Steganographic techniques

A brief survey

CS349 Cryptography

Department of Computer Science
Wellesley College

Steganographic taxonomies

- Stenanographic systems can be grouped by the type of covers used (graphics, sound, text, executables) or by the techniques used to modify the covers
  - substitution,
  - transform domain techniques,
  - spread spectrum techniques,
  - statistical method,
  - distortion techniques
  - cover generation methods

Using grammar for mimicry*

It's time for another game between the Whappers and the Blogs in scenic downtown Blovonia . I've just got to say that the Blog fans have come to support their team and rant and rave . Play Ball ! Top of the inning. Yup. What a game so far today . The pitcher spits. Frank Gavi adjusts the cup and enters the batter's box . Yeah. It's a rattler . He swings for the stands, but no contact . Here's the pitch It's a bouncing knuckleball . . .

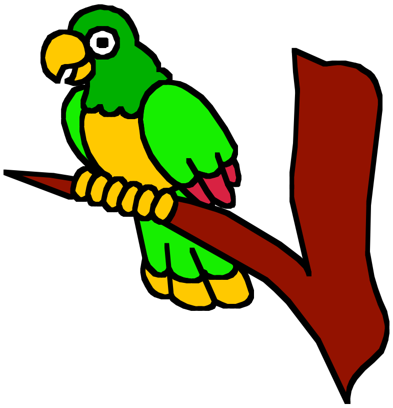

*http://www.wayner.org/texts/mimic/

Computer graphics

- A digitized photograph is represented by a matrix of numbers that stand for the intensity of light emanating from a particular place at a particular time.
- The array of charged-coupled devices that convert photons to bits are not perfect.

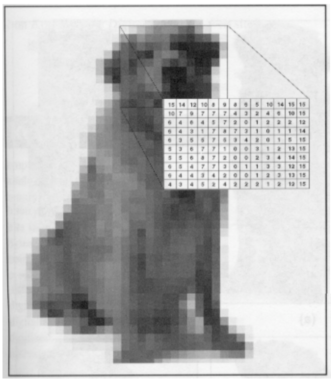

Sound

- Digitized sounds are lists of numbers representing pressure hitting a microphone at a sequence of time slices.
- Like digitized graphics, these numbers are imprecise and there is lots of room to hide information in the noise.

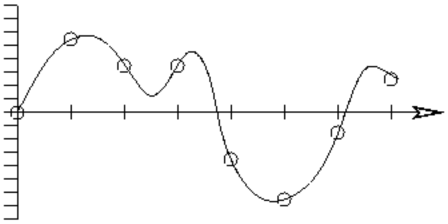

Hidden volumes

- At top scale, a Kodak photo-CD images is 3072 by 2048 pixels.
- Eight bits are used to encode each of the amount of red, blue, and green (or the amount cyan, magenta, and yellow).
- Borrowing the LSB of each color of each pixel yields about 1.8 megabytes of storage.

Where to hide

- The least significant bits seem close to random, but often contain hidden patterns.
- Destroying these subtle statistical correlations may well give the show away.

Peter Wayner's desk
Disappearing Cryptography.

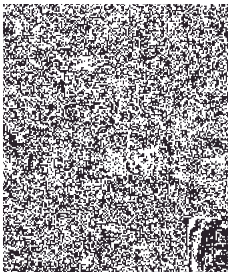
Least significant bits of same

Another problem

- Compressions artists often beat stenanographers to the punch.
- JPEG can often get by with one or two bits per pixel and save a factor of ten in file size.

Peter Wayner's desk again

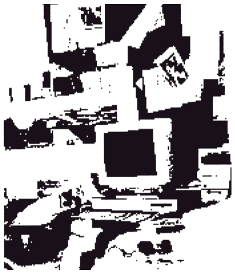
Most significant bits of same

Hiding in a GIF file

- GIF compression maintains a table of up to 256 colors that best represent the image.
- The color of each pixel is described by indexing into this color table.
- Changing the least significant bit may significantly change the image.

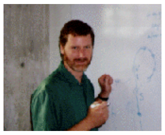

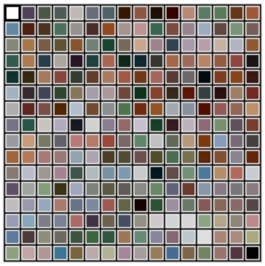

A possible solution

- Hide and Seek 4.1 converts the color table from 256 colors to 128, then duplicates each of these colors so that adjacent entries in the color table are duplicates of each other.
- Unfortunately, this technique leaves a large red flag for anyone scanning GIFs for hidden information.

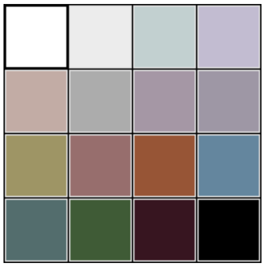

Another solution

- EzStego sorts the color palette so that the colors flow smoothly into each other.
- The hope is that changing the least significant bit doesn't drastically change the color.

