CERIAS Tech Report 2004-84

Benchmarking of Image Watermarking Algorithms for Digital Rights Management

by Benoit Macq and Jana Dittmann and Edward Delp

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907-2086

Benchmarking of Image Watermarking Algorithms for Digital Rights Management

BENOIT MACQ, SENIOR MEMBER, IEEE, JANA DITTMANN, AND EDWARD J. DELP, FELLOW, IEEE

Invited Paper

We discuss in this paper the issues related to image watermarking benchmarking and scenarios based on digital rights management requirements. We show that improvements are needed in image quality evaluation, specially related to image geometrical deformation assessments, in risk evaluation related to specific delivery scenarios and in multidimensional criteria evaluation. Efficient benchmarking is still an open issue and we suggest the use of open-source Web-based evaluation systems for the collective progresses in this domain.

Keywords*—*Benchmarking, perception and capacity, robust watermarking, Web-based evaluation tools.

I. Introduction: Benchmarking of Robust Watermarking for DRM

Digital rights management (DRM) systems are built from several components that allow setting efficient electronic commerce of intangible goods. A DRM system has to compromise between the security threats of the content owners, the privacy of the end user and the cost of the components that will be used to establish trust between the parties. In the digital world, security and privacy are implemented through the use of cryptographic algorithms and protocols. In the case of multimedia intangibles, the digital content has to be provided at the end point in an analog form: the digital image is transformed into light through a display, the digital sound is transformed into acoustic waves. Capturing and redigitizing these analog signals for illegal redistribution is always possible. This is the reason why the authors are convinced that robust watermarking, the digital insertion of marks to individualize, trace, and control usage of a digital copy, even when it is transformed into analog signals, will be one of the pillars of future DRM systems. Facing the problem of protecting a specific version of a waveform, including its analog representation, leads to complex criteria. Simplistic arguments [1] inspired from digital cryptography are not valid any more in this case. DRM systems have to be described and evaluated as a game theory problem [2]: the payment of online content by the end user will be obtained if his/her gains are worthwhile, while his/her privacy is sufficiently preserved. The watermarking process in this respect has to be evaluated in a risk-analysis (the provider point of view) and a gain analysis (the end-user point of view) perspective: what are the risks for the provider that watermarking protection could be circumvented, what is the gain for the end user to attempt to remove watermarks from the media? In this respect, the two main considerations for watermarking evaluation are a list of key parameters that should be taken into account in benchmarking metrics and methodology for determining their relative importance for specific DRM scenarios.

The goal of the DRM scenario is to analyze the potential security weaknesses in the distribution chain and identify at each point of the chain what tools have to be implemented as countermeasure. In this paper we will address several scenarios related to image distribution.

- The digital cinema distribution: the content is exhibited in a theater room. In this case, watermarking allows tracing the room identification and time of a projection, which should be rescanned by a camera during the exhibition. In this scenario, the retrieval of the parameters of an unauthorized copy can be done using the original version of the content.
- The contribution links, suited for distribution of content and archives between studios before packaging into programs.
- The broadcast scenario, in which a specific content is broadcast to set-top decoders: two features are addressed by watermarking in this scenario, i.e., tracing

Manuscript received October 25, 2003; revised December 19, 2003.

B. Macq is with the Department of Electrical Engineering, Université Catholique de Louvain, Louvain, Belgium (e-mail: macq@tele.ucl.ac.be).

J. Dittmann is with Otto-von-Guericke University, Magdeburg, Germany (e-mail: jana.dittmann@iti.cs.uni-magdeburg.de).

E. J. Delp is with School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: ace@ecn.purdue.edu).

Digital Object Identifier 10.1109/JPROC.2004.827361

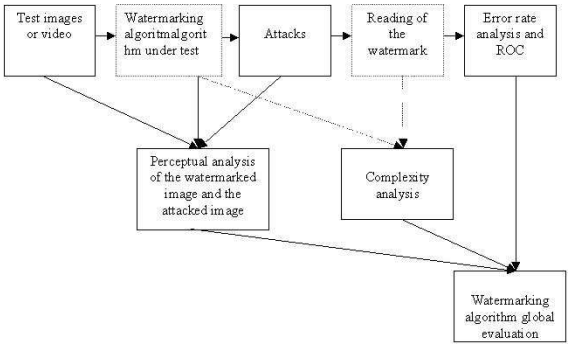

Fig. 1. A synoptic view of the benchmarking of watermarking algorithms.

of content and copy control. We include in that class of scenarios the copy control for the DVD or for PC.

- The Web publication of digital images.

We will show that each of these scenarios have specific requirements regarding the watermarking technology to be implemented. The most important properties of digital watermarking techniques are robustness, security, imperceptibility/transparency, complexity, capacity, and possibility of verification as well as invertibility [3]. The optimization of the parameters is mutually competitive and cannot be clearly done at the same time. If we want to embed a large message, we cannot require large robustness simultaneously. A reasonable compromise is always a necessity. On the other hand, if robustness to strong distortion is an issue, the message that can be reliably hidden must not be too long.

While each of these envisaged application domains has developed its own approach to evaluate watermarking, a more global approach to benchmarking has been sought for many years. We will describe its main ingredients. We propose a synoptic view of what should, therefore, be a watermarking benchmarking system in Fig. 1.

The role of benchmarking of watermarking is to provide a fair and automated evaluation of these parameters. Watermarking benchmarking research has been initiated by the pioneer work of F. Petitcolas, R. Anderson, and M. Kuhn [4] related to the StirMark system. The StirMark system introduced random bilinear geometric distortions as an innovative attack against image watermarks. The StirMark system is today a public automated benchmark evaluation service. The attacks which are included in the benchmark are cropping, JPEG, median cut, add noise, remove lines, affine transform, self-similarity, convolution, and random bilinear geometric distortion. StirMark uses Windows INI format to describe the evaluation profile. For each attack, it tests whether a message was correctly decoded or not. The StirMark benchmark has now been established as an evaluation tool for image and audio watermarking robustness. The general design concept as described in [5] and [6] (and shown in Fig. 6) is divided into three main parts: 1) the test library with the evaluation algorithms, evaluation profiles for the different requirements from the applications, and the multimedia database (the left side of the figure); 2) the benchmarking application with the marking scheme library and the quality metrics in the center of the figure; and 3) the results database with a Web server as Web interface for Web-based evaluation (the right side of the figure). The main idea is to encapsulate the test algorithms from the benchmarking to allow continuous development of new attacks independent of the actual available profiles integrated into the whole application. Furthermore, for offline testing the test library can be used as a stand-alone evaluation tool without using the Web evaluation service. The actual implementation for image benchmarking covers the test library as a stand-alone tool [7] and the Web service in the overall architecture design [8]. The audio evaluation also covers the test library as a stand-alone tool [9]. The Web service implementation for audio benchmarking is not yet available.

OPTIMARK has been developed by Solachidis *et al.* [10]. It runs the benchmarking on the Windows operating system. In OPTIMARK, a graphical user interface allows to provide interactively several statistical characteristics of the watermarking software under test, mainly receiver operating characteristic (ROC) curves. The statistical dependency of the keys and of the watermarking messages is also carefully tested. The envisaged attacks are similar to those of StirMark.

CHECKMARK [11] is very close to the European CERTIMARK system. Particularly, it uses, as CERTIMARK, a collection of intermediate results into XML files. A careful design of the attacks has been included in this system. It includes all StirMark 3.1 attacks, JPEG2000, projective transformations, warping, copy attack, template removal, denoising, nonlinear line removal, collage attack, down/up sampling, dithering, and thresholding. For each attack, it tests whether a message was correctly decoded or not.

CERTIMARK [12] is a project funded by the European Union. It includes all the features of CHECKMARK and OPTIMARK. The results of the project are unfortunately neither open source nor Web accessible.

In Section II we describe scenarios for DRM applications that would require efficient watermarking benchmarking. Section III shows current benchmarking problems: 1) perceptual quality evaluation and 2) capacity issues; and we give a short introduction to general attacks by indicating which attacks are already implemented. Section IV shows the general design of benchmarking suites with respect to management of benchmark complexity and evolution, by introducing a Web-based open-source benchmarking approach, which appears like an open working space instead of a certification authority for watermarking assessment.

II. Scenarios and Review of Specific Benchmarking Approaches

The scenarios described here give some illustrative applications in which watermarking can strongly improve the DRM system. The benchmarking of a watermarking algorithm has to be performed in the light of the attempted effect of the transmission chain, the potential attacks, and the required quality of the watermarked signal and the minimum acceptable quality of the attacked signals.

