On the Limitations of Digital Watermarks: A Cautionary Note

Stephen Wolthusen
Fraunhofer-Institute for Computer Graphics
64283 Darmstadt, Germany

ABSTRACT

The technology of digital watermarking has quickly become a respected concept for the solution of the copyright protection problems in the emerging global digital network. Content creators and vendors are increasingly aware of the pressing issues involved in protecting their intellectual property rights – which are crucial in establishing viable business models.
However, the author deems it necessary that the technical and practical limitations of this technology be shown so that all parties involved are endowed with a more realistic understanding of the capabilities of digital watermarking, the current limitations of the technology, as well as the limits of security that are achievable under realistic assumptions and the architectural restrictions deriving from this.

1. INTRODUCTION TO DIGITAL WATERMARKS

Digital watermarks are descendants of the ancient technology of steganography which has gained new important applications with the rise of digital data representation and automated techniques for data hiding in digital multimedia data (see [23] for a historical perspective, [47] for a general high-level introduction and the excellent bibliography [4] for further reading and new papers; much of the recent interest in watermarking and steganography was kindled by the Info Hiding Workshop in 1996[2]).
Digital watermarks are an extension of steganography in that they impose an additional robustness criterion on the embedded information, i.e. the watermark must be difficult or impossible to remove without noticeably degrading the original data and survive modifications that do not significantly degrade the quality of the data. At the same time the watermark must be imperceptible to all but the most discerning observers yet it must also be easily detectable by proper authorities, if at all possible without knowledge of the original data.
Various techniques have been suggested for steganographic and watermarking purposes – particularly for image data – ranging from simple modifications in the spatial domain [45,28] over DCT constraints and [25,26,49,41,14] to fractal compression schemes[39]. Besides images, models have been proposed for structured textual data[10,29], audio signals[44,9,19] and even 3D modeling data [32].
It should be noted that the difficulties and at the same time scientific challenges arise mainly from contradictory requirements imposed on digital watermarks. A common operation performed on multimedia data due to their redundancy is lossy compression; such algorithms are designed so that data imperceptible to humans is removed[22,46,21,33]. This caused Cox et al.[14] to postulate the requirement that watermarks need to be placed in perceptually significant portions of the data so that they will survive such compressions – in direct contradiction to the original requirements stated above. This conundrum is further deepened by the need for the watermark to be embedded in such a way as to be indistinguishable from normal background noise since otherwise it might be possible to derive the original data from the watermarked.

Digital watermarking applications and their requirements
Digital watermarking can be used for a variety of purposes; the original thrust of the body of literature as well as that of commercialization is in the area of copyright protection. Other possible applications include embedding hidden annotations and authentication of data and documents; the use of watermarks as a covert communications channel in situations where cryptography is not permissible and robustness is also an issue. Each of these applications has different requirements (e.g. higher than usual amount of embeddable data for hidden annotations as opposed to a high degree of robustness for limited data amounts in the case of copyright protection.

2. CONVENTIONAL THREATS TO DIGITAL WATERMARKS

The following discussion concentrates on the field of image watermarks used for copyright protection. This is largely due to the overwhelming majority of watermarking schemes being available for images and the already established commercial utilization of watermarks for intellectual property protection by such companies as MediaSec, Digimarc, and Highwater Signum.
Conventional threats to digital watermarks as may occur in innocent image manipulation include among others lossy compression (e.g. JPEG), cropping, color variations, rotations and local changes to the image (e.g. blackening out a logo).
Current systems are highly resilient against lossy compression such as according to the JPEG standard and slight modifications introduced in common image editing processes as may be found in both the traditional paper and web-based publishing industries (the photo in Figure 1 shows an original image (an image of a wolf scanned at a resolution of 1152x900); the photo in Figure 2 shows the same image after being watermarked and compressed by 90%; the watermark can still be recovered with a certainty of 99.999% even with posterization of full 8x8 pixel blocks that are plainly visible). Certain watermarks such as those produced by SysCoP[51] even survive the distortions found after printing and re-scanning digital images. Most systems survive modifications such as Gaussian blur, cropping, gamma correction, greyscale

conversion, hue changes, noise addition, rotations, changes in color saturation and scaling fairly well, even if noticeable changes and degradations are introduced into the image. From this it appears that at least the primary requirements for successful applications of digital watermarks are met by current systems, i.e. unlawful use of intellectual property is detectable even after modifications have taken place. While that much is true and is promoted by the watermarking companies, there are still a number of problems and issues to be resolved before this research area can be considered closed. However, the current systems reach their limits when confronted with a well-chosen blend of such image modifications (which has to be determined by experiments for each individual image).

