Global Information Assurance Certification Paper

Steganography and Steganalysis: An Overview

by Joshua Silman

gsec 1.2f (august 2001)

1 Introduction

What is steganalysis? Although my 11 year old daughter insists that steganalysis is the study of a two ton, thirty foot long dinosaur with a brain the size of a walnut, a more modern definition would include the discovery and destruction of hidden information. In order to understand steganalysis it is essential to have an understanding of steganography (stego), the techniques used to hide information.

2 Steganography

Steganography is a Greek word which means "covered writing" and can trace its origins as far back as 440 B.C.. In *Histories* written by Herodotus[1], he gives two examples of steganography. The first is of Demeratus, a Greek in the Persian court who sent warning of a forthcoming invasion by Xerxes by writing a message on a wooden pallet and then covering it in wax. The messenger was able to successfully smuggle the "blank" tablet to Sparta. A second example was that of Histiaeus who shaved the head of his most trusted slave and tattooed a message on his head. After the slave's hair grew in he was dispatched with the "hidden message".

As technology has evolved so has steganographic technique. Along with the printing press came the use of "invisible inks", usually crafted from organic materials such as milk, juices or urine. When heat is applied to the document the hidden writing becomes visible. Photography provided the opportunity to create microfilm(s) which could be smuggled in secret compartments in clothing and luggage. Microfilm was a popular medium during the Franco-Prussian War (1870 - 1871). By the turn of the century, photographic reductions made it possible to produce microdots, a picture that could be reduced to the size of a period. In the 21st century the governments began to use steganography to protect their currency from being counterfeited. They have employed special inks, dyes, Embedded threads and microstrips which denote the face value of the bill. Steganography has seen its greatest growth and use with the growth of the Internet. The power of the Internet lies in its ability to transmit large quantities of data, very quickly, to a large audience.

Why do we need to hide information? There are two major issues that drive the technology to hide information. In the first group are those who are trying to protect their intellectual property rights. With the high availability of information via the Internet it is becoming more difficult to protect intellectual property and enforce copyright laws. The use of digital watermarks provides a way to insert a copyright notice into a document or image. The watermark is often a small image or text that is repeated frequently through out the document or image. A similar technique is to Embed a digital fingerprint or serial number. The advantage of a fingerprint is that it can be used to trace the copy back to the original and is a powerful tool for prosecuting copyright violators.

The second group of people who are interested in hiding information are those who wish to convey information in a covert manner and avoid observation by unintended recipients. In this case the hidden message is more significant than the "carrier" object that is used to transport it. Steganography is often compared to cryptography in its ability to restrict unauthorized access

to information. Cryptography is used to encrypt or scramble the data in such a fashion that only the intended recipient can decrypt it. When transmitting an encrypted message it is obvious that some form of communication has occurred, even if the message cannot be read. Steganography is used to hide the very existence of the message.

How do we hide information in the electronic age? At the most fundamental level computers use binary, a combination of zeros and ones to represent text and graphics. The American National Standard Code for Information Interchange (ASCII) is the de facto standard for representing text and certain control characters. ASCII uses one parity bit and seven data bits to represent each character in the English language. For example an uppercase "A" is represented by 1000001. A digital image is composed of picture elements or "pixels." Each pixel contains information as to the intensity of the three primary colors, red, green and blue. This information can be stored in a single byte (8 bits) or in three bytes (24 bits). For example, in an 8 bit image white is represented by the binary value of 11111111 and black is 00000000. Current information hiding techniques rely on the use of a cover object (image, document, sound file, etc.) sometimes known as a carrier. The secret message is then broken down to its individual bits by a steganographic tool (stego-tool) and Embedded in the cover object. Many tools will utilize a password or passphrase which is necessary to extract the hidden message and is referred to as a stego-key. The result of this process is known as the stego-object.

