

Assuring Ownership Rights for Digital Images

Germano Caronni

Computer Engineering and Networks Laboratory

Swiss Federal Institute of Technology Zurich

E-Mail: caronni@tik.ethz.ch

Abstract

The use of digital data has become more and more commercialized. This is especially true for digital images, where proofs of origin and of content integrity are an important issue. This paper describes a problem related to 'proof of origin' and proposes a possible solution to it. After a discussion of the solution, possible extensions and related areas of work are addressed.

1 The Problem

Until now, digital data which was disseminated had no 'unique' features. Everybody received an identical copy of the data. Thus, if one of the copies was illegally distributed, it was impossible to determine the initiator of the unauthorized distribution. Typical effects are software piracy, the unauthorized distribution of vector fonts for printers and the distribution of certain digital images, such as art collections and satellite data. The same holds true for the distribution of confidential texts or images.

All possible kinds of digital data, such as computer software, fonts, texts, images and sound suffer from this problem. Only digital data in form of images¹ will be discussed here. Although related solutions for other types of digital data might be found, they have not yet been considered and would exceed the limits of this paper. A possible solution for formatted text may be found in [9] or [16].

A distributor of digital images of commercial or confidential nature usually is interested in detecting the source of illegal copies of his data. To do this, he has to provide each recipient with a different copy of his data. A process called **tagging** will be described, which includes hidden information in images, and thus makes distributed instances of an image different from each other. 'Hidden' here means that the inclusion of the data into the image causes quality degradation which is not perceivable by human eyes, and a receiver of the processed image is not able to detect or remove the included tags. As soon as the distributor of the original image

¹ Only digital (or digitized) images are considered, which contain a certain amount of noise, or variance in brightness. Thus images of 'Roger Rabbit' may not be acceptable, but a copy of Tizians 'Pietà' is.

somehow receives an illegal copy of it, he should be able to identify the original receiver of this particular image with high probability, even if the image suffered from some loss of quality.

Naturally, the distributor has to decide if the cost (time and effort) of tagging is adequate to achieve the intended results. If the distributed images have a short lifetime and are spread to a large audience, as with Reuters news images, tagging might be less adequate than in an art catalogue. At the same time, secure means for distribution and storage of tagged images have to be used, e.g. by applying commonly known cryptographic techniques, such as DES[11] or IDEA[12] for storage and additionally RSA[10] for transmission. Otherwise, a tagged image might be stolen from a legal customer, causing him to be accused for illegally spreading this image.

2 Requirements for successful tagging of images

The fundamental solution to the problem of detecting the distribution path of each image is to provide each recipient of an image with a different copy. The difference in the distributed images will allow the distributor to identify a certain recipient, by determining to whom he has given this instance of the original image.

As soon as a recipient, from now on dubbed **enemy**, wants to illegally spread his image, he will use countermeasures like the addition of noise, stretching of the image in one axis, or any other change which does not destroy the semantics of the image. This makes it more difficult for the distributor to identify him and has to be taken into account when looking for solutions to the following requirements:

- A **tag**² introduced into an image should have maximal information content to allow a good differentiation between different recipients.
- The tag should destroy as small as possible an amount of original information in the image. This guarantees high acceptance of the modified image by the recipient.
- The distributor should be able to easily separate the tags from the original image to allow detection of tags when an illegal copy of an image returns to him.
- There should be no possibility to separate the tags from an image without having access to the original untagged image.
- Removing or hiding the tags in the image should imply a maximum loss of quality in the image.

Some of these requirements work against each other, so a balance has to be found in order to get an optimal result. This balance depends on the actual needs of the distributor, and is influenced by e.g. the number of recipients or the fact if the distributor wants to recognize printed copies of the image.

3 Technical Approach

The issue of tagging images was partitioned into interdependent problems. Possible solutions to these problems are examined in the following sections. The approach presented here is partially based on heuristics, as formal models and methods have yet to be defined. To do this, information theoretical and statistical arguments have to be combined and discussed together. No tightly related work has been found. Although [18] pursues the same goals as this paper, the chosen approach is strongly related to DCT compression of an image, and has not been considered further. Loosely connected previous and related work is referenced.