Figure 1: Original image

Some limitations of steganographic technologies (and, more specifically, watermarking systems) are discussed in [5,18] The main issues to be considered are the constraints imposed by the imperceptability requirement (since such images are supposed to be sold for commercial purposes, degrading their quality would significantly affect the market value of any such images), the need for the embedded signal to be sufficiently indistinguishable from random noise (since otherwise it would at least be possible to find and correlate statistically significant signatures in a spectral analysis of the image, providing a powerful tool for the removal of the watermark), leading among others to the spread-spectrum approach[14,15] to digital watermarking. At the time of this writing it appears unlikely that a totally robust watermarking system resilient to all but the most egregious modifications can be constructed.
The abovementioned combination attacks on digital watermarks of images can to some degree be automated; examples of such programs are StirMark[27] which uses a combination of stretching, shearing, shifting, and rotating the watermarked image followed by interpolating the new pixel values and introducing minor errors, and UnZign [6] (no information were available on the techniques used in the UnZign process). Both systems have some success in making the embedded watermarks unreadable - yet at the cost of introducing visible artifacts into images (these are more pronounced in the case of UnZign). A properly calibrated digital watermark containing highly robust (but limited) copyright information does stand a good chance of surviving such attacks. For a more detailed discussion of possible automated attacks see for example [34].

3. UNFAIR ATTACKS

Attacks on fortresses have often succeeded not because fortifications were not designed strong enough or because the masonry was executed badly, but rather because of something the fortress builder had not taken into consideration. A similar situation is present in any case where digital data is to be protected and one cannot rely on cryptography alone to provide security. Even if the ciphers used are in themselves correct and secure, there are a multitude of systems that have failed the test of time because of failures in protocol and implementation.
While section 2 has concentrated on the strength of the actual watermarking mechanism, this section is concerned mainly with protocol attacks (specifically, attacks against copyright protection mechanisms and the consequences for the creation of secure digital watermarking systems.

Figure 2: Watermarked image after 90% compression

One of the earliest proposals for attacks on copyright protection watermarking systems (although predated by a report from NEC[42], presenting attacks on watermarks based on sample differences) was discussed in [16], this has become known as the „IBM attack“. In it the authors propose an attack based on the creation of counterfeit originals, relying on the invertibility of some watermarking schemes. Besides the proposition made in [16] to use a non-invertible watermarking scheme, another obvious solution is the use of time stamps. The former method – while in itself an useful characteristic of a watermarking scheme – still must be considered only a plausible conjecture.
Opportunities for attacking watermarking schemes also arise from the way the watermarking is used. Consider the application of publicly readable watermarks for the identification of the intellectual rights holder (each watermark contains an identification number which can be resolved into contact information when sending this number to a clearinghouse database); this is the model pursued by some commercial watermarking applications (e.g. Digimarc's PictureMarc, Highwater Signum's SureSign). While the

watermarks embedded by these products are remarkably robust (which in part is due to the fact that the volume of data is very small), there are other ways besides those discussed in [27]. While that approach is ignorant of the algorithms used in the watermarking software (and is therefore applicable to almost every watermarking software), there may be ways to overcome this particular class of watermarks. The algorithms used to embed the actual watermarks are kept secret – but since they are symmetrical (i.e. the same technique is used for embedding and retrieving the watermarks) and the retrieval software is distributed freely by a number of graphics software vendors and also by the companies themselves, it will ultimately reach the hands of determined adversaries. Disassembling and reverse-engineering software is – while prohibited by licensing agreements in some cases and also a relatively uncommon skill – very much a fact of the security software industry and can be used to find out and understand the no longer secret algorithm.

