NAGRA

TECHNOLOGY SERIES - WHITE PAPER

CAS
IPTV
ON DEMAND
INTERACTIVE
DRM
MOBILE
CMS

Forensic Watermarking Architectures in a CAS/DRM Environment

NAGRAVISION
KUDELSKI GROUP

1. FORENSICS IN A MEDIA DELIVERY ENVIRONMENT

Securing the value chain in an end-to-end media delivery environment typically requires a set of complementary technology tools designed to address the threats specifically raised within the corresponding ecosystems [1].

In this white paper we focus more specifically on the technology required to address the emerging threat of unauthorized content redistribution beyond the CA/DRM control, and how such a technology may be integrated with other security technologies such as content encryption, conditional access and rights management within an end-to-end content/service protection architecture.

What are the content/service protection threats calling for forensics?

A number of solutions have been put in place today in order to protect digital content assets along the value chain from content management in the B2B environment to a plethora of content distribution B2C models, including theatres, packaged media, TV broadcasts and VOD.

However, the evolution of digital networking technology leads now the consumer market to:

- Widely benefit from increase in domestic internet bandwidth access, both downstream and upstream;
- Adopt convenient media management tools in the form of media managers in personal computers, in addition to more conventional consumer electronics equipment dealing with digital video as well as music - the Apple iPod(tm) success perfectly illustrating the convergence between the two worlds.

This evolution introduces a new category of threats that were by far not as concerning in the analogue world: the C2C threats.
Indeed, as TV place shifting and peer-to-peer media exchange become convenient enough to be part of everybody's culture just as the practice of emails, chats, blogs and internet phone calls have become widely spread, illegal content sharing will become a real threat to the content distribution value chain.

Why would a service provider need to address them?

This C2C threat concern has been long identified by the content owners, in particular the main video studios learning lessons from how peer-to-peer has affected the music industry revenue. Forensic tracking technologies based on watermarking have been tested initially on Hollywood screeners in order to better track early content asset leakages still at the B2B stage, and were proved useful on a very popular 2004 case [2]. Forensic marking requirements are also explicitly described in the digital cinema specifications the movie industry is currently putting in place through their Digital Cinema Initiative [3].

In the meanwhile, beyond the movie and music industry, major owners of sport events distribution rights have recently started to sue C2C video sharing platforms like Google's YouTube for copyright infringement.
So there are basically two motivations why a service provider would need to consider the integration of forensic technologies into its end-to-end security:

- Access to early release windows for some top ranking HD content may be subject to the ability by the distributor (typically VOD or PPV service provider) to enable digital forensics in accordance with the content owners' requirements.
- Protection of the service provider revenue itself.

Indeed, the C2C threat will hurt not only the content owners revenue, but also the service providers', as soon as this technology becomes popular enough (due to the ease of the underlying technology combined with the lack of effective legal circumvention means) to provide C2C payTV subscription sharing.

What does forensic tracking add to a CA/DRM solution?

Conditional access and digital rights management solutions primarily focus on ensuring that only the end users' devices with the appropriate entitlements can make a specific usage of the content – such as view, record, copy, transfer, etc.

In order to keep close control on the content usage, the content is protected from unauthorized access by encryption (called scrambling in conventional payTV applications). The secure handling of content encryption keys along the content distribution chain is a key feature of any CA/DRM scheme, as those keys may also be subject to a C2C key sharing threat in particular in a broadcast environment. A dedicated security framework has therefore been developed by Nagravision to ensure that the operator STB matches the required level of security against this particular threat **[12]**.

However, the protection brought by the original content security stops as soon as the original CA/DRM has to hand over the content control to a third party technology, such as for instance in a Set-Top-Box HDCP on DVI or HDMI interfaces, or Macrovision on analogue outputs. The latter technologies provide some protection against copying, but do not address the C2C content sharing threat.

Is forensic tracking the holly grail of content protection?