² The sum of hidden information introduced into the image is named tag.

3.1 Information that Constitutes the Tags

To allow the distributor to differentiate between multiple instances of the same image, information has to be included into them. In its most abstract form, this information is a sequence of bits. Experiments have shown that, using the method presented in section 3.2, an image usually contains some hundred tag bits. Depending on the expected strategies of the enemies, different usage and interpretation of these bits should be chosen. Under the assumption that enemies do not cooperate (see section 3.3), the tag bits may provide maximum difference between different image instances. Principles applied to the construction of error correcting codes[1] (ECC) can be used to construct highly individual tag sequences. Under other circumstances, random bit sequences[13] may be used. They are easier to construct than ECCs, and give a better possibility to detect groups of cooperating enemies (see section 3.3).

3.2 Integrating the Tags into the Image

A mechanism has to be found to integrate the above defined tag bits into the image in a non-localizable manner. The distributor may not simply append the tags to the image, or place them in well-defined locations of the image, as an enemy might then just remove the tags, without suffering a loss of quality.

The idea of hiding information in an image to provide means of transferring the information without detection by an enemy is not new [2][3]. For example, a bitsequence could be directly integrated into the image by setting the least significant bit of the color values of a pixel to the value of one bit in the sequence. Nevertheless, currently known mechanisms are not fault tolerant, even slight distortion of the image makes the hidden information unrecoverable³, as no redundancy is provided.

If the tagging procedure were to be executed by a human he could modify some picture elements manually, thus minimally changing the semantics of the image. By introducing these modified elements (such as additional leaves of a depicted tree, a change in a shadow or a shift in the position of the sun) depending on the chosen bit sequence, a corresponding tag sequence would be produced. A similar but automated method for tagging purposes could shift borders detected in the image, replace homogenous areas by slightly different shades or change line widths of lines detected in the image. These two approaches (the manual and automatic change of image semantics) were not examined further, but still remain interesting, as they represent a near-optimal fulfilment of the requirements stated in section 2.

The approach taken in this work modulates the brightness of chosen rectangles in the image to hide its tagging information. Independent modulation of RGB color values is not suitable, as greylevel images are deemed to be of quite good quality, and the transformation from color to greylevel causes an extremely high information loss. Figure 1 illustrates the method.



Figure 1: Example on rectangular tags

To the left, an unmodified section of the image is displayed. The section in the middle is

³ The approach of image tagging might even be used to convey small amounts of information between communication partners in a unrecognizable and fault-tolerant way.

tagged with a modulation of 2% of the maximal brightness, allowing the recovery of most of the tags even after printing and rescanning the image. Finally, the section to the right is tagged with a modulation of 15%, giving the possibility to actually see the embedded rectangles.

Using rectangles introduces a high amount of redundancy for the tag information, allowing the detection of tags even after strong distortions of the image. Special considerations taken when placing the rectangles in the image cause them to disappear behind the ‘natural’ noise in the image. No rectangle is placed in a region which is too homogenous, or contains a sharp break, such as an edge. Homogenous regions have to be avoided to prevent enemies from extrapolating the state of the tag by analyzing the surroundings of the tag, edges have to be avoided to maintain image quality.

3.3 Recovering Tags from Distorted Images

To recover the tags from a distorted image, the possible actions of the enemies have to be considered: An enemy can try to work alone, having access to only one tagged image, or a group of enemies can work together, and devise strategies which use their differently tagged images to defeat the distributor.

An enemy who has access to only one tagged image is not able to detect the tags, as they are hidden behind the ‘natural’ noise in the image. He can distort the whole image or regions of it. This may be a change of contents, like adding noise, quantifying the colorspace of the image, applying dithering or a change in the form of the image such as stretching it, slightly rotating it, etc.

Unless this solitary enemy degrades the quality of the image by an amount which makes a future exploitation unlikely, the redundancy of the tags which were introduced by the distributor allows a good (> 90%) detection of the tag sequence. Methods to compensate for a change in form are known (e.g. [4],[5] and [6]), but have yet to be applied.