That knowledge of the algorithm can – in the vast majority of cases where there is natural noise in the image such as photographs and related artwork – be used to create another image, interpolating the data modified by the watermarking algorithm (this was also suggested by [18]) and removing any tell-tale modifications performed by the algorithm (such as modifications only in places where a watermark would have left traces) while retaining the image quality. Such images cannot be discerned quality-wise from the originals, even though it will most likely differ from them. The situation in the case where no natural noise is present (i.e. computer-generated imagery) requires a slightly different approach but has an even greater chance of recreating the original. After identifying the modifications performed on such imagery (which are notoriously difficult to perform without visibly degrading the image in the first place) one can use patches of unmodified portions of the image or recreate gradients found elsewhere to paste over modifications in the image. Alterations in the geometry of the image can also be detected.

If one wishes to further confuse the situation, such an image can then be used to embed one's own digital watermark (obviously using the same algorithm which is usually image-dependent is best for this purpose). A proof of ownership in this case is close to impossible.

Does this mean that digital watermarking is obsolete before it has achieved the breakthrough predicted by many experts[7]? Not necessarily. The true implications of this are that one needs to take a careful look at the infrastructure used for digital watermarking and the resolution of ownership rights questions so as to thwart potential attacks.

First, one must assume the algorithms used for watermarking to be common knowledge. Violating Kerckhoffs' principles[24] has cost cryptographers dearly in the past and would undoubtedly also hurt the vendors and users of digital watermarking systems.

Starting from this assumption one is faced with the need both providing a watermark readable by users (while this can also be achieved using a label attached to an image, a solution based on digital watermarks has the distinct advantage of being retained even when the data representation is changed or slightly modified, giving potential users more freedom to use the data to suit their needs) and one readable only by the creator with a publicly available software package. It is important to note that the public watermark can ultimately be removed given the current technology and foreseeable developments in the field.

Kerckhoff stated that the security of a cryptographic system should depend only on the key itself. That is also the case with secret digital watermarks. To briefly sum up the requirements imposed on a watermarking scheme, it should be dependent on the data being watermarked, on a secret key known only to the watermarker (and possibly trusted third parties) and be indistinguishable from random noise found in a spectral analysis of the data[14,15]. Obviously the scheme should also meet essential robustness criteria[26].

The problem of duplicate creation should also be addressed conservatively. Depositing the results of a one-way hash function on the original data along with the secret key used to watermark the image, possibly protected by a secret-sharing scheme [43,40]) with a trusted third party (TTP) which is required to perform a digital signature upon the data in conjunction with a time stamp once a watermarking user submits a data obviates the IBM attack as well as the one outlined above. If the legitimate owner of the watermarked work has identified one of his creations in the possession of a third party with which no licensing agreement exists (in most cases by means of the public watermark), he can prove his legitimate claim of ownership by asking the TTP to verify the secret watermark (after doing so himself to ensure that no false accusations are raised). Claims by the holder of the counterfeit that the disputed data are in fact original data can be refuted by referring to the time stamp of the original submission of the claimant although it should be noted that there is a small grey area when it comes to images that have been altered significantly – whether such modified images constitute original art is certainly beyond the scope of a watermarking algorithm.

4. CHALLENGES TO DIGITAL WATERMARKS AS EVIDENCE

Proving a negative

So far the discussion of watermarking systems has concentrated on proving that an image originated from a given source, it is conceivable that the opposite will be necessary at some point[50]. One scenario results from the concern that image processing has advanced sufficiently that image manipulations that introduce different semantics are almost indistinguishable. That, combined with a trend towards the use of digital technology in all stages from cameras to electronic distribution presents a grave danger to the notion of photographic evidence in courts of law. Unprotected digital imagery can be used to prove anything, leading to the consequence that any such evidence must be banned and courts are once again to rely on witnesses only. Such a situation is clearly unacceptable.

