Digital Watermarking

Introduction

[1]More information is transmitted in a digital format now than ever, and the growth in this trend will not plateau in the foreseeable future. Digital information is susceptible to having copies made at the same quality as the original.

There are many types of digital information and data. The types concentrated on in this report are:

- Digital Images
- Digital Audio, and
- Digital Videos

A watermark is [2]A pattern of bits inserted into a digital image, audio or video file that identifies the file's copyright information (author, rights, etc.). The name “watermark” is derived from the faintly visible marks imprinted on organisational stationery.

Unlike printed watermarks, which are intended to be somewhat visible (like the very light compass stamp watermarking this report), digital watermarks are designed to be completely invisible, or in the case of audio clips, inaudible.

In addition, the bits representing the watermark must be scattered throughout the file in such a way that they cannot be identified and manipulated. And finally, a digital watermark must be robust enough to survive changes to the file its embedded in, such as being saved using a lossy compression algorithm eg: JPEG.

Satisfying all these requirements is no easy feat, but there are a number of companies offering competing technologies. All of them work by making the watermark appear as *noise* - that is, random data that exists in most digital files anyway.

Digital Watermarking works by concealing information within digital data, such that it cannot be detected without special software with the purpose of making sure the concealed data is present in all copies of the data that are made whether legally or otherwise, regardless of attempts to damage/remove it.

---

1 http://www.itacs.uow.edu.au/research/smicl/IVM/Digital_watermarking_tutorial.html

2 http://www.webopedia.com/TERM/D/digital_watermark.html

Digital Watermarking

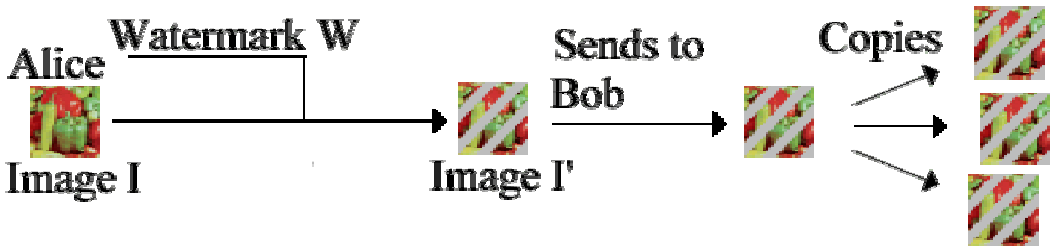

The purpose of digital watermarks is to provide copyright protection for intellectual property that is in digital format.

[3]As seen above, Alice creates an original image and watermarks it before passing it to Bob. If Bob claims the image and sells copies to other people Alice can extract her watermark from the image proving her copyright to it.

The caveat here is that Alice will only be able to prove her copyright of the image if Bob hasn't managed to modify the image such that the watermark is damaged enough to be undetectable or added his own watermark such that it is impossible to discover which watermark was embedded first.

[3] Secure Multimedia Information Communication Research Labs, University Of Wollongong
http://www.itacs.uow.edu.au/research/smicl/IVM/Digital_watermarking_tutorial.html

Digital Watermarking

Technical Details

Digital watermarking technology makes use of the fact that the human eye has only a limited ability to observe differences. Minor modifications in the colour values of an image are subconsciously corrected by the eye, so that the observer does not notice any difference.

While vendors of digital watermarking schemes do not publicly release the exact methods used to create their watermarks, they do admit to using the following basic procedure (with obvious variations and additions by each vendor).

A secret key (string or integer) produces a random number which determines the particular pixels, which will be protected by the watermarking. The watermark is embedded redundantly over the whole image, so that every part of the image is protected.

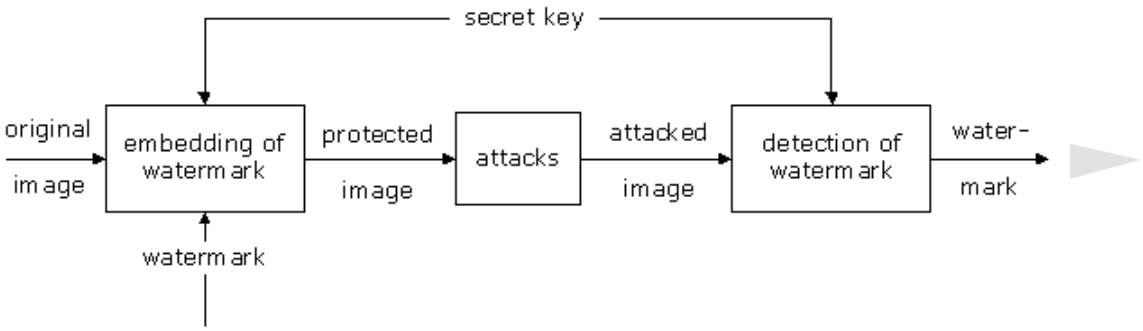

One way of doing this is by “Patchwork”. This technique uses a random number generator to select n pairs of pixels and slightly increases or decrease their luminosity (brightness level). Thus the contrast of this set is increased without any change in the average luminosity of the image. With suitable parameters, Patchwork even survives compression using JPEG.

