



RHOZET.

Content Control: Digital Watermarking and Fingerprinting

White Paper

2008

Table of Contents

| | |
|--|-----------|
| <u>Classic DRM</u> | <u>3</u> |
| <u>Watermarking vs. Fingerprinting</u> | <u>3</u> |
| <u>Watermarking</u> | <u>4</u> |
| <u>Watermarking Applications</u> | <u>6</u> |
| <u>Fingerprinting</u> | <u>6</u> |
| <u>Fingerprinting Applications</u> | <u>8</u> |
| <u>Watermarking and Fingerprinting in the Transcoding Workflow</u> | <u>9</u> |
| <u>Fingerprinting in Action: YouTube</u> | <u>11</u> |
| <u>Summary</u> | <u>11</u> |
| <u>Digital Watermarking and Fingerprinting Solutions Providers</u> | <u>12</u> |
| <u>References</u> | <u>13</u> |

Content producers and publishers are continually seeking more reliable ways to control access to their valuable media assets while, at the same time, making them available to generate revenue. Meanwhile, professional and casual content pirates are becoming ever more resourceful.

This white paper provides a high level overview of digital watermarking and fingerprinting and examines how these two technologies can be integrated into workflows for automatically tracking, protecting and monetizing content.

CLASSIC DRM

The classic digital rights management (DRM) approach, which attempts to authorize a specific user to watch a specific video on a specific device, can be very effective for some situations, such as allowing an account owner to rent a particular movie and watch it on their computer during a designated time period.

There are many situations, however, where classic DRM is neither appropriate nor effective — situations in which two emerging technologies come into play: digital watermarking and fingerprinting. These technologies promise to facilitate critical business functions, such as:

- Content Identification
- Copyright Control
- Copy Protection
- Viewer Tracking
- Forensics

WATERMARKING VERSUS FINGERPRINTING: WHAT'S THE DIFFERENCE?

Sometimes you will hear the terms “watermarking” and “fingerprinting” used interchangeably. While both the watermark and fingerprint uniquely identify a particular piece of video, they are very different in both purpose and execution.

A watermark is like a tattoo, permanently added to every frame of the digital media file. Watermarks may be visible to viewers (perceptible) or invisible (imperceptible).

Fingerprinting digital media works very much like fingerprinting people, relying on innate characteristics of the subject. Human fingerprints are characterized by patterns of loops or whorls; digital media fingerprints are comprised of clues such as audio waveforms and/or video characteristics. In both cases, a database of fingerprints must be maintained, against which to compare fingerprints “found in the field,” and technology is needed to rapidly match fingerprint “evidence” to fingerprints on file.

Many identical versions of the same piece of video can be created, each with its own unique watermark. If you sent those individually watermarked videos to 10 different people and one of them was illicitly uploaded to the web, the watermark would tell you which copy was posted and could, therefore, be traced back to the person responsible.

Video fingerprints are stored in a database. Any video clip can be compared with fingerprints on file, to see if there is a match. Imagine that you are running a user-generated video website. You want users to be able to upload their own videos, but you don't want to publish movies pirated from major content producers (for example Disney, HBO or CBS), because you don't want to be sued. So when a user submits a video to your site, you generate its fingerprint and send it to a service that compares it to the fingerprints of thousands of copyrighted movies on file. The service responds by telling you whether or not you should publish the video. And, of course, all of this is fully automated.

WATERMARKING

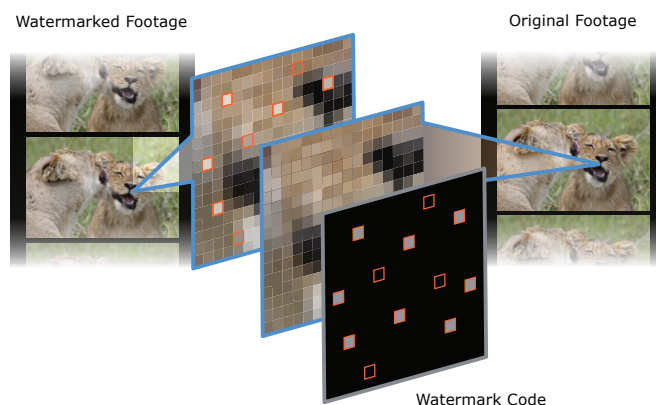
Watermarks are useful for:

- Tracking individual assets
- Helping to identify who created a particular piece of content
- Determining whether content was obtained by legitimate means

Visible watermarks, such as network logos or station IDs in the lower third, are perceptible — useful as branding tools. Content protection, on the other hand, relies on invisible, or imperceptible watermarks, in which the original source file and its watermarked counterpart are visually indistinguishable from one another. You could make a hundred copies of the same video and, with imperceptible watermarks, uniquely identify each. Anyone just viewing the videos would not be able to tell that they were different, but your watermarking system could identify each of them individually.

