



Article received
25 September 2006
Reviewed by EuDiRights
Workshop Committee
Approved for publication
10 October 2006

Complementing DRM with digital watermarking: mark, search, retrieve

Patrick Wolf, Martin Steinebach and Konstantin Diener
Fraunhofer IPSI, Darmstadt, Germany

Abstract

Purpose – The purpose of this paper is to show how digital watermarking can be applied to assist and improve cryptography-based digital rights management (DRM) systems by allowing the protection of content beyond the domain protected by the DRM system.

Design/methodology/approach – Digital watermarking is a passive technology, not allowing the active prevention of copyright violations. But it allows the irreversible linking of information with multimedia data, ensuring that an embedded watermark can be retrieved even after analogue copies. Therefore watermarking can be used where DRM fails: whenever content needs to be moved out of the protected DRM domain, e.g. when playing back content via analogue output channels it can mark the content with information that would help to identify its origin if it is used for copyright violations. The remaining challenge now is to find the marked content within the channels regularly used for copyright violations. The paper therefore introduces a concept for scanning file sharing networks for marked content.

Findings – The vast number of files present in the file sharing networks prohibits every approach based on completely scanning and analysing each file. Therefore concepts for filtered search queries where only potentially watermarked files are downloaded are discussed.

Originality/value – The paper shows how watermarking can be applied as a technology to allow active content protection beyond the limitations of current DRM systems.

Keywords Data security, Digital storage, Mark scanning equipment

Paper type Conceptual paper

Security challenges for digital media

Recent advances in multimedia document production, delivery and processing, including the wide availability of increasingly powerful devices for the production, communication, copying and processing of electronic documents, have made a large number of new opportunities available for the dissemination and consumption of multimedia content. At the same time these rapid developments have raised several important problems regarding intellectual property, digital rights management, authenticity, privacy, conditional access and security that risk impeding the diffusion of new services. Issues such as Internet piracy or electronic documents forgery may have a substantial negative impact on the commercial exploitation of the possibilities already offered by multimedia technology. Resolving these issues would provide new impetus to a more rapid introduction of products and services, benefiting both industry and consumers. Digital rights management (DRM) aims to be such a solution.

First attempts at DRM relied on cryptography, particularly since cryptography proved so useful for secure transmission of digital data through public channels. Unfortunately, when cryptography alone is applied to protect multimedia applications,



it often fails. The biggest challenge is the so-called analogue gap or analogue hole: content needs to be decrypted and is presented to the user in an analogue environment, as humans cannot perceive digital information. This enables an attacker to re-record the signal in its unprotected state. Schneier (2001) offers additional reasons for the failure of cryptography-based copy protection, including the easy distribution of tools to break the protection.

Driven by the necessities outlined above, the last few years have seen the development of new tools to tackle the problems encountered in media security applications, leading to a broader view of multimedia security, according to which security primitives and multimedia signal processing work together to ensure the secure access to the data. As it is now evident, multimedia security encompasses a wide range of diverse technological areas working together to cope with the complex problems characterising this rapidly evolving field (Lu, 2004). Enabling technologies include watermarking, data hiding, steganography, cryptography, perceptual hashing, biometrics, active fingerprinting, network security and digital forensics. Some of these technologies are described below in more detail. As an example, to meet the challenge of the analogue gap mentioned above in the case of music, both digital audio watermarking and perceptual audio hashing can be applied (Cano *et al.*, 2002a). The first approach adds information appropriate to analogue conversion to the musical piece which can be retrieved, for example, to identify the content. The latter approach stores a description of the musical piece appropriate to analogue conversion in a database which helps to identify the content at a later time.

Searching for content

The interest of the media industry in fighting illegal file sharing has inspired various approaches and business models to searching for copyright violations within file sharing networks. This area includes specialised agencies scanning the networks with a low degree of automation (for example, Partners for Management www.p4m.de/index.htm). While this approach can successfully identify copies within networks, it is time-consuming and therefore affordable only for media content where illegal copies cause huge financial losses, as in the film industry. As an alternative, Digimarc Image Search www.digimarc.com/mypicturemarc/imagesearch is an automated system currently addressing only Web pages and limited to single images to be protected. Other companies like GridPatrol www.gridpatrol.de offer a more complex service, including not only searching for content in various places like Web pages, FTP directories or file sharing networks, but also initiating legal measures when illegal content is found. Identification of illegal content here is done by logical analysis and statistical calculations.