One possible approach to this problem is to embed digital watermarking at the acquisition stage (i.e. digital cameras and audio recording equipment) directly in hardware. This is to ensure that only data which has not been tampered with significantly can be identified as original photographs and audio recordings. Since one of objectives here is to prevent fabrication of evidence, the photographer or the person recording audio data cannot be presumed trustworthy in this scenario.

It is therefore necessary to embed the secret key used in watermarking the digital data (as discussed in section 3)

directly and without the possibility for retrieval in the hardware used for creating the data. This is well within the capabilities of the technology and should impose no significant cost overhead except possibly for the need to make each recording device unique. The information making each device unique constitutes the secret key and must be known only to a trusted third party (to alleviate privacy concerns the use of a secret sharing scheme in which a customer is given a required datum for the identification process can once again be considered). If a dispute over the authenticity and source of an image arises, the TTP can easily identify the device used to create the data as well as ensure that no significant alterations of the image have taken place. (For an introduction to tamper-proof systems, see for example [20,13,48]).
This does by no means constitute a perfect scheme [3], but should serve to thwart all but the most determined attackers. This situation should be seen in the context of other types of evidence commonly accepted in courts of law all over the world. None of these types of evidence is in itself invulnerable to tampering and fabrication, yet the knowledge that such actions are beyond the reach of the vast majority of presumed perpetrators is considered sufficient.

Collusion attacks

Another possible threat to digital watermarking schemes arises when the same data (the following discussion is - without loss of generality - based on images) is watermarked multiple times and then distributed. This scenario is quite common when the goal is to link an image to a particular licensee and to identify the particular person once a breach of a licensing agreement occurred.
This method of identifying perpetrators has a high price in the form of a class of attacks known as collusion attacks.
Collusion attacks by definition cannot occur when only a single watermarked copy of an image is ever distributed and the original image is kept secret. The price for this is that only illegal users can be found and prosecuted while counterfeiters selling the images as their own cannot be traced.
There are a number of protocols that have been proposed to allow the distribution of multiple differently signed copies of watermarked images [12,8,37,38,35,36]. While these research papers provide a solution to the problem of collusion attacks, there are several problems involved which preclude these schemes from widespread use and applicability. First among them is an assumption that is common to all schemes proposed so far called the Marking Assumption, first proposed in [8]. It states that The main property the marks should satisfy is that users can not change the state of an undetected mark without rendering the object useless.
Obviously, given the limits of digital watermarks of today and the foreseeable future this assumption cannot hold for real watermarking schemes.
A second problem involving such codes is that for acceptable security given a larger distribution scale (for a maximum of c traitors out of N given an error probability ε), the size of the codewords to be embedded into the digital watermarks are $O(c^4 \log(\frac{N}{\varepsilon}))$ [8]. Embedding such volumes of data is certainly a hard problem in any case but should prove extremely difficult when coupled with the need to satisfy the Marking Assumption since such markings need to be highly robust and therefore redundant. Therefore, sadly, unless significantly better algorithms can be found, traitor tracing is of academic interest only.
This leads to the conclusion that the desirable feature of tracing individuals that have breached licensing agreements by technical means is as of yet not feasible to implement in watermarking schemes. To achieve a functional equivalent one is forced to resort to organizational means as well as elaborate licensing agreements. It also means that the only safe watermarking systems are those where only one single watermarked copy is released and the unwatermarked copy is kept safe.

