USER REQUIREMENTS FOR WATERMARKING IN BROADCAST APPLICATIONS

A.J.Mason[1], R.A.Salmon[1], O.H. Werner[2], J.E.Devlin[1]

[1]BBC Research and Development Department, UK and [2]WDR, Germany

ABSTRACT

Digital watermarking of audio-visual signals is a recent and rapidly developing field of research. There are several potential applications ranging from authentication to identification. For example, imperceptible identification marks can be used to help copyright protection, distinguish copies or monitor audience behaviour.

There is already a considerable amount of published work that deals with the design of algorithms. However, there is little work that addresses watermarking from the user's point of view. Potential uses, consequent system design issues in a broadcasting environment, performance specification and evaluation techniques are all crucially important. A broadcaster needs to know what a watermarking system can achieve, both theoretically and in practice. An understanding of the issues involved, the conflicting demands and, ultimately, a means of evaluating competing systems are required.

INTRODUCTION

Watermarking is already well known in the authentication of banknotes. In the field of audio and video signals watermarking is increasing in importance with the advent of digital audio and video media and transport mechanisms which allow perfect copies to be made. In this context a watermark is an imperceptible modification to the signal that can be detected on request and that conveys hidden data. The reasons why a watermark might be embedded into a signal and the benefits that might be had from detecting it are numerous. This paper concentrates on applications that are relevant to broadcasters.

There are several applications for watermarking in a broadcast scenario. A clear distinction must be made between watermarking for the purpose of authentication and watermarking for the purpose of identification. Authentication establishes the credibility of an audio-visual signal, that is, whether the signal is what it claims to be or not. Identification associates an audio-visual signal with descriptive information. This association can be made by using the watermark to convey a unique identification number pointing to a record in a database.

The properties required by authentication and identification are contradictory. An authentication mark must be difficult to create and sensitive to distortions. An identification mark should be difficult to remove and robust to distortions.

This work focuses mainly on identification. The possible uses of a watermark along the broadcast chain are discussed and the required properties for the watermark are derived. Test methods are proposed for checking a watermarking system against the requirements.

USES OF WATERMARKING IN BROADCASTING

There are several reasonably obvious uses of watermarking in broadcasting:

Rights and Metadata

Identifying programme material within the studio environment helps us to ensure that we meet contractual agreements concerning rights. Each item of programme material can carry an identifier that links it to a record in a database that enables the broadcaster to keep track of the use of performing rights.

As a programme is being made, the information about the rights being used can be collated automatically and information in the database updated to help ensure that rights owners receive the proper reward.

Rights and Copyright

Copyright protection can be assisted by use of watermarks. In the broadcasting context, the watermark in a programme would identify the original broadcaster and enable the detection of unauthorised re-use of programme material.

Programmes licensed to broadcasters might be broadcast more often than was agreed. Programmes licensed for broadcast in one geographical area might be broadcast in other areas. Programme material might be recorded off-air and re-broadcast by another broadcaster without authorisation. Illegal tape recordings of programmes might be sold on the market.

This raises the question of what to do when detection takes place. One course of action is to prosecute the person infringing the rights and take them to court. However, in practice, it might be advantageous to ask them if they would consider legitimising the use they were making of the rights. This might well bring in more revenue over a longer period.

Audience Research and Advertising

All broadcasters, whether public service or commercial, are interested in audience size. Knowing accurately how many people are tuning in is very important if you want to justify your advertising rates or licence fee. Rather than relying on people completing a questionnaire to record their viewing and listening habits a small, wearable, audio watermark detector could do the same job.

For the commercial broadcaster, verifying that the commercials that are contractually required to be broadcast are actually being broadcast is a clear case where watermarking might be employed to good effect.

Authenticating Web Site Content

The BBC has a well-known Online presence. A huge amount of material is stored and distributed this way. There is a concern about the security of the information stored on the servers in that it might be possible for a malefactor to exploit some software loophole and put material onto our servers. The possible consequences, in terms of loss of reputation, are enormous. It would be preferable to have systems in place that ensure that only authorised users have write-access to the servers, but it is also possible to use an authentication watermark and a web-crawler to monitor the state of the servers.