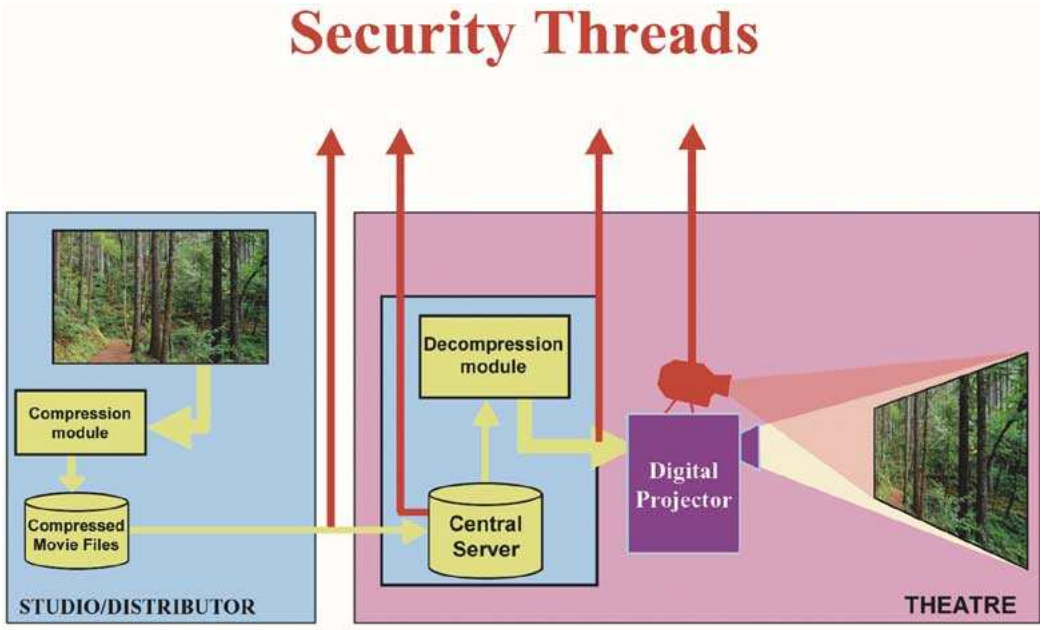

Fig. 2. Security threats in the digital cinema chain.

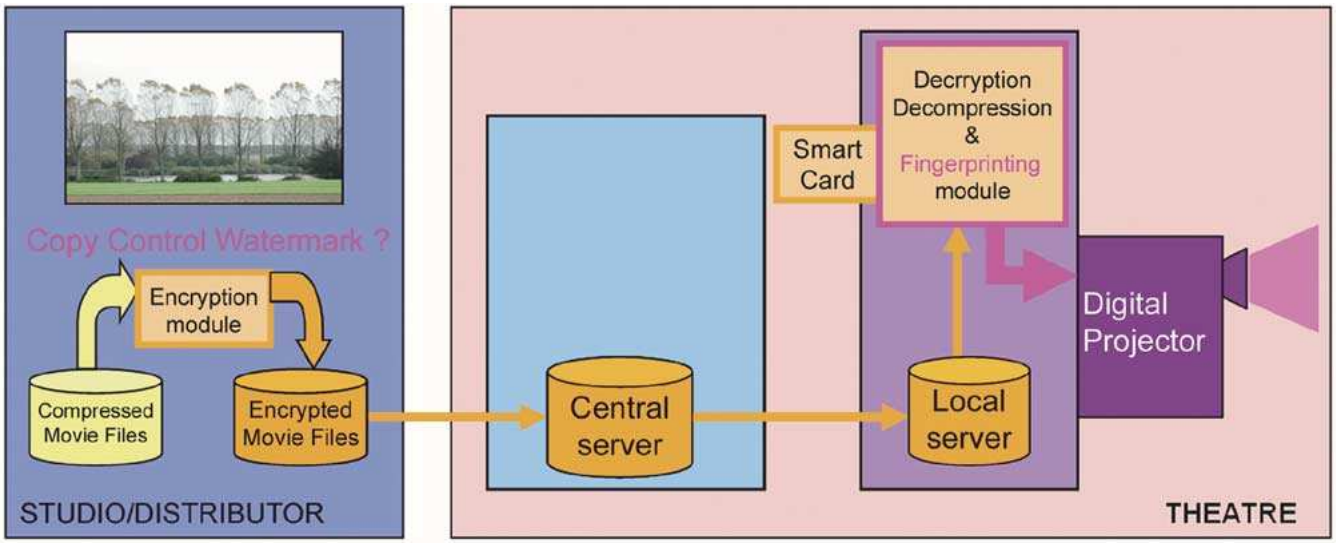

Fig. 3. Exhibition fingerprinting in digital cinema theater.

A. Digital Cinema

Digital cinema can be defined as the digital electronic distribution and display of theatrical film content or live material to the theater. However, this simple definition does not reflect the complex mesh of balanced business relationships between the different parties involved in worldwide cinematic distribution. The transition of the distribution business from analog to digital will be progressive, starting with only content being digitally distributed and going until the business is all digital. Both of these facts impose strong requirements on the DRM system for digital cinema.

Movie theaters receive content (movies) from national distributors. They store them and project the movie in one or more theaters under some contract conditions.

Piracy happens at two levels. The first one is obvious and consists in direct bit-to-bit copies done in the storage device. This kind of piracy can be totally solved by proper uses of conditional access systems, including encryption and reliable key distribution. The second one is also the responsibility of the movie theater owners. It consists in letting a spectator filming the projected movie with a handicam at the back of the theater. Security threats in the digital cinema chain are depicted in Fig. 2.

In this scenario, the most relevant watermarks are forensic tracking watermarks that we call "exhibition fingerprints." As a matter of fact the amount of high-value cinema content is generally not very high and one should directly be able to recognize a specific content (if one finds an illegal copy of a movie, the issue is not to identify the movie itself but the origin of the forgery). In the comparison to audio, the number of works is much higher, but solutions exist to automatically retrieve original works (see [13]). Exhibition fingerprints are applied during each exhibition (see Fig. 3). Exhibition finger-

prints identify the circumstances of the exhibition. The fingerprint should include identification of the content as well as the exhibition. The fingerprint should be upgradable. To be effective, the means of application of the fingerprint should be resistant to attempts to disable it. This may require placing the implementation within a secure perimeter.

This fingerprint should be resistant to the "handicam copy." This means severe image distortions, such as scaling, cropping, affine transforms, but also nonlinear geometrical transform due to optics.

Generally the exhibition fingerprint will not be directly attacked by people distributing the illegal handicam copy, since the only one responsible for the forgery is the exhibition theater room owner. Moreover, one could consider that the original movie material (or an excellent copy of the original) is available for the reading of the exhibition fingerprint. This watermarking scenario is, therefore, characterized by very high quality requirements for the watermarked signal (total invisibility of the watermark) and a very difficult transform (digital cinema to handicam shot). The worst case is the redistribution through DivX format. That means global projective geometrical transform, digital-to-analog-to-digital (DAD) conversion, global illumination change, and high compression.

Specific malevolent attacks could include any of the classical StirMark attacks. Collusion attacks are less practical. However, the reading process benefits from more favorable conditions for the reading process (no collusion attacks and availability of the original copy for decoding the watermark).

B. Contribution Links

The contribution links are the liaisons between content providers [like news operators or broadcasting unions (BUs)—the European Broadcasting Union (EBU) is one among them] and studios. The providers are multicasting content, which are remastered at each studio to be redistributed in secondary links. Generally, the studios are fair players and their main DRM concern is to identify the copyright owner (CO) of content when it has gone through several postproduction processes. In this case, a watermark with a payload containing the content owner's identity is a good solution to prove to a legal authority the ownership of a work.

It is a good means to solve conflicts because it is very robust and not easily removable and because the inscription of the watermark is made with the use of a secret key. Only the owner of this secret key can read or detect the watermark; at the same time, he is the only person having been able to produce that watermark. During an appearance before a court, the key owner will be the single person capable to produce the watermark in question; see such early papers as [14]).

Fig. 4 illustrates one possible generic scheme for a BU network's use, which is proposed by the EBU. A sequence is originated outside the BU network.

This sequence's originator watermarks the sequence W1. The sequence is securely delivered (scrambled) to BU members through the BU network. The sequence is then securely dispatched to the BU members. These two steps are probably

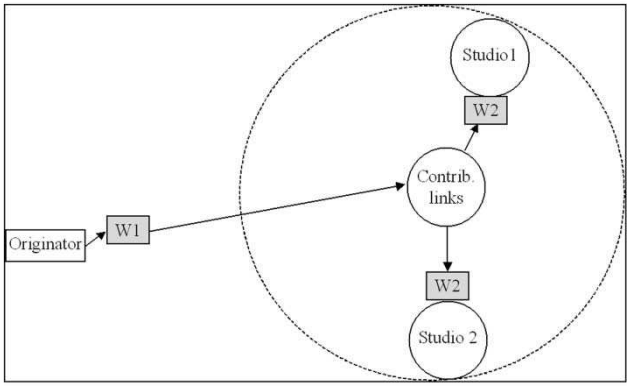

Fig. 4. Watermarking in a BU network—generic scheme [14].

achieved simultaneously. When one of the EBU members accesses the sequence, it is automatically watermarked (W2), fingerprinting the receiving service.

