

Technology

White Paper

Contents

About Dremedia	2
Company Background	2
Contacting Dremedia	2
Dremedia Benefits	3
Dremedia Application Platform	5
Processing Back-End Overview	5
Filters	6
DRE/AXE	7
XML Architecture	7
Core Technology	8
Video	8
Audio	8
Text	9
Performance and Scalability	10
Limitations of Alternative Approaches	11
Managing Assets Using Metadata	11
Video Indexing Systems	12
Limitations	12
Dremedia's Approach	12
Keyword Searching Systems	12
Phonetic Speech Recognition Systems	13
Conclusion	14

About Dremedia

Dremedia is one of the driving forces behind the changing face of the broadcast industry. Founded in 2001 by senior executives from the broadcast and technology industries, Dremedia is dedicated to addressing the challenges faced by the broadcast industry in the new digital age. Dremedia has now also applied this experience to the enterprise to enable Global 2000 companies, who deal with large amounts of rich media content on a day-to-day basis, to leverage those same technologies being employed by leading companies in the broadcast industry.

Dremedia products allow broadcasters, media companies and enterprises to improve the efficiency of their operations by increasing productivity, decreasing costs and reducing time-to-air. Dremedia cuts to the heart of these issues by delivering a technology platform that can automatically analyze, understand and manipulate audiovisual content. By forming an understanding of the actual content of audiovisual material, Dremedia products are able to automate the most manually intensive production tasks across all sources of material.

Dremedia's software products are based on robust and patented technologies provided by Autonomy, Softsound and Dremedia, which have years of academic and commercial development behind them. Dremedia has unrivalled access to these technologies and to the research departments who continue to pioneer their development.

Softsound's acclaimed speech recognition technology is fully integrated into all of Dremedia's products. This technology has consistently out-performed other systems for real-time, anonymous speech recognition, as benchmarked in US and European evaluations. The precision and utility of this technology, combined with Autonomy's advanced automation technologies (for concept-based search, retrieval and categorization) and Dremedia's synchronization technologies, delivers unparalleled automation to the production and management of audiovisual content.

Dremedia is a privately held company whose principle investor is Autonomy Corporation, which was founded in 1996 and has offices in Boston, Chicago, Dallas, San Francisco, New York, and Washington, D.C. in the United States, as well as offices throughout EMEA, including Amsterdam, Brussels, Cambridge, Frankfurt, Milan, Paris, Oslo, as well as in the Far East. In July 1998, Autonomy floated as a public company on the Nasdaq Europe exchange (AUTN), The NASDAQ National Market (NASDAQ: AUTN) in May 2000, and on the London Stock Exchange (LSE: AU.) in November 2000.

Contacting Dremedia

Dremedia Limited

63 St James's Street
London SW1A 1LY
United Kingdom
Tel: +44 (0) 20 7970 5670
Fax: +44 (0) 20 7970 5683

Dremedia

301 Howard Street
22nd Floor
San Francisco CA94105
USA
Tel: (415) 243 9955
Fax: (415) 243 9984

E-mail: for information: info@dremedia.com
Website: <http://www.dremedia.com/>

Dremedia Benefits

With the massive increase in the volume of audiovisual content coming into enterprises today, there is a urgent requirement for systems that deal with the key operations automatically – from ingest of source material, right through to transmission or play-out. Always under pressure to improve the efficiency of their operations and to increase the productivity of their employees, some of the world's leading broadcasters have recognized that this challenge can be met with Dremedia software installed in their facility. Now, they have access to the only fully automated, production system available today. This system is now also available to enterprises, who are coming to terms with the same problems that broadcasters are facing up to.

Borne of technology experience gained in broadcast organizations worldwide, Dremedia software utilizes advanced video, speech and text analysis technologies to allow that content to be quickly and easily accessed and manipulated, without first having to rely on costly manual processing and preparation.

