DIGITAL WATERMARKING:

Mark This Technology!

Have you pondered over how to protect ownership rights over digital data that's exclusively yours? You could look towards digital watermarking for a solution. The techniques currently available, though, are effective only in a limited way and susceptible to attacks. But a beginning has been made, and digital watermarking could be the best way to protect your digital data from illegal replications in the near future.

By: NEHA SINGH AND ARNAB NANDI

With the recent spate of cases involving fake currency, no one needs to be reminded of the importance of watermarking. A watermark is a form, image or text that is impressed onto paper, which provides evidence of its authenticity. Digital watermarking is an extension of this concept in the digital world. In recent years, the phenomenal growth of the Internet has highlighted the need for mechanisms to protect ownership of digital media. Identical copies of digital information, be it images, text or audio, can be produced and distributed easily. In such a scenario, who is the artist and who the plagiarist? It's impossible to tell — or was, until now. Digital watermarking is a technique that provides a solution to the longstanding problems faced with copyrighting digital data.

What are digital watermarks?

Digital watermarks are pieces of information added to digital data (audio, video, or still images) that can be detected or extracted later to make an assertion about the data. This information can be textual data about the author or its copyright or it can be an image itself. The information to be hidden is embedded by manipulating the contents of the digital data, allowing someone to identify the original owner, or in the case of illicit duplication of purchased material, the buyer involved. These digital watermarks remain intact under transmission/transformation, allowing you to protect your ownership rights in digital form.

Watermarks may be visible, in which case their use is two-fold—to discourage unauthorised usage, and also act as an advertisement. However, the focus is on invisible watermarks, as they do not cause any degradation in the aesthetic quality or in the usefulness of the data. They can be detected and extracted later to facilitate a claim of ownership, yielding relevant information as well. Watermarks may also be classified as robust or fragile.

Robust watermarks are those which are difficult to remove from the object in which they are embedded, despite various attacks they might be subjected to, discussed later. Fragile watermarks are those that are easily destroyed by any attempt to tamper with them. Absence of a watermark in a previously watermarked document would lead to the conclusion that the data has been tampered with.

For a digital watermark to be effective for ownership assertion, it must be robust, recoverable from a document, provide the original information embedded reliably, be non-intrusive, and also be removable by authorised users.

There are three main processes involved in watermarking—insertion of a watermark, detection of a watermark, and removal of a watermark.

Inserting a watermark

A general block diagram for the insertion of a watermark is shown, which provides a generic approach to watermarking any digital data. It consists of a watermark insertion unit that uses the original image, the watermark, and a user key to obtain the watermarked image.

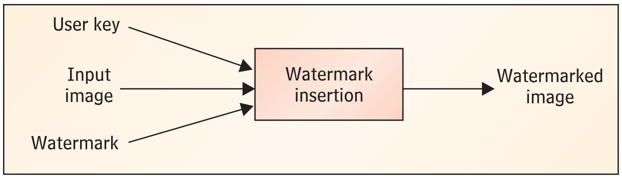

Figure 1: Watermark Insertion Unit

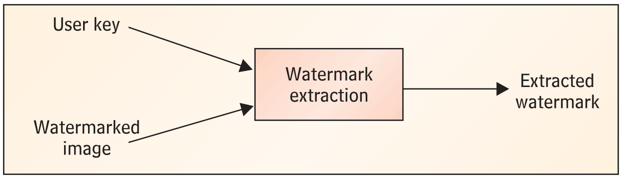

Figure 2: Watermark Extraction Unit

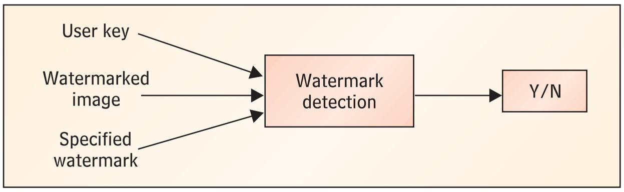

Figure 3: Watermark Detection

Similarly, watermark extraction and detection can also be performed using the units as well as the user key, as shown in Figure 2.