W1 implies that the BU commits to have some responsibility in the sequence use. Irregular sequences with only W1 imply that they have been pirated inside the BU network.

W2 implies the BU member responsibility in the sequence use. Irregular sequences with (W1 and) W2 imply that these sequences have been hacked after reception (and descrambling) by an identified BU member.

W2 implies that the sequence have been scrambled when broadcast to the BU members.

W2 must be applied at the reception point, just after descrambling and decoding.

In this scenario, the perceptual quality of the watermarked images must stay very high. Resistance against postprocessing in studios, like cropping, temporal cuts, zooming, and compression/decompression processes have to be very high. Resistance against malevolent attacks (collusion and intricate image processing) is less important. The EBU has been pioneering a benchmarking process dedicated to this scenario, which is described in [15]. The specific benchmarking approach in this case is summarized in Table 1.

C. Broadcasting of Images

The DRM cases, which are related to this theme, take place in the context of wide distribution and are mainly turned toward the tracking piracy over a media of distribution—mainly concerned with the defense of property rights, but also to copy control, which is probably the most difficult and controversial scenario to implement through the use of watermarking.

1) Tracking Piracy Over a Media of Distribution: Content providers over broadcast channels are wary of any breach of contract, either because content is shown more often or at other channels than has been agreed upon (e.g., news clips by Reuters) or less than has been agreed upon (commercial verification). In many cases, the channels over which such content is distributed to some degree are uncontrollable. The last resort for verifying the proper showing of content is by verifying in the field, that is, by having monitoring stations in every major region where verification is required. In order to reduce the complexity

Table 1
Robustness Requirements for Television Broadcast Monitoring

No	Attack	Equipment	Description
1	Digibeta	Sony AVW-A500P	Source format
2	Aspect ratio	DVE:	16:9 to Letterbox on 4:3
3	conversion	CO6S/514 BBC / Radamec	4:3 to centre of 16:9 with side curtains
4	Shift	DVE	4 pixels right
5			8 pixels left and 4 lines up
6			64 pixels right and 26 lines down
7			200 pixels left and 150 lines down
8			360 pixels right and 288 lines up
9	Scaling	DVE	200% (Zoom in)
10			50% (Zoom out)
11	Bend/shear/ rotate un-noticeable amount	DVE	less than 2° rotate
12	Bend/shear/ rotate noticeable amount	DVE	less than 10° rotate and shear
13	Re-sampling	D-A + A-D conversion	via RGB; also increase in brightness
14	White Noise	TriMedia	Gaussian noise -30 dB, via PAL
15	Sony IMX	c/o Sony Broadcast	MPEG 2 4:2:2 50Mb/s recorder
16	Sony DV	Sony DV	Low-end production (4:2:0)
17	Pal and Beta SP	Beta SP	Analogue input
18	Panasonic DV or JVC Digital-S	Panasonic DVC Pro	Low-end production (4:1:1, 25 Mb/s)
19	Beta SX	Beta SX	
20	Avid	AVR 77	Desktop editor - approx. 3:1 compression
21	Slow Motion	Digibeta	1/3 rate (0.34)

and security issues of such a monitoring station (having all the originals at a monitoring station is a bad idea, as well as high-bandwidth pipes to central video servers with original content), content is being stamped with an invisible marker that cannot easily be retrieved from the content after distribution. In other words, a robust watermark. In a typical broadcast verification scenario, the watermark carries an index into a large database where the associated broadcasting rights and permissions are stored. On a more abstract level, this application is very similar to the people metering application, where consumer TVs have been replaced by professional monitoring stations.

In both cases, the following requirements are encountered:

- robustness to nonintentional attacks related to usual manipulations: MPEG compression, transcoding, analog to digital and digital-to-analog conversions, standard conversions (PAL-NTSC), change of geometry;
- high probability of detection and high probability of correct extraction when the watermark is present, low false detection probability when not present;
- real-time extraction for reasonable complexity both for embedding and detection;
- blind extraction;
- invisibility (studio level);
- granularity less than 1 s;
- payload between 64 and 72 b.

For tracking piracy (i.e., illegal redistribution of content) one must add the following robustness requirements on embedded watermarks:

- robustness to overwatermarking up to three other watermarks (from pirates or distributors);
- robustness to any intentional attack aims at breaking the synchronization of the watermark (making it undetectable) or removing it.

For example, in the CERTIMARK project, lists of robustness requirements are summarized separately for different attacks.

2) Copy Protection: The primary purpose of a copy-control watermark is to prevent content that has lived in an uncontrollable environment from reentering the compliant world of devices compliant with the copy-control procedures. For that purpose a copy bit is needed which is unremovably tied to the content. This copy bit is to be implemented as a robust watermark. The set of requirements for this copy control watermark are quite severe.

For the DVD case, the payload of the watermark is a byte, where two bits are used to indicate copy-free (CF), copy-never (CN), copy-once (CO), and copy-no-more (CM). The copy protection system implemented by the watermark must allow changing the state from CO to CM (for example, by remarking).

The watermark may not perceptually degrade the content, i.e., the watermark has to be below the visibility threshold (VT) with very severe dual-stimuli tests (expert viewers, freeze-frames modes, etc).

The watermark must have an extremely low false positive rate of less than 10^{-12} per basic detection. The granularity of the watermark detection is 10 s, with a reliable payload of 8 b.

The watermark must be robust to all common processing, including MPEG compression down to 2 Mb/s, DAD conversions, standard conversions, and zooming of the images. The watermark detection procedure must be simple enough to detect real time both in baseband and in MPEG bitstreams and within cheap consumer devices (for example, in a DVD drive in a PC) without unnessesarily burdening the total costs of those devices.

The watermarking technology must resist against easily available hacking tools, such as frame deletion/duplication, cropping, and grayscale conversions. Watermark detection must be performed without the original (blind detection).

One could imagine that the DVD copy protection system could be directly expanded to a set-top box containing local storing devices.

For musical content, such an approach was attempted by the Secure Digital Music Initiative (SDMI) without being successful, the removal of the copy control watermarks being too easy for usual hackers.

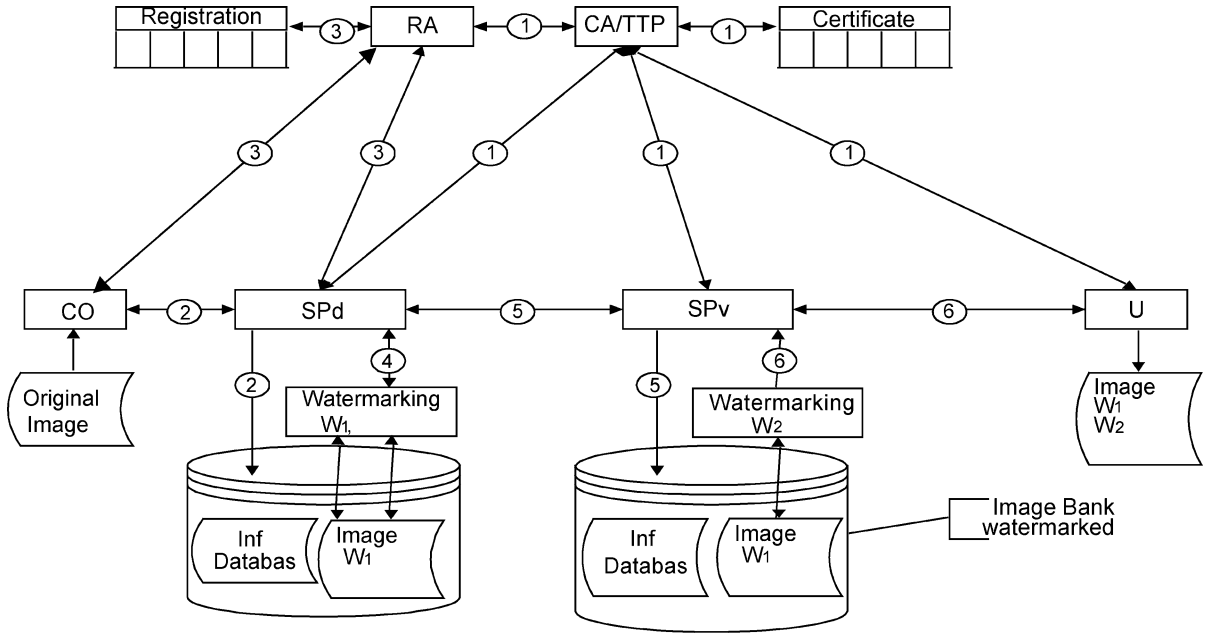

Fig. 5. Creations' secure trading over the Internet: common functional model for the use of watermarking.