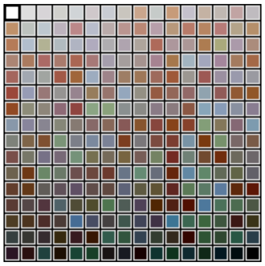

The Traveling Salesperson Problem

- Ordering colors in one dimension is easy. Ordering a three dimensional color space is not.
- EzStego treats the colors as cities in RGB space and and tries to find the shortest path through all of the stops.

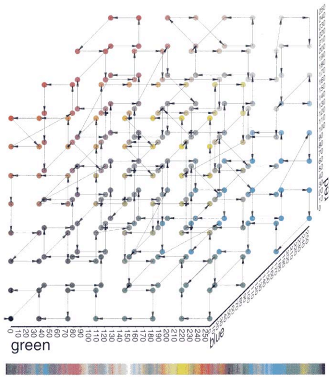

A heuristic for TSP

Begin with a list of two cities, {c_1 c_2}. Set $C = c_1$.

1. Find city, *d*, that is farthest from *C*.
2. Scan the list to find i such that

   □(d, c_i) + □(d, c_{i+1}) is minimized.
3. Insert *d* between *i* and *i*+1 and set *C* = *d*.

EzStego covers its tracks

1. Sort the palette so closest colors fall next to each other.
2. Encode the encrypted message by twiddling the least significant bit.
3. Unsort the palette by renumbering all of the colors with their original values.
4. Ship the image.
5. Receiver resorts palette using same algorithm and extracts bits by using the sorted palette.

GIFShuffle

- A palette containing N colors can be permuted in N! different ways. Each permutation may encode $\log_2(N!)$ bits of information.

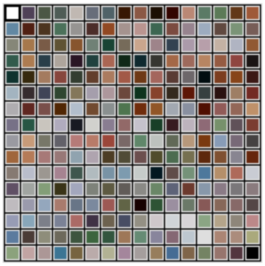

Original palette

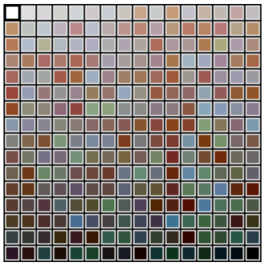

Sorted palette

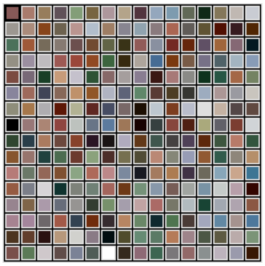

Permuted palette

Transform domain techniques

- LSB modification techniques are highly vulnerable to even small cover modifications.
- Transform domain methods hide messages in significant areas of the cover image making them more robust modification than LSB methods.
- JPEGs are perfect candidates for hiding information using such methods.

Fourier series

- Fourier series are expansions of periodic functions f(x) in terms of an infinite sum of sines and cosines of the form

$$f(x) = \sum_{n=0}^{\infty} a'_n \cos(nx) + \sum_{n=0}^{\infty} b'_n \sin(nx).$$

- Successive partial sums produce increasingly accurate approximations to the f(x).

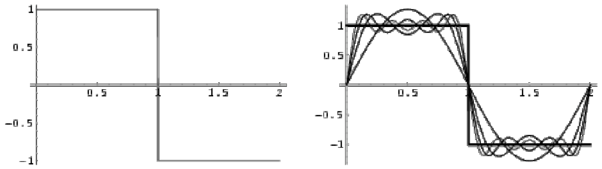

Cosines suffice

- Cosine functions similar to those shown below are used to model sound waves in MP3 music compression.

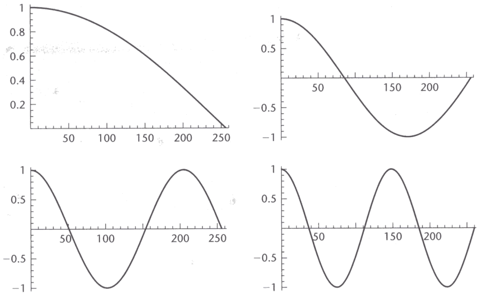

Gimme a W

- The Wellesley W was recreated using the four cosine functions shown on the previous slide:

  1.0 cos(x) + 0.5 cos(2x) - 0.8 cos (3x) + 0.3 cos(4x).
- MP3 compression uses Discrete cosine transforms to find cosine coefficients given sequences of sound samples.

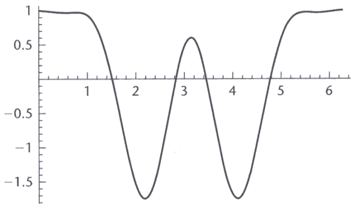

A two-dimensional discrete analogue

- The discrete cosine transform of a two-dimensional N x N array, s, of real numbers is given by

$$S(u,v) = \frac{2}{N}C(u)C(v)\sum_{x=0}^{N-1}\sum_{y=0}^{N-1} s(x,y)\cos(\frac{\pi u(2x+1)}{2N})\cos(\frac{\pi(2y+1)}{2N})$$

- The inverse cosine transform is given by

$$s(x,y) = \frac{2}{N}C(u)C(v)\sum_{u=0}^{N-1}\sum_{v=0}^{N-1} S(u,v)\cos(\frac{\pi u(2x+1)}{2N})\cos(\frac{\pi(2y+1)}{2N})$$

where C(u) equals $2^{-1/2}$ if u is 0 and 1 otherwise.