Signal processing collusion

Otherwise collusion attacks (with the modified meaning that it is not the intention of the traitors to create marked copies that cannot be traced to an individual but rather to create copies that are indistinguishable in terms of quality from the original image) are feasible. One such collusion attack is outlined below.
The idea behind the attack is that digital watermarks can usually be interpreted as noise over a spread spectrum (usually after some appropriate transformations such as a DCT [1], wavelet transform [11,17] or FFT [30,31] have been performed). If multiple watermarked images or other data, again, assuming images without loss of generality, (NB not multiple watermarks within the same image — this is perfectly acceptable) exist and a number of traitors conspire by combining their images, this will result in an image of a quality deemed acceptable for general distribution by the owner of the original image. The spread-spectrum noise embedded by the digital watermark has a certain energy that is deemed an acceptable degradation of the image quality by the creator of the digital watermark. Since the noise can, by most watermarking system definitions, reside anywhere within a predetermined frequency band, it must be presumed that adding noise energy anywhere within this frequency band while retaining the original noise energy will also result in an acceptable quality signal.
One can easily see that by averaging the signals that represents the different copies of the image one obtains a combined signal in which for a sufficient number of watermarked images the watermarking noise is effectively cancelled and the original image (with a smaller amount of noise present than in any of the watermarked images) will emerge. This effectively defeats individual watermarking schemes proposed for use tracking, although it does not necessarily affect a watermark common to all distributed copies which may still be used to track illicit use of the digital data.

5. CONCLUSION

This paper has attempted to present the challenges faced by digital watermarking technology when entering the application domains envisioned for it. While much progress has been made in recent years in terms of the achievable robustness of embedded watermarks and their perceptability, dedicated attackers can still outmaneuver these algorithms in most cases, even if doing so requires manual intervention.

At the same time it should have become clear that possible attacks impose certain restrictions on every watermarking algorithm; in particular, the notion that one can uniquely identify the purchaser of digital content (and therefore prosecute that particular user along with unlicensed users in case illegal use is detected) must be given up since even a small collusion (depending on the quality and robustness of the watermark) can create fraudulent data by combining differently watermarked copies.

Similarly, the introduction of a Trusted Third Party that serves as a clearinghouse for watermarked data and a timestamping service for content providers is virtually inevitable for reasonably secure systems since only timestamping can provably be used to counter the „IBM attack"[16]; at the same time a TTP is required in case digital watermarks do not work as a deterrent alone but their existence in an illegal copy must actually be proven in a court of law.

Finally, for digital watermarking to be useful in the authentication of original imagery and audio recordings, an infrastructure and set of standards for digital watermarks and their tamper-proof embedding into recording devices and cameras needs to be created. Only the creation of such an architecture effectively ensures the continued viability of digitally stored and processed environmental data as evidence.

REFERENCES

[1] Ahmed, N., Natarajan, T., and Rao, K. R. On image processing and a discrete cosine transform. *IEEE Transactions on Computers C-23* (1974), 90–93.

[2] *IEEE Transactions on Computers C-23* (1974), 90–93. Anderson, R., Ed. *Information hiding: first international workshop, Cambridge, U. K.*, May 30-June 1, 1996: proceedings (New York, NY, USA, 1996), vol. 1174 of *Lecture Notes in Computer Science*, Springer-Verlag Inc.

[3] *Information hiding: first international workshop, Cambridge, U. K.*, May 30-June 1, 1996: proceedings (New York, NY, USA, 1996), vol. 1174 of *Lecture Notes in Computer Science*, Springer-Verlag Inc. Anderson, R., and Kuhn, M. Tamper resistance – a cautionary note. In *Second USENIX* Workshop on Electronic Commerce (Oakland, California, Nov. 1996), USENIX, pp. 1–11.

[4] In *Second USENIX* Workshop on Electronic Commerce (Oakland, California, Nov. 1996), USENIX, pp. 1–11. Anderson, R. and Petitcolas, F. Information Hiding: An Annotated Bibliography. Tech. rep., Computer Laboratory, University of Cambridge, U. K., January 1998.

[5] Tech. rep., Computer Laboratory, University of Cambridge, U. K., January 1998. Anderson, R. and Petitcolas, F. On The Limits of Steganography. *IEEE Journal on Selected Areas in Communications: Special Issue on Copyright & Privacy Protection* (April 1998).

[6] *IEEE Journal on Selected Areas in Communications: Special Issue on Copyright & Privacy Protection* (April 1998). Anonymous. UnZign. Available for download at http://altern.org/watermark/, 1997. The author can be contacted at unzign@hotmail.com.