All of these approaches share one important weakness: they cannot identify the sources of illegal copies, but only the current point of distribution; or, in the worst case, only the existence of illegal content within their search space.

This paper takes up the challenges just presented by describing how watermarking can be actively used to complement current DRM systems. First, we will provide the technical background of watermarking. It will become apparent that watermark embedding and retrieval alone will not suffice in advancing DRM systems. What is missing is an active component searching for potentially watermarked files on the Internet, which we will develop as a meta-search framework in the next section. Finally, we will take a closer look at what can be achieved by the results of such a meta-search.

Technologies for multimedia protection

Because of the massive illegal copying of media files, content owners utilise various protection technologies. In online media stores normally digital rights management (DRM) and/or digital watermarking is applied. DRM hopes to prevent the illegal copying by restricting the usage of the content. As long as a DRM system has control over the media file, it cannot be moved to a file-sharing network. But often users circumvent these mechanisms and remove the restrictions. The result is an unprotected and freely distributable file. Digital watermarking embeds individual information within the content sold by the online store. This allows tracing illegal copies back to the original buyer of the content. The customers are granted freedom to use the content in any way they want, including burning a CD or copying the audio file to an MP3 Walkman. Misuse such as illegal P2P distribution is discouraged by the embedded information about the customer, which stays in the content after format changes or even in analogue copies. One important challenge here, among others, is to find marked content after illegal usage.

Perceptual hashing

Perceptual hashing algorithms have also been called robust hash algorithms, passive fingerprinting as well as audio IDs (when designed for audio data) in the literature. The concept here is to derive a robust content-dependent description from media data to be able to identify the media data by comparing the stored and a newly calculated description. This description aims to be sensitive to modifications of the media data (e.g. lossy compression). Known concepts include the inherent periodicity of signals, the time-frequency landscape given by the frame-by-frame mel-frequency cepstral coefficients, principal component analysis, adaptive quantisation and channel decoding (Özer *et al.*, 2005). Various other approaches for deriving a robust content description have been introduced for acoustic (Cano *et al.*, 2002a; Cano *et al.*, 2002b; Allamanche *et al.*, 2001; Mihcak, 2001) or visual data (Burges *et al.*, 2002; Skrepth and Uhl, 2003; De Roover *et al.*, 2005; Mihcak, 2005; Oostveen *et al.*, 2001). Snocap www.snocap.com offers a content-recognition system based on passive fingerprinting enabling reliable filtering of illegal content in file-sharing clients or playback devices.

Passive fingerprinting, one popular synonym for perceptual hashing, is called “passive” to distinguish it from the usage of specially designed binary sequences for customer identification, also called fingerprints. The act of embedding such fingerprints into a media file via digital watermarking is called active fingerprinting (Lu, 2004).

Digital watermarking

Digital watermarks are one of the various strategies which should help to make the distribution of digital material more secure. A distinction can be made between active and passive strategies:

- active methods, such as cryptography, directly prevent unauthorised distribution of data; and
- passive methods, such as digital watermarks, serve more as a method to provide proof of ownership rights.

Watermarks have been, and continue to be, used in other physical media such as banknotes or paper. In order to use watermarks for digital media, they are adapted for the relevant types of multimedia data. Watermark solutions are thus specially developed for areas such as audio and video.

Watermark information is integrated with the data so that, unlike information merely attached to the data, it cannot simply be removed. In most cases it has the appearance of a noise pattern spread transparently across the medium. The noise pattern here needs only to be adapted, depending on the multimedia data format used. In audio data, for instance, individual frequencies are altered, and for image data colours of individual pixels are altered. Special conditions arise in the area of video, since the temporal sequence of pictures has to be taken into consideration as well. Watermarks have a variety of properties which receive various weightings depending on the way in which the watermark is used, including robustness, transparency, payload and security.

Usage of digital watermarking

In order to use digital watermarking for assisting DRM systems, two general strategies can be followed:

- (1) *Public watermarking.* The watermark is embedded by a DRM system when the protected content leaves the DRM system, e.g. when sending it to an analogue output channel. It includes information about the usage rules for the content within the DRM system. If the file re-enters the DRM domain, the watermark is retrieved and the DRM rules are re-assigned on the file.
- (2) *Private watermarking.* Here the watermark is embedded and retrieved by the owner of the protected content or a service provider. In most cases this is done at the sales stage in an online shop to include an individual customer-dependent identifier of a sold file to be able to trace the customer in case of copyright violations. There is no interaction between the customer and the watermark; the embedding process is transparent to him.

