

White Paper

**Embedding and
Retrieval-in-Order
of Multiple
Watermarks in
PACS Images**

Published: January 2011

www.astoninventions.com/WM

White Paper

Embedding and Retrieval-in-Order of Multiple Watermarks in PACS Images

Contents

1	Executive Summary	1
2	Introduction	2
3	DICOM Security	3
4	Digital Watermarking	5
5	A Novel Multiple Watermark Embedding Scheme	9
6	Applying the Technology.....	11
7	Conclusion	12

1 Executive Summary

With the rapid diffusion of Picture Archiving and Communications Systems (PACS) throughout the healthcare industry, and the increased sharing of medical images between healthcare institutions, interoperability is now a key requirement. Interoperability has however led to increased security concerns, particularly relating to compromised patient anonymity. While the Digital Imaging Communications in Medicine (DICOM) interoperability standard allows for various security technologies, the biggest threat to data security is often posed by disgruntled or negligent employees—those with authorised access to confidential data, but who deliberately or accidentally distribute medical images. Even highly trusted employees can pose a significant threat to data. The use of digital watermarking technology developed at Aston University in Birmingham, England can prevent such activity by enabling traitor-tracing and by acting as a deterrent to would-be traitors. Other benefits include storing an embedded audit trail within each image.

An in-depth technical description is available from <http://www.astoninventions.com/WM>.

2 Introduction

Modern society increasingly relies on digitized information that can be easily accessed, copied and transmitted. Numerous technologies have been developed to facilitate this reliance: the floppy disk, email, FTP, the Zip drive, recordable CDs and DVDs, the flash memory drive, and P2P, among many others. These technologies have however also enabled the dramatic rise of copyright infringement seen in recent years.

In parallel, the medical field has seen the broad diffusion of Picture Archiving and Communications Systems (PACS), with tens of thousands of installations now in place worldwide. The benefits of such systems are immense. Each year, up to 20 percent of traditional medical images are estimated to go missing, leading to cancelled clinics and procedures, inconvenienced patients, increased costs, and often compromised diagnoses. PACS ensures that medical images of patients are easily and instantly available, so hospitals and clinicians can save time and money and increase their efficiency.

With dozens of firms developing PACS applications, interoperability is a key issue. And, with the rapid growth of telemedicine, it's increasingly important for a specialist in one facility using a particular PACS application to readily access medical images stored in a different facility that uses a different PACS application.

The Digital Imaging Communications in Medicine (DICOM) standard was developed to facilitate interoperability of PACS applications. DICOM dictates how objects (X-ray, CT, MRI or PET images, reports, measurements, etc.) should be structured, and how they can be exchanged. To facilitate this exchange, these systems usually utilise the same data sharing technologies used in copyright infringement.

Interoperability therefore inevitably leads to increased security concerns. In particular, interoperability increases the risk that patient anonymity will be compromised, potentially leading to a weakening of the patient-doctor relationship. Although DICOM has been updated frequently since it was formulated in 1985, security and confidentiality are only guaranteed insofar as the authorised end-users of DICOM-based PACS systems are themselves trustworthy and diligent in adhering to process and procedure.

This white paper argues that the use of digital watermarking technology can diminish the threat posed by insiders, by providing imperceptible, robust, and traceable evidence of unauthorised activities by particular individuals. This paper also proposes a patent-pending scheme that meets this need.

The technology described here can also be used in military and intelligence-gathering operations, in particular to protect sensitive data, such as satellite images or video recordings.

3 DICOM Security

The advent of picture archiving and communications systems (PACS) in radiology brings its own challenges as regards security but confidentiality in PACS must be maintained on the same basis as any other aspect of the practice of medicine—
The Royal College of Radiologists.

Unfortunately, the DICOM security model is relatively weak on matters of security, and only assumes that implementers of DICOM systems will utilise appropriate security policies. For instance, DICOM files are unencrypted and open to all users, whether authorised or not.

