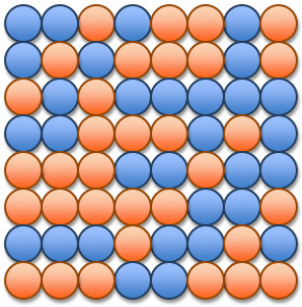

UNIT 7C

Data Representation: Images and Sound

Pixels

- An image is stored in a computer as a sequence of *pixels*, picture elements.

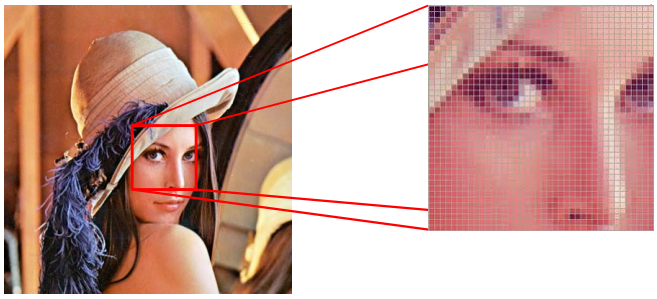

Resolution

- The resolution of an image is the number of pixels used to represent the image (e.g. 1024 X 768).
- Each pixel represents the average color in that region.
- The higher the resolution, the more accurate the image will appear.

Vector vs. Bitmap

- There are two major ways to store images:
  - Vector (a series of line segments)
  - Bitmap (a series of pixels)

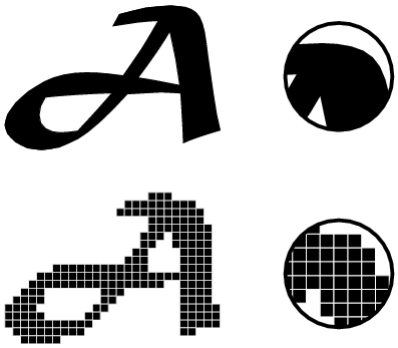

Storing Bitmap Images

- In bitmapped images, each pixel is represented in computer memory in binary, just like other data types.
- If pixels of an image are black or white only, then we only need 1 bit per pixel to store the image.

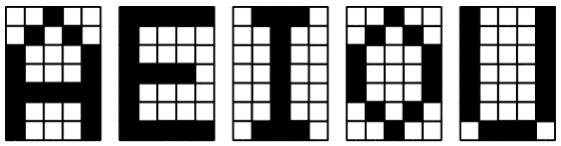

Grayscale Images

- Grayscale images contain pixels that are various shades of gray, from black (maximum gray) to white (minimum gray).
- If there are 256 levels of gray for pixels, we can represent each pixel using 8 bits.
  11111111 = white
  : (shades of gray)
  00000000 = black

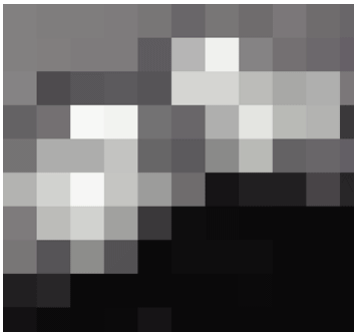

256-color images (8-bit color)

- Each pixel is represented with a 8-bit value that is an index into a table of 256 colors.

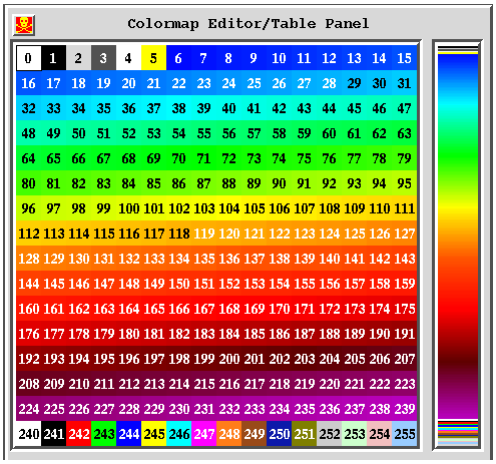

RGB-color images (24-bit color)

- Colors are represented as mixtures of red (R), green (G), and blue (B).
- Each pixel is represented using three 8-bit values, one for each color component.
- This representation allows for 2^{24} = 16,777,216 different colors.
- This representation is also called “true color”.