A trusted PC could also use watermarking for copy control. Copy control implies strict compliance between the media and the targeted hardware, which is subject to some privacy concerns (Trusted PC are discussed by R. Anderson in [16]).

D. Internet Distribution of Images

The alternative to copy control and trusted computers relies on the responsibility of the content user and tools to mark this responsibility. Legal actions against copyright infringements on the Web have already decreased the amount of peer-to-peer redistribution of content. Therefore, means for increasing user awareness in his responsibility is precisely watermarking, combined with registration authorities and transaction certification. Such an approach is described in Fig. 5.

Functions are classical and have already been defined in the current literature. Three function classes can be identified: management functions, production functions, and diffusion functions.

Management functions correspond to the upper layer of the figure. The certification authorities (CA) and the trusted third parties (TTP) are responsible for the actors' certificates management (registration, revalidation, and blacklisting). It corresponds to the first process. The registration authorities are responsible for the intellectual property rights (IPR) registration or the legal deposit. Production functions correspond to the path toward the final, consumable, product. On one side, the creator (the CO) will generate original creations. On the other side, service producers (SPd's) will give some shape to this creation. By the way, he can generate a composite work composed with shaped creations and original ones.

Diffusion functions correspond to the path toward the end-consumer. When entering into this phase, creations should not be modified anymore. In the figure above, it corresponds to the service provider (SPv) and the user (U).

First of all, an initiation stage is necessary. All actors are registered to a CA or TTP. They obtain certificates including their status (through a set of rights), a validity period, and a public key. Ideally, the corresponding private key is distributed on a smart card to protect its access. It corresponds to the first process.

For the first implementation of this business scenario, we only considered a linear distribution, from the creator to the user.

- Once an original creation is achieved (process 2), a registration or a legal deposit is mandatory for IPR protection. It can be done by the CO or by the SPd. It is best to do it as soon as possible.
- In return, the creator receives a unique identifier for its work (process 3).
- To definitively guarantee its ownership, a secure and secret binding of the creation with the unique identifier must be realized. Applying a first watermark (called W1 at this stage) as soon as it is in a digital format does this. This is process 4.
- At this stage, the work may enter the distribution phase. Therefore, it is transmitted to a reseller, also called SPv (process 5). IPR management between those CO, SPd, and SPv has to be done offline. Usually, it consists of business contracts established prior to the transactions. The SPv will propose thumbnails of the creations on a Web site.

U browses the thumbnails freely. He can express his interest by clicking one of those. At this stage, secure communications are necessary because we enter into a business process and, thus, into the sixth process. Both parties are authenticated through an exchange of certificates

and a form is proposed to U. This form will be the basis for a trading contract. Its content depends on SPv. A minimal set of information is the identity, the localization, and the purpose of U. This form is encrypted and sent to SPv. SPv can generate a contract, sign it electronically, and propose it to U. If U accepts the terms of the contract, he signs it electronically too and send it back to SPv and waits for the creation delivery. SPv will probably sell the creation to other users. In order to track the use of it, he must be able to differentiate those. Therefore, a watermark is applied before the delivery. This secret watermark (called W2 at this stage) establishes a link between the creation and the trading contract.

Such watermarks have very high constraints in term of perceptual quality of the watermarked content and have to face very difficult attacks. A big advantage in this case, compared to the copy control approach, is that the hacker will be never sure that the watermark has been removed, since the effect of the watermark is only accessible to the authorities and COs.

Also, in the CERTIMARK project, the main robustness requirements are shown against common processing and attacks that the watermarking algorithm has to satisfy, by identifying parameter ranges for the different manipulations and the importance of robustness for each manipulation.

III. In Search of an Automated Benchmarking Suite

Several projects have attempted to provide fair benchmarking suites for watermarking evaluation. However, these attempts are far from a definitive and global solution. The main difficulties rely in perceptual evaluation, coding capacity modeling, and basic classification of attacks. We review some potential paths for these three challenges.

A. Automated Perceptual Evaluation

The perceptual evaluation has two purposes.

1) The first question is whether or not the watermarked data is perceptually different from the host data. The referenced host data can be data at study-quality level (for professional, contribution, and high-end applications) or data at consumer-quality level (typically JPEG or MPEG-2 compressed quality).
2) The second question is to determine the perceptual degradations resulting from attacks from which the original message can no longer be retrieved (i.e., what is the remaining perceptual value of an attacked image, for which it is probable that the watermark is removed). The perceptual model for attack effect evaluation should deal not only with filtering and additive noise but also geometrical transforms. The perceptual evaluation of additive noise has been widely studied, particularly in the field of image compression. Therefore, a key question is to determine a method for geometrical transform characterization from a perceptual point of view.

1) Subjective Quality Assessment: There are two ways of measuring the perceptual quality of the data, namely, subjectively and objectively. The subjective method involves a panel of viewers to whom host, watermarked, and received data are presented for grading. This panel is asked to rate the visual quality of the data, using procedures such as the single/double-stimulus methods and the two alternatives forced choice. ITU Recommendation 500 gives recommendations for standardized subjective image quality evaluation procedures. Ratings are typically done on a five-point (continuous) scale such as in the ITU-R BT.500 double stimulus continuous quality scoring (DSCQS). From the set of measurements, the mean and standard deviation are calculated, which then serve as the final numerical result for the subjective quality metric. Since subjective quality assessment methods involve human intervention (i.e., test panels), they are obviously not suitable for automatic benchmarking. One could, however, imagine an online Web-based system in which a panel of online experts would evaluate the data.

A particular quality level is the so-called VT, which is the strength of a watermark above which the watermarked material is determined as (even slightly) impaired compared to a nonwatermarked one. This threshold is also measured with a double stimulus method with a forced choice. When the choice between watermarked images and nonwatermarked ones is completely random, the watermark is below the VT. The VT is generally arbitrarily set when the forced choice gives 75% of well-classified images (watermarked and nonwatermarked ones). The VT is also very useful for dealing with attacks. Attacks above the VT introduce artifacts, which may be not acceptable.

For an automatic benchmarking of the visual quality of the data, we have to rely on numerical measurements. In the next section, we will discuss different metrics that qualify for inclusion in the benchmark. It is, however, noted beforehand that there is no single universally accepted objective visual quality measure that correlates well with the outcomes of subjective quality measurements for all applications. Therefore, we have to determine as how well the selected metric correlates with subjective quality assessment and to explore possibilities for watermarking-specific objective quality measurements. Subjective ratings of watermarked and received data in a common test set will be used as input for these research issues. There are, however, some hopes to directly use perceptual quality criteria in the case of grayscale images. This is discussed in the next section. In this respect, one of the key challenges in watermarking is the evaluation of the perceptual degradations of a watermarked image attacked by geometrical transforms.

2) Objective Visual Quality Metric: The PSNR Approach: The most well known and widely used quality measure is the global mean-square error. By normalizing this metric on the signal's variance and taking the 10-logarithm of this ratio, the signal-to-noise ratio (SNR) metric is obtained. If the normalizing takes place on the signal's (squared) peak value, the peak-SNR (PSNR) metric is obtained.

Although it is known that this criterion may not correlate too well with the subjective ratings, PSNR gives significant

indication of fixed image quality in case of filtering or in case of additive random noise (or spread watermark). Since the difference between host and watermarked data will be small in general, we expect that reasonable correlations between these initial ratings and the subjective quality of the watermarked data will be obtained in most cases.

3) Advanced Objective Visual Quality Metric Issues:

a) Weighting and Masking: The WMSNR: The PSNR metric is not an acceptable quality metric in a final benchmarking system. Advanced objective visual quality criteria have to take into account the fundamental aspects of the visual perception of moving pictures. Human visual quality models account for the multiresolution structure of the early stages of human vision, sensitivity to contrast, visual masking, color perception, and interactions between spatial and temporal perception (see [17]).

The models of spatio-temporal vision are then parameterized by psychophysical experiments on human subjects so as to obtain estimation of the human spatio-temporal sensitivity to contrast. The experiments have been carried out with synthetic signals modeling coding noise. The resulting spatio-temporal contrast sensitivity function especially characterizes sensitivity to video additive coding noise.

A hierarchy of models has been described in the literature, each one corresponding to a finer modeling of human vision. This ranges from a simple multichannel model for video that combines essential features of visual perception to a model that accounts for normalization of the cortex receptive field responses and interchannel masking.