Forensic tracking is primarily a dissuasive content protection measure, to be used in association with the appropriate legal framework. Forensic tracking neither address the access and usage control functionality nor the rights management that are essential components of an end-to-end content delivery architecture. It is rather a complementary tool that may be useful to specifically address emerging C2C content sharing threats, depending upon the service providers and content owners needs to protect their revenue flows after the initial content purchase, the latter remaining under close control by CA/DRM solutions.

2. OVERVIEW OF FORENSIC WATERMARKING TECHNOLOGIES

What do we mean by forensic tracking?

In order to trace illegal digital content leaking sources from broadcast PayTV services, a forensic tracking[1] mechanism is desirable that will enable to embed a Set-Top-Box (or subscriber) identifier into the content at the time it leaves the CA/DRM control.
Retrieval of the identifier in illegal content, typically traced by dedicated monitoring services from the black market or the Internet, will then enable:

- To target the necessary technical countermeasures though the CA system controls to prevent content access from the «leaking» decoders;
- To provide proofs to the content owners to prosecute the «leaking» subscribers in accordance with the US DMCA or EU Copyright directive legislations, etc.

[1] Forensic tracking is also sometimes called fingerprinting. However the latter terminology also refers to the traceability of content signatures, rather than end-user signatures, which is a different application, so we prefer the former, more accurate wording.

The simplest forensic tracking solution simply consists in adding the forensic identifier into the digital content format itself, as a simple tag in one or the other private extensions placeholders enabled by the content representation syntax (e.g. in some side fields of MPEG streams). However this approach is very easy to detect and circumvent, and it does not survive to format conversions (e.g. analogue).

What is watermarking?

Watermarking has therefore been identified by the academic as well as industrial players as the technology offering this forensic tracking capability in a fairly robust way. It consists in embedding a signal directly into the content (audio or video) in such a way that it is invisible to the human, detectable by computing, while remaining robust to further content transformations (including digital-analogue-digital conversions and compression).

Watermarking originates from signal processing research and is typically fine-tuned to match the signal characteristics (e.g. hiding the information in the form of discrete variations into the signal noise and redundancies, so that it is imperceptible to the human but remains detectable by a machine... while remaining robust to compression, although the latter process typically removes the noise and redundancies from the original signal!).

Example of signal processing attacks a watermarking technology shall be robust against include (for video):

- Digital to analog to digital conversion

– Format ratio change

– Decompression-recompression in lower rate (e.g. 1.5Mbits/s or 3Mbits/s instead of 5)

– Geometrical transforms (rotate, translate, re-scale)

– Various filtering, esp. noise removal, image sharpening, color adjustment

Advanced watermarking technologies also combine the signal processing algorithms with a number of additional improvements ranging from cryptography add-ons (used to control the insertion/detection algorithm or the watermark information itself) to error detection and correction codes (used to optimize the watermarking detection robustness).

How is a watermark embedded into the content?

The watermark embedder can be configured in various ways, but typically takes as input the original content, the information to embed (a.k.a. watermark payload), and possibly a number of configuration parameters such as a cryptographic key, a target signal strength or redundancy (watermarking-technology dependent), etc *(Figure 1)*.
The watermark information itself is usually very limited in payload, typically in the order of magnitude of 32 to 64 bits to warrant a robust detection out of a few minutes of typical MPEG compressed video or audio.

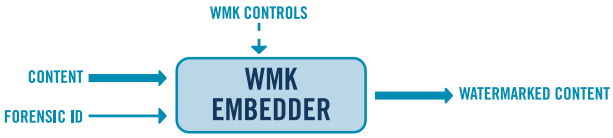

Figure 1 – Watermark embedding in content

Is it possible to embed a watermark in a compressed stream?

In general, a watermark embedder needs to operate on clear content – audio stream or video frame. Depending on the underlying signal processing approach, some video watermarking solutions, as developed for instance by Thomson **[5]** and Philips **[6]**, can also be configured to embed watermark at the MPEG internal representation level (DCT texture coefficients, frame-to-frame motion vectors, displaced frame difference…). However, in the latter case, the embedder still requires to at least partially decode the content (e.g. entropy coding by means of variable length codewords in MPEG-2 video stream syntax), and partially re-encode it, while maintaining a constant bitstream size throughput to the extent possible *(Figure 2)*.