(image from Wikipedia)

RGB example

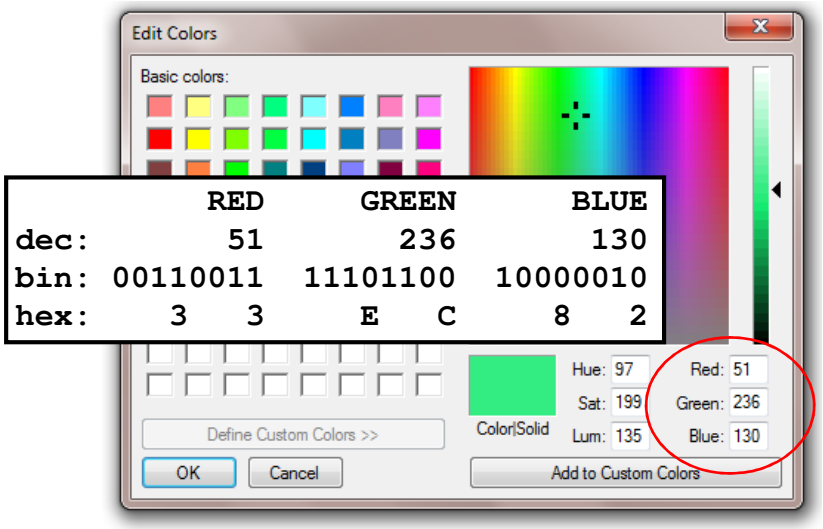

Comparing Representations

- If an image has a resolution of 640 X 480 (307,200 pixels), how many bytes does each representation require?
  - B&W 38,400 bytes
  - 8-bit grayscale 307,200 bytes
  - 256-color (8-bit color) 307,200 bytes
  - True color (RGB) 921,600 bytes
- A single RGB screen image of size 1600 X 1200 requires over 5.76 million bytes!

GIF

- Developed by CompuServe in the late 1980s to represent 8-bit images efficiently.
- Each pixel is an 8-bit value, mapping to a table of 256 24-bit RGB colors.
- Useful for representing images with fewer colors or large areas of color like company logos.

GIF compression

- A lossless compression technique is used for GIF files called run-length encoding.
- Instead of storing the 8-bit value for every pixel, we store an 8-bit value along with how many of these occur in a row (run).
- This saves a lot when there are large runs of the same color.

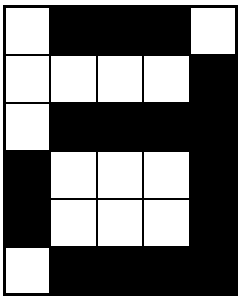

Color, Run, Color, Run, ...

255,1,0,3,255,1

255,4,0,1

255,1,0,4

0,1,255,3,0,1

0,1,255,3,0,1

255,1,0,4

(Colors: 0=Black, 255=White)

Comparison

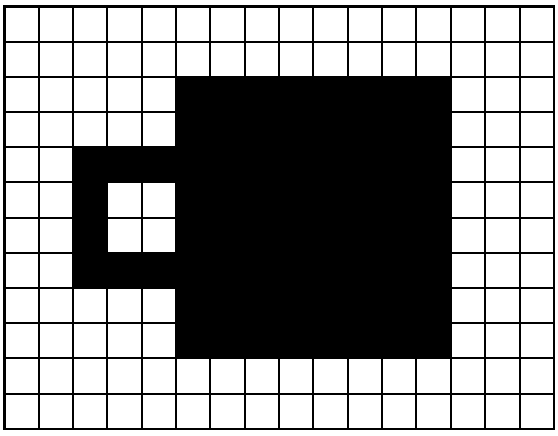

GIF	BMP
2 bytes	16 bytes
2 bytes	16 bytes
6 bytes	16 bytes
6 bytes	16 bytes
6 bytes	16 bytes
10 bytes	16 bytes
10 bytes	16 bytes
6 bytes	16 bytes
6 bytes	16 bytes
6 bytes	16 bytes
2 bytes	16 bytes
2 bytes	16 bytes
64 bytes	192 bytes

JPEG (JPG)

- A lossy compression technique used generally for photographic images.
  - Uses a variant of Huffman encoding.
- Supports varying levels of compression.