JPEG compression

Convert image into YCbCr color space and break up each color plane into 8x8 blocks.

Compress the quantized DCT coefficients using entropy coder.

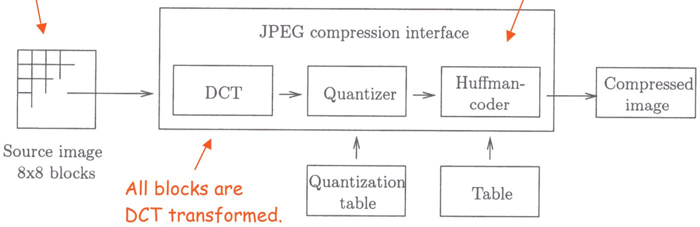

All blocks are DCT transformed.

Divide DCT coefficients by predefined quantization values to modulate influence of spectral components on image.

Quantization factors

- DCT coefficients corresponding to highest frequency cosines are likely to be dominated by noise.
- Values in the quantization table are calculated to reduce the influence of these "noisy" coefficients.

(u, v)	0	1	2	3	4	5	6	7
0	16	11	10	16	24	40	51	61
1	12	12	14	19	26	58	60	55
2	14	13	16	24	40	57	69	56
3	14	17	22	29	51	87	80	62
4	18	22	37	56	68	109	103	77
5	24	35	55	64	81	104	113	92
6	49	64	78	87	103	121	120	101
7	72	92	95	98	112	100	103	99

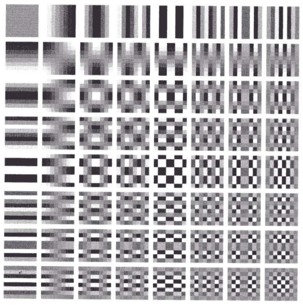

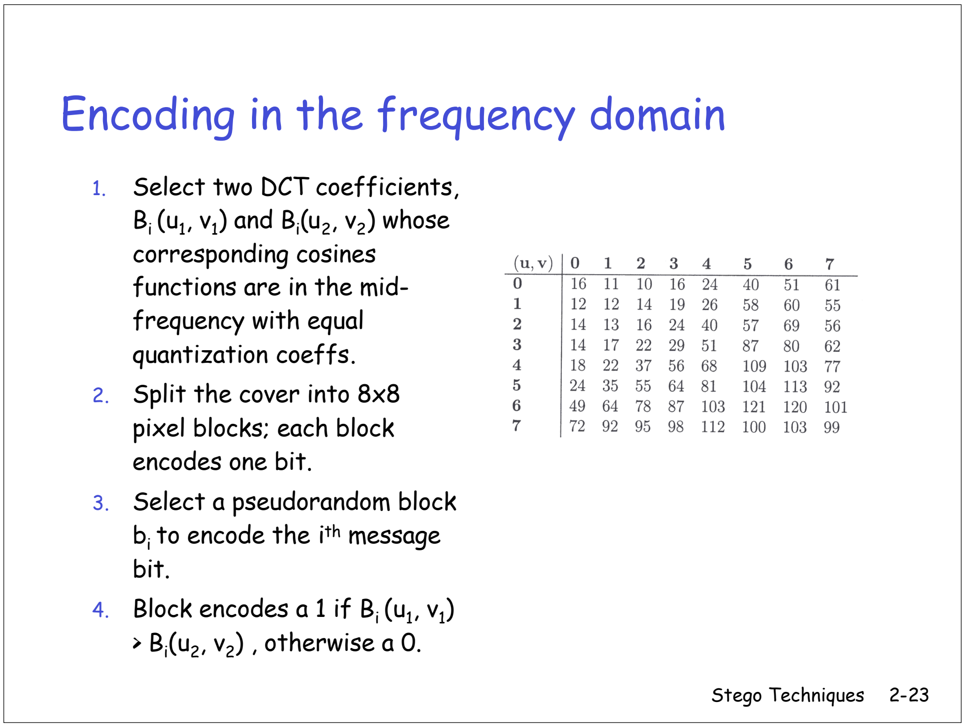

Encoding in the frequency domain

1. Select two DCT coefficients, $B_i(u_1, v_1)$ and $B_i(u_2, v_2)$ whose corresponding cosines functions are in the mid-frequency with equal quantization coeffs.
2. Split the cover into 8x8 pixel blocks; each block encodes one bit.
3. Select a pseudorandom block b_i to encode the i^{th} message bit.
4. Block encodes a 1 if $B_i(u_1, v_1) > B_i(u_2, v_2)$, otherwise a 0.

(u, v)	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**
0	16	11	10	16	24	40	51	61
1	12	12	14	19	26	58	60	55
2	14	13	16	24	40	57	69	56
3	14	17	22	29	51	87	80	62
4	18	22	37	56	68	109	103	77
5	24	35	55	64	81	104	113	92
6	49	64	78	87	103	121	120	101
7	72	92	95	98	112	100	103	99