SYSTEM DESIGN CONSIDERATIONS

A generic watermarking system consists of an embedder using a secret key, a channel which may contain systems that attack the watermark, and a detector using the same secret key. Once the watermark is detected successfully, the key can additionally provide access to a hidden data channel from which one can retrieve a user-defined payload.

The embedder and detector can be deployed at different points along the broadcast chain depending on the specific purpose of the watermark.

Embedding

Figure 1 shows a simple model of a broadcast chain. It consists of four areas with different functions: signal acquisition, programme production, broadcast, and archive. Acquisition is where sounds and pictures are captured, or, alternatively, they are synthesised. Production is where these signals are edited, mixed and variously processed to produce a programme. Broadcast is where programmes are sent to the network for distribution and emission by analogue, digital, terrestrial, satellite, cable, or whatever other distribution mechanism, for example, the Internet. In this model, programmes get into the archive either as records of what has been broadcast or as partial or complete programmes from the studio (not necessarily broadcast).

Some points in this model at which a watermark might usefully be applied are as follows:

1) Acquisition: each camera or microphone could apply its unique identifier consisting of time, place, equipment serial number, and so on.

2) Production output: each finished programme, or useful parts thereof, could contain a programme identifier.

3) Network input: each network could be marked to identify the broadcaster's network, the time and date of transmission, and the programme.

4) Archive input: as each programme item is catalogued and stored in the archive, the catalogue number could be included in a watermark.

5) Archive output: as each request for material from the archive is fulfilled, a watermark identifying the recipient (a so-called 'fingerprint') could be added to the material.

Detection

A point that is overlooked more often than one might expect is that a watermark is only useful if one has a detector and somewhere to use it. Unless some monitoring is in place the watermark is worthless.

Figure 2 shows some places in the broadcast chain where one might monitor for watermarks and the benefit that this can have.

1) Input to production: if one needs to keep track of what shots are being used for what purpose, or identify the source of material, we can monitor all that comes in to the studio. This could act simply as a backup system for the correct labelling of tapes together with some of the meta-data that would be associated with the incoming material. For this purpose the watermark would have to have a relatively high data capacity because the length of segments of material that get used might be short and the information that is useful is quite long.

2) Output of production: There are many rights issues associated with a programme. If the watermark identifies the source and provides a pointer to information in a database then the rights issues can more easily be resolved. The data-rate depends on whether one needs to identify a source, or a segment of material from a particular source.

3) Input to presentation, or play-out: it would be useful in a presentation area to be able to identify incoming programme feeds to ensure that the right programme is being selected at the right time.

4) Network input: as programmes are sent to air, automatic logging of what material has been broadcast could be done using watermarks. The use of rights can then be monitored precisely, and not be dependent on meticulous record keeping throughout the production process.

5) Input to archive: the process of cataloguing what gets put into the archive can be simplified by monitoring the information that might be carried by a watermark. The hidden data conveyed by the watermark can point to a record in a production database from which information is automatically copied into the archive database when material is brought into the archive. Again, the data is low for long programmes and high for short items.

6) The whole world: finally, the entire world of broadcast media can be monitored to trace signals and to check whether rights are being infringed.

This arena presents the most stringent requirements for a watermarking system. Whilst the data rate required to identify the rights owner might be relatively low, the attacks that might be made on the integrity of the watermark, with the deliberate aim of attempting to render it undetectable, can be substantial.

It is this final detection option that has attracted much of the attention in recent years. The costs and benefits do merit careful examination. Monitoring all broadcasts in all parts of the world is expensive and time-consuming. Far more likely is the possibility that most of the benefit can be gained by monitoring a small proportion of broadcast networks in a few carefully chosen locations. The rest simply are not worth the effort.

