

Data Compression

G. W. Cox -- Fall 2007

Compression

The idea

- Say we have to send a 1K x 1K pixel, 24-bit-color image = 3MB.
- As-is, the transmission delay for a 64Kbps link will be ~7min.
- If we can compress it so that it is 10 times smaller ("10:1 compression ratio"), it will be 300KB, which takes ~42sec to send.

G. W. Cox -- Fall 2007

Compression

Types of compression algorithms

cs570

- Lossless
 - The original data can be perfectly recovered by decompression.
 - Needed for text files, code, executables...
- Lossy
 - The original data cannot be exactly recovered.
 - OK for photos, video, voice...

cs570

Lossless algorithms

Run-Length Encoding

cs570

- The idea: Replace runs of a symbol with one copy of the symbol + the count.
- Example:
 - “AAAABCCCCDD” (10 symbols) → A4BC3D2 (7 symbols)
- Good where there are long runs of a symbol (faxes, .bmp images...)
- Typical Performance:
 - For scanned text, 8:1 compression is common

G. W. Cox -- Fall 2007

Compression

Huffman encoding

cs570

- The idea: Instead of using the same number of bits for every symbol in a string, we encode each symbol with a number of bits that is inversely proportional to the symbol's frequency in the string.

- Example:

1000-symbol string With 4 symbols		Normal encoding		Huffman encoding	
Frequency:		A	00	A	01
B	30%	B	01	B	0
C	50%	C	10	C	011
D	15%	D	11	D	0111
	5%		—		
		Total bits	2000	Total bits	1750

*Performance varies with string and frequencies.
Not good for even distribution of symbols*

G. W. Cox -- Fall 2007

Compression

Pulse-Code Modulation (Differential)

cs570

- The idea: Symbols are encoded based on their “distance” from a reference symbol.
- Example:
 - ASCII Source string = “AAABCCDDEGF” (88 bits)
 - We choose to encode symbols with 8 bits and distances with 2 bits (when the distance becomes greater than 3, choose a new base symbol).
 - Encoded string = “A0012233E21”
 - Number of bits = 82222222822 (35 bits)
- Effective when adjacent symbols are “similar” (e.g, imagery)

G. W. Cox -- Fall 2007

Compression

Dictionary methods (Lempel-Ziv) - 1

cs570

- The idea: Build a dictionary of the “phrases” in the source. Assign each phrase a number and xmit the numbers. (phrases can be anything – fixed # characters, words in a language...)
- Note: you either must have a standard dictionary, or you must send the dictionary along with the encoded string
- Effective when the number of phrases is small relative to the length of the string.
- Used in Unix “Compress” and Windows “Zip”

G. W. Cox -- Fall 2007

Compression

Dictionary methods (Lempel-Ziv) - 2

cs570

Example:

"one small step for man one giant leap for mankind" (392 bits in ASCII)

Dictionary	
one_	000
small_	001
step_	010
for_	011
man_	100
giant_	101
leap_	110
mankind_	111

Encoded form: 000 001 010 011 100 000 101 110 011 111 → 30 bits

If you must send the dictionary, you have an additional 42 characters → 336 bits

So, for standard dictionary = 30 bits, ~ 13:1 compression

for non-standard dictionary = 366 bits, ~ 1.1:1 compression

G. W. Cox -- Fall 2007

Compression

Graphical Interchange Format (GIF)

cs570

- Used for image compression – only lossless when original image has fewer than 256 colors
- Algorithm:
 1. If necessary, reduce # colors to 256 (8 bits per pixel)
 2. Run LZ on the result
- Works best for images that have repeating patterns (e.g, photo of a picket fence or tiled floor).
- Typical: 4:1 – 10:1 compression ratio

G. W. Cox -- Fall 2007

Compression

GIF example (lossy)

cs570



Original JPEG (24-bit color)



GIF (8-bit color)

G. W. Cox -- Fall 2007

Compression

cs570

Lossy algorithms

G. W. Cox -- Fall 2007

Compression

JPEG

cs570

- JPEG = ISO "Joint Photographic Experts Group"
- Algorithm:
 1. Do a spectral analysis of each 8x8 pixel block (lossless)
 2. Quantize (round) the spectral results (lossy)
 3. Run RLE and Huffman on the result (lossless)
- User controls quality of reconstructed image by choosing how much rounding is done in step 2 → trade quality for size
- Typical:
 - HQ: 10:1 → 20:1
 - MQ: 30:1 → 50:1
 - LQ: 60:1 → 100:1

G. W. Cox -- Fall 2007

Compression

JPEG example

cs570



High Quality
34KB



Medium Quality
6KB



Low Quality
2KB

G. W. Cox -- Fall 2007

Compression

Video streams

cs570

- Uncompressed VCR-class video
 - 352 x 240 pixels, 24 bit color, 25 frames/sec
 - ~51Mbps
- Uncompressed Broadcast-quality standard video (NTSC)
 - 720 x 480 pixels
 - ~207Mbps
- Uncompressed HDTV-quality video
 - 1920 x 1080 pixels
 - ~1.2Gbps
- Significant compression needed to carry these at reasonable bandwidths

G. W. Cox -- Fall 2007

Compression

MPEG

cs570

- MPEG = Motion Picture Experts Group
- Standards:
 - MPEG-1: VCR-quality video, 1.2Mbps
 - MPEG-2: HDTV-quality video, 4-8Mbps
 - MPEG-3 ("MP3"): CD-quality audio, 96-128Kbps
 - MPEG-4: MPEG-2 + VRML, DRM, Interactive video

G. W. Cox -- Fall 2007

Compression

MPEG video concepts

cs570

- Based on JPEG
 - We could just do a JPEG image of each frame of video
 - But we can improve on that by taking advantage of common features across frames



Same image, translated

*There is generally
A significant
amount
of repeated
Sub-images
from
one frame to
another*

G. W. Cox -- Fall 2007

Compression

MPEG-2 Approach

cs570

- Each frame in the video stream are encoded as one of the following types:
 - I-frame:
 - A JPEG-encoded still image
 - I-frames are included periodically as a reference
 - P-frame:
 - An encoding of the differences between the last frame and this one
 - Image broken up into 16x16 subimages that are similar in both frames, maybe in different places ("macroblocks")
 - Differences expressed in terms of movement, transformation of macroblocks.
 - B-frame:
 - Like a P-frame, but based on either previous or next frame

G. W. Cox -- Fall 2007

Compression

Audio streams

cs570

- Uncompressed POTS-quality sound
 - 4Khz Bandwidth
 - 8bit sample / 125usec → 64Kbps
- Uncompressed CD-quality audio:
 - 44Khz Bandwidth
 - 16-bit sample / 23usec x 2 channels for stereo → 1.4Mbps

G. W. Cox -- Fall 2007

Compression

MPEG-3 audio compression

cs570

- The idea: don't transmit the sounds people don't hear
 - Strong sounds at one frequency tend to mask softer sounds at nearby frequencies – do not transmit the masked sounds
- Algorithm:
 - For every 26msec, measure power in 32 frequency bands
 - Allocate higher bps to bands with most power, less bps to bands with less power
 - Huffman code the allocated bits
- Performance:
 - CD quality at 128Kbps – 12:1 compression

G. W. Cox -- Fall 2007

Compression