Both strategies have drawbacks and advantages. Private watermarking requires the complete computational costs of watermark embedding and retrieval to be paid by the content owner or the service provider, while for public watermarking the DRM system of the customer's computer is responsible for watermarking. For public watermarking, the embedded information becomes an active part of the DRM system. All content to enter the DRM environment is first scanned for embedded watermarks, disabling attacks to remove DRM policies via the analogue gap. The private watermarking strategy requires additional efforts to enforce the content protection it is aiming at. After embedding the customer identifier, the watermark rests passively within the content bought by the customer. It does not hinder the customer in further distributing the content, for example, by placing the files into file-sharing networks. Only if the content owner uses active measures to search for illegal copies of his content, retrieve the watermark from it and address the identified customer does the watermark offer active protection.

The most important difference between both strategies is the inability of users to check whether attacks against the embedded watermarks are successful. This is a major challenge for watermarks used in public scenarios: if the watermark is detected

by the DRM system and DRM regulations are issued to it based on the retrieved information, attackers can use the system to verify the success of attacks against the watermark. They can try an arbitrary number of operations on the protected file, move it into the DRM system and see if DRM rules are re-assigned to it. As soon as this is no longer the case, they have identified a successful attack on the watermark and can try to duplicate it for other protected files. This is especially critical, as one common watermark with respect to secret key and algorithm would need to be used for all users of a DRM system; otherwise files copied from one DRM system would not be recognised by another system, and DRM rules could not be issued for them.

The private approach does not share this weakness. Customers may try to remove the watermark, but only the reaction of the content owner would indicate success or failure of the attack. In this case the number of attacks that could be tried by attackers is minimal, as legal action would occur soon after the first copyright violations were identified due to failed attacks.

The weakness of the public approach therefore lies within the retrieval of the watermark and the feedback to the customer. A scenario where the watermark is only embedded by a DRM system and then detected by a private party would not allow attack optimisation. If this is the case, both strategies share the challenge of finding potential watermark content. In the following, we propose a strategy for automated scanning of file-sharing networks for marked content to meet this challenge.

Internet media search

Potentially marked media on the Internet can be accessed in several ways. Most of the files are published in file-sharing systems, but the Web, FTP servers or News Forums are used as publication platforms as well. The search for potentially marked media on the Internet means scanning all these different sub-spaces for the media files; and it therefore requires an Internet-wide meta-search engine (Figure 1).

This meta-search engine must provide a generic search interface to the client issuing a search request as well as simple access to all possible sub-spaces. This is best realised by an extendable framework that allows definition and registration of specialised components for individual sub-spaces. Thus the media meta-search framework covers two aspects:

- (1) *Search interface.* This part contains the definition of complex search queries and representation of generic search results.




Figure 1.
Internet meta-search engine

- (2) *Search execution architecture*. This part describes the extendable architecture for search execution on the concrete Internet sub-spaces and the platform dependent transformation of generic search queries.

Search interface

Internet sub-spaces that allow searching explicitly for digital media files like file-sharing clients or the Google image search use a metadata-based search mechanism. Unlike the classical text document search engines like Google, Lycos or Yahoo no content-based search mechanism is provided. All these media search interfaces hardly differ. Figure 2 shows the common functionality.

The common media search interface accepts a set of (textual) metadata as input (keywords, file type, media type, etc.). This search query is processed on the underlying network (file sharing) or database (Google Image Search). From every file that matches the search query a search result entry is created. This result contains metadata about the file and provides an access to the medium's essence – a file-sharing client offers the file for download and Google Image Search provides a link to the file.

As all the Internet sub-space's search interfaces based on a metadata search the meta-search, search interface will do so as well. But the search interfaces support different sets of metadata for search parameter definition and search results. Because of this, the framework will define generic query and result metadata.

Queries

The set of metadata defined for meta-search query definition is reduced to minimum of items. This ensures that all sub-spaces will be able to support the metadata-based media search. Table I contains the metadata items supported for query.