A group of enemies working together is able to initiate a much stronger attack by mixing or comparing their differently tagged images. This way, they can reduce the detectability of tags or even localize a certain amount of them. Estimates on the strength of such attacks may be found in section 5.2. To solve the problem of cooperating enemies in a better fashion, special tag sequences or even a different tagging method have to be developed. A possible approach to do this might be derived from [17].

After the tag sequence is retrieved by the distributor, it is compared with all generated tag sequences. The ones that are most similar represent the enemy or group of enemies who has distributed the image.

4 Realisation

In this section, the proposed simple tagging mechanism and the detection of tags shall be examined in greater detail, after discussing some preliminaries.

The tagging process introduces noise into an image, thus degrading its quality. This quality degradation (and the degradation that occurs when enemies apply countermeasures to a tagged image) has to be measured. This may be done by some humans, stating their subjective impression about the image. Preferring more objective data which may be collected in an automated way another approach has been taken. The correlation coefficient between original and modified image is measured. This coefficient is calculated on the brightness of each corresponding pixel in the two images ($b_o(x, y)$ for the original and $b_m(x, y)$ for the modified image respectively). It is defined as:

$$R = \frac{v_{om}}{v_o v_m}.$$

$$v_{om} = \frac{1}{(X \cdot Y) - 1} \sum_{x=1}^X \sum_{y=1}^Y (b_o(x, y) - m_o)(b_m(x, y) - m_m)$$

is the covariance between original and modified image, where m_o and m_m represent the mean brightness of either one. v_o and v_m are the variances of the two images, v_o is defined as

$$v_o^2 = \frac{1}{(X \cdot Y) - 1} \sum_{x=1}^X \sum_{y=1}^Y (b_o(x, y) - m_o)^2 .$$

When comparing two identical pictures, $|R|$ will have the value of 1, the more differences the pictures show, the more $|R|$ will decrease towards 0. This method for comparing images can only be applied to images having the same size, which sometimes might require the pre-processing of images

4.1 How to Integrate the Tags

In this tentative realisation of the tagging mechanism, the bitsequence which constitutes the tags is generated by a simple random number generator[14]. For more serious applications better generators have to be chosen to disallow attacks based on this information.

Tags are represented by rectangles which get modulated onto an image. The more geometrical deformation of the image is expected, the bigger a tag should be. They have a fixed size of $2 \cdot 2$ up to $2n \cdot 2n$, ($n < \min(X, Y)/2$) pixels, which is chosen at program start. Tags of 4×4 up to 16×16 pixels have been examined in [8] and in section 5 of this paper. In a first step, all locations in the image where a tag could possibly be placed are identified by calculating the variance of regions of size $n \cdot n$ in the image and comparing it against a upper and a lower limit. These limits were empirically defined. After having located all possible positions, some of these positions are randomly chosen, keyed by a so called **group identification** and a probability for each possible position to be actually used. Care is taken to provide each rectangle with a border of n unmodulated pixels. This is needed for a later detection of the tags. At the same time, the direction in which a future tag may get modulated (brighter/darker) is randomly chosen.

The location and possible modulation of tags in an image is the same for all customers who receive this image, as long as the group identification is the same for all customers. To differentiate between customers, a *serial number* is used, again keying a random generator. The thus generated bitsequence triggers the actual modulation of the tags, and is at the same time used to add some noise (currently 0.5% of the maximal brightness) to each pixel of the image. The activation of a tag alters the brightness of a corresponding rectangle in the image by e.g. 1%. Again these values are hardcoded. Figure 2 illustrates the different modulations which are superimposed on top of the original image.

Actual data on some examples (number of tags and correlation coefficient) may be found in section 5. Adapting the variance in brightness to the actual variance of the local region might lead to a noticeable increase in tag detection by the distributor, and will be subject to further study.

As tag rectangles are placed only in regions with a minimal variance, it is expected that the 'additional' information added by the tag disappears behind the image noise. Tags introduced in an image usually are not visible to a careful observer.