Enterprises are now able to achieve real-time completely automated indexing, handling and management of any audio and video material. The dynamically created index consists of a time-ordered set of time-coded statements, describing attributes of the source content. Because these statements are time-ordered and have millisecond-accurate time-codes, they can be used to manipulate the source material trans-modally - i.e. allowing the editing of the video, by synchronistically manipulating the text, video and audio components.

Dremedia software produces frame-accurate output marked up with:

- speech
- speaker IDs
- accent
- language
- shot-change
- fade (first frame time-code, last frame time-code)
- zoom (first frame time-code, last frame time-code, direction)
- pan (first frame time-code, last frame time-code, direction)
- searchable image (i.e. shot) attributes

By having the above information automatically available, enterprises now have the unique opportunity to enhance existing and develop new applications in areas that would previously have been too costly to pursue.

Dremedia's products facilitate, automate and enhance process in the following areas:

- Asset Management
- Newsroom Systems
- Post Production Systems
- Film & Video Editing
- Interactive TV
- Syndication
- Content Creation
- Content Repurposing
- Media Monitoring

By utilizing the Dremedia technology the following unique applications can be achieved.

Integrated within an AV edit suite, it can allow an editor to:

- jump to particular words
- edit a clip by selecting text, speaker or image
- jump to next speaker
- jump to next instance of current speaker
- search for named speaker
- search on accent or language
- view key-frame of shot
- extract pans, fades etc
- find visually similar material

In real-time multi-media production, the system can automate:

- the association of hyperlinked documents with real-time AV
- instant cross-referencing of live material with archived material
- triggering of events by attribute (e.g. show name when speaker X is talking)

For AV archives, the system can:

- provide automatic categorization of live material
- automatically re-categorize multiple archives
- make archives searchable from any production system
- enable advanced concept-based retrieval as well as traditional keyword or Boolean methods
- automatically aggregate multiple archives
- automatically extract and append meta data

For Media Monitoring, the system can:

- automatically monitor any number of feeds across TV, radio and print
- provide real-time alerts via email, SMS or pager about subject-matter, speaker or individual words
- provide regular (daily, weekly or monthly) reports as well as statistical representations
- automatically archive content and make it searchable

Dremedia Application Platform

Processing Back-End Overview

The function of the processing back-end, on an AV asset or stream, is to produce an XML index containing a rich variety of AV metadata. This index may include:

- a transcript of the spoken language
- speaker segmentation and/or identification
- language, accent, dialect identification
- shot change information
- image association signatures
- any other auxiliary data available from legacy systems or source meta-data
- once processed, the XML can be written to a file system, or automatically indexed into the Autonomy XML engine (AXE™)

A highly modular approach has been adopted for the design of the XML and the processing system that creates it. This modular approach extends the front-end Graphical User Interface (GUI) components. Generally there is a one-to-one relationship between the marked up fragments contained in the XML output, the processing modules (filters) contained within the Dremedia back-end and the GUI components.