Three main effects can lead to a more refined objective measurement than the simple PSNR. These effects are quite well determined for grayscale pictures.

1) Weber–Fechner's law states that if the luminance of a test stimulus is just noticeable from the surrounding luminance, then the ratio of the luminance difference to the surrounding luminance is approximately constant. Thus, the VT of a noise is larger for bright areas than for dark ones.
2) Contrast sensitivity functions (CSFs) express contrast sensitivity, i.e., the reciprocal of the just visible contrast, by an expression which depends on parameters such as mean luminance, spatial frequency, and orientation. This CSF dependency can be modeled as a bandpass filter, and the SNR should be weighted to take this effect into account.
3) Masking is generally defined as any interference between two or more visual signals or stimuli that results in an increase or, more often, a decrease of their visibility. Masking is generally modeled as a decrease of the CSF depending either on spatial activity or in an oriented spatial frequencies domain. The SNR should, therefore, be decreased to take into account the activity in the original image.

Those phenomena allow improving the fidelity of a quality criterion by introducing a weighted masked (WM) SNR. This approach is valid in case of additive random noise and is widely used by watermarking algorithms to spread more energy in masked or less visible areas of the images. It is, however, not valid to use it as a fidelity criterion because any slight desynchronization or small geometrical deformation should increase largely any weighted or masked SNR while the so-deformed image is perceptually very close to the original one. Before measuring the WM-PSNR, image resynchronization is first mandatory.

Alternatively simpler analytical approaches like the one proposed in [18] offers a very simple mean to measure the structure of the noise and gives surprisingly good results. A Matlab code of this approach is available freely on the Web.

b) Quality of Image Distorted by a Geometrical Transform: Geometrical transformation attacks provide a unique challenge to a watermark algorithm developer. On one hand, geometrical attack are generally difficult to model and, thus, difficult to anticipate (in opposite to, for example, a lossy compression attack). On the other hand, a geometrical transformation attack usually only affects the synchronization between the embedded watermark and the watermark detector. The watermark itself (or a major part thereof) is usually still present in the data. The overall problem is by using an exhaustive search space that watermarks may always be recovered, but the false positive rate increases too. We attempted to assess quality of geometrically deformed images in [19]. A geometrical transformation attack can take many forms, from relatively simple to complex. One of the simplest forms of geometrical transformation attacks is the rotation, scaling, and translation (RTS) transformation. In this transformation, shapes and angles are preserved. This transformation has four degrees of freedom and can be described using the following formula:

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = s \begin{pmatrix} \cos(r) & -\sin(r) \\ \sin(r) & \cos(r) \end{pmatrix} \times \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} t_1 \\ t_2 \end{pmatrix}. \quad (1)$$

We can assume that a global uniform affine RST transform does not degrade the visual quality of an image.

Examples of more complex transformations include the bilinear and curved transformation, shown in (2) and (3), respectively. It should also be noted that a transformation might be a combination of two or more transformations

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \times \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} e \\ f \end{pmatrix} xy + \begin{pmatrix} g \\ h \end{pmatrix} \quad (2)$$

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} \{(1-\beta)a + \beta b\} \sin(\alpha\pi) \\ \{(1-\beta)c + \beta d\} \sin(\alpha\pi) \end{pmatrix}. \quad (3)$$

The previous examples shows geometrical transformation applied to the entire image area. In other words, the geometrical transformation of the entire image can be described using a single mathematical expression and a single set of parameters. A geometrical transformation attack can also operate locally, i.e., in several small locations of the image. In this case, the mathematical expressions are different for each particular location. The perceptual effect of this kind of geometrical transformation attacks are more difficult to model, even when the individual local transformation is relatively simple (e.g., an RTS transformation).

For real applications, geometrical transformation attacks can happen either due to physical manipulation of the

watermarked material or to the digital manipulation of the watermarked image. Examples of the physical manipulations include the printing and scanning attack and the digital cinema scenario (the handicam attack). In these examples, the geometrical transformations that the watermarked material underwent are usually a by-product of the attack and are, therefore, usually limited in the sense of variation of transformation and transform parameters. The transformations that are usually involved are complex global geometrical transformations. Examples of the digital manipulations of the watermarked image include the common procedure of resizing or cropping the image and the random bending attack (RBA) incorporated, for example, in the StirMark benchmark. Unlike the previous manipulation, digital manipulation can involve many exotic transformations and is very difficult to predict. The resizing or cropping operation can be classified as simple global transformation while the RBA attack can be classified as a complex global transformation involving many local transformations.

The global geometrical transform of the image can be approximated as local RST transforms applied on blocks of the image: the image is segmented into blocks of variable size, until the evaluated RST on each of this block gives a result close enough to the transformed image. One way to proceed, in a similar way to [18], is to measure the perceptual degradation of the attack as a function of the locality of the equivalent affine transform and the variance of the local affine transform parameters. The locality of the affine transform is defined as the largest image block size in which the distortion can still be approximated using affine transform that yields a residual error below a certain threshold. The smaller the necessary block size, the more complex the attack transforms. The variance of the affine transform parameters of the blocks also determines the perceptual degradation of the transform. A large parameter variance means large degradations, and vice versa. The procedure of the approach described in [19] is as follows. The measurement is performed in a top-down fashion using a quad tree decomposition of the image. First, a global affine parameter estimation is performed to the whole image to find the parameter set with the smallest residual error. This residual error is compared to a certain predetermined threshold. If the residual error is still above the threshold, a quad tree decomposition of the image is performed and the parameter estimation for each block is repeated. The residual error of each block is compared. Any block with residual error above the threshold is further quad tree decomposed and processed. The blocks with residual error below the threshold is not processed any more. This whole process is repeated until a minimum block size is reached or until all blocks have a residual error above the threshold. The final output of the aforementioned procedure is the segmentation of the image into blocks, each block having its own set of local affine transform parameters.

The perceptual degradation of the geometrical transform is now determined by looking at the average block size of the final segmentation result and the variance of the local affine transform parameters contained in the parameter sets.

After geometrical transform compensation, the quality assessment can be complemented with the WMSNR.

B. Automated Coding Capacity Evaluation

In watermarking optimization, there is a tradeoff between the visual quality of the watermarked signal, the robustness of the watermark and the capacity, i.e., the maximum payload of the system. Each algorithm claimed to offer a certain payload as an intrinsic characteristic, choosing its size *ad hoc* for empirically accomplishing certain robustness or application requirements (typically, watermarks of 64 b are admitted as sufficient to give pointers to copyright or fingerprint databases).

Most of the studies related to watermarking capacity determine the bit-error rate for several attacks, for several payload lengths and for several visibility levels of the watermark. Conversely some studies present the amount of watermarked data that is required to retrieve the payload with a bit-error rate below a given threshold. In most cases, it is sufficient to assume that a given application requires a certain payload size; a random payload with this size is generated and the probability of decoding error is computed when varying the intensity of an attack when the quality of the watermarking medium is fixed.

C. Capacity

In practice, a benchmark should be able to evaluate the capacity of an algorithm presented to it as a black box. Fixing a given host signal, embedding strength, and attack, we may try to estimate capacity empirically using its definition; for this purpose, we should fix a sufficiently low probability of error for which we must determine the maximum payload yielding lower probabilities of error than the threshold. This amount would be in principle a fair approximation to capacity in the context where the test takes place.

An algorithm similar to the payload checker described by Lagendijk and Setyawan in [21] or to that described by Solachidis *et al.* in [22] can be applied: if a payload variation is allowed by the algorithm, it is used as a "loop control," i.e., the payload size is increased step by step and the probability of error measured until the stopping condition is verified. Notice that the approach is valid irrespective of what happens inside the black box, e.g., it does not matter if it uses side information or not.

Once the watermark is decoded, it is important to determine the reliability of the received message. As most of the watermarking are based on spread-spectrum, the decoding is based on a correlation detector. A threshold is fixed for the message detection. The lower the threshold, the higher the probability of message detection. Lowering the threshold may also create some nonexistent messages to be decoded from the noise or from the host image. The ROC curve is the mean chosen to determine this tradeoff between the hit rate (true positives) and the false alarm rates (false positives) [23].