4.2 Recovering the Tags

The algorithm which recovers the tags is designed to exploit the fact that image distortion introduced by an enemy or e.g. lossy compression algorithm usually are not localized exactly on the effective tag rectangles. Distortion is expected to equally spread on the rectangles (or

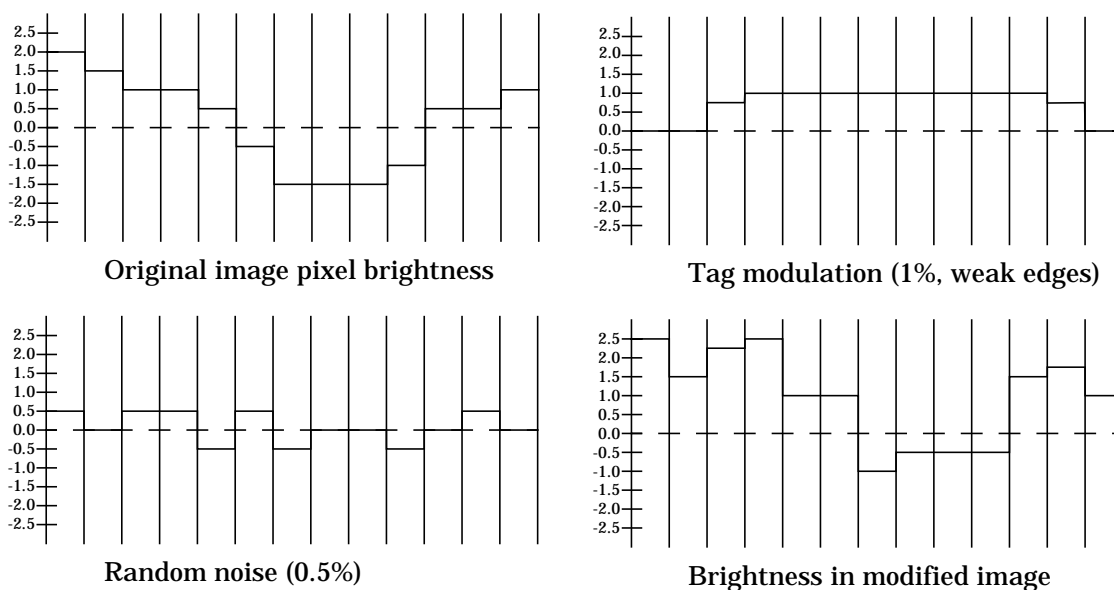


Figure 2: Modulation of an image by tagging information

part of them) and their unmodified surroundings. It is a precondition that the image to be processed has the same size as the original image, and that geometrical distortions (like rotation) have been eliminated from it.

In a first step, the brightness of each pixel in the received image is subtracted from the original one. Now, having knowledge of possible tag positions, the algorithm tries to recover the original modulation of the rectangle, thus identifying the state of the corresponding bit in the tag sequence. Around the original tag with size $2n \cdot 2n$ an unmodified region of size n should exist. After the subtraction, the mean brightness of the border region should be 0. The actual value is calculated, and the so won offset used to correct the mean value for the brightness in the tag rectangle. This is done separately for each quarter of the tag rectangle, allowing a future balancing of the four mean values extracted from the rectangle on a nonlinear base. Currently, just the arithmetic mean of the four values is taken and compared with a threshold. If the mean value is higher than $1/2$ of the modulation strength of the rectangle, the corresponding tag bit is taken as '1' in the other case as '0'.

After this has been done for each tag rectangle in the image, the distributor is now in possession of a recovered tag sequence. By comparing it with the stored tag sequences of all customers the enemy may be identified. If a group of enemies shall be detected, groups of different tag sequences have to be generated, and just the bits in each sequence which are equal to all customers in the assumed group have to be checked.

5 Evaluation

To substantiate some of the claims in this paper, data has been collected. The main purpose of this data is to show the detectability of tags in distorted images on the one hand, and on the other hand give some hints on how strong the quality degradation of the images in the course of tagging actually is.