So how do imperceptible watermarks work? All watermarking techniques are examples of steganography — the process of hiding secret messages in a document or picture so that only the sender and intended recipient can recognize the message and, more importantly, know how to retrieve it. The classic example is the spy who sends a long letter to a “friend,” describing her travels. The document appears innocuous to anyone intercepting it. But the recipient knows that every 10th letter in the document spells out a secret message. Similarly, watermarking subtly modifies each image in a video so that no one can tell it has been altered.

Just like the example of the letter, a very simple way to watermark images would be to either increase the brightness of every 100th pixel by a small amount or leave it alone, depending on whether we wanted to encode a “1” or a “0.” The recipient of our watermarked image could simply subtract the original image from the watermarked image and would be left with a series of grey dots on a black background. The grey dots would be our hidden binary message. This is an extremely simplified example; actual methods currently in use are much more complex.

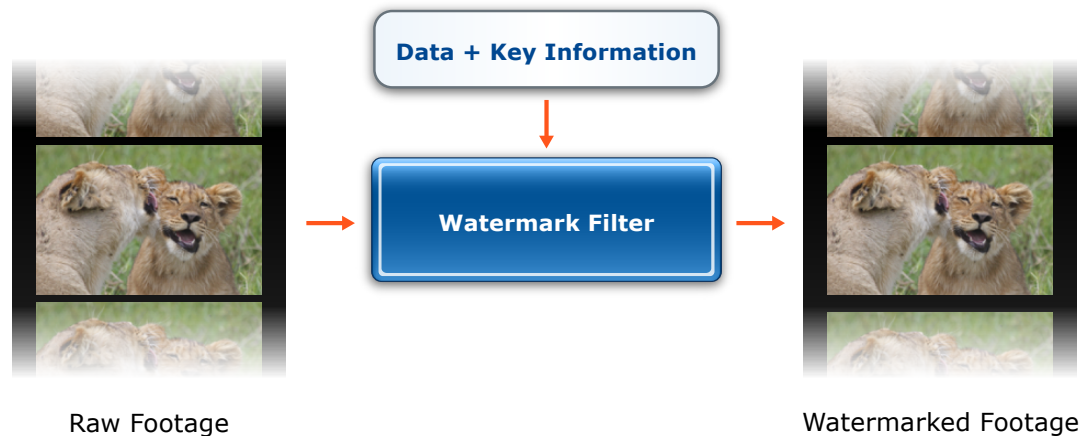


Many current digital watermarking methods embed codes so that the image can be altered (transcoded, cropped, scaled, etc.) without losing the ability to extract the watermark. Current methods also do not require that you have the original media for comparison, in order to extract the watermark data. Many techniques are similar to those used in compression technologies. The watermark ends up as very subtle color variations in the final image. If you don't have the mathematical "key" revealing where the watermark data is hidden, you cannot find it.

Watermarking systems and techniques are not generic or standardized — a watermark generated by one technology cannot be read by a system using a different technology. And even when two systems use the exact same technology, one customer would not be able to read another's watermarks without the secret "key" that reveals where to find the watermark and how to decode it.

Regardless of how complicated the math, the basic process of adding a watermark is fairly simple. The watermark is typically executed as a "filter" applied to an uncompressed video frame, resulting in an uncompressed frame that contains the embedded information. The watermarking filter must be programmed with the data to be embedded, as well as with the "key" that enables that data to be hidden.

Since digital watermarking is usually performed on uncompressed frames, it is typically carried out as part of a transcoding process. The transcoder demultiplexes the video, decodes it into uncompressed frames, feeds the frames through the watermarking filter, then compresses the resulting frames and multiplexes them into the final format or formats.



While there are some watermarking techniques that can be performed directly on compressed video, they are most often limited to maintaining the same compressed format (MPEG-2, for example).

For a watermark to be useful, there must be a way to extract it and compare it with known watermarks. Some resources that provide watermarking technology also provide tracking services, which analyze content on the web and search for watermarked content. Other resources provide just the technology, assuming that only the customer will embed and detect watermarks. Your requirements and budget will dictate the type of watermarking you do.