Extracting the watermark can be divided into two phases—locating the watermark and recovering the watermark information. Two kinds of extraction are available—using the original document and in the absence of the original document.

A watermarked detection unit consists of an extraction unit to first extract the watermark, and later compare it with the original watermark inserted. The output is 'Yes' or 'No' depending on whether the watermark is present.

Image watermarking depends on the domain in which the watermarking is done—the spatial and frequency domains. Watermarking in the spatial domain involves selecting the pixels to be modified based on their location within the image and is very susceptible to cropping and the mosaic attack, discussed later.

A simple spatial watermarking algorithm—the LSB technique

The LSB technique is the simplest technique of watermark insertion. If you specifically consider still images, each pixel of the colour image has three components—red, green and blue. Let us assume you allocate 3 bytes for each pixel. Then, each colour has 1 byte, or 8 bits, in which the intensity of that colour can be specified on a scale of 0 to 255.

So a pixel that is bright purple in colour would have full intensities of red and blue, but no green. Thus that pixel can be shown as

$$X_0 = \{R=255, G=0, B=255\}$$

Now let's have a look at another pixel:

$$X_1 = \{R=255, G=0, B=254\}$$

We have changed the value of B here. But how much of a difference does it make to the human eye? For the eye, detecting a difference of 1 on a colour scale of 255 is almost impossible.

Now since each colour is stored in a separate byte, the last bit in each byte stores this difference of one. That is, the difference between values 255 and 254, or 127 and 126 is stored in the last bit, called the Least Significant Bit (LSB).

Since this difference does not matter much, when we replace the colour intensity information in the LSB with watermarking information, the image will still look the same to the naked eye.

Thus, for every pixel of 3 bytes (24 bits), we can hide 3 bits of watermarking information, in the LSBs.

Thus a simple algorithm for this technique would be:

```
Let W be watermarking information
For every pixel in the image, $X_i$
Do Loop:
    Store the next bit from W in the LSB position of $X_i$[red] byte
    Store the next bit from W in the LSB position of $X_i$[green] byte
    Store the next bit from W in the LSB position of $X_i$[blue] byte
End Loop
```

To extract watermark information, we would simply need to take all the data in the LSBs of the colour bytes and combine them.

A modification of this method would be to use a secret key to choose a random set of bits, and replace them with the watermark. This technique of watermarking is invisible, as changes are made to the LSB only, but is not robust. Image manipulations, such as resampling, rotation, format conversions and cropping, will in most cases result in the watermark information being lost.

Other algorithms

In another technique, the pixels are divided into 2 equal sets A & B randomly by a secret key. A small integer k is added to the intensity of each pixel in set A, and subtracted from each pixel in set B. Since the integer k is small, changes are imperceptible.

To detect whether the image is watermarked, we simply calculate the average intensities of the two areas. If the two values differ by 2k, the image is watermarked. If they differ by 0, the image is not watermarked.

In the superimposition technique, a watermark symbol is selected and either scaled, or the canvas is enlarged, so that the two images, the watermark and the original image, have the same dimensions. The two images are then added together as follows. For each pixel making up the watermark symbol, a fixed intensity is added to the corresponding pixel in the original image. The resulting watermark may be visible or invisible, depending on the intensity chosen. Robust in the face of most common geometric transformations, simple implementations may be defeated by rotations.

Frequency-based watermarking

Watermarking in the frequency domain involves selecting the pixels to be modified based on the frequency of occurrence of that particular pixel. This is to overcome the greatest disadvantage of techniques operating in the spatial domain, i.e. susceptibility to cropping. The mosaic attack (discussed later) defeats most implementations of digital watermarking operating in the spatial domain but the frequency domain watermarking is less susceptible. The LSB technique can also be applied in the frequency domain, by selecting the pixels according to frequency, though this is not a robust way. Common transforms, such as Fast Fourier Transforms, alter the value of pixels within the original image, based on their frequencies.

The watermark is more commonly applied to the lower frequencies within an image, as higher frequencies are usually lost when an image is compressed, or to frequencies considered to contain perceptually significant information. Frequency-based techniques result in a watermark that is dispersed throughout the image and are, therefore, less susceptible to attack by cropping. However, these techniques are susceptible to standard frequency filters and lossy compression algorithms, which tend to filter out less significant frequencies.