Nevertheless, DICOM does suggest the use of various security mechanisms, such as access control, anonymization, cryptography, firewall, intrusion detection and virtual private network software. When used together as a package, these mechanisms provide a powerful toolkit for ensuring confidentiality as well as reliability and availability of medical images.

Nonetheless, even well-implemented security schemes are not flawless. Cryptographic systems provide a case in point. Such tools preserve an encrypted copy of data, rather than the original data. To view the encrypted data, authorised end-users utilise a cryptographic key to restore the data to a viewable state. Each person who needs access to the data therefore requires the key.

However, encrypted data can still be copied and distributed widely using various Internet applications or portable storage media. Crucially, encryption keys can also be distributed along with the encrypted data. A skilled person who receives both the encrypted data and the encryption key can therefore view the data with relative ease.

To resolve this significant threat, a number of vendors have devised schemes that use a digital signature stored on a hardware-based cryptographic device. Such schemes essentially use microprocessors embedded within USB sticks or smartcards to encrypt and decrypt the data. These provide improved security because the physical device is needed to view the image. These devices can nevertheless be physically removed by a determined attacker, and can reduce the availability of an image.

Other vendors have developed Data Loss Prevention (DLP) applications, which monitor or prevent transfer of data. These applications tend not to be useful however when the sharing of confidential information is an essential part of the job, when confidential information is viewed over the Internet, or when the data remains encrypted during transfer.

3.1 Internal Security Threats

No security scheme can moreover prevent an insider with full authorisation to access and share the data from illegitimately distributing copies of the data, either intentionally or accidentally.

How big is this threat? It's well documented that authorised personnel pose the greatest threat to data security. A study conducted on behalf of IT-giant Cisco found that most IT professionals believe their company's employees pose a greater threat to data security than any outside source. The biggest threats were posed by disgruntled and negligent employees.

Another study found that 48% of over 900 data breaches investigated by Verizon and the US Secret Service during 2004-9 were caused by insiders. Of these, 90% were thought deliberate, with the insiders often working on behalf of external agents. Yet another study, of 57,000 internal security incidents during 2008, found that 19% of firms thought the incidents to be deliberate. Still another study found that 51% of data breaches were thought intentional while 43.5% were thought accidental.

Take for example the recent leaks to the Wikileaks website of 500,000 classified documents regarding the wars in Iraq and Afghanistan and 250,000 cables between US embassies around the world and the State Department. The alleged perpetrator of these leaks was a US Army intelligence analyst with high-level security clearance.

The cost of such data breaches is thought to be massive. The authoritative Ponemon Institute's recently published *Benchmark Study on Patient Privacy and Data Security* report found a total economic burden of data breaches on the healthcare industry of nearly \$6 billion annually, or \$1 million per healthcare organization. The typical organization suffered about 1.2 data breaches per annum, with an average of 1,769 lost or stolen records per breach. Some 15% of breaches were caused by malicious insiders, with a further 10% caused by intentional but non-malicious employee actions. Some 52% of breaches were unintentional.

Therefore, it is not unreasonable to suggest that such data breaches pose a significant threat to the revenues and profits of healthcare organizations, as well as a heightened risk of litigation by patients suffering from compromised confidentiality and even prosecution in countries with strict data protection legislation.

4 Digital Watermarking

Digital watermarking is a technique for embedding digital identification codes into data files, whether the file is a still image, a video stream, or an audio stream. Watermarking involves making subtle changes to the original data, known as the *cover*, so that the watermark is imperceptible to all, or so that a visible identification watermark cannot be removed without detailed knowledge of the creation process and other know-how. For instance, identification data is embedded by increasing the level of noise in a particular region of the cover data, or by rearranging existing noise.

In a generic watermark encoding scheme, the cover data, a watermark and a key are fed into a watermark encoding algorithm (Figure 1). This algorithm then outputs the watermarked data, consisting of the cover data with an embedded watermark.




Figure 1: Generic watermark encoder

In a generic watermark decoding scheme, the same components, along with the watermarked data, are fed into a watermark decoding algorithm (Figure 2). This algorithm then outputs the watermark. The same key is used in both the encoding and decoding process. When published to aid in data authentication, the key is known as a *public key*. When kept private to increase the integrity of the watermark, the key is known as a *private key*.