Although the amount of secret information has no direct impact on the visual fidelity of the image or the robustness of the watermark, it plays an important role in the security of the system. The key space, that is the range of all possible values of the secret information, *x* must be large enough to make exhaustive search attacks impossible.

In the process of extracting the watermark, the secret key is used to identify the manipulated pixels and finally to decode the watermark.

As an example of poor engineering, an early version of Digimarc’s watermarking software gave each licensed user an ID and a two-digit

numeric password, which were issued when she registers with Digimarc and pays for a subscription.

The password checking mechanism could easily be removed by flipping a particular “flag” bit and the passwords had only 99 possibilities so it was short enough to be found by trial and error.

A deeper examination of the image also allowed a villain to change the ID and thus the copyright of an already marked image as well as the type of use (such as adult -> general public content).

Before embedding a mark, watermarking software usually checks whether there is already a mark in the picture, but this check can be bypassed fairly easily with the result that it is possible to overwrite any existing mark and replace it with another one.

The quality of digital watermarks can be judged in two ways; firstly it must be able to resist intentional and unintentional attacks and secondly the embedded watermark must not detract from the quality of the image.

The higher the resistance of a watermark against attacks, the higher the risk of the quality of the image being reduced, and the greater the chance of obvious visual artefacts being created.

[4]Methods used to test Watermark Robustness

Images

These are some of the methods that can be used to test whether a watermark can survive different changes to the image it is embedded in.

Compare this Original Image with the attacked images below, and see if you can spot any changes in quality.

Horizontal Flipping

Many images can be flipped horizontally without losing quality. Few watermarks survive flipping, although resilience to flipping is easy to implement.

[4] Attacks on Copyright Marking Systems, University of Cambridge http://www.cl.cam.ac.uk/~fapp2/publications/ih98-attacks.pdf

Rotation & Cropping

A small rotation with cropping doesn't reduce image quality, but can make watermarks undetectable as rotation realigns horizontal features of an image used to check for the presence of a watermark. The example at left has been rotated 3 degrees to the right, and then had its edges cropped to make the sides straight again.

JPEG Compression/Re-compression

JPEG is a widely used compression algorithms for images and any watermarking system should be resilient to some degree to compression or change of compression level e.g. from 71% to 70% in quality like the example at left.

Scaling

Uniform scaling increases/decreases an image by the same % rate in the horizontal and vertical directions. Non-uniform scaling like the example at left increases/decreases the image horizontally and vertically at different % rates. Digital watermarking methods are often resilient only to uniform scaling.

Dithering

Dithering approximates colors not in the current palette by alternating two available similar colors from pixel to pixel. If done correctly this method can completely obliterate a watermark, however it can make an image appear to be "patchy" when the image is over-dithered (as in the elbow area of the image at left).

Mosaic

[5]A mosaic attack doesn't damage the watermarked image or make it lose quality in any way, but still enables the image to be viewed in eg: a web browser by chopping the image into subsections of equal size and putting it back together again.

To the viewer a "mosaic" image appears to look the same as the original but a web crawler like DigiMarc's MarcSpider sees many separate images and doesn't detect that these separate images are parts of a watermarked image.

This means that the watermark cannot be detected, as a problem common to all image watermarking schemes is that they have trouble embedding watermarks into small images, (less than 256 pixels in height or width).

Stirmark

StirMark is the industry standard software used by researchers to automatically attempt to remove watermarks created by Digimarc, SysCoP, JK_PGS (TALISMAN project – É.P.F.L. algorithm), Signum Technologies and EIKONAmark.

Stirmark attacks a given watermarked image using all the techniques mentioned in this report as well as more esoteric techniques such as low pass filtering, gamma correction, sharpening/unsharpening etc.

All vendors of digital watermarks have their products benchmarked by Stirmark and as of August 2001, no watermark from any vendor survives the test, ie: the watermarks are all removed without degradation to image quality occurring.

