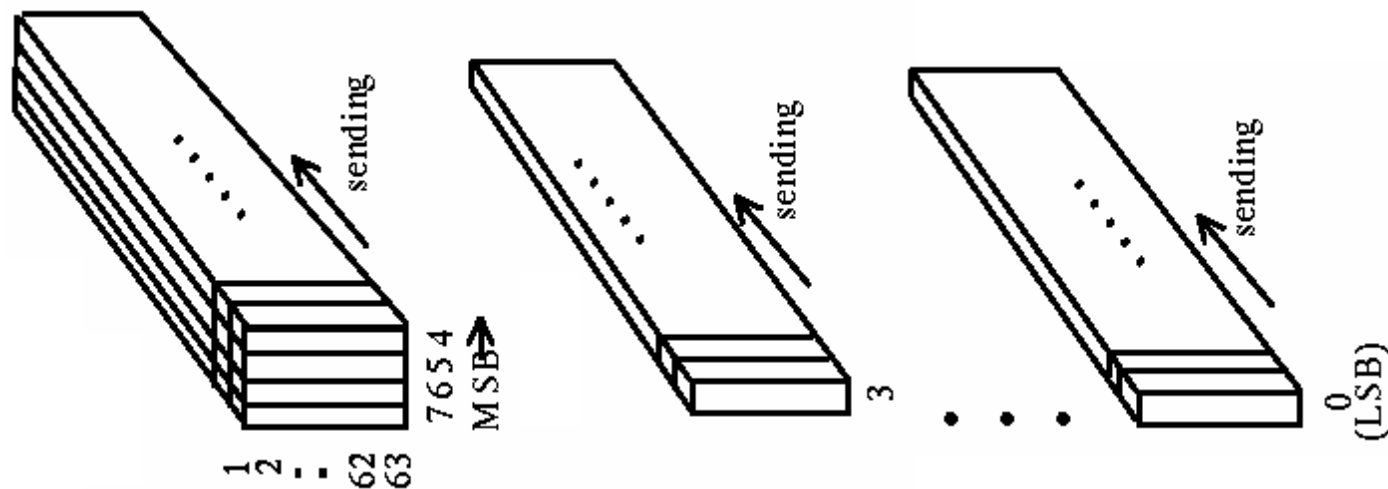
- When there are just DCs, each 8x8 block is filled with equal pixels.





Successive Approximation

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




Successive Approximation

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



JPEG quality

● 1%
3203
bytes



● 5%
11832
bytes



● 20%
32155
bytes



● 100%
284179
bytes