As most of the time one single criterion will not fit the programmer's needs, there must be an option to combine multiple criteria with a complex one. Therefore the framework contains a set of logical operator criteria: NOT, AND, OR, and XOR. All operators are criteria objects themselves and can be nested elements of a logical




Figure 2.
Common media search
interface

Keywords	The simplest way is to map the keywords to the media file's name. This option is supported by all Internet sub-space. Other possible mappings are to an MP3 file's ID3 tag or metadata embedded in an image
Medium type	The current design provides three types of media: image, audio, and video. The easiest way to determine the medium type is to check the media file extension. Some of the sub-spaces allow obtaining the media type more distinctly. In addition they often allow specifying the requested media type together with the keyword-based query definition
Medium file size	The framework provides the search criteria for the media file size: greater than, equal, and smaller than

Table I.
Metadata items for
meta-search query
definition

operator as well. The criteria objects grouped by logical operators create a tree structure that can be modelled using the composite pattern.

Results

The search process creates asynchronously a set of search results. In addition to the media file or media essence every search result contains a set of metadata that describe the found medium. This set of metadata differs for every Internet sub-space. Since the generic meta-search interface must provide a unique description for all search results, the framework defines an interface for all search results. Each result contains a relationship to the original criterion it matches. Thus a search result that possibly fits with more than one search request could easily be associated with the correct search criterion.

Each search result also contains a set of media file locations for requested media files. This is due to the peer-to-peer basis of some media publication platforms on the Internet. One result of the peer-to-peer architecture is a non-guaranteed availability of the network nodes or the media files stored on them. In order to make the published media files as available as possible, the files are not stored on only one single node. Nearly all published media files are stored on multiple nodes. Therefore a search query to a file-sharing platform will create a set of search results. The media file's metadata and the access to its content are not modelled directly in the search result's interface. This is coupled to a medium descriptor, which provides access to the media file and its metadata.

The framework's search process will work asynchronously, as most of the Internet sub-spaces for the media search process search queries asynchronously. This is motivated mainly by the peer-to-peer-concept on which many media publication platforms are based. The platform client sends a message containing its query to the network and triggers a search result if another node answers this query. The framework provides an observer pattern-based publish/subscribe mechanism for search result delivery. Every client interested in a search result has to implement a subscriber interface and to register at the framework's meta-search interface that is described in the following section. After registration the client receives all triggered search results from the meta-search interface.

Search execution

As mentioned earlier, our framework divides the Internet into sub-spaces. File-sharing, World Wide Web and FTP servers are three examples of sub-space offering different opportunities for media search and download. But all these spaces can be divided in different spaces themselves. File-sharing is a highly volatile field with very heterogeneous sets of publication platforms and access tools, and it is constantly changing. New tools like torrent networks appear in the space, and older tools like Napster or KaZaA are displaced. This brief look at only one of the potential spaces for media publication on the Internet shows the requirements which arise for the search execution architecture: an extendable and easily modifiable environment. A single search component is not able to fit the described requirements to access the Internet as a whole. Therefore the search execution logic must be delegated to specialised search components. The search component interface is shown in Figure 3 (*NetworkConnector*). The diagram also contains the meta-search interface – a manager class for the search





Figure 3.
Search component
interface

components (*NetworkConnectorManager*) that delegates all search requests to the registered search components.

Every component encapsulates one pre-defined sub-space and hides its specific attributes and constraints to the caller. Two of these specific attributes are the concrete sets of metadata supported for query definition and result description. Beside the interaction with the target space's search interface the component's task is to transform the generic search criterion to a target space-dependent query statement and convert the target space's search results to generic search results and vice versa. The component's whole search process is shown in Figure 4.

- (1) The search component transforms the generic meta-search query to a search query supported by the target space. All generic metadata are mapped to metadata accepted by the target space's search interface. This mapping ignores all metadata that are not accepted by the search interface. These metadata are relevant in Step 4.
- (2) The search component sends the search request to the target space's search interface. This could mean to call functions of a network client API.
- (3) The search component receives a search result synchronously or asynchronously from the target space's search interface.

Figure 4.
The component's internal
search process



- (4) The target space's results metadata are mapped by the generic medium descriptors. As most of the time the set of metadata provided in the form of search results is larger than the set accepted for query definition, the search component checks the results conformity with the generic query again. The generic search result is triggered only if this evaluation succeeds; otherwise the target space's search result is discarded.