When one has decided for what purpose a watermark is used one can decide whether it is permissible for others to have access to the watermark and the information that it conveys. Theoretically, if someone can detect a watermark they can work out how to remove it. In the limit, with a detector that says, "watermark present" or "watermark absent" an attacker has only to tweak pixel values until the detector changes from "present" to "absent". So, in some circumstances it is highly undesirable that others should be able to detect our watermarks. This leads on to the selection of systems that rely on some concept of a 'secret key' rather than a 'public key' and all the issues of key management that go with them. Standardisation often has benefits, but in watermarking it brings with it the need for trusted third parties or registration authorities to manage the secrecy.

Broadly speaking, all the application scenarios that one could derive from the points at which one could embed and detect watermarks in the broadcast chain, fall into two categories are far as the parameters of the watermark technology are concerned. It is taken for granted, for several reasons, that imperceptibility is required. As a result, we find that we need two different kinds of watermark, or one kind with two different sets of operating parameters. One set of applications calls for a high data capacity but can have relatively low robustness to attacks (unintentional). The other calls for a high level of robustness to attacks both unintentional and intentional, but can have a lower data rate.

PROPERTIES REQUIRED OF A WATERMARK

To consider what properties the watermark should have it is useful to focus on one particular application within the broadcast chain, namely the watermarking of the signal at the input of the transmission network. We will only consider video watermarking but similar requirements apply to audio watermarking systems.

There are three principal requirements, which in any practical system have to be traded off one against the other:

Perceptibility.
That the watermark should be invisible on the picture is an over-riding requirement, taking precedence over all others. If the watermark produces a visible degradation of the received signal then the entire system is likely to be a non-starter. Unfortunately, this requirement is not as clear-cut as it might be: there is the vitally important effect of the watermark on the coding performance of MPEG-2 coders in the subsequent digital transmission chain. Adding a noise-like signal to video may well cause the MPEG coder to have to work harder, thus degrading the picture quality at a given MPEG-2 bit-rate.

Robustness.
The watermark must still be detectable after a series of "attacks" on it. These attacks can be categorised as either "friendly" or “hostile”. Friendly attacks would be those to be expected in the broadcast environment, such as standards conversion, MPEG-2 coding at bit-rates as low as 2Mbit/s, or PAL transmission. Hostile attacks are those that are applied during the course of unauthorised use of the material, such as deliberate geometric distortion, noise coring or VHS tape recording.

Data capacity.
There is a requirement for the watermark to convey some hidden data channel such that the source, ownership, time, and method of transmission can be determined on detection of a pirated broadcast. A data capacity of a packet of 64 bits of information is usually considered appropriate. The other factor to be considered is the "watermark minimum segment" which is the minimum length of video sequence in which one wishes the watermark to be detected. This is the length of video which the broadcaster considers to be of value, and something between 1 and 10 seconds is usually considered appropriate, implying that the data packet remain unchanged for this period.

For a given watermarking technique there is a trade-off among the above three requirements, since changing any one of them has an effect on the other two. Thus fixing imperceptibility and a data-rate has the effect that the level of robustness cannot then be chosen freely.

Cascadability of the watermark.
As indicated earlier, watermarks might be applied at various stages of the broadcast chain, so multiple watermarks may have to exist on the same video signal. One requirement is that the invisibility of the watermark should not be compromised by multiple applications.

Orthogonality.
Many watermarks make use of some form of spread spectrum technique employing a pseudo-random sequence generated from a secret key. In order to be able to detect and separate the payload of cascaded watermarks, the different watermarks should be mathematically orthogonal. This means that the addition of a second watermark using a different key does not compromise the detectability of the first, and vice versa.

False positive probability.
The probability of detecting a watermark in an un-watermarked video stream must be sufficiently remote that the cost of investigating occurrences of false positives is not excessive, and, in the event of a case coming to court, the lawyers could not argue that the watermark could have been apparently detected even though it were not there. A false positive probability of the order of 10^{-12} may well therefore be required.