5.1 Tagging and Quality Loss

Depending on the size and the 'noisiness' of the image, and on the tag size, a different number of tags can be placed in the image. Table 1 enumerates the number of tags which was measured on a variety of randomly collected pictures. At the same time values of $|R|$ are dis-

played, giving a hint on quality loss introduced by the tagging process.

Image:	#Tags 4x4	#Tags 8x8	#Tags 12x12	#Tags 16x16	R 4x4	R 8x8	R 12x12	R 16x16	R Ref. $\pm 1\%$ Noise
bud (640x480)	690	427	254	156	.9998552	.9998131	.9997896	.9997647	.9988916
zurlim (512x512)	1593	606	282	156	.9999024	.9998786	.9998695	.9998585	.9994244
pic3 (502x900)	614	445	293	204	.9998595	.9998270	.9997997	.9997749	.9988591
ystone (1152x779)	1208	1076	683	453	.9995562	.9994302	.9993338	.9992625	.9964358
lake (512x512)	1530	608	299	175	.9998826	.9998515	.9998394	.9998352	.9993038

Table 1: Number of tags and value of correlation coefficient (tagging with 1.2%)

5.2 Countermeasures

As stated in section 3.3 enemies might apply different kinds of modifications to a tagged image to make it harder for the distributor to recover the tag sequence. The list of possible modifications and attacks on tagged images in this paper represents in no way an exhaustive overview, nor does it prove anything. It just gives a hint on the possibilities of the enemy⁴.

A group of enemies working together is able to initiate a strong attack. They may simply mix their images, giving each pixel of their ‘output’ image the value of the mean of all the corresponding pixels in the different images. This way, they can reduce the detectability of some of the tag bits by flattening the profile of the corresponding tag rectangles. Additionally they may compare their images, thus detecting differently modulated tags (see figure 3).




Figure 3: The detection of differing tags by enemies (20 tags detected)

They are then capable of falsifying their tag sequence. Assuming a randomly constructed bit sequence as identifier for each customer, N enemies may detect a fraction of $1 - 2^{1-N}$ of all tags. As long as the number of enemies is small, the distributor may still identify them by checking the bits they were not able to detect; if the number of enemies gets larger ($2^N \geq \text{Number of Tags}$) it is impossible to detect them.

A solitary enemy is not able to gain any information on the tags in the image. Thus his possible attacks are of two distinct classes:

1. Modification of image geometry

The enemy may slightly rotate, shrink, stretch, shift, etc. the whole image, or parts of it. This causes the locations of tags to be shifted, making it difficult for the distributor to (automatically) check the tags.

Just to give an example, some images have been shrunk by 50%. About 2/3 of all tags were still detectable, while $|R|$ dropped to about 0.85 and the images were subjectively severely degraded. The main problem here is to undo the geometrical distortion introduced by an enemy to allow the subsequent detection of tags. The application of [6] will at least partially solve this problem.

⁴ Usually it is very difficult for the designer of a cryptography or protection related algorithm to prove the strength of his algorithm, or assess all possible methods to counter it.

2. Modification of image content

The goal of content modification is to ‘remove’ the tags from the image, or at least distort the brightness of tag rectangles as much as possible, thus disallowing the distributor to successfully recover the bit sequence hidden in them. Image content modification comprises many possibilities. The following mechanisms have been employed to gain some data:

- Noise has been randomly added to the tagged image. The noise has been added to the brightness of each pixel, changing it by $\pm 2\%$, respectively $\pm 4\%$ of its maximal value.
- The JPEG lossy image compression algorithm[15] has been employed on the tagged images. The quality of the image was reduced to 75% and 30% respectively, where a quality of 30% represents a rather degraded picture.
- The colorspace of the tagged image has been reduced to 32 colors. At the same time dithering with Floyd-Steinberg error diffusion has been employed. The output of this step is in the range of a very sophisticated color printer.