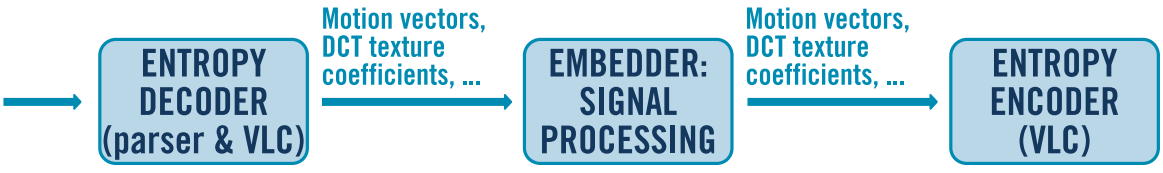

Figure 2 – Conventional watermark embedding in compressed streams requires partial decoding and re-encoding of the bistream: syntax parsing, extraction of compressed features (motion or texture) from Variable Length Codewords, signal processing to identify and mark the compressed signal, and reconstruction of the bitstream now carrying the hidden watermark payload

This conventional watermarking approach consequently raises a number of issues in conventional CA/DRM protected delivery environment, where the content is usually distributed and stored in compressed and scrambled form, as will be further detailed in the remainder of this paper. In particular:

– Forensic tracking requires individual embedding of the watermark, possibly into the STB itself, which adds complexity to the end device and may not always be realistic, for instance in a low consumption mobile receiver.

– Forensic tracking requires embedding of the watermark into the content itself while obviously requiring secure processing of clear content, possibly into the STB itself, which requires careful security design for content workflows.

– Just as any security technology, watermarking must be thought to be renewable. Integrating a specific algorithm into the silicon of the STB chipset to enhance the efficiency and security for instance makes it hard to renew afterwards.

Is it possible to pre-watermark the content?

When watermarking is applied to carry forensic tracking information, it is obviously not possible to embed the forensic identifier in the content before knowing it in the first place. However, it may be possible to pre-process the content in order to ease the actual embedding process afterwards. Such an approach has been specifically developed by CINEA whose watermarking technology splits the actual embedding into 2 stages **[7]** *(Figure 3)*:

1. Signal processing, representing the heavy computational part, is achieved in a pre-processing step once at authoring stage, possibly offline.

2. The watermark payload itself is embedded using a so-called «replacement model» at the place where forensic tracking is needed, typically into the STB. This process has low complexity and can be applied on the compressed stream right after descrambling without the need for partial transcoding. It can therefore also benefit from the descrambler control architectures put in place by the CA/DRM scheme.

The main drawback of this approach is an extra bandwidth requirement to transmit dedicated watermark control streams. In addition, this technology necessitates close integration with the CA/DRM scheme itself as the watermark control streams need to be protected (integrity and confidentiality) along the content distribution chain just as the content keys and rights control streams under CA/DRM protection.

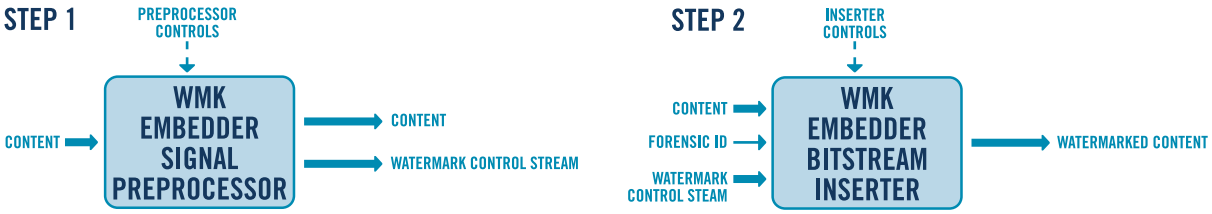

Figure 3 – Workflow of 2-stages replacement watermark model from CINEA

How is the watermark detected?