Disclosure of algorithm.
It is a likely requirement that during court action the watermarking algorithm used would have to be disclosed. This is also needed to enable the user to confirm the required figure for the false positive probability, since it is not feasible to test the system for a long-enough period to confirm this requirement by experiment alone. It is therefore essential that the security of the watermark resides in some secret key rather than in the knowledge about the algorithm.

Delay.
The smallest possible video delay in the watermark embedder is required, of the order of 2 fields (for interlaced raster scan) being reasonable.

Detector complexity.
Detection must be done in real time and must be low cost.

TESTING AGAINST REQUIREMENTS

With different manufacturers now offering watermarking systems it is necessary to consider methods of testing those systems against the requirements. The two areas requiring particular attention are perceptibility and robustness. Again these comments are aimed directly at video watermarking, but the principles are much the same for audio.

Testing for perceptibility.

In the first instance an indication of the perceptibility of the watermark would be judged informally by experts on properly lined-up picture monitors. We have also found that it is important to view material on plasma panels, particularly ones with a direct digital input, since such displays seem even more critical than the traditional grade 1 monitor, and are now available to the general public. A wide range of test material must be viewed to determine what is critical for the particular watermarking system under test.

For a more formal evaluation the ITU-R BT.500-8 "Double-stimulus, continuous quality-scale method" of subjective testing is preferred, since this method is very suitable for testing systems introducing a minimal amount of distortion. Care must be taken in the choice of test material, since only a limited range can be presented to the test subjects if the tests are to be kept down to a reasonable duration.

It is often suggested that an objective quality measure be used to evaluate such systems, either using a simple signal-to-noise measure or a more sophisticated perceptual model.

There is an inherent danger that a system optimised to such a measure may appear to perform better than a system that has been optimised in another way that may in fact produce artefacts that are actually less annoying. Formal subjective tests must therefore remain the principal measure of picture quality. As an example, consider a watermarking system that worked by displacement of edges in the picture by a very small amount. Whilst being unnoticeable to the viewer, this might score very badly according to a perceptual model assuming additive noise effects.

Robustness testing.

The system must be tested to ensure that the watermark can survive the normal operations carried out on the picture in a studio/broadcast environment, and also that it has a measure of resistance to attacks such as VHS recording and duplication, and at least two generations of MPEG-2 coding and decoding with different group of picture (GOP) alignments, at bit-rates down to 2 or 3 Mbit/s. In the case of VHS recording it is important to reproduce the real-life scenario where the replay is made on a different machine from the recording.

The resistance of the watermark to geometrical attacks must also be tested, whether they be straight-forward aspect ratio conversion, or more severe attacks such as might be applied by a pirate, including line or column dropping, shear or rotation. Cropping and picture shifts horizontally and vertically also need to be considered; they can be observed to occur routinely in signal distribution, quite apart from any hostile attack of this sort. The system should also be capable of recovering the watermark after noise-reduction techniques have been applied to the signal, since such techniques are on occasion applied to the input of MPEG-2 coders to improve the coding quality.

CONCLUSION

This work discusses the user requirement of watermarking in a broadcast scenario. As a consequence of the imperceptibility constraint, the possible applications fall broadly into two categories: one set of applications requires a high data capacity but compromises on the robustness. The other set requires a high level of robustness but compromises on the data capacity. Possible applications of the first kind are typically encountered in a controlled environment where watermarking can assist in the automation of digital content and asset management. Watermarking to help copyright protection and to trace signals outside the broadcaster’s own network are possible applications of the second kind. Test methods are proposed, suitable for verifying whether emerging watermarking systems fulfil the required imperceptibility and robustness requirements.