Back End Topology

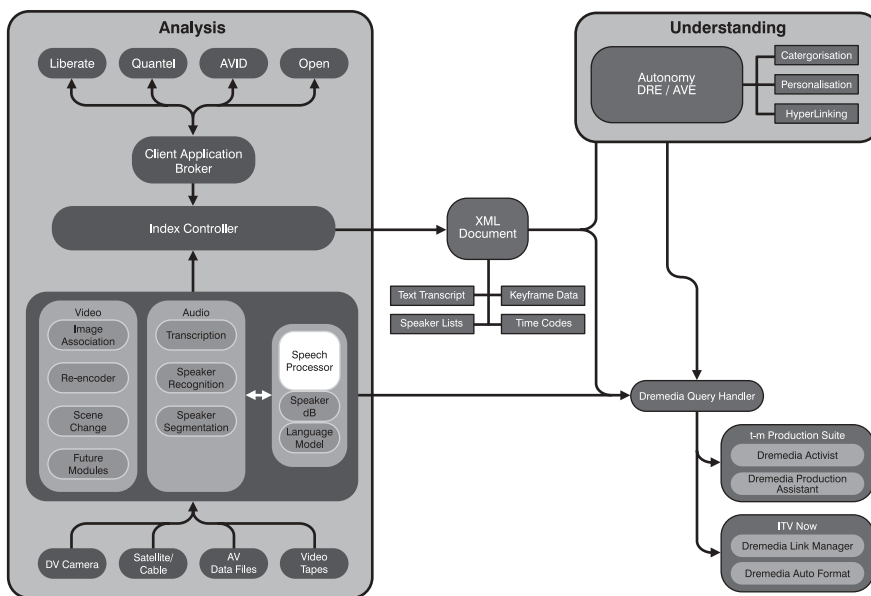


Figure 1 : Dremedia Architecture

Index Controller

The Index Controller is an HTTP- controlled server that coordinates AV processing jobs. It performs no processing itself, but instead distributes the processing load over a number of slave Filter-Boxes. The Index Controller is responsible for listening for processing requests and passing them on to the Filter-Boxes. It then manages the XML results they each stream back. It also takes care of control commands that enable the management of processing jobs and enquiries into job statuses.

There are a number of ways of issuing processing jobs requests and monitoring their progress. At the highest level there is the Indexing GUI, which is suited to one-off processing requests, since it requires the manual entry of various processing parameters. At a lower level there are scriptable command line utilities that enable the more technical user to run batch-processing jobs. At a lower level still, there are the HTTP, C, C++ and COM APIs that enable the application programmer to control the processing back-end directly and programmatically. In keeping with current market trends there is also a SOAP interface, which enables the control of the system from any platform/language for which a SOAP API has been written.

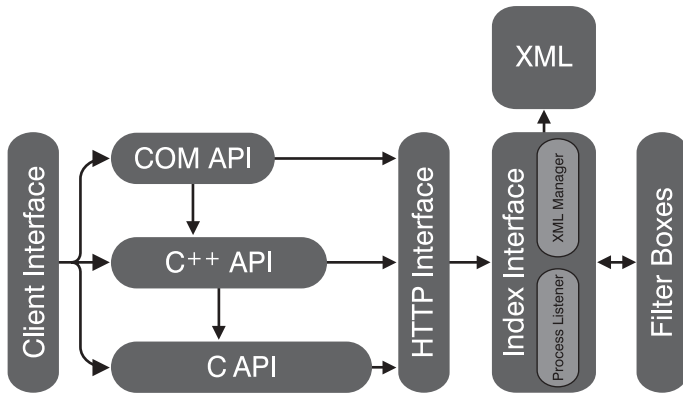


Figure 2 : Index Controller Architecture

Because the Index Controller is controlled using HTTP commands, processing requests can be issued via an HTTP proxy server and so can cross network boundaries even with a firewall in place.

Once a processing job is complete the Index Controller can perform one or both of the following actions:

- index the results into AXE™
- save the results into a local directory where it can be accessed by external processes

Filter-Box

The Filter-Box is a slave process to the Index Controller, and is the functional unit that performs the AV processing. Multiple Filter-Boxes can be installed on different physical machines and the Index Controller configured to distribute processing requests to each of them. A large enough cluster of these processing machines can ensure real-time or better processing rates even when the processing load is very high. Each filter performs a different processing task on either the audio or video stream. Alternatively, a filter may accept as its input the output of another filter.

The processing tasks performed by an individual Filter-Box are determined by the plug-in filters configured. This modular architecture allows filters to be added as new ones become available or are required/requested by clients.

Filters

Dremedia Speech Services Filter

This filter extracts the audio from AV content and streams it across a network to the Dremedia Speech Services module. Dremedia Speech Services is the distributable platform on which the key Dremedia and Softsound speech analysis processes are implemented. At the current time the key services provided are:

Transcription Filter

This process transcribes speech into text in a variety of different languages. Dremedia's unique relationship with the Softsound speech scientists means that Dremedia has a direct input into the development cycle and can drive the production of new language models as required.

