A Metadata Model for Multimedia Databases

Cristina Ribeiro [*] *and Gabriel David* [*]

[*] Faculdade de Engenharia da Universidade do Porto/ INESC Porto

E-mail: {mcr,gtd}@fe.up.pt

ABSTRACT

Largely due to the facilities provided by the technology, there is currently a large effort put into the organization of large electronic repositories of documents, covering both present-day materials and those of historic interest. It is also technologically feasible to gather heterogeneous documents, and the same collection may include text, images, video and audio items, either independently or as multimedia objects. Professional and lay users have great expectations on the possibilities of searching such materials, partly due to the success of search engines on the Web. Content-based retrieval on multimedia objects is a hard task, however, and results similar to those obtained with text retrieval are not foreseeable in the near future for images or video. Retrieval of multimedia objects must therefore rely on the wealth of associated information, available as metadata of diverse nature.

The paper describes a metadata model where contextual information, in an archival perspective, is combined with the description of contents, as required for content-based search. Contextual information can be valuable for search, and is important when communicating search results, where issues such as date, access rights or age rating may condition the usage of retrieved documents. The metadata model adopts a flexible multi-level approach to both context and content description, to allow for different granularities in the description of materials. It may be impracticable to provide descriptions at the individual item level in the digitization of an historic photo archive. Instead, a top-down approach to the description of collections can be a good compromise. For the video materials of a broadcast station, on the other hand, there is automatically produced information for each recorded shot that can be also accommodated in the model.

The model has been tested in a prototype database system where specialized interfaces are offered for cataloguing and for search. The database contains data from historic, photo and video archives. The prototype also provides tools for filtering information into standard formats for interchange. The relations of the proposed model with existing and emerging metadata standards are presented. An application based on the prototype will be used for testing user behavior with data from an historic archive.

KEYWORDS: multimedia databases, metadata, archival description

INTRODUCTION

Multimedia archives both in the sense of local repositories and, in a broader view, as a fraction of the Internet, are increasing in importance and availability. However, there is a wide

469

recognition of the difficulties in producing accurate and trustworthy search results. The addition of metadata [2] to archival objects enables informed and efficient search, provided it includes both contents and context descriptions. Intense research is being carried out in multimedia content description standardization [5][6] as well as in the extension of traditional context description to the world of electronic documents [9].

The paper starts by laying out some basic ideas on multimedia archives, arguing in favor of the definition of a widely acceptable metadata scheme for organizing documents and associated descriptions. The incorporation of these ideas into constructs is presented with examples taken from several categories of documents. The Metamedia scheme is then assembled, showing how different metadata categories are integrated and the devices available for controlling the structure of archives. Conformity to standards is discussed. The conclusion puts this work in the context of the projects where it has been developed.

BASIC IDEAS

Building a multimedia archive involves collecting large quantities of documents in different media and formats, storing the documents and any associated information and being able to search the whole set based on features of its content. It is also essential to be able to provide the set of documents obtained as the answer to a search in some standard formats, in such a way that they can be included in some collection elsewhere or presented to the final user in a convenient display format. Additionally, it is important to have appropriate interfaces both for the archiving and the search activities. In the following, the ideas on which the proposed metadata

standard is based are put forward.

Archival Description and Web Search

Traditionally, archives have organized collections of documents with great concern with their authenticity, organization and preservation. The organization has been mainly built on the provenance of the individual documents; it is expected that documents resulting form the activity of a specific office are kept together and grouped according to the functional divisions existing in the office. The physical support for the documents is typically intended for their long-time preservation. The ability to search the archival materials has not been a requirement of classical archival systems, where search is supposed to be an accessory function, carried out by specialized personnel.

A very different perspective exists on the side of the institutions and individuals who are disseminating information on the Web. Electronic documents on the Web are unorganized, arbitrarily linked and hard to authenticate. They are also highly volatile: a server may go down making a whole repository of documents unavailable and thus increasing the miss rate of the information provided by the search engines, which are therefore forced to implement refresh mechanisms to maintain their own data quality. The success of the huge archive of materials that the Web constitutes comes from the ease with which lay users browse it and use sophisticated search engines to find information that suits them. The advances in indexing and searching text have made retrieval based on the textual content of documents feasible and usually quite effective.

But materials on the Web are also quite

470

heterogeneous: the same document may include text, image, audio or video, in several formats and with complex links. The digital formats pose new problems both for search and for preservation. Effectively searching audio, images or video will continue to be a hard task in the near future. Guaranteeing that the document file that is being displayed today will continue to be available in 10 years time may require that a safe storage system will be kept and that the application used to open it will still be available, or else that the document content is preserved throughout technological changes.