Where can information be hidden? Almost anywhere on the Internet! The standard protocol suite used on the Internet is the Transmission Control Protocol / Internet Protocol (TCP/IP). The headers used to transfer data between computers allow the use of flags and certain reserved fields. With the appropriate tool, information can be inserted into these fields. The advantage of this technique is that headers are rarely read by humans and thus makes an ideal place to hide data. The disadvantage of this method is that firewalls can be configured to filter out packets that contain inappropriate data in the reserved fields, thus defeating the steganographic transmission. Another popular technique for hiding information is to include extra spaces in documents. These spaces may contain hidden characters. Again this is a simple technique for hiding information and consequently is easy to detect and defeat. By opening such a document in a word processor the unusual spacing becomes readily apparent. Reformatting the document can remove the hidden message. The use of audio files can provide a good carrier for hidden messages. By their very nature sound files tend to be large in size and thus do not attract attention. In particular MP3Stego, a tool available from [2], can be used to hide information and maintain nearly CD quality sound.

The most prevalent cover objects in use today are digital images because of their potential payload (hidden information). A typical image with 640 x 480 pixels and 256 colors (8 bit) can hide approximately 300 Kilobytes of information. A high resolution image, 1024 x 768 pixels and 24 bit color could hide approximately 2.3 Megabytes worth of data. Due to the potential large size of such files compression algorithms are used to reduce the image to a suitable size for sending across the Internet. There is a wide variety of compression algorithms available, but the three most common are Windows Bitmap (BMP), Graphic Interchange Format (GIF) and Joint Photographic Experts Group (JPEG). When choosing a cover image for use in steganography the first two compression algorithms, BMP and GIF are preferred because they offer "lossless" compression. The compressed image is an exact representation of the original. The JPEG compression algorithm uses floating point calculations to translate the picture into an array of integers. This conversion process can result in rounding errors which may eliminate portions of

the image which are not visible to the naked eye. Although this rarely causes a noticeable change to the image it can significantly alter or destroy any information that was hidden in the image.

Embedding data into an image can be accomplished by either of two techniques, Image Domain tools or Transform Domain tools. Image Domain tools, also known as Bit Wise Methods, manipulate the Least Significant Bit (LSB) of the cover image. In this method the leftmost bit of each pixel in the cover image is replaced with one bit from the secret message. Because the LSB can only contain zeros and ones, approximately half the time the bit does not need to be altered in order to Embed the data from the secret message. In a low resolution (small number of pixels) image with 8 bit color the effects of manipulating the LSB can cause noticeable shifts in colors. As the resolution and depth of color increase in an image the impact of manipulating the LSB becomes less noticeable. Thus high resolution images are preferred for use as cover images. One exception to this rule are gray scale images. A gray scale image uses 8 bits to define 256 shades of gray between white and black. In a gray scale image pallet each shade represents an increment (or decrement) of 1 bit from the previous shade. Thus when the LSB is manipulated it is less likely to create a "new" or previously unused shade within the pallet. Most of the stego-tools available today utilize bit wise methods for hiding information. Some of the more popular Image Domain tools include; Hide and Seek, Mandelsteg, Steganos, StegoDos, S-TOOLS, and White Noise Storm.

Transform Domain tools utilize an algorithm such as the Discrete Cosine Transformation (DCT)* or wavelet transformation to hide information in significant areas of the image. Stego-tools which utilize one of the many transform domain techniques are more robust, have a higher resilience to attacks against the stego-image such as compression, cropping and image processing[3]. As of this writing all of the stego-tools which can manipulate JPEG images are transform domain tools such as; Jpeg-Jsteg, JPHide, Outguess, PictureMarc and SysCop.

*Hiding Information in JPEG Images

"The JPEG image format uses a discrete cosine transformation (DCT) to transform successive 8x8 pixel blocks of the image into 64 DCT coefficients each. The least-significant bits of the quantized DCT coefficients are used as redundant bits into which the hidden message is embedded.

In some image formats, *e.g.* GIF, the visual structure of an image exists to some degree in all bit-layers of the image. Steganographic systems that modify the least-significant bits of these formats are often susceptible to visual attacks.