WATERMARKING APPLICATIONS

Content Control. The ability to trace the source of leaked classified information, proprietary corporate research or unauthorized copies of movies is one of the primary benefits that digital watermarking offers content owners. Companies are able to keep tabs on confidential recordings and videos. If something secret is leaked, the person responsible could be traced using digital watermarks.

Digital watermarks can also be used to monitor distribution of sensitive material. For example, studios can use watermarks to track whomever has access to unreleased work prints and dailies. If someone makes and distributes illegal copies, they can be traced. During the Oscar season in Hollywood, for example, preview copies of new movies are individually watermarked before they are sent to voting members of the Academy of Motion Pictures. Anyone who leaks a movie can be prosecuted.

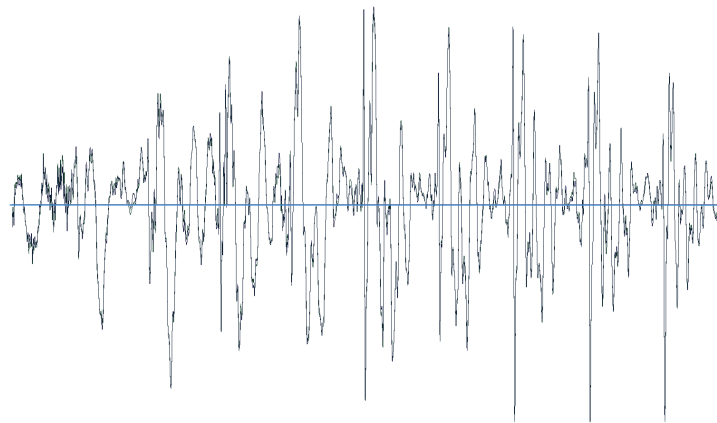
Digital Cinema. Anti-piracy is a major concern for digital cinema. If films were sent to theaters in perfect digital form, they could be easily copied and illegally distributed. So digital watermarks identify each individual copy released, facilitating tracking and piracy prevention.

Content Identification. By embedding watermarks in footage, content creators can identify the sources of specific content. For example, news organizations can determine whether footage retrieved from an archive came from a staff reporter, a freelancer or a third party.

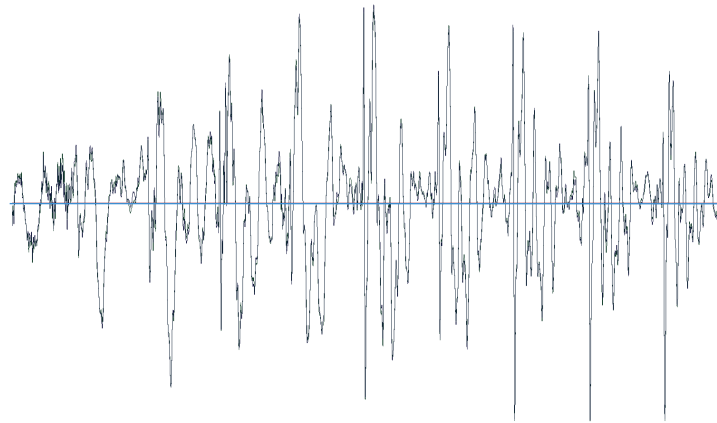
FINGERPRINTING

In purpose, value and execution, fingerprinting is quite different from watermarking. Watermarking relies on embedding information into the video and/or audio, then uses that information to identify the piece of content. Fingerprinting does not embed any information; it analyzes the video and/or audio to determine the unique characteristics of the content. The identified pattern is stored in a database and can be used for recognizing the content in the future.

The graphic below represents the audio waveform for a specific piece of music.