Speaker Segmentation / Identification Filter

This process analyses the audio stream to identify the speaker against its database of known people. Even if the speaker is unknown, the changes in voice can be detected and the time offsets at which these changes occur can be used for speaker segmentation and navigation through the original AV asset. Custom speaker databases can be produced for a variety of environments using Dremedia's speaker database management tools.

Scene Change Filter

The Scene Change filter analyses the video stream and detects scene and shot changes. It can distinguish between fades, cuts, zooms and pans. Upon detecting a shot change it outputs a thumbnail picture and a frame-accurate time code. A front end GUI can access this information, display the thumbnails and permit the user to navigate through the AV clips e.g. by clicking on the thumbnail.

Image Association Filter

The Image Association filter analyzes the output of the Scene Change filter and extracts the thumbnail image from its results. It applies a series of algorithms to images to produce a characteristic image signature. This signature can be indexed into a DRE so that the image can be dynamically associated with any other images in the database from any media asset processed by the system.

DRE/AXE

AXE™, the latest XML enabled version of Autonomy's patented DRE (Dynamic Reasoning Engine) technology is fully integrated into Dremedia's products and its proprietary synchronization architecture. AXE™ provides storage and search facilities for the XML results. Once the XML has been ingested by AXE™ a variety of queries can be made upon the information therein. In addition to using Autonomy's powerful natural language queries against any part of the XML, AXE™ also can be queried to provide suggestions of similar content.

Queries from the front-end are handled by the Query Controller and passed on to AXE™. The Query Controller server process allows for greater flexibility in the storage mechanisms used and allows the integration of 3rd party storage solutions such as XML repositories and relational databases.

Using AXE™ also allows a broad variety of Autonomy's infrastructure tools to be brought to bear on the information stored inside. AXE™ is also used to store and associate image signatures for the image association features of the Dremedia platform.

XML Architecture

Dremedia's technology utilizes XML as the central communication mechanism between the core technology modules, between all production systems and across all data sources. This provides a flexible framework and eases future expansion.

Dremedia software enables the automatic operation, creation and exchange of XML metadata and delivers cross-platform flexibility within a highly scaleable architecture. Dremedia's XML modular architecture provides accessibility to existing legacy and production systems as well as allowing third party developers to integrate with Dremedia's software. Dremedia's use of XML also guarantees both backward and forward compatibility and allows for remote configuration.

Dremedia's technology avoids traditional problems associated with XML because Autonomy's core automated text-handling technology enables computers to understand any form of unstructured data. This enables Dremedia's software to automatically extract meta-data for integration into existing media databases and to automatically attach the appropriate XML tags, based on the textual output extracted from video and audio streams as well as the entire content of text-based material. These XML tags enable material to be identified and distributed and the system to automatically capture, categorize, link or make material available via the appropriate system. Editorial staff or administrators no longer need to waste time manually inserting XML tags, thus reducing labor costs, and removing the inefficiencies of human error.

Dremedia's software uses its core technology to address XML's reliance on manual processes, adding a layer of intelligence to the management of XML, by understanding the content and purpose of either the tag itself, or related material, or both.

Core Technology

Dremedia's core video analysis technology allows the content of digital images to be managed in a manner previously impossible. Dremedia's proprietary technology – using advanced signal processing techniques – enables an image to be segmented into its component parts, such that they may be identified, understood, labeled and indexed. Such tasks have previously been performed by journalists, interns or researchers viewing hours of tape and then manually performing these operations. Dremedia's technology enables the automation of these operations, decreasing time-to-air and delivering immediate returns on investment.

Dremedia's video analysis technology uses a variety of techniques to provide automatic segmentation, cueing, manipulation and object recognition within all formats of video. The system technology is based around condensation tracking and temporal fusion.