[7] The author can be contacted at unzign@hotmail.com. Berghel, H. Watermarking Cyberspace. *Communications of the ACM* 40, 11 (November 1997), 19–24.

[8] *Communications of the ACM* 40, 11 (November 1997), 19–24. Boneh, D., and Shaw, J. Collusion-secure fingerprinting for digital data. *Lecture Notes in Computer Science 963* (1995), 452–??

[9] *Lecture Notes in Computer Science 963* (1995), 452–?? Boney, L., Tewfik, A. H., and Hamdy, K. N. Digital watermarks for audio signals. In *1996 IEEE Int. Conf. on Multimedia Computing and Systems* (Hiroshima, Japan, 1996), pp. 473–480.

[10] In *1996 IEEE Int. Conf. on Multimedia Computing and Systems* (Hiroshima, Japan, 1996), pp. 473–480. Brassil, J., Low, S., Maxemchuk, N., and O'Gorman, L. Electronic marking and identification techniques to discourage document copying. In *Proceedings of Infocom '94* (June 1994), pp. 1278–1287.

[11] In *Proceedings of Infocom '94* (June 1994), pp. 1278–1287. Chan, Y. T. *Wavelet Basics*. Kluwer Academic Publishers, Boston, 1995.

[12] Kluwer Academic Publishers, Boston, 1995. Chor, B., Fiat, A., and Naor, M. Tracing traitors. In *Advances in Cryptology—CRYPTO* '94 (21–25 Aug. 1994), Y. G. Desmedt, Ed., vol. 839 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 257–270.

[13] In *Advances in Cryptology—CRYPTO* '94 (21–25 Aug. 1994), Y. G. Desmedt, Ed., vol. 839 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 257–270. Choudhury, A. and Maxemchuk, N. and Paul, S. and Schulzrinne, H. Copyright Protection for Electronic Publishing over Computer Networks. *IEEE Network Magazine* (June 1994), 18.

[14] *IEEE Network Magazine* (June 1994), 18. Cox, I. J., Kilian, J., Leighton, T., and Shamoon, T. Secure spread spectrum watermarking for multimedia. Technical Report 95-10, NEC Research Institute, 1995.

[15] Technical Report 95-10, NEC Research Institute, 1995. Cox, I. J., and Miller, M. L. A review of watermarking and the importance of perceptual modeling. In *Proc. of Electronic Imaging '97* (February 1997).

[16] In *Proc. of Electronic Imaging '97* (February 1997). Craver, S., Memon, N., Yeo, B.-L., and Yeung, M. Can invisible watermarks resolve rightful ownerships? Tech. Rep. RC 20509, IBM Research Division, July 1996.

[17] Tech. Rep. RC 20509, IBM Research Division, July 1996. Edwards, T. Discrete wavelet transforms: Theory and implementation. Department of Statistics, Stanford University, 1991.

[18] Department of Statistics, Stanford University, 1991. Fridrich, J. and 2LT Baldoza, A. and Simard R. On Digital

Watermarks. In *2nd Information Hiding Workshop* (April 1998), Springer Verlag.

[19] In *2nd Information Hiding Workshop* (April 1998), Springer Verlag. Gruhl, D. and Lu, A. and Bender, W. Echo Hiding. In *Info Hiding 96* (1996), vol. 1174 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 295–315.

[20] In *Info Hiding 96* (1996), vol. 1174 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 295–315. Hauser, R. *Control of Information Distribution and Access*. PhD thesis, Department of Computer Science, University of Zurich, Winterthurerstr. 190, CH-8057 Zurich, Switzerland, tel: +41-1-257-4311, fax: +41-1-257-4505, Sept. 1995.

[21] PhD thesis, Department of Computer Science, University of Zurich, Winterthurerstr. 190, CH-8057 Zurich, Switzerland, tel: +41-1-257-4311, fax: +41-1-257-4505, Sept. 1995. ISO/IEC-JTC1/SC29. Very-low bitrate audio-visual coding. Tech. rep., ISO, Mar. 1993. New Work Item Proposal.