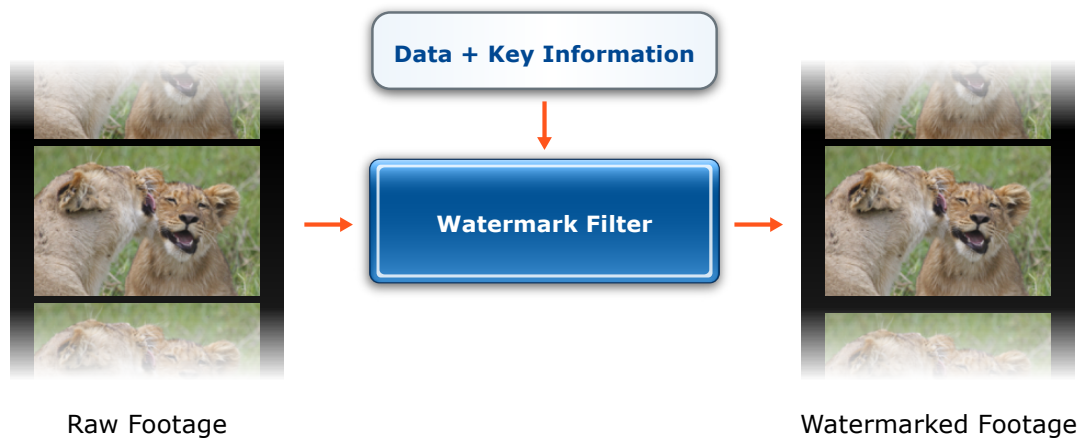


A fingerprinting system needs to create a database that can be used to identify an exact piece of content, should the system ever encounter the same piece of content again. One way of doing that would be to store the entire piece of music in the database and then compare every new piece encountered with the entire original. This would work, but it would be slow, and the database would become very large. Imagine doing the same thing with video — the database would be unmanageably large! So instead of storing the entire piece of audio, only a statistical sample is stored. For example, four samples could be taken every 10th of a second, as shown below.



Now imagine that sampling continues in this manner through the entire audio waveform. If the source file is a two-minute song, there would be over 4,000 samples in the fingerprint, but the actual fingerprint file would still be a thousand times smaller than the original audio file. Why is it so small? The original audio might have had a sampling frequency of 48kHz (a common standard), which means 48,000 samples per second. The fingerprint is only being sampled 40 times each second (four samples every 1/10th of a second), which means there is a thousand times less data. These samples comprise a unique signature that corresponds to the specific piece of music. Any other piece of music would have a different pattern of samples. Just as human fingerprints don't reveal what a person looks like, the samples don't really indicate what the music sounds like. They just allow it to be uniquely identified.

The example above is merely a simplified example of how audio fingerprinting might work. Modern fingerprinting algorithms actually sample a wide variety of audio and/or video characteristics to produce the fingerprint. To thwart pirates, sophisticated techniques are used to hide the exact nature of the properties on which fingerprints are based. For example, random temporal and spatial sampling is used to hide information about which properties are used to produce the fingerprint.



A fingerprint is generated from a series of uncompressed frames. The fingerprint can incorporate metadata about the media, along with the fingerprint pattern.

Unlike watermarking, where application of the watermark filter produces a file of uncompressed video frames (which still yields a relatively large file even after compression), the fingerprint exporter does not generate a viewable media file at all — it creates a much smaller fingerprint file, documenting inherent characteristics of the media content.

Video fingerprints are resolution and format independent. They can be used to identify complete videos, portions of videos and, more important, very short snippets of videos. Fingerprints can even be used to spot video content that has been manipulated, as it might be if included in a video mash-up.

One of the ways pirates steal films is by secreting a video camera into a theater and illegally taping the movie. Video fingerprints cannot be defeated in this manner.

Audio fingerprints, unfortunately, aren't as resistant to manipulation as video fingerprints. It's relatively easy to alter or completely replace audio tracks. For example, completely new soundtracks usually accompany video mash-ups, or pirated films might be shown in foreign countries with new dialogue tracks. So identifying content based solely on an audio fingerprint can be problematic.

FINGERPRINTING APPLICATIONS

Broadcast and General Media Monitoring. Fingerprints can be used to track when and where a video has been shown. This capability will be useful to advertising agencies and their clients who want to monitor their media activities. Content syndicators will use it to track when and where programming has appeared. Talent agencies will use it to monitor activities for which their clients are owed performance royalties. Organizations such as the AP and Reuters will monitor how their content is being used by broadcast and online news operations, as well as bloggers.

Copyright Control. Cross-referencing actual usage rights and permissions in a fingerprint database will facilitate monitoring of authorized and unauthorized uses of content. Content producers and distributors will use fingerprints to determine whether a database contains unauthorized content. Stock footage providers will use fingerprints to spot the clips they license in commercial programming.

Metadata. Metadata allows content creators to store all sorts of useful tracking information associated with content. For example, metadata can be used to supplement video fingerprints by storing vital information about users (such as who created the content and who has modified it), a history of use, which operating systems a video has been played on and by what version of which player technology, which networks it has traveled on and more. This information will be invaluable during forensic investigations to trace pirating operations.