Figure 2: Generic watermark decoder

Other attributes can be assigned to the watermark before it is embedded within the cover data. These consist of transparency, robustness, and information capacity (Table 1). These conflict with each other, and create a three-dimensional trade-off relationship of perceptibility versus imperceptibility, robustness versus fragility, and high payload versus low payload. For instance, an increase to the payload of a watermark will decrease the watermark's transparency and robustness.

<i>Transparency</i>	The degree to which a watermark is perceptible to the human audio/visual system.	<i>Perceptible</i>	A visual pattern such as a logo or copyright information that is visible in the cover data.
		<i>Imperceptible</i>	An invisible watermark hidden within the watermarked data, with very high similarity between the cover data and the watermarked data.
<i>Robustness</i>	The degree to which a watermark can survive various distortions or attacks. Malicious attacks aim to remove the watermark; non-malicious attacks result from common transmission and signal processing practices (e.g. compression, resizing, cropping, rotation, noise addition, quantization).	<i>Robust</i>	A watermark designed to survive heterogeneous attacks, which assures the highest levels of security, including copyright protection.
		<i>Fragile</i>	A watermark embedded with very low robustness, which is altered or destroyed by even the slightest manipulation. Useful for checking the authenticity and integrity of watermarked data.
<i>Information capacity</i>	The degree to which a hidden payload can be embedded within a watermark. The payload can comprise such information as serial numbers, copyright messages, ownership data, and transaction dates.	<i>High payload</i>	Fragile watermarks can contain high payloads.
		<i>Low payload</i>	Robust watermarks typically contain low payloads.

Table 1: Attributes of watermarks

4.1 Digital Watermarking of PACS Images

Digital watermarking techniques have been proposed lately as a valuable technique for resolving the security and privacy issues affecting PACS applications. In particular, watermarking can be used to ensure patient anonymity, to verify the authenticity of the image, to certify the image has not been modified, and to record the history of the image (Table 2). These techniques in particular offer a highly effective solution to insider threats.

<i>Anonymity</i>	Hide patient data, annotations and other confidential information within the cover image.
<i>Authenticity</i>	Verify that the cover image is authentic, for instance, that it was produced by a specific medical facility.
<i>Integrity</i>	Show that a cover image has not been tampered with or otherwise distorted subsequent to its production.
<i>History</i>	Trace the different processes to which the cover image has been subjected by different service providers.

Table 2: Potential benefits of digital watermarking of PACS images

Various research groups have looked at digital watermarking of PACS images, and three types of watermarking method have been suggested for PACS images (Table 3). These types include assigning a Region of Non-Interest (RONI) within the image, embedding a reversible watermark, and embedding a classical watermark with minimum distortion.

<i>Region of Non-Interest (RONI)</i>	Techniques that embed a watermark in a region of little interest to the medical practitioner, such as the black background or an area of solid colour within the cover image.
<i>Reversible watermarking</i>	Fragile watermarking techniques that enable removal of the watermark and exact retrieval of the cover image once the embedded content has been read.
<i>Classical watermarking</i>	Techniques that embed a watermark in different image domains. One type of technique embeds watermarks in the least significant bits of a cover image—the spatial domain. Another embeds watermarks in the transform coefficients of a cover image—the transform domain. In both cases, the intention is to minimize any distortion of the cover image.

Table 3: Types of watermarking method for PACS images

4.2 Embedding Multiple Watermarks

While individual watermarks can provide proof of image authenticity or origin, they cannot robustly preserve the *temporal* history of end-user activity in relation to a particular image. Preserving the temporal history of an image in watermarks would be useful, as it would show the identity of the last end-user to access a suspect image prior to its unauthorised distribution, and hence the identity of who might have misused the image.

The embedding of multiple watermarks within an image—with each watermark representing an end-user event—has been proposed as a secure method of preserving the temporal history of an image.