A major difficulty relies on the probability estimation for very low error rates [generally, the targeted false alarm

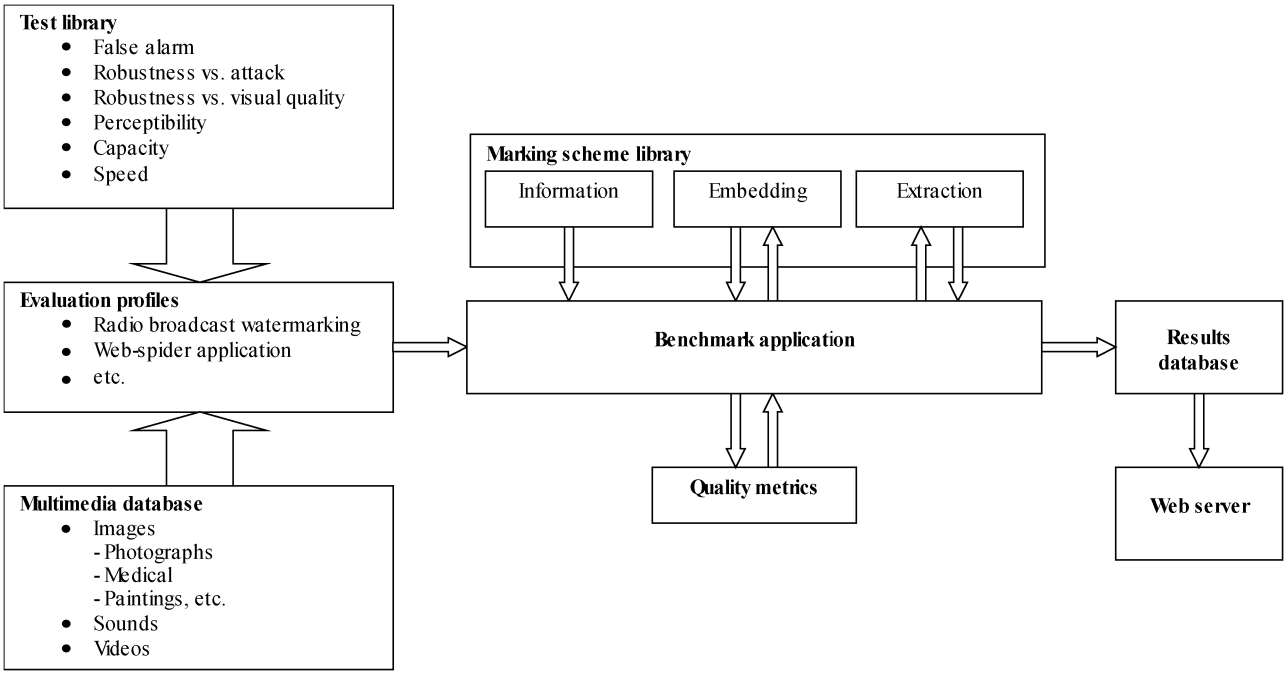

Fig. 6. Dataflow for the watermarking evaluation service from [10].

detection for copy control watermarks should be very low (below 10^{-10}), and, therefore, very difficult to estimate experimentally].

D. Attacks

A wide range of attacks has been described in the literature. Successful attacks on images protected by watermarks can be divided into four large categories. [24] distinguishes between removal attacks, geometrical attacks, cryptographic attacks, and protocol attacks. If somebody tries to remove the watermark from the data, this is called a removal attack. The means employed most frequently are filter models taken from statistical signal theory. Denoising the marked image through median or highpass filtering as well as nonlinear truncation or spatial watermark prediction are methods considered very likely to succeed. Contrary to this, geometrical attacks are not aimed at removing the watermark, but try to either destroy it or disable its detection. Both the removal and the geometrical attack are mostly aimed at the robustness of the watermark. Cryptographic attacks cover, for example, direct attacks to find the secret key or attacks called collusion attacks. The attacks in the last group, the protocol attacks, neither aim at destroying the embedded information nor at disabling the detection of the embedded information (deactivation of the watermark). Rather, they take advantage of semantic deficits of the watermark's implementation. Consequently, a watermark must not be invertible or to be copied. A copy attack, for example, would aim at copying a watermark from one image into another without knowledge of the secret key.

The actual benchmarking suites as summarized in Section I consist already of a variety of attacks and allow combinations of removal and geometrical attacks. The actual main challenges in the benchmarking design are to find the appropriate attack strength and relevant attack combinations depending on the image characteristics, the actual watermarking algorithm and the application scenario. In the field of cryptographic and protocol attacks, the benchmarking suits offer fewer approaches and are widely neglected. Most cryptographic attacks are based on an exhaustive key space search or collusion attacks, while protocol attacks are not considered in most benchmarking systems.

IV. Benchmarking Suites

We start by giving the design philosophy of watermark benchmarking, as we have to meet special requirements: management of benchmark complexity and evolution.

As introduced in [6] for StirMark, a generic watermarking evaluation and benchmarking services should be implemented with different reusable modules and integrated in a general dataflow model as shown in Fig. 6. As explained in [10], the marking scheme is provided by the user as a library of functions (binary). This library exports in particular an information function, which is used to select which evaluation profile (simulating the scenarios) has to be used. The evaluation profile is composed of a list of tests or attacks to be applied and a list of multimedia objects required for the test and sorted by types and categories. All tests results are uploaded to an SQL server connected to a Web server.

The benchmark we have to implement is a quite complex tool: several parties will provide parts of the benchmark. Many tests need to be performed within a benchmarking session; this is the other side of benchmark complexity. Since the main idea is to simulate a complete watermarking chain, we can assume data will pass through this chain more than

once, with various parameters. The benchmark can, therefore, be seen as a pipeline. Nevertheless, as parts of this pipeline are not our concern (the watermarking algorithms we want to test), we choose to make it modular so that its pieces can be easily exchanged.

Then, modularity also rules our general approach: every basic operation that is part of the classic watermarking chain will be separated from the others into one isolated module. These modules will have to comply with a general detailed interface. This philosophy induces strong advantages: modules can be exchanged easily; they can be developed separately (given the interface) and upgraded when needed. Although this flexibility allows a good understanding of every separate part of the benchmark, one has to interface them correctly together with some sort of consistency all along the benchmarking process. More generally, one has to be confident in the results of the benchmarking session. This implies that every module has to conform to certain rules (e.g., with regard to file management, error handling and reporting, etc.) for interoperability, and that the benchmark framework needs to guarantee control and integrity of the benchmark as a whole. As stated, there is a need for evolution (new attacks will appear, progressive support of new media types, etc.) that implies the ability to replace any particular part of the benchmark. But the core benchmark process should not be modified when one module changes. We will, therefore, tie the modules together with a script that launches the execution of the modules in a consistent way. This script will allow benchmark global tuning from application-dependent benchmarking to future evolution of watermarking algorithms.

We would like the benchmark to be as flexible as possible, thus allowing common data management. That is, we need to use data types for intermodule communication that are standardised and easily extensible. Give the previous remarks; we focus on the resulting basic needs.

Intermediate files: Many intermediate multimedia files will be created and modified during a benchmarking session. They will need to be managed carefully: given an intermediate file, we must be able to place it in its correct context and at the correct pipeline stage. This control will be fully automated.

Error handling: We have to carefully design the error handling part of the benchmark, as we will have to handle errors that will not come from the benchmarking platform itself. For example, a watermarking can suffer exceptions (insufficient memory, access violation, floating-point exception, etc.) that should not interfere with the overall benchmark process. Standard error reports will be defined for common cases, along with common return values.

Benchmark consistency: we emphasise on benchmark modularity for it allows easy development and maintainability of the whole benchmark. Nevertheless, we have to ensure global consistency throughout the modules that will be used for performing a particular session. This will help to achieve benchmark reproducibility.

Result output: Another critical issue is that of accurate result output when all tests of a session have been performed. Depending on the type of user, we will have two different kinds of output. For developers, the report writer module will provide raw results (for instance, for optimization purposes or statistical analysis), whereas choice makers will receive a certification of the submitted algorithm from the certification module with respect to the specific application performance specifications they have defined.

A. Open-Source and Web-Based Benchmarking

Collaboration and publications constitute the foundations of academic research. Information, knowledge are actually well and easily distributed thanks to electronic journals or forums. A second aspect is the algorithmic and programming collaboration. Unfortunately for the heterogeneity of the programming languages, the operating systems used by researchers are major problems to design a common testing platform. Current solutions exist to develop collaborative work. Generally, these solutions impose a specific programming language and/or operating systems to developers. Some others specific rules have to be respected. These heavy constrains slow down the utilization of collaborative programming platform. The OpenWatermark [25] and the Watermark Evaluation Testbed (WET) [26] projects propose a modern architecture for cooperative programming exchange that takes all these aspects into account. Developers work on their preferred programming language and operating systems.

OpenWatermark is a distributed application system whose initial purpose is to allow the execution and the comparison (i.e., benchmarking) of programs uploaded by the user.

The user first logs into a Web site using her/his preferred Web browser, fills a form where she/he is asked to explain some characteristics (such as the programming language used and the syntax of its command-line arguments) of her/his program and to upload it. Those characteristics, programs as well as the input data sets, are stored into an SQL database.