Gracenote's TuneUp Companion for iTunes is one of the first applications to combine metadata with fingerprinting. The application's primary use is to clean up metadata associated with recorded music. It will also be used to retrieve album cover art, tour schedules, YouTube videos, news about recording artists and more.

Behavioral Modeling Advertising. Interest-based or behavioral advertising matches ads to individuals based on a user's past online activities, such as visiting a website or searching for information on a particular subject. Video fingerprinting extends that capability to marketers who want to reach consumers based on their video viewing interests. It also brings behavioral advertising models to new domains, such as VOD services and cable television where signals must pass through a set-top box.

Copy Protection. Video fingerprints can be used as a copy protection tool. For example, both a video fingerprint and an authenticating signature could be required before a file could be copied or replicated.

Forensics. Another promise of video fingerprinting is in the area of information forensics where fingerprints could be used to detect whether video footage has been manipulated. Research is ongoing in this area.

Additional Business Opportunities. Digital fingerprints must be matched against extensive content ownership databases to be effective. It stands to reason that maintaining, licensing and managing access to large-scale finger-print databases is a potential revenue opportunity. Audible Magic, for example, has already established a subscription-based business leveraging its extensive ownership database.

WATERMARKING AND FINGERPRINTING IN THE TRANSCODING WORKFLOW

When you add the enormous accumulation of video content that resides in the archives of media organizations to all of the new content that's constantly being produced, you realize that sheer volume poses a significant challenge to utilizing both watermarking and fingerprinting technologies.

At the very minimum, implementing an anti-piracy solution around watermarking requires the following capabilities:

- Method to embed watermarks
- Database to track watermarks
- Method to detect watermarks

For watermarking, consider whether to handle detection in-house or whether to employ an outside service. The choice depends on how the watermark will be used. For example, to trace a single video back to its source should it be leaked, detection could be easily handled in-house. If your preference is to let someone else hunt down leaks, perhaps because they could come from a large number of unauthorized sources, consider an external detection service.

Fingerprinting requires the following capabilities:

- Method to generate fingerprints
- Database to store metadata relating fingerprints to originals
- Third-party service (or multiple services) that tracks fingerprints and provides access control information

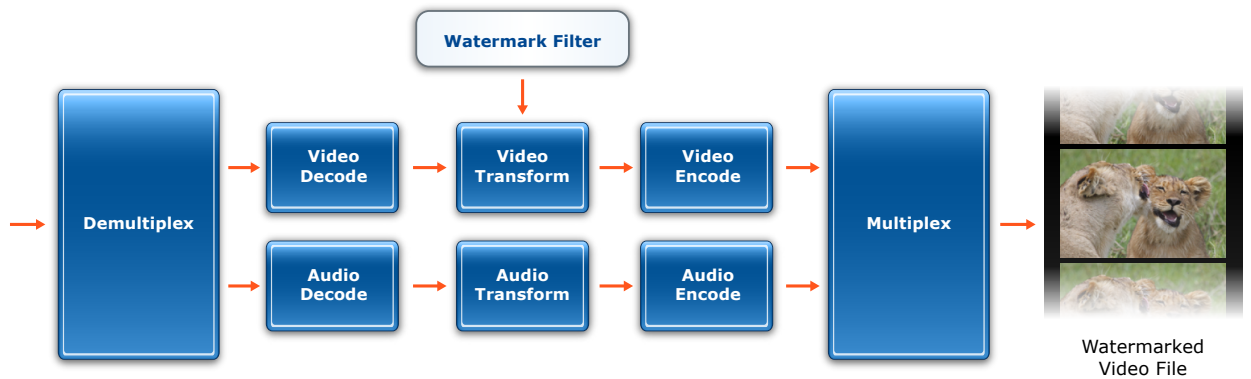
One of the problems with current watermarking and fingerprinting technologies is that they only accept a limited number of input formats. And, in the case of watermarking, they only generate a limited set of output formats. This is where Rhozet™ and transcoding come into the picture. Rhozet Carbon Coder™ is a general-purpose transcoder that handles dozens of different format types. Plugging a particular watermarking or fingerprinting technology into Carbon Coder allows you to handle any media type. Carbon Coder can be run as a stand-alone transcoding engine or as part of a larger transcoding farm for higher volume workflows.