The video stream is broken down into frames, which are analyzed and compared in order to identify objects within the complex image, to identify changes of shot and to reconcile unsynchronized soundtracks.

The major benefits of the approach to image analysis taken by the Dremedia audio analysis technology are:

- format independence
- frame-accurate analysis of any input format - digital or analogue
- browse in Real, Windows Media, QuickTime
- image analysis
- identify and track objects and find similar using condensation tracking
- locate shot changes, including fades, pans and zooms

Resynchronization

Many video library archives are too costly to reuse because the audio and video signals are out of synchronization. Combining image analysis of speaking lips with speech recognition and frame accurate signal manipulation allows resynchronization to be achieved at the press of a button.

Audio

Softsound's audio analysis technology, with its heritage from Cambridge University, uses advanced statistical methods to deal with all aspects of processing the digital audio signal. It employs a wide range of recognition technologies - from keyword spotting to small and large vocabulary continuous speech recognition and information retrieval from spoken documents - to enable applications in speech and audio processing, such as: speech recognition, speaker recognition, language recognition and speech enhancement. Because of the accuracy and utility of the speech recognition, Dremedia's products are the only products that enable the frame-accurate selection of video via an interface onto an automatically generated transcript.

Softsound's technology has consistently out-performed other systems for the precision of its real-time anonymous speech recognition tasks as benchmarked by:

- DARPA and NIST (both US Government agencies)
- THISL (an EU research initiative)
- SQUALE (a pan-European evaluation)

To achieve anything near the level of precision that Softsound's recognition engine achieves, the closest of the other systems required a massive increase (around ten times) in hardware investment and their associated running costs, to get close to the precision that the Softsound technology delivers on a standard Pentium box.

Softsound's audio analysis technology uses neural network technology and "Hidden Markov Models" (HMMs) to construct a highly efficient acoustic model, which is able to provide a fast, accurate and dynamic solution within variable and rapidly changing acoustic environments.

The system technology is based around decomposing digitized speech into its phonetic constructs. The phonetic sequence is then analyzed in conjunction with acoustic model and statistical improbabilities to calculate which is the most probable sequence of hence words and utterances.




Figure 3 : Recurrent network used for phone probability estimation

The major benefits of the approach to speech processing and recognition taken by the Dremedia audio analysis technology are described below.

Large Vocabulary Recognition

Using patented predictive technology the speech analyzer is able to provide the benefits of a large vocabulary speech recognition system without the overhead of a vast search space when considering sample audio.

Inter-Speaker Independence

Transcription of speech and segmentation by speaker requires no initial per-speaker training. Whereas many other approaches require training data from specific speakers to realize their full potential, the Dremedia system performs consistently well across a wide variety of previously unencountered speech sources. This is because the underlying technology was initially developed to maintain inter-speaker independence and not as a single user transcription tool.

Non-Dictated and Noisy Speech

Information feeds such as news broadcasts are often intrinsically difficult to transcribe due to noisy conditions and less than perfect articulation. Dremedia's sophisticated signal processing and statistical techniques enable the transcription engine to filter out extraneous noise, compensate for low volume levels and probabilistically predict intended dialogue.

Text

Dremedia utilizes Autonomy's technology to automatically handle large amounts of structured and unstructured text, including the textual output from our AV analysis technologies. Based on advanced pattern-matching technology that exploits high-performance Bayesian probabilistic techniques, this proprietary technology enables computers to form a conceptual understanding of text in any format and automate key operations and processing tasks upon it. Given that such tasks have previously been performed by expensive and slow manual labor, by automating the key tasks that are performed every day, Autonomy's technology delivers immediate return on investment.

At the heart of Dremedia's technology is Autonomy's XML Engine which performs the core information operations. The AXE™ is based on advanced pattern-matching technology that exploits high-performance probabilistic modeling techniques to accept a piece of content and decide which operation to automatically perform on that information. A piece of content refers to a sentence, paragraph or page of text, meta-data, a record containing human readable information, or the derived contextual information of an audio or speech snippet. These automated operations are:

Retrieval

The AXE™ accepts high-precision natural language queries and returns a list of documents containing the concepts looked for, ordered by contextual relevance to the query. The AXE™ also supports Boolean and structured field queries as standard.