The generic search criteria are grouped by logical operators. The result of this grouping is a tree structure modelled through the composite pattern discussed earlier. The search component has to traverse this tree structure and to transform every tree node. In order to provide various ways to traverse the tree structure the iterator pattern is used (Gamma *et al.*, 1996).

The transformation algorithm has to map the generic search criteria to the metadata supported for search query definition by the target space. That means that the node transformation component depends on the constraints of the target space. On the other hand all transformation objects of search components must have the same interface. Otherwise it would not be possible for the iterator object to transport them from one node to the other. The visitor pattern provides a way to extend tree structures with new functionality without changing the node objects (Gamma *et al.*, 1996). These visitor objects have a generic interface, and in our case they implement the transformation or mapping logic. They are also used to implement the evaluation algorithm required in Step 4 of the process described above.

Searching in multiple sub-spaces for potentially watermarked media might provide the same files (bitwise). Since checking for watermarks is a costly operation (in terms of computational cost), it is desirable to eliminate as many duplicates as possible. This can be achieved by grouping network connectors themselves in a hierarchical structure, so that a "redundancy removal connector" would receive a search query, delegate it to its child P2P network connectors, which eventually return some results. These results are then checked for redundancies, and only the really different files are returned by the redundancy removal connector.

But such a strategy is helpful only if a duplicate version of the same work really exists on the Internet. In order to verify this, we have downloaded MP3 copies of the song, "Hips Don't Lie", by Shakira, a Top 10 hit at the time, via the four file-sharing networks Gnutella, Bittorrent, eDonkey and DC++. Our goal was to see:

- (1) if we could find files which were available in more than one network that were identical hash-wise; and
- (2) if we could find files that differed hash-wise when looking at the MP3 but resulted in identical.wav files after MP3 to.wav conversion.

Point (1) motivates the usage of hash comparison before watermark retrieval when downloading and checking copies from different networks. While all four networks use hash methods to identify files, a cross-identification of files between networks is impossible as each network uses an individual hash algorithm.

Point (2) motivates the usage of hash algorithms on pure media data. As most media file formats allow adding meta-data to the multimedia content, changing this metadata also changes the file hash. With respect to watermark retrieval, the added metadata is irrelevant as the information is only embedded into the multimedia data. Therefore using a hash algorithm, taking into account only the pure multimedia data, may identify redundant files before watermark retrieval.

As hash calculation is much faster than watermark retrieval, filtering already checked copies before the detection process improves performance of the overall system. For our experiment, we downloaded 102 copies of the song. In six cases we found hash-wise identical files which were present in more than one network. After converting all MP3 files to.wav, 47 files were hash-wise identical and therefore redundant with respect to watermark retrieval. In this case, hashing the pure multimedia data would reduce by 46 per cent the number of files where a watermark would be retrieved.

Results of search

What does a search for watermark information in media on the Internet truly reveal? First of all, it gives information about when and where the file was shared. Depending on the sub-space, the IP addresses of the source may be revealed as well (not true for some P2P networks). This is general information that every investigation can reveal. More importantly, the search for watermark information reveals the information that was originally embedded into the medium. In our case this is a direct link to the original buyer of the medium (either through a unique transaction number or a customer ID). Direct means that the instant that it is able to interpret the watermark, it can use this information to match the medium to a real-life person. No other parties like Internet Service Providers or even authorities need to be involved. This is true even for “stealth networks”, where uploads can be performed completely anonymously. The link to the original buyer resides in the medium itself. As long as the medium can be accessed (i.e. downloaded or even simply listened to), the potential source of the illegal distribution can be identified.

A valid question about using this approach is: What does a copyright owner gain by being able to identify the original buyer? First of all, online stores selling digital goods download the general terms and conditions of the contract with the original buyer which usually not only forbid publishing the work but also the sharing of it with close friends, which in some countries might be otherwise legal (*CNET News*, 2004). Finding such a link in a medium does not necessarily incriminate the original buyer, but with the link investigations at least have a starting point – this is not the case for cryptography-based DRM systems. Recent cases in the film industry show that

watermarking can be successfully used as a first hint at copyright violation (*The Guardian*, 2004). Still, the media could have been stolen, and in court neither the watermark nor the IP address alone will suffice to sentence the original buyer. Additional proof will be needed. Thus the challenges described at the beginning of this paper have turned into legal problems, since the supporting technology is already there.