[22] New Work Item Proposal. Joint Photographic Experts Group ISO/IEC, JTC/SC/WG8, CCITT SGVIII. JPEG technical specifications, revision 5. *Report JPEG-8-R5* (Jan. 1990).

[23] *Report JPEG-8-R5* (Jan. 1990). Kahn, D. *The Codebreakers*, 2nd ed. Scribner, 1996.

[24] Scribner, 1996. Kerckhoffs, A. La cryptographie militaire. *Journal des Sciences Militaire* 9th series (February 1883), 161–191.

[25] *Journal des Sciences Militaire* 9th series (February 1883), 161–191. Koch, E., Rindfrey, J., and Zhao, J. Copyright Protection for Multimedia Data. In *Proc. of the International Conference on Digital Media and Electronic Publishing* (December 1994).

[26] In *Proc. of the International Conference on Digital Media and Electronic Publishing* (December 1994). Koch, E., and Zhao, J. Towards robust and hidden image copyright labeling. In *Proc. of 1995 IEEE Workshop on Nonlinear Signal and Image Processing* (Halkidiki, Greece, June 1995), pp. 452–455.

[27] In *Proc. of 1995 IEEE Workshop on Nonlinear Signal and Image Processing* (Halkidiki, Greece, June 1995), pp. 452–455. Kuhn, M. Stirmark. Available for download at http://www.cl.cam.ac.uk/ mgk25/download/, November 1997. Security Group at the Computer Lab, Cambridge, U. K.

[28] Security Group at the Computer Lab, Cambridge, U. K. Kutter, M., Jordan, F., and Bossen, F. Digital signature of color images using amplitude modulation. In *Proc. SPIE Storage and Retrieval for Image and Video Databases* (San Jose, California, 1997), vol. 3022, pp. 518–526.

[29] In *Proc. SPIE Storage and Retrieval for Image and Video Databases* (San Jose, California, 1997), vol. 3022, pp. 518–526. Low, S., Maxemchuk, N., Brassil, J., and O'Gorman, L. Document marking and identification using both line and word shifting. In *Infocom '95* (Boston, MA, April 1995).

[30] In *Infocom '95* (Boston, MA, April 1995). Maslen, D. K., and Rockmore, D. N. Generalized FFTS - A Survey of Some Recent Results. Tech. Rep. PCS-TR96-281, Dartmouth College, Computer Science, Hanover, NH, Apr. 1996.

[31] Tech. Rep. PCS-TR96-281, Dartmouth College, Computer Science, Hanover, NH, Apr. 1996. Nussbaumer, H. J. *Fast Fourier Transform and Convolution Algorithms*. Springer, Berlin, 1981.

[32] Springer, Berlin, 1981. Ohbuchi, R. and Masuda, H. and Aono, M. Watermarking Three-Dimensional Polygonal Models. In *ACM Multimedia '97* (Seattle, WA, November 1997), ACM Multimedia.

[33] In *ACM Multimedia '97* (Seattle, WA, November 1997), ACM Multimedia. Pan, D. Y. Digital audio compression. *Digital Technical Journal of Digital Equipment Corporation 5*, 2 (Spring 1993), 28–33 (or 28–40??).

[34] *Digital Technical Journal of Digital Equipment Corporation 5*, 2 (Spring 1993), 28–33 (or 28–40??). Petitcolas, F. and Anderson, R. and Kuhn, M. Attacks on Copyright Marking Systems. Unpublished., 1998.

[35] Unpublished., 1998. Pfitzmann, and Waidner. Anonymous fingerprinting. In *EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT* (1997).

[36] In *EUROCRYPT: Advances in Cryptology: Proceedings of EUROCRYPT* (1997). Pfitzmann, B. Trials of traced traitors. *Lecture Notes in Computer Science 1174* (1996), 49–??

[37] *Lecture Notes in Computer Science 1174* (1996), 49–?? Pfitzmann, B., and Schunter, M. Asymmetric fingerprinting (extended abstract). In *Advances in Cryptology—EUROCRYPT* 96 (12–16 May 1996), U. Maurer, Ed., vol. 1070 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 84–95.