Embedding additional watermarks in a way that preserves the legibility of preceding watermarks is not trivial however. Multiple watermark embedding schemes must ensure that all watermarks are fully and permanently legible—that is, subsequent watermarks must not overwrite preceding watermarks. Moreover, such schemes must ensure that the original image is not degraded beyond a fit-for-purpose limit.

One earlier proposed scheme aimed to embed watermarks in multiple transform frequency bands within an image, with each end-user adding a unique watermark to a different frequency band. Another earlier scheme aimed to divide an image into multiple spatial domains, with each end-user inserting their watermark into a different domain. Both of these schemes therefore provide solutions for preserving the legibility of preceding watermarks and images.

Neither of these schemes however aimed to securely preserve the embedding order of each watermark. While adding a timestamp to each watermark could preserve the embedding order, timestamps require a relatively high resolution and a correspondingly high degree of watermark perceptibility. These characteristics make such watermarks more susceptible to an attacker attempting to render the timestamp illegible, make it more difficult to add multiple watermarks to an image, and make it relative easy to forged a watermark.

Moreover, neither of these schemes offers a comprehensive scheme for embedding watermarks and checking their temporal history.

5 A Novel Multiple Watermark Embedding Scheme

Researchers at Aston University in Birmingham, England have devised a reliable patent-pending method of determining the *order* in which multiple watermarks are embedded in a cover image. This capability is a major advance over existing methods.

The technology essentially uses the embedding of multiple watermarks as a log of the end-users who have accessed a PACS image. This provides an effective method of not only embedding and retrieving the log but also the order of the end-users who have accessed the PACS image.

Each time an authorised end-user—that is, an individual who has logged into the PACS system—accesses a medical image, an imperceptible robust watermark is automatically added to the image. If that person then illegitimately distributes the image, via whatever means available, the watermark will travel with that image.

The technology potentially also offers an embedded audit trail within each image, showing who has accessed the file, when, and what tasks were undertaken.

The proposed architecture comprises a Content Management System, which in turn comprises four core components (Figure 3):

- Content Database
- Watermark Database
- Content Provision System
- Watermark Analysis System

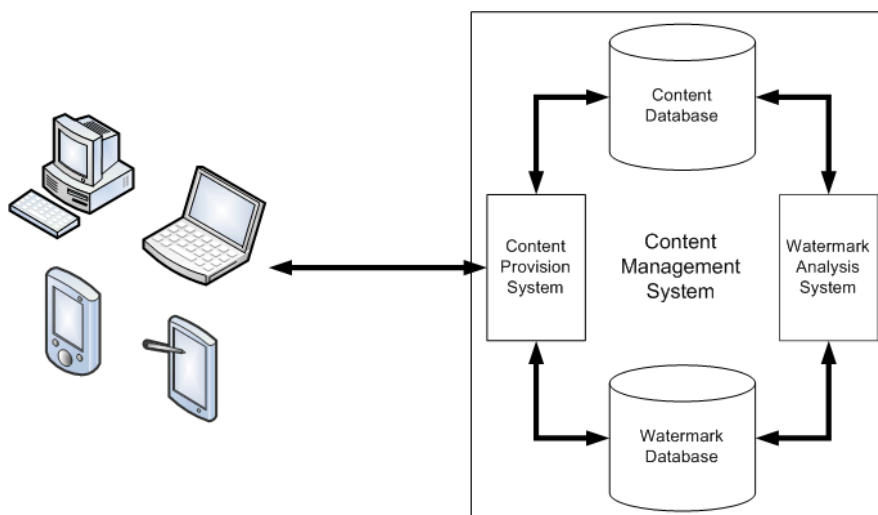


Figure 3: Content Management System Architecture

End-users will be able to connect with the Content Management System using a diverse array of computing devices.

The Content Database stores the cover images. It also contains a table of records, with each record representing a particular stored cover image.

The Watermark Database stores records that correspond with each cover image in the Content Database. Each record comprises data about the process used for adding a watermark to its corresponding cover image.