Higher quality
Compression 2.6:1
(images from Wikipedia)

Medium quality
Compression 23:1

Lowest quality
Compression 144:1

Sound

- When an instrument is played or a voice speaks, changes occur in air pressure which our ears interpret as sound.


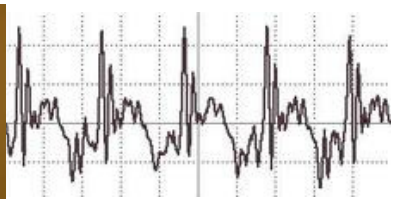


Sampling

- A sound is represented digitally by sampling an electronic version of the audio signal.
- The amplitude of the signal is measured (sampled) at equally-spaced time intervals.
- The amplitude axis is divided into equally-spaced intervals depending on how many bits we want to store per sample.

Sampling

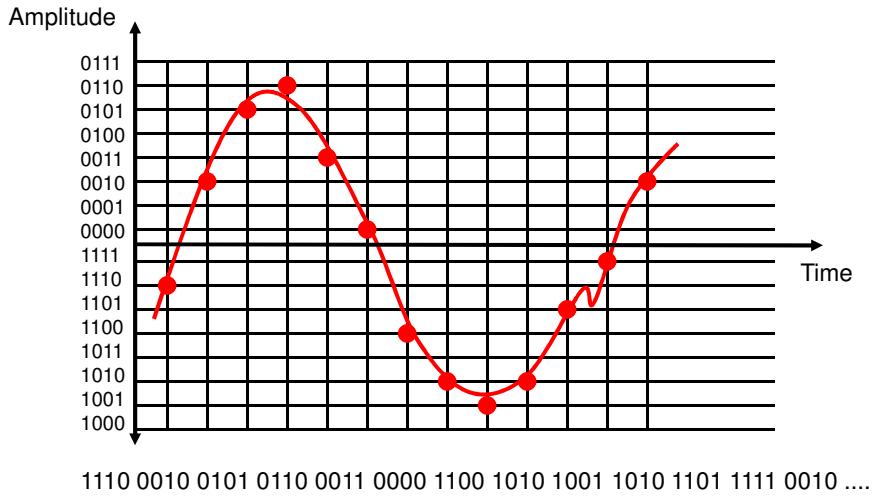

Sampling

- In order to reproduce the audio waveform as accurately as possible, we need to increase the sampling rate (samples per second) and the number of amplitude levels (bits per sample).
- Note in the previous picture how some of the samples had to be moved up or down to match an amplitude level and some finer changes in the sound signal could be missed if the sampling rate is too low.

Sampling

- Digital recordings are typically sampled at 44,100 samples per second (Hertz – Hz).
  - This is due to the "sampling theorem" that states that the sampling rate must be at least twice the highest frequency in the sound, and humans can hear up to approx. 20,000 Hz.
- For accurate amplitude readings, sound is often sampled at 16-bits per sample (so there are 65,536 amplitude levels that can be measured).
  - Some systems sample at finer amplitude levels (e.g. 24 bits per sample)

Example

- For 1 second of 16-bit stereo sound sampled at 44,100 Hz (samples per second), we would need:

  16 bits/sample * 44100 samples/second
  * 2 channels * 1 second/channel

  = 1,411,200 bits = 176,400 bytes
- WAV files and CD audio files store uncompressed samples.

MP3

- MP3 is a lossy compression technique.
- This format takes advantage of some facts about human hearing.
  - We can't hear certain sounds (very low or very high frequencies).
  - Our hearing is better in mid range frequencies than on the low and high ends.
  - If a loud and soft sound play at the same time, we can't hear the soft sound.
- MP3 filters the audio signal based on these properties.

MP3 Compression

- Like JPEG, MP3 has various levels of compression:

Bit Rate	Compression Ratio	Comments
256Kbps	5:1	Supreme (near best)
192Kbps	7:1	Excellent (better)
128Kbps	11:1	CD quality (good)
96Kbps	19:1	Near CD quality (fair)
64Kbps	22:1	FM quality (poor)

- MP3 also has Variable Bit Rate (VBR) since compression ability can vary at different segments of the digital recording.