Conclusions

It seems that DRM systems that are based on cryptographic mechanisms alone will not be able to fully close all the holes where digital works can leak out of the systems. Digital watermarking is the most promising technology to complement cryptographic mechanisms for closing the analogue hole. By embedding a link to the original buyer into a bought work (such as a customer identifier) digital watermarking is able to identify the potential source of illegal distribution. Since digital watermarking is a passive technology, it needs to be supplemented by an active component. The key element is to find suspect works that can be checked for watermarks.

We have presented a framework that serves as a meta-search engine to find potentially watermarked works on the Internet and retrieve the watermarks. The meta-search engine is fed with abstract metadata-based search criteria and is able to search arbitrary search spaces (P2P networks, "Google space", eBay, etc.) by using space-specific network connectors. Each network connector behaves as any other client of the search space and is thus indistinguishable from regular clients, so that search spaces cannot take active measures against being searched. The search reveals – even for stealth networks – a direct link to the original buyer of the work; thus misuse is strongly discouraged.

In order to close the analogue hole in cryptographic DRM systems the workflow of digital watermarking needs an active component that finds potentially marked works. This summarises the digital watermarking workflow as mark, SEARCH and retrieve.

References

- Allamanche, H., Helmuth, F. and Kasten, C. (2001), "Content-based identification of audio material using MPEG-7 low level description", *Electronic Proceedings of the International Symposium of Music Information Retrieval*, available at: <http://ismir2001.ismir.net/papers.html>
- Burges, J., Platt, J. and Jana, S. (2002), "Extracting noise-robust features from audio data", paper presented at the International Conference on Acoustics, Speech and Signal Processing (ICASSP 2002), Orlando, FL, 13-17 May.
- Cano, P., Batlle, E., Kalker, T. and Haitisma, J. (2002a), "A review of algorithms for audio fingerprinting", paper presented at the International Workshop on Multimedia Signal Processing, US Virgin Islands, December.
- Cano, P., Gomez, E., Batlle, E., de Gomes, L. and Bonnet, M. (2002b), "Audio fingerprinting: concepts and applications", paper presented at the International Conference on Fuzzy Systems Knowledge Discovery (FSKD'02), Singapore, November.
- CNET News (2004), "File sharing legal in Canada", *CNET News*, available at: http://news.com.com/2100-1027_3-5182641.html

-
- De Roover, C., De Vleeschouwer, C., Lefebvre, F. and Macq, B. (2005), "Key-frame radial projection for robust video hashing", paper presented at the Workshop on Image Analysis for Multimedia Interactive Services (WIAMIS), Montreux.
- Gamma, E., Helm, R., Johnson, R. and Vlissides, J. (1996), *Design Patterns: Elements of Reusable Object-Oriented Software*, Addison-Wesley Longman, Boston, MA.
- (The) Guardian (2004), "Oscars piracy man fined £165,000", *The Guardian Digital Edition*, available at <http://film.guardian.co.uk/piracy/story/0,1397436,00.html>
- Lu, C-S. (2004), *Multimedia Security*, Idea Group Publishing, Hershey, PA.
- Mihcak, K. (2005), *Robust Image Hashing*, Microsoft Research, Polytechnic University, Brooklyn, NY, available at: <http://cis.poly.edu/seminars/abstract6fall05.shtml>
- Mihcak, V. (2001), "A perceptual audio hashing algorithm: a tool for robust audio identification and information hiding", *Proceedings of the 4th International Information Hiding Workshop*, Pittsburgh, PA, April 2001.
- Oostveen, J., Kalker, T. and Haitsma, J. (2001), "Visual hashing of digital video: applications and techniques", *Proceedings of the SPIE 24th Conference Applications of Digital Image Processing*, San Diego, CA.
- Özer, H., Sankur, B. and Memon, N. (2005), "Robust audio hashing for audio identification", *European Signal Processing Conference (EUSIPCO)*.
- Schneier, B. (2001), "The futility of digital copy prevention", *CRYPTO-GRAM*, 15 May.
- Skreph, C. and Uhl, A. (2003), "Robust hash functions for visual data: an experimental comparison", *Proceedings of IbPRIA 2003*.

Corresponding author

Martin Steinebach can be contacted at: martin.steinebach@ipsi.fraunhofer.de