This is not true for the JPEG format. The modification of a single DCT coefficient affects all 64 image pixels. For that reason, there are no known visual attacks against the JPEG image format."
(Provos, CITI Technical Report 01-11) [4]

3 Steganalysis

With careful selection of an appropriate cover image and a stego-tool it is possible to create a stego-image that does not appear to be different within the limits of human perception. However, electronically each of these tools leaves a fingerprint or signature in the image that can be used to

alert an observer to the presence of a hidden message. Discovering a hidden message is the first step in steganalysis and is considered an "attack" on the hidden information. Attacks may come in several different forms depending on what information is available to the steganalyst (see table 1).

There are two other types of attacks against steganography. The first is the known message attack. In this case the steganalyst (one who does steganalysis) has a known hidden message and the corresponding stego-image. In this case the objective is to determine patterns that result from hiding the message. These patterns can then be used to analyze other stego-objects in the future. The second attack is the chosen-message attack. In this case the steganalyst will create a message and use a known stego-tool to create a stego-image. This known stego image is then analyzed to determine patterns for later use against other stego-images.

Table 1.	
stego-only attack	Only the stego-object is available
chosen stego attack	The stego-tool (algorithm) is known and the stego-object is available
known cover attack	The stego-object and a known original copy of the cover object are available.
known stego attack	The stego-tool (algorithm) is known and both the stego-object and the original cover are available.

In order to be effective at steganalysis one must have good pattern recognition skills. In some instances comparing stego-images prepared with Image Domain tools and their original cover images will result in detectable visual noise. Noise is defined as a pixel that stands out from the other pixels in its area or "neighborhood." For example a lone red pixel on a white field. Another visual clue to the presence of hidden information is padding or cropping of an image. The Hide and Seek tool can only produce images of a fixed size [5]; 320x200, 320x400, 320x480, 640x400 and 1024x768. If an image does not fit into one of these sizes it is cropped or padded with black spaces. StegoDos has a similar problem.

The majority of stego-images do not reveal visual clues when compared with their cover image and thus require a more detailed analysis in order to determine that information has been concealed. In their work with current steganographic tools [5], Johnson and Jajodia discovered several possible electronic signatures. The simplest signature is an increase in the file size between the stego-image and the cover image. Most of the other signatures manifest themselves in some form of manipulating the color palette of the cover image. These fingerprints can include a large increase or decrease in the number of unique colors. Another fingerprint is colors in a palette which increase incrementally rather than randomly. The exception, of course, is gray scale images, which do increase incrementally. The presence of a disproportionate number of shades of black in a gray scale image is another strong indicator.

Once a stego-image has been discovered there are several steps that can be taken to disable or destroy the hidden message. Stego-images created with an Image Domain tool can be rendered useless (the hidden message can not be recovered) by simply converting the image to a JPEG format [3]. Images created with Transform Domain tools require a more aggressive

approach in order to disable the hidden information. Although they can survive any single image manipulation, multiple manipulations on the same image have defeated all of the known tools [3]. Image manipulation includes techniques such as: cropping, removing portions of the image; rotating the image; blurring, decreasing the contrast between pixels; sharpening, increasing the contrast between pixels (opposite of blurring); adding or removing noise; resampling; converting between bit densities (gray scale, 8 bit, 24 bit); converting from digital to analog to digital (print the image then rescan it); adding bit wise messages; adding transform message.

4 Current Events

If steganography is so easily detected and defeated who would use it? According to Ross Anderson, of Cambridge University, "There are about three or four generations of stego software. The stuff you can download is first generation and easily defeated."[6] It is important to remember that steganography is only a single tool. Repeated use of the same tool will provide an unintended recipient with a large body of stego-objects which can facilitate the cracking of the stego-system, thus revealing all of your communications. There are several steps which can be taken to improve the security of your data, including encrypting messages before applying steganography and changing stego-tools periodically.