Transcoding a media file from one video format to another involves a number of steps:

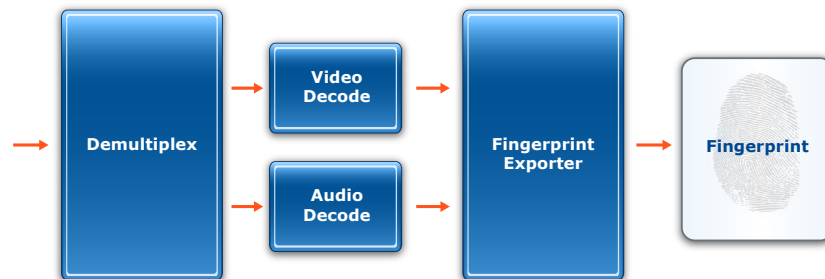
- Demultiplexing the original into its constituent audio and video components
- Decoding the audio and video using the appropriate codecs
- Transforming the audio and video (e.g., frame size, frame rate, graphic overlays, filters, etc.)
- Encoding the audio and video with the target codecs
- Multiplexing the audio and video into the final wrapper (e.g., MPEG-2 Transport Stream)

For the fastest execution, all of this occurs on-the-fly, in memory.

For watermarking, the watermarking filter is plugged into the Carbon Coder pipeline and is applied during the transform stage. The result is a media file, in whatever format is required for distribution.



For fingerprinting, the fingerprint technology is embedded into Carbon Coder as an exporter. It analyzes the uncompressed audio and video frames and generates a fingerprint file, which can be used to recognize the original media in whatever format it is found. There are no transform, encode or multiplex steps, because the output is only a fingerprint; not a media file.



Many watermarking and fingerprinting companies are now integrating their technology with Rhozet transcoding systems to take advantage of the Rhozet transcoding workflow. You can actually have multiple technologies from different vendors all plugged into the same Carbon Coder system. This allows you to use a single transcoding network to create and detect a variety of watermarks and fingerprints.

FINGERPRINTING IN ACTION: YOUTUBE

YouTube provides an excellent example of a real-world fingerprinting application. Google has created a sophisticated mechanism for helping to control piracy on the YouTube site. YouTube allows content owners to fingerprint their original media and then upload the fingerprints to YouTube's central database. Every video that is uploaded to YouTube is first compared with the central fingerprint database, before it is displayed.

The content owner can specify usage policies that govern how YouTube should handle matched content, including monetization and blocking. This means that a content owner can share in the advertising revenue from his video, regardless of who actually uploaded it.

If someone uploads a piece of content that its original owner requested be blocked, then YouTube will not allow this video to be shown. This is an extremely effective way to prevent piracy. Rather than having to request the removal of content after it has been posted (which is like shutting the barn door after the horse has left), this technique actually prevents the video from being posted in the first place.

Google and Rhozet have integrated the YouTube fingerprint creation software right into Carbon Coder software. This free-of-charge component allows content owners to create fingerprints as part of their standard workflows. Thus fingerprinting can happen much earlier in the production pipeline. For example, a customer could fingerprint dailies of a movie 18 months before its release, in order to guarantee that no leaks get shown on YouTube.

YouTube fingerprinting is unique to the YouTube website. Blocking something on YouTube does not make it automatically blocked on some other site. But since YouTube accounts for approximately 40% of all video viewed on the web, its fingerprinting system provides a crucial component in the fight against piracy.

SUMMARY

Watermarking and fingerprinting are exciting technologies that can facilitate critical business functions, including content identification, copyright control, behavior tracking, copy protection and forensics. In addition, these technologies provide the means for content owners who embrace emerging distribution platforms, such as user generated content websites, to develop new revenue models based on what they previously might have considered pirated content.

Rhozet does not sell these anti-piracy technologies, but instead partners with their developers to allow them to be integrated into Rhozet transcoding solutions.

WATERMARKING AND FINGERPRINTING SOLUTIONS PROVIDERS

Audible Magic

<http://www.audiblemagic.com/products-services/contentsvcs/>

Audible Magic provides sophisticated content identification services that allow content owners to track and manage their copyrighted work in electronic form. Using patented CopySense identification techniques, the service recognizes content based on digital fingerprints. The approach is highly accurate and does not rely on metadata, watermarks or file hashes. The technology is immune to compression and distortion as well as codec choice, file type and streaming format. An efficient and compact API library facilitates easy implementation of CopySense technology. Audible Magic maintains a growing ownership database that currently includes over 6 million works.