Concept Matching

The AXE™ accepts a piece of content or reference (identifier) as an input and returns references to conceptually related items ranked by relevance, or contextual distance. This is used to generate automatic hyperlinks between pieces of content.

Alerting

The AXE™ accepts a piece of content and returns similar pieces of content ranked by conceptual similarity. This is used to deliver real-time alerts via any digital device or to discover users who are interested in similar content, or find users with common interests.

Categorization

The AXE™ accepts a piece of content and returns categories ranked by conceptual similarity. This is used to discover which categories the content is most appropriate for, allowing subsequent tagging, routing or filing.

Summarization

The AXE™ accepts a piece of content and returns a summary of the information containing the most salient concepts of the content. In addition, summaries can be generated that relate to the context of the original inquiry - allowing the most applicable dynamic summary to be provided in the results of a given inquiry.

Clustering

The AXE™ can organize large volumes of content or large numbers of profiles into self-consistent clusters. Clustering is an automatic agglomerative technique, which partitions a corpus by grouping together information containing similar concepts.

Active Matching

The AXE™ can accept textual information describing the current user task and returns a list of documents ordered by contextual relevance to the active task.

Performance and Scalability

Dremedia's AV analysis technology is highly scalable and can be tailored to many different applications, from small to full-scale broadcast channel monitoring.

A system for the monitoring of a multiple real-time broadcasts is possible using standard entry-level hardware with the filters being able to process incoming audio and video feeds in excess of twice real-time.

When working in RAID and extended storage provider environments, such as EMC, Dremedia leverages disk optimization and caching features to ensure the optimal performance.

Communication between the Dremedia components is consistent with all other Dremedia modules, using HTTP over TCP/IP allowing for a highly scalable distributed system to be constructed easily.

Limitations of Alternative Approaches

Many companies claim to have workable solutions for solving the problems associated with dealing with large amounts of audiovisual content. However, most of these solutions have a fundamental problem – they are heavily reliant on human processing, input and integration and are inherently manual.

Such approaches increase the costs associated with handling large amounts of audiovisual content and restrict the ability to scale operations, as the volume of content being handled increases.

Only Dremedia's approach removes the need for those manual, expensive tasks. Dremedia's proprietary technology uses advanced, high-precision speech recognition and video analysis to actually understand the content of the broadcast stream and locate the specific segment without searching, logging, time coding or creating meta data.

Managing Assets Using Metadata

The use of metadata – or the tagging of important or relevant information – to manage information content has become increasingly popular over the past few years, so much so that Asset Management systems now place a heavy reliance on metadata to bring structure to their solutions.

Limitations

Reliance on human behavior and user experience

In order to achieve maximum efficiency from a meta-tagging system, every item of importance would need to be tagged correctly. This relies on the user doing the meta tagging having knowledge of what is important in every piece of content and also being able to accurately describe these items of importance. Increasingly the user's ability to describe what is important is dependant on their personal experience, knowledge and opinions. These "elusive" variables vary between data and meta-data loggers and have a dramatic effect on both the quality and accuracy of the meta-tags.

In practice, users – especially when under time and deadline pressure - consistently fail to tag content and adhere to meta data schemes. Attempts to simplify or standardize meta data creation or generation through the imposition of meta data 'standards' also fail victim to inconsistent user behavior, creating a backlog of content that needs to be re-categorized or is simply 'lost'. Notoriously, in one of the largest meta data schemes - run by US government departments - over three-quarters of all content that had been tagged, was tagged with the category "General".