The Content Provision System receives and fulfils all end-user requests for images. The system identifies the specific end-user making the request, and then retrieves the requested image from the Content Database. It then uses a Watermarking Module to embed a unique watermark into the requested image. The newly watermarked image is then stored in the Content Database, where it replaces the original image. Only then is the watermarked image provided to the end-user who made the original request.

The Watermark Analysis System receives all suspect images that need to be analysed, such as a confidential image found circulating on the Internet. This is achieved by decoding and analysing all watermarks embedded within the suspect image, and determining the order in which the watermarks were embedded.

An in-depth technical description is available from www.astoninventions.com/WM.

6 Applying the Technology

The previous sections have discussed a critical need faced by PACS users and proposed a technical architecture that solves this need. How then can the proposed technology be applied?

When it is first installed, the Content Provision System analyses each cover image stored in a pre-existing Content Database. A preliminary ownership watermark is then (optionally) applied to each image. Cover images in the original Content Database will be overwritten with the newly watermarked images. If copies of the original images are required, to satisfy corporate policy or national legislation, a backup should be stored in a secure facility.

We envisage that these watermarks will be imperceptible, robust and carry a low payload, to maximise the security provided by the system. Alternatively, a relatively higher payload might be used—one that contains hidden confidential information—but this involves a trade-off with watermark robustness.

Each time an end-user requests access to a PACS image, the system would require them to enter their unique user ID and password, something already required by many PACS systems. This would ensure that the correct identity is embedded within each watermark. In the event that an end-user's ID and password are stolen and abused by end-users, a watermark will still be embedded, allowing appropriate counter-measures to be taken.

Ideally, the system would conceal that a watermark is embedded after each such request. We envisage the Aston technology will be completely hidden from end-users, to ensure that deliberate abusers of the data will be unaware that their activities are traceable. Alternatively, the presence of a watermarking scheme can be advertised, as a deterrent to would-be traitors. Nonetheless, if the end-user discovers the watermark, attempts to remove it will render the cover image highly distorted and unusable.

We further envisage that the Aston technology will be supplied as an integral component of an established PACS system. When such a PACS system is installed or upgraded, it is proposed that the Aston technology will be included within the package. Alternatively, an ISV might offer the Aston technology as an add-on to a variety of PACS systems.

Aston researchers will be available to offer advice on all aspects of implementing the technology.

7 Conclusion

In the wake of the Wikileaks scandal and numerous other leaks of confidential data, an imperative now exists to protect data from disgruntled or negligent employees. This paper reveals how multiple embedded digital watermarks can be used to trace traitors and to deter potential traitors, whether the activity is intentional or accidental.

The proposed system traces traitors, by providing an imperceptible watermarking scheme that preserves the temporal history of end-user access to the watermarked image. If a watermarked image is illegitimately distributed by an end-user, the watermark reveals the identity of the perpetrator.

The proposed system could also act as a deterrent to potential traitors, by providing them with the knowledge that their identity is automatically associated with unauthorised activities.

Other benefits of the technology include the potential to store an embedded audit trail within each image.

Aston University is now actively seeking commercial partners in the PACS industry to help us bring this innovative and highly useful technology to market.

About Aston University

Aston University is a long established British university known for its ground breaking research and strong links to industry. Research at Aston makes a real difference to individuals, organisations and society—we're advancing scientific and theoretical knowledge, and working to share and implement the outcomes of our activities.

Aston has a policy of actively creating, protecting, and commercializing its intellectual property. We currently manage some 50 patent families, and we're seeking licensees who can help bring these technologies to market. We also have a number of spin-out opportunities of interest to entrepreneurs and investors.

The technology described in this white paper is the subject of an international patent application (PCT/GB2010/000308).

Further information about this technology—including a technical description, a copy of the patent application and a podcast—is available from www.astoninventions.com/WM.

References used during the creation of this document are available upon request.

Aston University
Business Partnership Unit
Aston Triangle
Birmingham B4 7ET, United Kingdom
+44 (0)121 204 4242
www.astoninventions.com

Copyright © 2011 Aston University. All rights reserved. Aston Inventions, Aston University and the Aston University logo are trademarks or registered trademarks of Aston University. Other names may be trademarks of their respective owners.