Table 2 depicts the quality loss experienced when employing above methods on the original images (col: number of colors in the original image):

	Noise 2%	Noise 4%	JPEG Q75	JPEG Q30	FSQUANT 32
bud 256 col	.9969303	.9879267	.9941969	.9749811	.9900836
zurlim >99999 col	.9983958	.9935527	.9971826	.9918425	.9949042
pic3 76540 col	.9968435	.9875711	.9984049	.9965283	.9725430
ystone >99999 col	.9901941	.9624366	.9959695	.9912676	.9583207
lake >99999 col	.9980696	.9923478	.9971620	.9942864	.9911683

Table 2: Quality degradation after distortion of original images

A very special kind of modification is the repeated tagging of an already tagged image. Some trials assuming the knowledge of the tagging algorithm and all its parameters except the group identification and the original picture have shown a quality degradation of about 0.0002 per tagging iteration, and a loss of 3-4% of the original tags per iteration. After about the fifth iteration the images subjectively become more and more distorted.

5.3 Success in Recovering the Tags

Having produced a variety of tagged images (tagged with different tag sizes and differing strength of tag rectangle modulation) the content distortions mentioned above have been applied. Afterwards the tag sequences were recovered and compared with the originally introduced tags. Table 3 enumerates the percentage of tags that were successfully detected in each case for different tag sizes and tag modulation strengths.

Using a modulation strength of 2% and a tag size of 16x16 pixels, it was possible to recover 75% of the tags from enlarged, (color-)printed and rescanned images.

6 Summary and Future Work

A new and interesting problem has been presented, and some basic approaches for a solution have been discussed. Although there is still a lot of work to do, the results are promising. Additional efforts on both the theoretical and the practical side need to be done on at least the following points:

- Explore other forms of tagging and modulation of tags, including ‘Adaptive Tagging’.
- Explore hierarchical distribution paths for the images (multiple tagging?).
- Apply ‘tagging’ to sound (Tagging text has in the meantime been done by [9])
- Prove the nondetectability of tags introduced into images.

		Noise 2%				Noise 4%				JPEG Q75				JPEG Q30				FSQUANT 32			
		4x4	8x8	12x12	16x16	4x4	8x8	12x12	16x16	4x4	8x8	12x12	16x16	4x4	8x8	12x12	16x16	4x4	8x8	12x12	16x16
bud	1,0%	81	98	99	100	68	83	90	99	82	99	100	100	63	83	93	100	76	91	94	98
	1,2%	84	98	100	100	70	85	93	99	85	100	100	100	65	86	96	100	72	91	94	99
	1,4%	88	99	100	100	73	90	97	100	89	100	100	100	68	92	99	100	82	96	95	98
zurlim	1,0%	81	98	100	100	68	87	93	97	82	99	100	100	65	83	96	98	75	90	92	94
	1,2%	85	99	100	100	70	89	96	99	86	100	100	100	66	87	98	99	78	93	94	94
	1,4%	88	100	100	100	73	92	99	100	89	100	100	100	69	90	99	100	81	95	95	96
pic3	1,0%	83	98	100	100	69	84	96	98	83	99	99	100	66	85	96	99	68	84	86	92
	1,2%	85	99	100	100	71	86	96	99	84	99	100	100	66	89	96	100	71	84	88	94
	1,4%	88	100	100	100	74	91	99	99	88	100	100	100	69	94	99	100	76	89	94	94
ystone	1,0%	82	97	99	100	68	83	94	98	85	99	100	100	67	89	96	99	72	86	87	90
	1,2%	85	98	100	100	70	86	95	99	85	100	100	100	68	91	98	100	76	88	89	90
	1,4%	89	99	100	100	73	90	98	100	89	100	100	100	71	94	99	100	79	90	90	92
lake	1,0%	80	98	99	100	67	88	94	98	83	99	100	100	68	86	96	99	69	85	88	94
	1,2%	83	99	100	100	69	90	96	100	86	99	100	100	69	89	98	99	71	87	90	93
	1,4%	87	100	100	100	72	94	98	100	89	100	100	100	71	93	99	100	73	89	92	94