Meta-tagging is expensive and time consuming

Creating meta-tagging solutions requires the consultation of all targeted users – archivists, producers and IT personnel. In addition tagging all historical and current assets is a manually intensive task which increases project costs and project completion times. The effort expelled in tagging assets never matches the desired levels of accuracy or usability.

No scalability

To develop a refined and searchable asset database, assets would require a high number of tags. As new assets become available, more tags would be required. In time as the number of tags increases, so do the resources required, which increases the risk of miscalculation. Equally, the manual input required is directly linked to the volume of content requiring to be tagged – the more content, the more people.

Dremedia's Approach

Dremedia's approach directly addresses the problems associated with manual meta-tagging by adding a layer of intelligence and automation to the management of XML by understanding the content and context of the either the tags themselves or the associated information. In effect, the Dremedia solution removes the need for meta-tags or explicit meta data. Metadata is implicitly (covertly) inferred through the installed layer of intelligence.

However, if metadata is required, intuitive user interfaces are provided to add reassurance and additional information. In situations where there are already large amounts of existing meta-data and/or established taxonomies, Dremedia is used to automatically add new content to these schemes and append the appropriate tags. Dremedia is also frequently used to deliver immediate benefits by automatically integrating disparate meta data schemes and provide a single, unified view of the content with no manual overhead.

Video Indexing Systems

Video indexing systems identify scene changes within a video stream and attempts to automatically associate the clips with metadata, thus making the video stream searchable, using basic keyword search engines.

Limitations

Manually intensive

The video indexing process is manually intensive in the extreme. Someone still needs to watch the video asset for scene changes and attach descriptive metadata to the clip. This increases project costs and the time taken to index the video asset.

Unsynchronized

As video indexing systems rely so much on human intervention, they become cumbersome and expensive to integrate into other production or asset management systems. So much user intervention takes place that the outputs from the video indexing components cannot be effectively or automatically exported to the next system in the production process.

Dremedia's Approach

Dremedia's proprietary video analysis technology automatically and seamlessly identifies the scene changes within a video stream. These scene changes are ordered by time code and using similar pattern matching technology as described above all clips can be "understood" by Dremedia's products. In addition by using advanced color and shape analysis algorithms it becomes possible to search the asset database for similar video clips, without relying on either metadata or human intervention.

These outputs are completely synchronized with all other outputs to the millisecond on a frame-accurate basis. This means that the images are synchronized with the relevant sentences within an automatically generated transcript, the words spoken are synchronized with the relevant speaker, the audio transcript is synchronized with the appropriate scene changes etc.

This unsurpassed level of synchronization enables users to simultaneously and inter-changeably navigate through large amounts of audio visual content by image, word, scene, speaker, offset etc., with no manual integration required to facilitate this.

Keyword Searching Systems

Historically, the only way to search for content within an asset management system was by using traditional keyword searches. This simple method involved asking a user to enter a search keyword or a combination of keywords into a query field. The search engine then searches through a list of documents or a table of metadata and returns a list containing the search keyword. The quality of results and user experience are identical to using an Internet search engine.

Limitations

No context

Using this method a user enters a keyword or term into a search box, for example "penguin". The search engine then searches for any assets containing the word "penguin" The fundamental problem is that the search engine is simply looking for that word, regardless of how it is used or the context in which the user requires the information. i.e. is the user looking for a penguin bird, a publisher or a chocolate-brand? Moreover, this approach requires the relevant word to be present and for the content to have been tagged with the word. Any new subjects, names or events will not be present and the system

Manual

Keyword searches do nothing more complex than look for the occurrence of the searched word or term. These processes require a significant amount of hardware resources, which increase systems overheads. In addition keyword search systems require a significant amount of manual intervention so that words and the relationship between similar words can be identified. (Penguin = flightless birds = fish eating birds).

No dynamic intelligence

Keyword search engines cannot "learn" through use, nor do they have any understanding of queries on specific words. For example when the word "penguin" is entered, keyword search engines cannot learn that the penguin is a flightless black and white bird that eats fish.