There is some concern that terrorists using steganography. In February 2001 Jack Kelly wrote two articles in USA TODAY which indicated the Osama Bin Laden and his organization, Al-Qaeda, as well as other known terrorist groups were using steganography to plan and implement terrorist acts. It was suggested that stego-images were being placed on auction sites such as e-Bay and Amazon as well as sports chat rooms and pornographic sites. Several other news agencies ran similar articles later in the month. None of the articles offered definitive proof, other than anonymous quotes from federal law enforcement agencies, indicating that stego-images had been found. There were several references to encrypted e-mail and files that had been recovered. Based on the allegations that terrorist organizations were using steganography, Niels Provos and Peter Honeyman, researchers at the University of Michigan, launched a project to determine the truth of the matter. In their technical report [4] published on August 31, 2001, Provos and Honeyman outline the tools they used (Stegdetect, Stegbreak, Crawl and Disconcert) to launch an automated, statistical analysis of over 2 million JPEG images found on the e-Bay web site. As of this writing only one stego-image has been discovered. The image, sovereigntime.jpg (below) contained a gray scale image (below) of the "B-52 graveyard" at Davis-Monthan Air Force Base. These images were part of an ABC interview with an Internet security consultant who was demonstrating steganography. The authors conclude that based on their statistical analysis of images there is a small chance that they have not yet detected the stego-images that the terrorists are using. They believe it is more likely that as August 2001, there are no stego-images on the Internet.

5 Conclusion

Steganography is a dynamic tool with a long history and the capability to adapt to new levels of technology. As the steganographic tools become more advanced, the steganalyst and the tools they use must also advance. Like any tool, steganography (and steganalysis) is neither inherently good nor evil, it is the manner in which it is used which will determine whether it is a

benefit or a detriment to our society.