Often, methods or a combination of methods, considered unintentional are used intentionally as an attack on a watermarked document in order to render the watermark undetectable. Compression is a common attack, as data transferred via a network is often compressed using JPEG (most commonly). High quality images are often converted to JPEG to reduce their size. Another method of attack is deletion or shuffling of blocks. In audio data, small blocks may be deleted or shuffled with no noticeable decrease in quality. In images, rows or columns of pixels may be deleted or shuffled without a noticeable degradation in image quality. Other common attacks include horizontal or vertical flipping, small angle rotation and cropping. These may render an existing watermark undetectable.

Often attackers are only interested in a small subsection of the image. A watermark at the edge of an image can often easily be cropped out of the picture without any significant loss. The mosaic attack is an extreme form of this method. In a mosaic attack, the attacker breaks up the entire watermarked image into many small parts. For example, a watermarked image on a Web page can be cut up and reassembled as a whole using tables in HTML. The only defence against this attack is to tile a very small watermark all over the image, and allow retrieval of the watermark from any of the small sub-sections of the fragmented image. However, the attacker can always create smaller blocks, and the watermarked image also has to be large enough to be distinguishable.

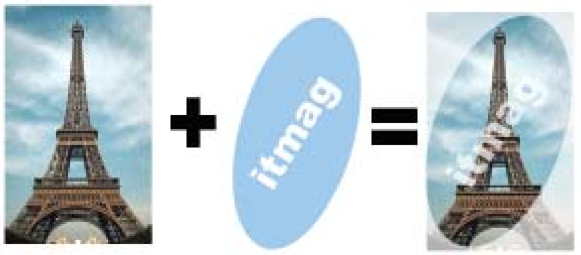

The watermarking of a digital image.

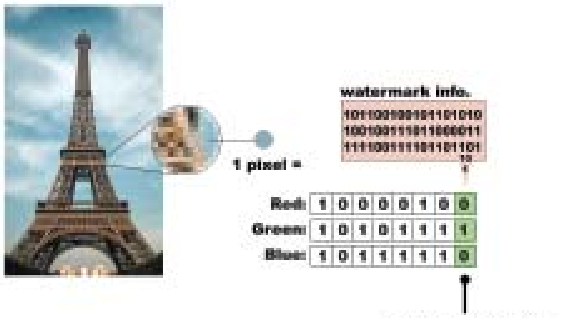

These LSB values are discarded and the watermark values are put in.

The LSB Technique

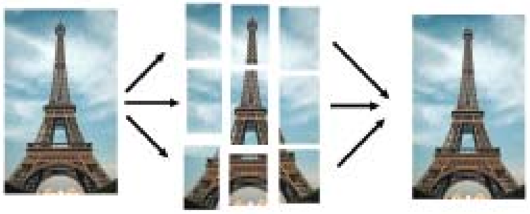

Mosaic Attack : Images are split and recombined using HTML, etc.

Audio watermarking

Watermarking is not restricted to just images. Audio watermarking uses the time and frequency masking properties of the human ear to conceal the watermark, and make it inaudible. One technique is echo-hiding, which involves hiding information within recorded sound by introducing very short echoes, relying on the fact that the human auditory system cannot perceive echoes shorter than a few milliseconds. Information is embedded into audio data by introducing two types of echoes, characterised by their duration and relative amplitude. This allows us to encode ones and zeroes within the audio data.

While digital watermarking holds much promise, current techniques are inadequate for general use. Watermarks are destroyed too easily to be used as evidence of copyright infringement in a court of law.

Despite the vulnerability of current techniques, watermarking remains important as long as it hinders the task of copyright infringement, and current tools offer this to a limited degree. Digital watermarking is an unproven technology today. Though in its infancy, it is all set to grow, and to make a large impact on the way in which digital media is distributed.

Neha Singh is a student of computer engineering at NSIT, Delhi. Arnab Nandi is a student of computer science at Delhi University.