The OpenWatermark system (see Fig. 7) determines on which machine the execution should be scheduled and the context of those executions, that is to say, mainly which data sets and parameters should be used and what kind of output should be expected. It then connects to the machines concerned and requests them to download the executable from the database as well as the associated data sets, run it using the previously specified command-line options, and upload back to the database the results of this execution. Finally, the results could be consulted by the user on the same Web site as soon as they become available. The user interface is entirely constituted of Web pages written in Java Server Page (JSP), and communicates with the Java application responsible for the execution of the tested program running on each of benchmarking hosts using Remote Method Interface (RMI) and with the SQL server using Java Database Connectivity (JDBC).

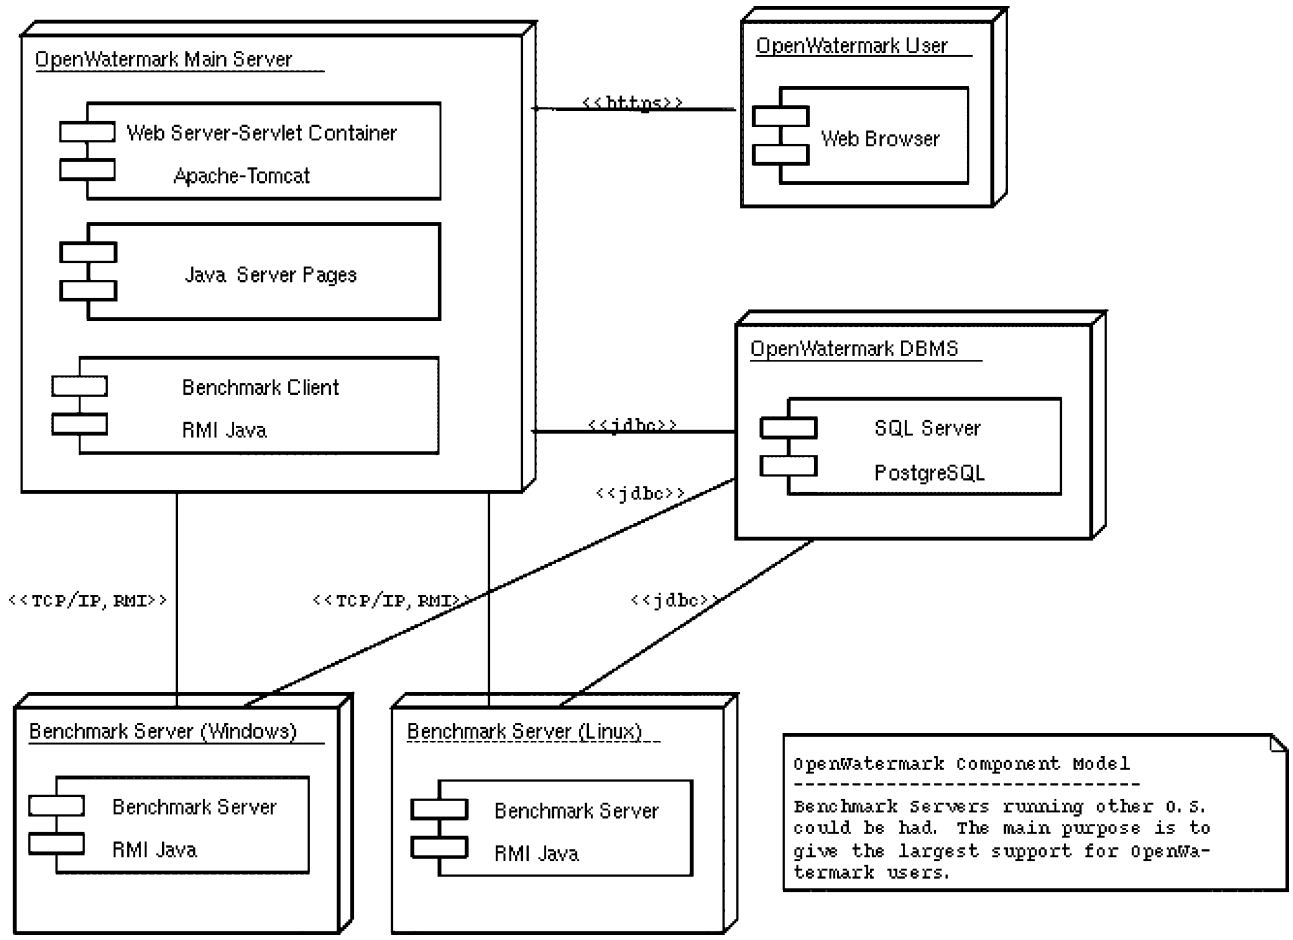

Fig. 7. OpenWatermark architecture.

This architecture is, therefore, independent of the operating system used and, hence, easily portable on any platform supporting Java, RMI, and JDBC.

In addition, to allow the upload of the programs and their characteristics, several forms allow the elaboration of benchmarking templates. Those templates are constituted of blocks and relations composing a block diagram. The blocks represent either a data set, an option (e.g., a compression or quality parameter) or a program.

When the user selects templates to be run, the JSP script launches several threads managing the required executions. Browsing the block template, these threads evaluate how many executions of each program (with specific data sets and parameters) are required, and what their dependencies are. They then request the benchmarking hosts through RMI calls to fetch the corresponding data from the database into files which will the be passed as command-line arguments (in addition to the other mandatory or optional parameters) to the program to be executed. As soon as the execution successfully ends, if it did generate an output file, then it is uploaded by the local Java server into the database using JDBC. The state of the execution (finished or not) could be seen from the HTML user interface. The already available results are directly accessible either for download or (in the case of images) view from the user's Web browser. This platform-independent architecture, thus, constitutes a generic framework, which could be used in order to comparatively evaluate the efficiency of test programs on various data sets with reusable.

WET is a Web-based system for evaluating the performance of watermarking techniques. WET provides a user interface which allows the user to perform various watermarking operations over the Internet. The image processing engine used for embedding, detection, attack, and performance evaluation is the open-source program Gimp. WET consists of four major components: the front end, the Web server, the database server, and the gimp-perl server. WET has two modes of operation: the initial version and the advanced version. The initial version has a very intuitive interface, but it allows only limited functionality. The advanced version gives the user complete access to the system. Currently WET has six embedding and detection methods, an attack suite that includes StirMark, four performance metrics (including mean-square error and execution time), and a database of more than 3000 images. WET also supports user provided techniques through the use of GIMP plug-ins.

V. Conclusion

New ways to represent efficiently the characteristics of the watermarking algorithms are under development.

Instead of delivering certification for watermarking, the authors are convinced that the benchmarking needs still further research. They promote the use of a cooperative Web-based approach because key factors like quality evaluation under geometrical transforms and capacity evaluation still need to be better understood. Furthermore we see that beside robust image watermarking evaluation (see also, for example, [27]–[31]), other media like audio raises importance and first approaches for Web-based audio watermarking benchmarking can be found in the StirMark Audio Benchmarking project; see, for example, [32] for compression robustness and [33] for perceptual quality evaluation of attack parameters. Therefore, a media overlapping benchmarking suite with a generic profile design for different attack types and attack combinations is one of the future challenges in benchmarking.

ACKNOWLEDGMENT

The authors would like to thank all involved persons for their stimulating discussions and their work. Some ideas presented here come from the previous works done in CERTIMARK and StirMark.