The forensic tracking process needs to include the search, identification and seizure of potentially illegally leaked content assets. This process needs to be put in place as part of a more global security reactive countermeasures/monitoring framework, to be discussed on a project basis with the service provider and/or content owners requiring the forensic tracking service. On the technology side itself, the seized streams need to be analyzed in search for forensic data (here, the watermark information) by a so-called watermark extractor.

The watermark extractor is meant to remain the highest protected equipment, to be operated at a secure premise under strict access control, to prevent so-called oracle attacks. The latter attacks become popular after they ruined the SDMI watermarking standardization efforts for digital music protection in 2001 **[8]**.

Indeed, without any watermark extractor available to test the watermark attacks efficiency, the attackers will always doubt whether they actually managed to remove or distort the mark, typically by using signal processing attacks **[9][10][11]**. This is a fundamental difference with cryptography, where getting content in the clear means a successful attack.

The watermark extractor may use the original material as input at the expense of extra logistics in order to enhance its robustness (e.g. to efficiently counter geometrical attacks). This specific model is called informed recovery, as opposed to the blind recovery model in which no original material is used by the extractor.
The watermark extractor outputs the retrieved forensics identifier (if any) as well as, in general, a reliability measurement *(Figure 4)*. The more attacks performed on the content, the less reliable the result, but it is expected that for most practical attacks the reliability threshold will remain high enough to validate the retrieval with the expected level of confidence.

Lastly, when applying watermarking to the specific application of forensic tracking, specific attention must be given to the robustness of the watermarking technology to the so-called collusion attacks. The latter consist in merging the outcome from several forensic marker devices, for instance by averaging the pixel values, substituting frames or parts of the video from one device to another, etc, in order to confuse the watermark extractor on which devices originally leaked the content.

Figure 4 – Watermark extractor architecture

3. APPLICATION TO VOD

How can I watermark VOD content?

In the specific case of unicast VOD, it is highly desirable to mark the content at the source (head-end), for the following advantages:

- Legacy, off-the-shelf watermark-agnostic receivers can be used, with neither impact on the watermark end-to-end system security nor on the receiver.
- The watermark technology does not need to be exposed to field attacks. In particular, the watermark embedder remains under strict control by the VOD operator as is the watermark extractor.

The main requirement in that case is that the watermark embedding process needs to be real-time and applicable to a compressed asset, in order to be applied to the content at the source at the time it is delivered to the end user, be it streamed or downloaded *(Figure 5)*.

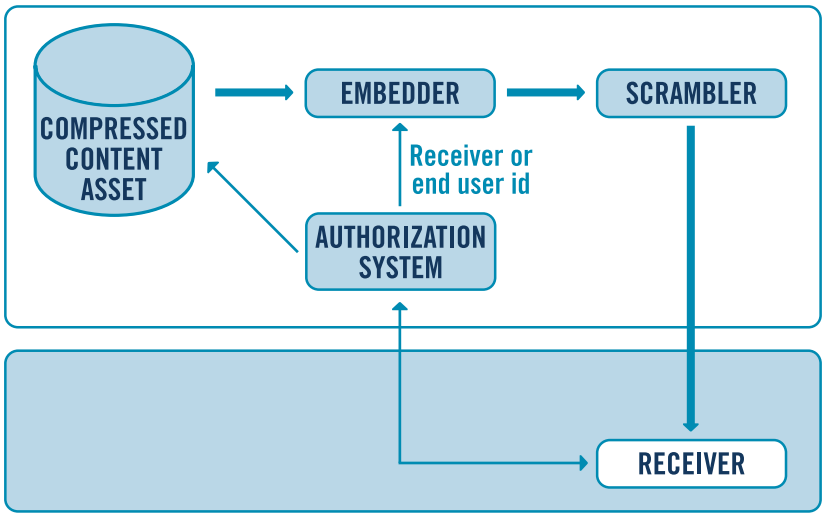

Figure 5 – Forensic tracking embedding in a VOD architecture with on-the-fly content encryption

Is forensic watermarking compatible with pre-encrypted VOD operation?