Images Discovered by Provos and Honeyman

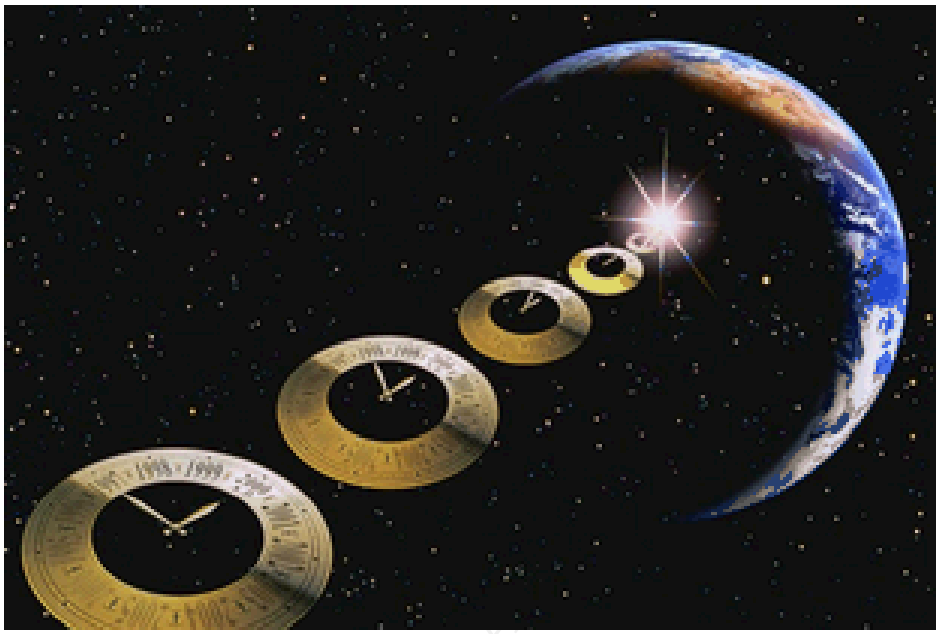

sovereigntime.jpg

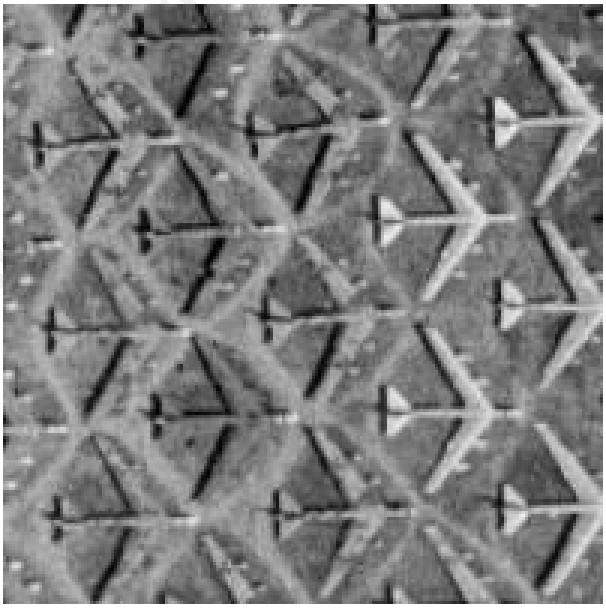

"B-52 graveyard" at Davis-Monthan Air Force Base

6 References

[1] Petitcolas, F.A.P., Anderson, R., Kuhn, M.G., "Information Hiding - A Survey", July 1999, URL: http://www.cl.cam.ac.uk/~fapp2/publications/ieee99-infohiding.pdf (11/26/01 17:00)

[2] An archive of steganography and steganalysis tools: URL: http://members.tripod.com/steganography/stego/software.html (11/26/01 17:00)

[3] Katzenbeisser, S., Petitcolas, F.A.P., *Information Hiding Techniques for Steganography and Digital Watermarking*, Norwood: Artech House, 2000, pg 56 - 92

[5] Johnson, N.F., Jajodia, S., "Steganalysis of images created using current steganographic tools", April 1998, URL: http://www.ise.gmu.edu/~njohnson/ihws98/jjgmu.html (11/26/01 17:00)

[4] Provos, N., Honeyman, P., "Detecting Steganographic Content on the Internet", August 2001, http://www.citi.umich.edu/techreports/reports/citi_tr_01-11.pdf (11/26/01 17:00)

[6] McCullagh, D., "Secret Messages Come in .Wavs", Feb 20, 2001, Wired News, URL: http://www.wired.com/news/politics/0,1283,41861,00.html (11/26/01 17:00)

[7] Artz, D., "Digital Steganography: Hiding Data within Data", *IEEE Internet Computing*, May-June 2001, pg 75-80

[8] Beyda, W.J., *Data Communications From Basics to Broadband 3rd edition*, Upper Saddle River: Prentice Hall, 2000, pg 38 - 40

[9] Kelley, J., "Terrorist instructions hidden online", *USA TODAY*, 06/19/2001, URL: http://www.usatoday.com/life/cyber/tech/2001-02-05-binladen-side.htm (11/26/01 17:00)

[10] Johnson, N.F., Jajodia, S., "Exploring Steganography: Seeing the Unseen", February 1998, URL: http://www.jjtc.com/pub/r2026.pdf (11/26/01 17:00)

[11] Johnson, N.F., Jajodia, S., "Steganalysis: The Investigation of Hidden Information", *IEEE Information Technology Conference*, September 1998, URL: http://www.jjtc.com/pub/it98a.htm (11/26/01 17:00)

[12] Kelley, J., "Terror groups hide behind Web encryption", *USA TODAY*, 06/19/2001, URL: http://www.usatoday.com/life/cyber/tech/2001-02-05-binladen.htm (11/26/01 17:00)

[13] McCullagh, D., "Bin Laden: Steganography Master?", *Wired News*, 07 Feb 2001, URL: http://www.wired.com/news/politics/0,1283,41658,00.html (11/26/01 17:00)

[14] Schneier, B., *Crypto-Gram Newsletter*, October 15 1998, URL: http://www.counterpane.com/crypto-gram-9810.html (11/26/01 17:00)

[15] Schneier, B., *Crypto-Gram Newsletter*, September 30 2001,
URL: http://www.counterpane.com/crypto-gram-0109a.html (11/26/01 17:00)

[16] Schneier, B., "War on Terrorism," *Crypto-Gram Newsletter*, October 15 2001,
URL: http://www.counterpane.com/crypto-gram-0110.html (11/26/01 17:00)