[38] In *Advances in Cryptology—EUROCRYPT* 96 (12–16 May 1996), U. Maurer, Ed., vol. 1070 of *Lecture Notes in Computer Science*, Springer-Verlag, pp. 84–95. Pfitzmann, B., and Waidner, M. Asymmetric fingerprinting for larger collusions. Research Report RZ 2857 (#90805), IBM Research, Aug. 1996. accepted for: 4th ACM Conference on Computer and Communications Security, Zurich 1997.

[39] accepted for: 4th ACM Conference on Computer and Communications Security, Zurich 1997. Puate, J., and Jordan, F. Using fractal compression scheme to embed a digital signature into an image. In *Proceedings of SPIE Photonics East'96 Symposium* (Boston, Massachusetts, 1996).

[40] In *Proceedings of SPIE Photonics East'96 Symposium* (Boston, Massachusetts, 1996). Shamir, A. How to share a secret. *Communications of the ACM 22* (1979), 612–613.

[41] *Communications of the ACM 22* (1979), 612–613. Smith, J. R., and Comiskey, B. O. Modulation and information hiding in images. In *Workshop on Information Hiding* (Isaac Newton

Institute, University of Cambridge, UK, May 1996), vol. 1174 of *Springer-Verlag Lecture Notes in Computer Science*.

[42] In *Workshop on Information Hiding* (Isaac Newton Institute, University of Cambridge, UK, May 1996), vol. 1174 of *Springer-Verlag Lecture Notes in Computer Science*. Stone, H. Analysis of Attacks on Image Watermarks with Randomized Coefficients. Tech. rep., NEC Research Institute Technical Report, 1996.

[43] Tech. rep., NEC Research Institute Technical Report, 1996. Sykes, D. The management of encryption keys. In *Computer Security and the Data Encryption Standard, NBS Special Publication 500-27*, Branstad, D., Ed. National Bureau of Standards, 1977, pp. 46–53.

[44] In *Computer Security and the Data Encryption Standard, NBS Special Publication 500-27*, Branstad, D., Ed. National Bureau of Standards, 1977, pp. 46–53. Tilik, J. F. and Beex, A. A. Encoding a hidden digital signature onto an audio signal using psychoacustic masking. In *Proc. 1996 7th International Conf. on Signal Processing Apps and Tech.* (1996), pp. 476–480.

[45] In *Proc. 1996 7th International Conf. on Signal Processing Apps and Tech.* (1996), pp. 476–480. van Schyndel, R., Tirkel, A., and Osborne, C. A digital watermark. In *IEEE International Conference on Image Processing (ICIP'96)* (Los Alamitos, CA, 1994), vol. II, IEEE Press, pp. 86–90.

[46] In *IEEE International Conference on Image Processing (ICIP'96)* (Los Alamitos, CA, 1994), vol. II, IEEE Press, pp. 86–90. Wallace, G. K. Overview of the JPEG (ISO/CCITT) still image compression standard. In *Proc. SPIE's Visual Communications and Image Processing* (1989). Paper distributed at conference.

[47] Paper distributed at conference. Wayner, P. *Disappearing Cryptography*. Academic Press, 1996.

[48] Academic Press, 1996. White, S. and Weingart, S. and Arnold, W. and Palmer, E. Introduction to the Citadel Architecture: Security in Physically Exposed Environments. Tech. Rep. RC 16672, IBM T. J. Watson Research Lab, May 1991. Version 1.4.

[49] Version 1.4. Zhao, J. Embedding Robust Labels Into Images For Copyright Protection. In *Proc. of the International Congress on Intellectual Property Rights for Specialized Information, Knowledge and New Technologies* (August 1994).

[50] In *Proc. of the International Congress on Intellectual Property Rights for Specialized Information, Knowledge and New Technologies* (August 1994). Zhao, J. Personal communication, 1997.

[51] Personal communication, 1997. Zhao, J. and Koch, E. A Digital Watermarking System for Multimedia Copyright Protection. In *ACM Multimedia '96* (Boston, MA, November 1996), ACM Multimedia.