REFERENCES

[1] C. Herley, "Why watermarking is nonsense," *IEEE Signal Processing Mag.*, vol. 19, pp. 10–11, Sept. 2002.
[2] M. Lesk, "The good, the bad and the ugly: what might change if we had good DRM," *IEEE Security Privacy Mag.*, vol. 1, pp. 63–66, May–June 2003.
[3] J. Dittmann, P. Wohlmacher, and K. Nahrstedt, "Using cryptographic and watermarking algorithms," *IEEE Multimedia*, vol. 8, no. Oct.–Dec., pp. 54–65, 2001.
[4] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Attacks on copyright marking systems," in *Proc. Information Hiding 2nd Int. Workshop*, 1998, pp. 219–239.
[5] F. A. P. Petitcolas, "Watermarking schemes evaluation," *IEEE Signal Processing Mag.*, vol. 17, pp. 58–64, Sept. 2002.
[6] F. A. P. Petitcolas, M. Steinebach, J. Dittmann, C. Fontaine, Raynal, and N. Fatès, "A public automated Web-based evaluation service for watermarking schemes: StirMark Benchmark," *Proc. SPIE: Security and Watermarking of Multimedia Contents III*, vol. 4314, pp. 575–584, 2001.
[7] StirMark [Online]. Available: http://www.petitcolas.net/fabien/watermarking/stirmark/index.html
[8] StirMark—WEB [Online]. Available: http://stirmark.kaist.ac.kr/
[9] StirMark—AUDIO [Online]. Available: http://amsl-smb.cs.uni-magdeburg.de/stirmark/index.php
[10] V. Solachidis, A. Tefas, N. Nikolaidis, S. Tsekeridou, and A. Nikolaidis, "A benchmarking protocol for watermarking methods," in *IEEE Int. Conf. Image Processing (ICIP'01)*, 2001, pp. 1023–1026.
[11] S. Pereira, S. Voloshynovskiy, M. Madueño, Marchand-Maillet, and T. Pun, "Second generation benchmarking and application oriented evaluation," presented at the Information Hiding Workshop III, Pittsburgh, PA, 2001.
[12] CERTIMARK [Online]. Available: http://vision.unige.ch/certimark/public/CMK_D22.pdf
[13] P. Cano, E. Batlle, T. Kalker, and J. Haitsma. (2002) A review of algorithms for audio fingerprinting. *Proc. International Workshop on Multimedia Signal Processing* [Online]. Available: http://www.iua.upf.es/mtg/publicacions2002.php3?lng=eng&authorlink=3&showpdf=1&urldocid=207
[14] D. Augot, J. M. Boucqueau, J. F. Delaigle, C. Fontaine, and E. Goray, "Secure delivery of images over open networks," *Proc. IEEE*, vol. 87, pp. 1251–1266, July 1999.
[15] L. Cheveau. Choosing a watermarking system for digital television—the technology and the compromises. *IBC 2002* [Online]. Available: http://www.broadcastpapers.com/asset/IBCEBUWatermarking03.htm
[16] R. Anderson, "Trusted computing and competition policy—Issues for computing professionals," *Upgrade*, vol. I4, no. 3, pp. 35–39, June 2003.
[17] S. Western, K. L. Lagendijk, and J. Biemond, "Perceptual image quality based on a multiple channel HVS model," in *Proc. ICASSP*, 1995, pp. 2351–2354.
[18] Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Processing Lett.*, vol. 9, pp. 81–84, Mar. 2002.
[19] I. Setyawan, D. Delannay, B. Macq, and R. L. Lagendijk, "Perceptual quality evaluation of geometrically distorted images using relevant geometric transformation modeling," presented at the SPIE/IST 15th Electronic Imaging, Santa Clara, CA, 2003.
[20] S. Baudry, J.-F. Delaigle, B. Sankur, and B. Macq, "Analyses of error correction strategies for typical communication channels in watermarking," *Signal Process.*, vol. 81, no. 6, pp. 1239–1250, June 2001.
[21] R. Lagendijk and I. Setyawan, "Digital watermark benchmarking: Requirements and implementation issues," Information and Communication Theory Group, Tech. Rep., 2002.
[22] V. Solachidis, A. Tefas, N. Nikolaidis, S. Tsekeridou, A. Nikolaidis, and I. Pitas, "A benchmarking protocol for watermarking methods," presented at the ICIP 2001, Thessaloniki, Greece.
[23] G. Voyatzis and I. Pitas, "The use of watermarks in the protection of digital multimedia products," *Proc. IEEE*, vol. 87, pp. 1197–1207, July 1999.
[24] M. Kutter, S. Voloshynovskiy, and A. Herrigel, "Watermark copy attack," presented at the IS&T/SPIE's 12th Annu. Symp. Electronic Imaging 2000: Security and Watermarking of Multimedia Content II, San Jose, CA, 2000.
[25] OPENWATERMARK [Online]. Available: http://www.openwatermark.org
[26] [Online]. Available: http://www.datahiding.org
[27] M. Kutter and F. Petitcolas, "A fair benchmark for image watermarking systems," *Proc. SPIE, Security and Watermarking of Multimedia Contents*, vol. 3657, pp. 223–226, 1999.
[28] S. Voloshynovskiy, S. Pereira, A. Herrigel, N. Baumgärtner, and T. Pun, "Generalized watermark attack based on watermark estimation and perceptual remodulation," presented at the Security and Watermarking of Multimedia Content II, San Jose, CA, 2000.
[29] J. K. Su and B. Girod, "Power-spectrum condition for energy-efficient watermarking," in *Proc. IEEE Int. Conf. Image Processing (ICIP99)*, pp. 301–305.
[30] J. K. Su, J. J. Eggers, and B. Girod, "Analysis of digital watermarks subjected to optimum linear filtering and additive noise," in *Signal Process. (Special Issue on Information Theoretic Issues in Digital Watermarking)*, vol. 81, 2001, pp. 1141–1175.
[31] F. Hartung, J. K. Su, and B. Girod, "Spread spectrum watermarking: malicious attacks and counter-attacks," presented at the Conf. Security and Watermarking of Multimedia Contents, San Jose, CA, 1999.
[32] M. Steinebach, A. Lang, and J. Dittmann, "StirMark benchmark: Audio watermarking attacks based on lossy compression," *Proc. SPIE, Security and Watermarking of Multimedia Contents IV*, vol. 4675, pp. 79–90, 2002.
[33] A. Lang, J. Dittmann, and M. Steinebach, "Psycho-akustsiche Modelle für StirMark Benchmark—Modelle zur Transparenzevaluaierung," presented at the Informatik 2003, Bellingham, WA.

Benoit Macq (Senior Member, IEEE) was born in 1961. He received the Engineer and Ph.D. degrees from Université Catholique de Louvain (UCL), Louvain, Belgium, in 1984 and 1989, respectively.

He is currently Professor at UCL, in the Telecommunication Laboratory. He has been Guest Editor for *Signal Processing*. He was Co-Guest Editor for a Special Issue on security for image communications of *Image Communications*. He teaches and does his research work in image processing for visual communications. His main research interests are image compression, image watermarking and image analysis for medical and immersive communications.

Dr. Macq received the Bell Telephone Award in 1990. He has been a Member of the program committee of several IEEE and SPIE conferences. He is a Member of the IEEE Technical Committee for Image and Multidimensional Signal Processing (IMDSP). He was Co-Technical Chair of the IEEE Conference on Multimedia and Expo (ICME02) and Member of the Board of EUSIPCO 2002 and ICPR 2002. He will be on the Technical Committee of ICASSP 06. He is Associate Editor of the IEEE TRANSACTIONS ON MULTIMEDIA. He was Co-Guest Editor for a Special Issue on watermarking for the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY and Guest Editor for the PROCEEDINGS OF THE IEEE.

Jana Dittmann was born in Dessau, Germany. She received the Engineer degree in computer science and economy and the Ph.D. degree from the Technical University of Darmstadt in 1990 and 1999, respectively.

Since 2002, she has been a Full Professor of multimedia and security at the Otto-von-Guericke University, Magdeburg, Germany. She has many national and international publications. She is Associate Editor for the *ACM Multimedia Systems Journal*. She specializes in the field of multimedia security. Her research is mainly focused on digital watermarking and content-based digital signatures for data authentication and for copyright protection.

Dr. Dittmann is a Member of the Association for Computing Machinery and GI Informatik. She is a Member of several conference Program Committees and organizes workshops and conferences in the field of multimedia and security issues. She was involved in all last five multimedia and security workshops at ACM Multimedia. In 2001, she was a Cochair of the CMS2001 conference. She was a Guest Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING.

Edward J. Delp (Fellow, IEEE) was born in Cincinnati, Ohio. He received the B.S.E.E. (*cum laude*) and M.S. degrees from the University of Cincinnati, Cincinnati, OH, the Ph.D. degree from Purdue University, West Lafayette, IN, and the Honorary Doctor of Technology from the Tampere University of Technology, Tampere, Finland, in 2002.

From 1980 to 1984, he was with the Department of Electrical and Computer Engineering at the University of Michigan, Ann Arbor. Since 1984, he has been with the School of Electrical and Computer Engineering and the Department of Biomedical Engineering at Purdue University. In 2002, he received a chaired professorship and is currently the Silicon Valley Professor of electrical and computer engineering and Professor of biomedical engineering. His research interests include image and video compression, multimedia security, medical imaging, multimedia systems, communication and information theory.

Dr. Delp is a Fellow of the SPIE, a Fellow of the Society for Imaging Science and Technology (IS&T), and a Fellow of the American Institute of Medical and Biological Engineering. In 1990, he received the Honeywell Award and in 1992 the D. D. Ewing Award, both for excellence in teaching. In 2001, he received the Raymond C. Bowman Award for fostering education in imaging science from the IS&T. In 2002, he was awarded a Nokia Fellowship. In 2000, he was selected a Distinguished Lecturer of the IEEE Signal Processing Society. He was the Program Cochair of the IEEE International Conference on Image Processing that was held in Barcelona in 2003.