Audio

The most common method of watermarking audio is to mark every x^{th} bit in an audio file depending on the random generator seed calculated from the watermarking key applied to the audio.

These are some of the ways watermarks can be removed from audio files.

---

[5] http://www.cl.cam.ac.uk/~fapp2/watermarking/2mosaic/

MPEG1 Layer III (MP3) audio compression

[6]A digital audio compression algorithm that achieves a compression factor of about twelve while preserving sound quality. What this lossy compression does is remove the frequencies not heard by the human ear from the audio. If a raw audio file is converted to MP3 at a bit-rate of 128kbps than roughly 90% of the frequencies are removed. This means that a search for the watermark needs to find an unaltered length of samples that contains at least 2 watermarked bits to prove the watermarks existence.

Audio Restoration programs

Audio restoration programs are designed to remove hisses, crackles and pops from audio recordings. They do this by searching through the wavelength, removing samples that don't "fit in" amongst neighbouring samples, and replacing them with an average of the two neighbour samples. Although the removal of digital watermarks is obviously not a purpose of these programs, they work remarkably well at doing so as the sample bits inserted to watermark the audio don't fit in with their surrounding pixels, and are therefore removed.

Echo Hiding Removal

Echo hiding relies on the fact that we cannot perceive short echoes, eg: 1 millisecond(ms) and embeds data into a cover audio signal by introducing an echo characterised by its delay and its relative amplitude compared to surrounding samples. The echo delays are chosen between 0.5 ms and 2 ms and the best relative amplitude of the echo is around 0.8 ms.

However specialised software which looks for echoes with a length between 0.5 ms and 2 ms [7](as seen below), can be used to detect and remove these echoes without effecting sound quality.

---

[6] http://www.dictionary.com/cgi-bin/dict.pl?term=MPEG-1%20audio%20layer%203

[7] Attacks on Copyright Marking Systems, University of Cambridge http://www.cl.cam.ac.uk/~fapp2/publications/ih98-attacks.pdf

Digital Watermarking

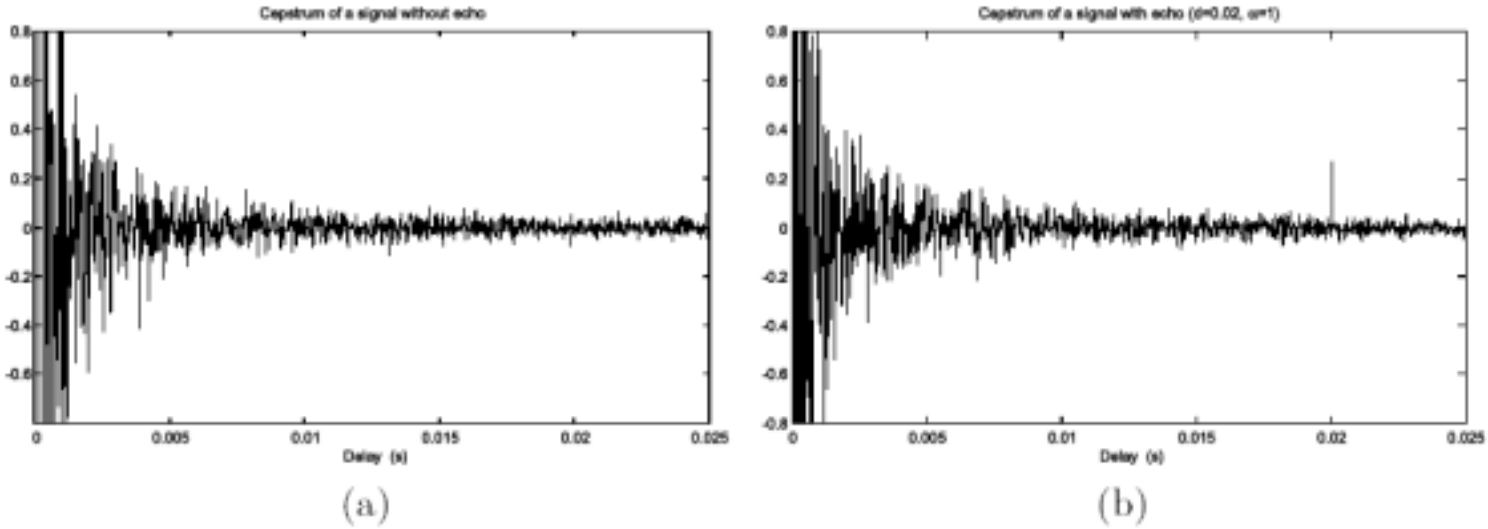

(a) (b)

Jitter

The simplest and most effective attack on any audio watermarking scheme is to add jitter to the signal. [8]"In our first implementation, we split the signal into chunks of 500 samples, either duplicated or deleted a sample at random in each chunk (resulting in chunks of 499 or 501 samples long) and stuck the chunks back together. This turned out to be almost imperceptible after altering, even in classical music; but the jitter prevents the marked bits from being located", and therefore the watermark is obliterated.