Significant user refinement required to boost accuracy

Keyword search engines rely heavily on the expertise of the end user to create queries in such a way that the results are most accurate. This requires complex and specific Boolean syntaxes, which the ordinary end-user would not be able to complete. i.e. to get an accurate result for penguins, an end user would have to enter the query as follows: "Penguin AND (NOT (Chocolate OR Clothing OR Publishing) AND Bird.

Dremedia's Approach

Dremedia's concept matching technology avoids these problems by matching concepts instead of simple keywords. The Dremedia search technology takes into account the context in which the search terms appear, thus excluding many inaccurate hits while including assets that may not necessarily contain the keywords, but do contain their concept. This also allows for new words or phrases to be immediately identified and matched with similar ones, based upon the common ideas they contain as opposed to being constrained by the presence or absence of an individual word; this equally applies to misspelt words.

In addition to the concept matching technology, the Dremedia engine accepts standard Boolean text queries or any combination of Boolean or concept queries.

Phonetic Speech Recognition Systems

Phonetic speech recognition systems identify and index the phonemes of words to transcribe the spoken work into text. Once the text is available the premise is that the audio itself is also searchable.

Limitations

Reliance on exact keywords and phrases

In order to for the search to be accurate the keyword being used needs to match the word exactly. If for example you were looking for a tape on the current "French downhill champion", you need to rely on that phrase appearing on the tape. "French+skiing+gold" or "France's first gold medal" would not return any results.

The system is also constrained by the limitations discussed in Keyword Search system.

No scalability

As more and more audio is ingested into the asset management system, the performance of the search engine becomes slower and slower.

Word level searching has a number of disadvantages

Word level searching is dependant on the exact word you are looking for being in the vocabulary database. This poses difficulties when new proper names appear in the public domain. In addition, word level searching is extremely processor intensive so has a significant hardware overhead attached to it.

No performance and accuracy benefits

Phonetic transcription systems claim that they are able to offer superior performance and the transcript of a higher accuracy as compared to their competitors. Industry benchmark tests have proved that the large vocabulary systems employed by Dremedia and other transcription systems are both faster and more accurate than phonetic transcription engines.

Dremedia's Approach

Dremedia uses a patented, proven large vocabulary recognition engine that guarantees accurate and scalable solutions with very little processor or manual overhead. Unlike similar products, Dremedia's speech recognition technology is not a plug-in; it is fully integrated into all of our products. The resultant transcripts are highly accurate and are editable within the production interface, any edits being automatically synchronized with the video. In addition, by using the concept matching techniques discussed above audio can now be searched and manipulated in any way.

Conclusion

- **Only Dremedia provides an only fully automated, synchronized platform for audiovisual content production and interactivity.**
- **Unlike any other alternative approach, Dremedia's technology is completely automated and does not rely upon or require any meta data or manual pre-processing for content acquisition, organization and management.**
- **Dremedia's integrated, scalable solution provides everything you need to automate audio visual content production and is accessible from and compatible with all existing production and legacy systems.**

The information contained in the document represents the current view of Dremedia Limited on the issues discussed at the date of publication. Because Dremedia must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Dremedia, and Dremedia cannot attest to the accuracy of any information presented after the date of publication.

This Document is for informational purposes only; Dremedia is not making warranties, express or implied, in this document.

Dremedia Limited
4th Floor
63 St James's Street
London SW1A 1LY
UK

Tel: +44 (0) 20 7970 5670

Fax: +44 (0) 20 7970 5683

Email: info@dremedia.com

Dremedia
301 Howard Street
22nd Floor
San Francisco CA 94105
USA

Tel: (415) 243 9955

Fax: (415) 243 9984

Email: info@dremedia.com



www.dremedia.com

Copyright © 2002 Dremedia Limited.
All rights reserved. Other trademarks
are registered trademarks and the
properties of their respective owners.
Product Specifications and features are
subject to change without notice. Use of
Dremedia software is under license.