Table 3: Measured success in detecting tags (in percent)

- Define probability limits for detecting enemies after receiving distorted images.
- Explore other geometrical shapes or overlapping shapes to carry tag information. Is spread spectrum technology applicable to the process of tagging?
- Adapt the ‘decomposition of deformation’ [6] to the analysis of tagged images.
- Develop better tag sequences for groups of enemies.
- Do extensive tests on different types of images.
- Find alternative methods to measure quality degradation of images.
- Analyze tagging in connection with confidential data and for steganographic purposes.
- Classify different possible types of tagging mechanisms, depending on the kind of document which is to be tagged.
- Study this approach in relation to the detection of covert channels [7].

Acknowledgements

The author would like to thank Bernhard Plattner and Ueli Maurer for their encouragement and support, which made this work possible.

References

- [1] Shu Lin, Daniel J. Costello jr., “Error Control Coding: Fundamentals and Applications”, Prentice Hall, 1983.
- [2] D. Kahn, “The Codebreakers”, Macmillan, New York, 1967, pp. 523.
- [3] Friedrich Bauer, “Kryptologie: Methoden und Maximen”, Springer-Verlag Berlin, 1993, pp. 5-20.
- [4] A.W. Gruen, “Adaptive Least Squares Correlation: A powerful image matching technique”, Report Number 115 of the Institute for Geodesy and Photogrammetry, ETH Zürich, 1986.
- [5] William K. Pratt, “Correlation Techniques of Image Registration”, IEEE Transactions on aerospace and electronic systems, vol AES-10, no 3, May 1974.

- [6] Fred L. Bookstein, "Principal Warps: Thin-Plate Splines and the Decomposition of Deformation", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol 11, no 6, June 1989, pp. 345-365.
- [7] National Computer Security Center, "A Guide to Understanding Covert Channel Analysis of Trusted Systems", (NCSC-TG-030), NCSC, National Security Agency, INFOSEC Awareness Division, Ft. George G. Meade, MD 20755-6000.
- [8] Germano Caronni, "Ermitteln unauthorisierter Verteiler von maschinenlesbaren Daten", in german only, Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology, August 1993.
- [9] J. Brassil, S. Low, N. Maxemchuk, L. O'Gorman, "Electronic Marking and Identification Techniques to Discourage Document Copying", *Proceedings of Infocom '94*, pp. 1278-1287, June 1994.
- [10] R. L. Rivest, A. Shamir, L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems", *CACM*, vol. 21, no. 2, pp. 120-127, Feb. 1987.
- [11] "Data Encryption Standard (DES)", NBS-FIPS Publication 46, National Technical Information Service, Springfield, VA, April 1977.
- [12] Xuejia Lai, "Detailed Description and a Software Implementation of the IPES Cipher", Institute for Signal and Information Processing, ETH Zürich, 1991.
- [13] M. Blum, S. Micali, "How to generate cryptographically strong sequences of pseudo-random bits", *SIAM J. Comput.*, vol. 13, no. 4, pp. 850-864, Nov. 1984.
- [14] Stephen K. Park, Keith W. Miller, "Random Number Generators: Good Ones are Hard to Find", *CACM*, vol. 31, no. 10, pp. 1192-1201, Oct. 1988.
- [15] Gregory K. Wallace, "The JPEG Still Picture Compression Standard", *CACM* vol. 34, no. 4, pp. 30-44, Apr. 1991.
- [16] J. T. Brassil, S. Low, N. F. Maxemchuk, L. O'Gorman, "Hiding Information in Document Images", Submitted to IEEE Symposium on Security and Privacy 1995.
- [17] Dan Boneh, James Shaw, "Collusion-Secure Fingerprinting for Digital Data", Technical Report at Princeton University (<ftp://ftp.cs.princeton.edu/reports/1994/468.ps.Z>), October 1994.
- [18] K. Tanaka, Y. Nakamura, K. Matsui, "Embedding secret information into a dithered multilevel image", *Proceedings of the 1990 IEEE Military Communications Conference*, pp. 216-220, September 1990.