The majority of CA/DRM protected VOD solutions are based on content pre-encryption today for a number of operational and cost-efficiency issues, such as the scalability of the VOD server. It is therefore counterproductive to envisage

the individual marking of content on the fly when it is served, as this would imply steps of (at least partial) decryption, decoding, re-encoding and re-encryption, thus annihilating the benefits of pre-encryption architecture optimization.

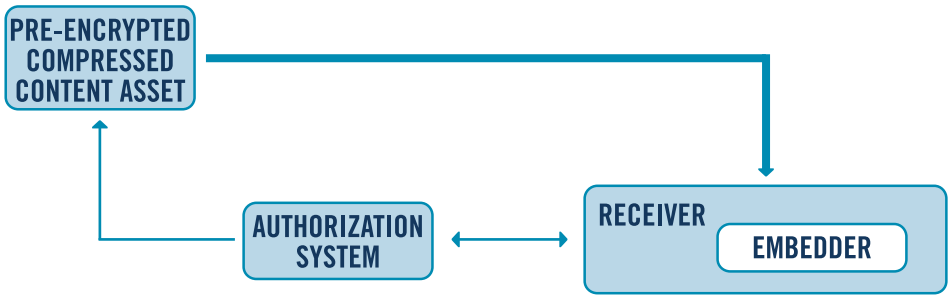

Figure 6 – Forensic tracking embedding in a VOD architecture with content pre-encryption

When pre-encrypted VOD is architected, the content has therefore to be marked by the receiver just as in the broadcast application *(Figure 6)*.

An alternative may be the use of partial pre-encryption, where part of the content would be kept in the clear and marked on the fly at the VOD transaction time, while the other part of the content would be pre-encrypted and left mark-free. However:

- Such an approach will be strongly watermarking technology-dependent, depending on which part of the content signal is processed by a specific technology.
- Depending on the content format, signaling which parts are kept in the clear as opposed to scrambled chunks may be tricky. For instance, in the MPEG-2 transport stream format, the finest granularity of scrambling on/off signaling is at the MPEG-2 packet level.

Is forensic watermarking compatible with pushVOD?

In pushVOD applications, the content is primarily pre-encrypted and broadcast to the whole set of decoders, so it cannot be marked at the source. The applicability of forensic watermarking to broadcast TV content, including pushVOD content, is discussed in details in the next section.

4. APPLICATION TO BROADCAST TV

How can a forensic watermark be embedded in a STB?

In a broadcast delivery framework, it is not possible to embed forensics prior to content delivery as this would require to individualize, thus unicast, the content streams. Therefore, forensic watermarking needs to be embedded into the content itself by the STB receiver *(Figure 7)* – either the audio or the video streams (or both) can be marked. Only video will be discussed here; similar principles are applicable to audio watermarking.

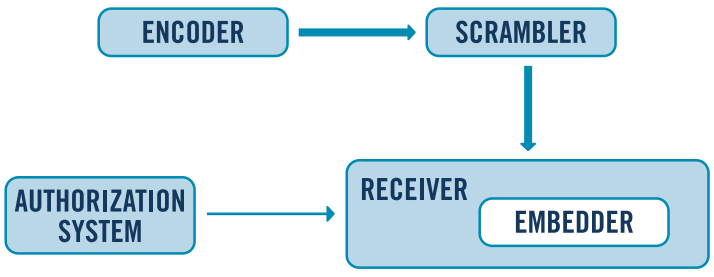

Figure 7 – Forensic tracking embedding in broadcast architecture

When the full watermarking processing takes place in the STB, it needs to operate on partially or fully decoded (depending on the technology) video stream. The broadcast descrambling, the transport stream demultiplexing and the initial video decoding stage (variable length decoding and signal features extraction, such as texture or motion representation in compressed video) need to be passed through before the watermark can be actually hidden into the video signal itself. Therefore, most conventional watermarking schemes have to be embedded into the decoder silicon - which requires a dedicated, watermark-embedding-capable chipset *(Figure 8)*.