Civolution

<http://www.civolution.com/>

Civolution offers an extensive portfolio of digital fingerprinting and watermarking technology solutions for IP licensing, Pay TV, forensic tracking and digital cinema applications. Civolution uses fingerprinting technology to automatically identify multimedia content. Fingerprints are extracted and stored in a database. Content fingerprints are matched against the contents of the database. Civolution watermarking technology embeds a unique, imperceptible watermark into video or audio material that can be used by content owners to identify and track their material. Civolution video watermarking technology was used to identify the source of illegal copies of the 2003, 2004 and 2005 Academy Award Screeners.

Fraunhofer

<http://www.sit.fraunhofer.de/EN/forschungsbereich/tad/index.jsp>

The Fraunhofer Institute for Secure Information Technology (SIT) researches, develops and deploys security solutions for IT environments. Along with a number of related initiatives, their Transaction and Document Security department is engaged in digital watermarking and cryptography. Fraunhofer is responsible for developing the MP3 codec.

YouTube

http://www.youtube.com/t/video_id_about?gl=GB&hl=en-GB

YouTube Video Identification, developed in conjunction with Google, uses video fingerprinting to recognize content by comparing its fingerprint against a fingerprint database. YouTube/Google can create the fingerprint when video is submitted to YouTube.com, or content creators can create a fingerprint independently and submit just the fingerprint to YouTube's fingerprint database. Content owners are able to specify rights and permissions that govern the use of their content, whether they themselves submit it to YouTube or users upload it. Content owners can block or promote their content or even derive revenue from it regardless of whether it was uploaded "officially" or by users.

Thomson

http://www.thomson.net/GlobalEnglish/Innovation/Innovation_centers/tracking%20and-security-technologies/Pages/default.aspx

Thomson NexGuard™ watermarking adds an imperceptible and indelible code to video sequences. When NexGuard is used in conjunction with Sapphire VOD servers, a unique watermark is generated each time content is requested, allowing the content to be tracked forensically to determine the source of illegally copied and distributed content. NexGuard has been implemented by studios, post-production facilities and in digital cinema environments as well as in set-top boxes, satellite and IPTV applications.

Vobile

<http://www.vobileinc.com/technology.html>

The Vobile Content Identification Platform is based on patented VideoDNA™ fingerprinting technology that is used to extract a fingerprint without altering the source content. Identifying an unknown video involves extracting its VideoDNA (a fingerprint) and matching it against entries in the Vobile Content Registry (a fingerprint database). The Vobile Content Identification Platform supports a variety of applications including media rights management, targeted advertising, business intelligence, metadata services, asset management and video search and categorization.

REFERENCES**Digital Watermarking**

http://en.wikipedia.org/wiki/Digital_watermarking

Digital Watermarking Alliance

<http://www.digitalwatermarkingalliance.org/>

Digital Video Fingerprinting

http://en.wikipedia.org/wiki/Digital_video_fingerprinting

Forensics

http://www.csl.uiuc.edu/research/info_foren1.asp

New Applications for Music Fingerprinting

<http://www.drmwatch.com/watermarking/article.php/3759456>

IBM Developer Works

<http://www.ibm.com/developerworks/power/library/pa-soc11/index.html>

Collusion-Secure Fingerprinting for Digital Data

<http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.40.2227>

Video Fingerprinting Based on Centroids of Gradient Orientations

http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&arnumber=1660364&isnumber=34758

Collusion-Resistant Fingerprinting for Compressed Multimedia Signals

Exploring QIM based Anti-Collusion Fingerprinting for Multimedia

<http://www.ece.umd.edu/~ashwins/fingerprinting.html>

Collusion-Resistant Video Fingerprinting for Large User Group

www-video.eecs.berkeley.edu/Proceedings/ICIP2006/pdfs/0002301.pdf

History of Steganography and Cryptography

<http://www.petitcolas.net/fabien/steganography/history.html>

Robust Video Watermarking of H.264/AVC

http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4100886

A Novel Scheme for Hybrid Digital Video Watermarking: Approach, Evaluation and Experimentation

<http://ieeexplore.ieee.org/Xplore/login.jsp?url=http%3A%2F%2Fieeexplore.ieee.org%2Fiel5%2F76%2F32993%2F01546010.pdf%3Farnumber%3D1546010&authDecision=-203>