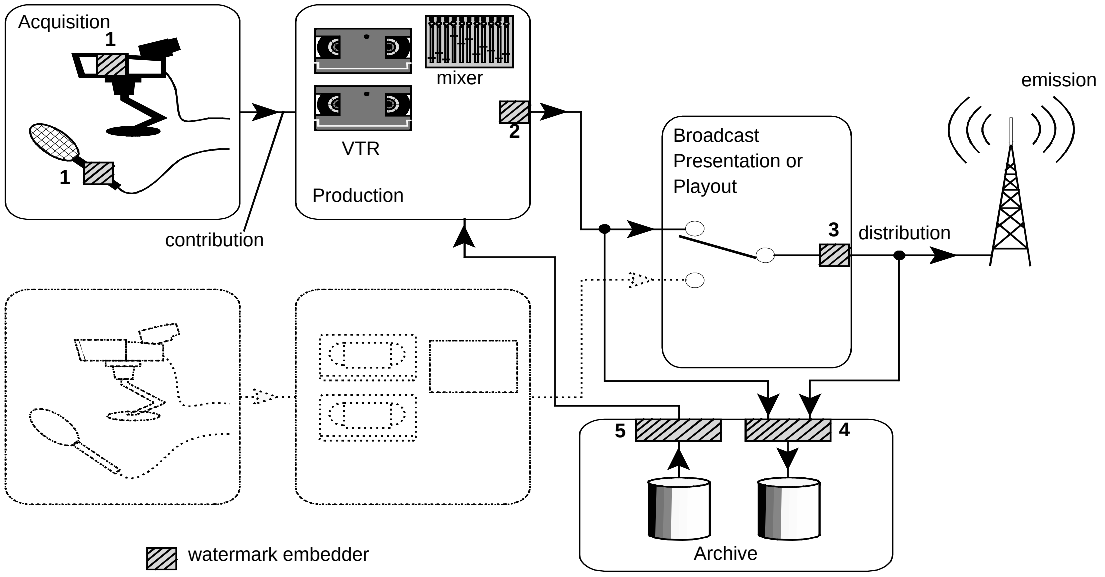

Figure 1 Watermark embedding points

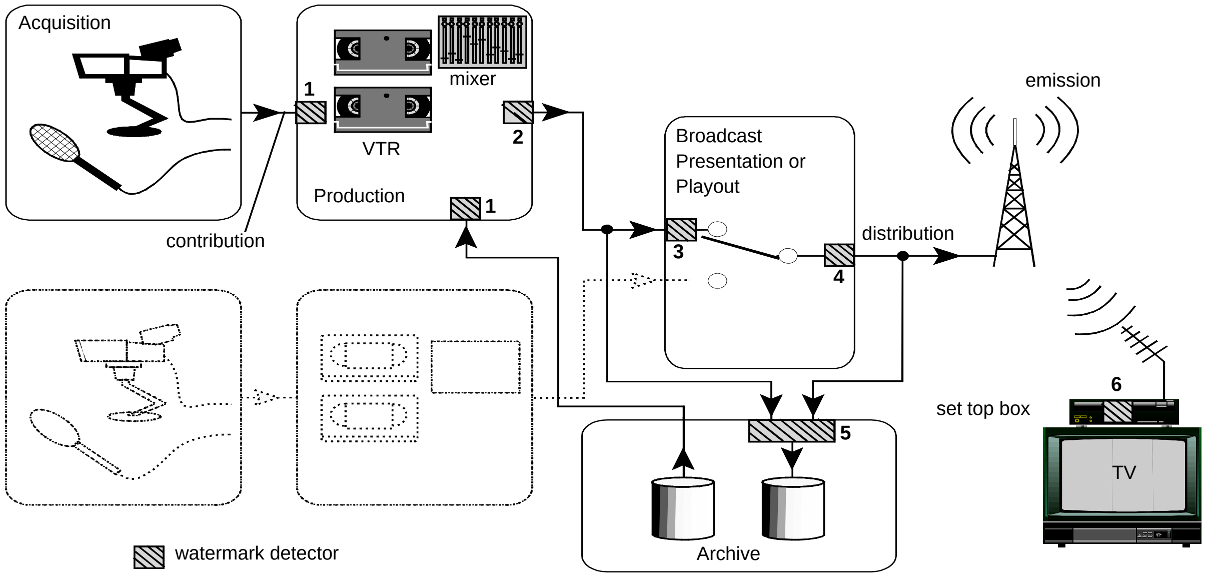

Figure 2 Watermark detection points