Figure 8 – Conventional watermark embedding technologies integration in the STB

Alternately, the watermark embedding may be inserted on the video frames themselves once they have been decoded, for instance by a dedicated chip in addition to the conventional chipset, or in the content handling area of some advanced chipsets *(Figure 9)*. However, this approach only makes sense when the decompressed outputs (analogue, DVI, HDMI) need to be protected. Whenever the content stream is processed, recorded or redistributed in compressed form, this method of watermark embedding requires an additional stage of decoding and re-encoding the video stream, which is obviously too cumbersome to be applicable even in a high-end STB.

Figure 9 – Conventional watermark embedding at picture level is not applicable to compressed content flows

Lastly, another way to perform efficient watermark insertion in the STB consists in applying of a two-stages watermarking embedding process, one stage being pre-processing on the head-end and only the second stage of actual insertion of the marked data taking place in the STB. This approach was chosen in particular by CINEA. In the latter case *(Figure 10)*, the insertion can apply to compressed streams (yet still in the clear, hence after the broadcast descrambler).

Figure 10 – Replacement model [7] embedding is applicable at descrambler level in a STB

What kind of forensic watermark can be embedded in a STB?

In a conventional STB supporting a CA solution, depending on the forensic tracking goals or STB architecture constraints, the following identifiers may be traced:

- STB identifier
- Smartcard identifier

The smartcard identifier can be linked to the end user thanks to the operator SMS.
The STB identifier used in CA can be linked to the smartcard, hence to the end user, when the controlling CA applies pairing of the smartcard to the STB.

In addition, depending on the CA countermeasures capability, either the STB identifier or the smartcard identifier may be used to apply dedicated countermeasures to «leaking» STB.

How can a forensic watermarking embedder be secured in a STB?

The security of forensic watermarking embedding in a STB is a key issue requiring careful STB security architecture. In particular, the following threats shall be addressed:

- It shall not be possible to bypass the embedder.
- It shall not be possible to falsify the forensic watermark payload (identifier).
- When extra watermarking controls are needed, they shall be protected against user access and modification.

The required expertise to this end is in fact very close to the one required to secure the CA operation inside a STB as developed by Nagravision **[12]**.

Is forensic watermarking compatible with the use of a CA descrambling module?

According to the DVB-CI or CableCard specifications, a CA descrambling PCMCIA module only implements the broadcast descrambler, and in the CableCard case, a local rescrambler, while the receiver implements the decoder. The module has limited CPU capability and no access to the decoding routines which are implemented into the host STB chipset itself. Therefore, full watermarking insertion processing is not applicable.

The only suitable technology may be the two-stages approach developed by CINEA, where pre-processing would be computed at the head-end (if applicable on-the-fly to a real-time broadcast environment) while the actual embedding would take place into the CAM *(Figure 11)*.

The advantage of this approach is that CA modules like the ones manufactured by the Kudelski group affiliate SmarDTV are under strict CA control and can therefore benefit from a closer CA security integration into the descrambler chip running into the CAM itself.

Moreover, in the absence of local rescrambling as in the DVB standard, forensic tracking adds the ability to trace possibly leaking content from the clear CI or unsecured standard host, with no impact on the latter performance or architecture.
The drawback of this approach is that it requires extra security bandwidth over-the-air to carry the embedder control data.

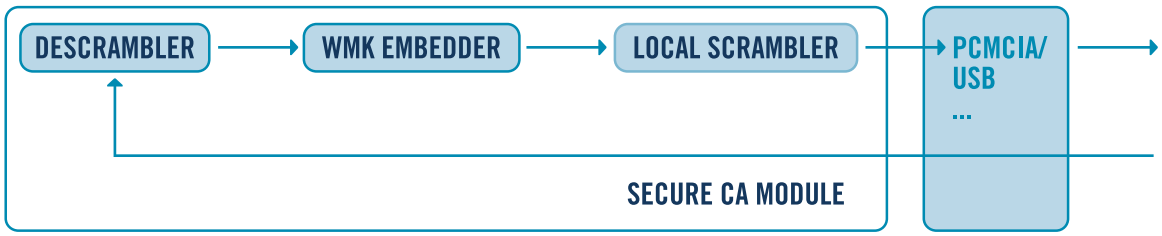

Figure 11 – Application of replacement model [7] embedding to a CA module

Is forensic watermarking compatible with recording on a PVR?