[9]In his paper titled "Audio watermarking: Features, Applications And Algorithms", Michael Arnold agrees with the Cambridge team stating that
"one of the greatest challenges [of watermarking] is the robustness against the so-called jitter attack".

Video

At present there is no known method to remove a digital watermark from a stream of video. This is probably because those who trade in pirated video, (especially in DivX format), store their pirated movies locally on their hard disk drives or on CD-R disks where they cannot be checked for watermarks by anyone.

---

[8] Attacks on Copyright Marking Systems, University of Cambridge
http://www.cl.cam.ac.uk/~fapp2/publications/ih98-attacks.pdf

[9]Audio Watermarking: Features, Applications and Algorithms, SysCoP
http://syscop.igd.fhg.de/Publications/Arnold00a.pdf

Digital Watermarking

Australian Digital Watermarking Scene

Across all the digital watermarking literature researched no mention was made of any Australian company implementing or selling a digital watermarking technique.

The only known group researching digital watermarks in Australia is the Secure Multimedia Information Communication Research Labs at the University Of Wollongong
http://www.itacs.uow.edu.au/research/smicl/

This research team is supported by the Motorola Australian Research Centre, Visual Information Processing Lab.

Analysis & Recommendations

The results of attacks on watermarks indicate that there are several problems which need to be solved by vendors before watermarks can become a viable option for those people/organizations who want to permanently embed proof of ownership or any other data into their audio or image creations.

The problems that must be resolved by vendors are:

- How effectively their image watermarking techniques can survive attacks by Stirmark, and other manipulation/transformation methods when applied with intent to remove the watermark or simply to edit it,

- [10]The basic problem that many schemes provide no intrinsic way of detecting which of two watermarks was added first: the process of marking is often additive, or at least commutative. So if the owner of the document *d* encodes a watermark *w* and publishes the marked version *d* + *w* and has no other proof of ownership, a pirate who has registered his watermark as *w*' can claim that the document is his and that the original unmarked version of it was *d* + *w* – *w*',

- How effectively their audio watermarking techniques can survive the jitter, MP3, echo removal and other methods.

Until these problems can be overcome, I don't believe digital watermarking techniques as used at present are robust enough for organisations and the general public to rely on.

[10] Weaknesses of copyright marking systems
http://www.cl.cam.ac.uk/~fapp2/papers/acm98-weaknesses.doc

Resources for Further Information

Internet resources used additional to those quoted (and many, many sub-links within followed!) were:

Watermarking & Stenography
http://www.cl.cam.ac.uk/~fapp2/index.html

Watermarking Stuff
http://www-sigproc.eng.cam.ac.uk/~pl201/watermarking/index.html

Watermarking of Video and Multimedia Data
http://www-nt.e-technik.uni-erlangen.de/~hartung/watermarking.html

Digital Watermarking
http://www.cosy.sbg.ac.at/~pmeerw/Watermarking/

Watermark links
http://www.elec.uq.edu.au/~ajantha/WmkLinks.html

Dipartimento di Elettronica e Telecomunicazioni- Università di Firenze
http://cosimo.die.unifi.it/~piva/Watermarking/watermark.html

Watermark Webring
http://nav.webring.yahoo.com/hub?ring=watermarking&list

Signum Technologies
http://www.signumtech.com/new/template7.asp?pageID=20&table=partners#%2012

DCT
http://www.dct-group.com/

Electronic Watermarks, Copy Control, fingerprints
http://buffy.eecs.berkeley.edu/~linnartz/water.html

Patrick's Watermarking Page
http://www-sigproc.eng.cam.ac.uk/~pl201/watermarking/

Watermarking and Data Hiding
http://www-nt.e-technik.uni-erlangen.de/~hartung/watermarkinglinks.html

The DICE Company
http://www.digital-watermark.com/

Peter Wong's Watermark Information Page
http://www.ee.ust.hk/~eepeter/watermark/