When recording on the PVR is done after the content has been broadcast descrambled, forensic watermarking may be applied before the content is recorded, typically under a local re-scrambling protection.
However, depending on the actual watermarking technology, specific chipset security architectures may be required to enable partial decoding and insertion of the forensic watermark into the video stream between the broadcast descrambler and the local re-scrambler *(Figure 9, Figure 10)*.

Is forensic watermarking compatible with pushVOD?

PushVOD content is recorded on the hard-disk drive prior to purchase, hence prior to broadcast descrambling. Therefore forensic watermarking can only take place at purchase time, when the content is actually descrambled and decoded.

Is forensic watermarking compatible with export to a Portable Media Player?

If the Portable Media Player records the content as a copy transferred from the STB PVR, the most advisable approach consists in forensic marking the PVR copy in the first place (at recording time), so that the content source is already marked.
When trans-coding, trans-rating or downsizing is required for the STB to export the content onto the PMP, the forensic mark shall be robust to those transforms, so the PMP export is agnostic to it.

Is forensic watermarking compatible with export to a home network?

If the STB exports the content from a copy recorded on the STB PVR, the most advisable approach consists in forensic marking the PVR copy in the first place (at recording time), so that the exported content source is already marked. When trans-coding, trans-rating or downsizing is required for the STB to export the content towards a home network output, the forensic mark shall be robust to those transforms, so the home network devices are agnostic to it.
Conversely, if the STB exports the content directly from the live feed, there are 2 cases to consider:

- The content remains under CA control (shared tuner use case): the original broadcast scrambling is not removed, so no watermark insertion can take place in the source STB. Forensic watermarking can only be embedded into the slave STB.

- The content is «bridged» from CA control to CP or DRM control: the original broadcast scrambling is replaced by a local scrambling matching the CP/DRM technology. Forensic watermarking can take place in between, yet very careful attention shall be given to the bridge security design, so that the content is not exposed in the clear and the forensic watermarking embedding takes place securely.

Note that if the forensic watermarking can only be embedded after decoding *(Figure 9)*, the direct export use case described above would require to go through descrambling, decoding and re-encoding, and re-scrambling of the stream before it can be exported, which makes no architectural sense.

5. REFERENCES

[1] "Pragmatic Content Security as a Business Model Enabler", Nagravision White Paper, IBC2007.

[2] http://www.cnn.com/2004/SHOWBIZ/01/23/oscar.arrest/index.html

[3] http://www.dcimovies.com/DCI_DCinema_System_Spec_v1_1.pdf

[4] http://www.reuters.com/article/internetNews/idUSN0445675320070505?pageNumber=1

[5] http://www.thomson.net/EN/Home/Press/Press+Details.htm?PressReleaseID=bc673dc4-0eee-4183-b03a-616bcbd95de5

[6] http://www.business-sites.philips.com/contentidentification/Products/Index.html

[7] http://www.cinea.com/whitepapers.html#

[8] http://www.theregister.co.uk/extra/sdmi-attack.htm

[9] http://www.petitcolas.net/fabien/watermarking/stirmark/

[10] http://www.certimark.org/

[11] http://www.ebu.ch/en/technical/trev/trev_286-cheveau.pdf

[12] http://www.nagravision.com/pdf/NV_TechSeries_STBsecurity.pdf

For more information on this WhitePaper, please contact one of the authors:

Corinne Le Buhan Jordan
PhD, Head of New Technology
Nagravision
corinne.lebuhan@nagra.com

Daniel Rufle
Marketing Manager
Nagravision
daniel.rufle@nagra.com

Philippe Stransky
Chief Technical Officer
Nagravision
philippe.stransky@nagra.com

Ivan Verbesselt
Senior VP, Head of Strategic Marketing
Nagravision
ivan.verbesselt@nagra.com

SECURE ACCESS. CREATE SUCCESS.

www.nagra.com

NAGRAVISION
KUDELSKI GROUP