The intention of the proposed metadata scheme is to have a model that can serve for describing a video document, a picture or a text document, exploring features that are meaningful in any of the media. The advantage of a common model comes from using a single formalism to deal with multimedia objects of different kinds, that actually do appear bundled in current documents. This model can be implemented in a database management system to efficiently store and retrieve documents and can also be turned into a descriptive interchange format to be communicated between applications. It is also envisaged that descriptions using this model might be turned into descriptions in other models, for the purpose of exporting documents into other systems, without significant information loss. The proximity to relevant standards has therefore been taken into account. The metadata scheme has put emphasis on descriptors that conform to the requirements of archival systems.

Hierarchical
Documents can have complex structure and diverse relationships with other documents. Capturing the relationships

between documents is essential both for their description and for search. Documents do not exist alone and their full meaning can only be grasped if the context of their creation, use, and ownership is known, as well as that of the documents they refer to or that refer them. This information should be recorded. In a similar way, relevance of retrieval requires knowledge about the role a document plays in a set where, for instance, one document presents the definition of a concept and others are detailed explanations or case studies.

In electronic documents, it is more and more the case that documents are composed of parts that may be reused to create new documents. In traditional archives, documents are structured to identify their relationships to the entities that originally created them.

Due to their diversity, the semantics of the relationships between documents can be involved and hard to represent. There is one kind of relationship that is given particular importance in the organization of archives: the hierarchy of containment.

This is mostly due to the fact that archives have traditionally had the *a posteriori* task of organizing materials based on features of their creation, therefore grouping together documents from the same origin an organizing them in a hierarchy where both the individual document and the collection are described. If we take a more dynamic document production environment, like a corporation or a TV production studio, we see that the rate of production of materials is not compatible with an *a posteriori* effort of description. But it is also a fact that the environment of creation is getting more sophisticated, and the production tools can be

471

customized to add a substantial amount of metadata to the items as they are generated. It is also desirable to have documents, as they are produced, organized in collections according to an *a priori* classification. As an example, a company may be organized in 3 branches, each having some departments and some services within the departments. This organizations might be reflected in the organization of documents originated in the company, so that documents from one service would be kept together as a set which in turn would be part of the documentation of the corresponding department and this would go up to the company level. A similar case might be identified in a governmental or administrative service. A video production studio is also organized in units that create what can also be classified as documents, and in the same vein their materials can be grouped at creation time in a hierarchy that reveals their natural organization. The proposed metadata scheme assumes that there can be varied ways for an archival item to be related to others, and that there should be room for such links to be established. The hierarchy of containment is however considered important enough to be built into the model, so that specialized constructs are provided to configure it and to control its use. It is supposed that the applications handling data according to this model take advantage of the hierarchy in the interrogation.

Semantic Networks Versus Flat Indexing

The era of Web information retrieval has brought about a big success based on the use of the abundant disk space and computing power to go through vast amounts of textual documents and build indexes that allow quick retrieval based on user-supplied terms. The representation of documents in Web search engines relies on collecting words from the text and *meta* components of Web documents. The most successful retrieval procedures, however sophisticated, are based on those words and overlook the information conveyed by the image and video components. The efforts towards a deeper understanding of documents have not, in its present state, contributed to significant improvement in retrieval quality. The success of brute-force indexing is as yet restricted to text search. With the huge amounts of audiovisual information available online, it is getting more and more interesting to be able to search the contents of such documents. But there is no easy syntactic channel between information and its user for an image as is the case of words in text. A word is easily identifiable in a text and is usually associated with a concept at the level of human communication. An automatically identifiable feature of an image, such as a color distribution, does not provide a retrieval system with a concept that is equally useful for user interrogation and therefore does not constitute a useful representation of the document from an indexing point of view. There are strong lines of research investing on higher-level representations both for text and for other forms of information such as image or video, but automatically searching multimedia materials is not expected in the foreseeable future. In fact, the current technology in image and video processing is not yet able to extract features abstract enough for direct user interaction. So metadata is regarded as a feasible way to attach meaning to documents in a machine understandable format.

A *metadata scheme identifies well-*

established parts of the information associated to documents and standardizes their representation and sometimes even the vocabularies used to instantiate them for the items being described. This can be very useful for search: not only the values of the descriptors associated to a document are directly usable, but the semantics of the descriptors themselves is fixed and therefore specialized search tools can take advantage of this fact.

Even if no interest were to be found in the use of metadata for search, it would still be invaluable in the interpretation of search results. When some document is to be included in the answer to a query, the user should be aware of items that belong to the document's metadata, such as date, creator or terms and conditions of use.

Libraries, Archives and Museums
A metadata scheme can incorporate several views of the same document. It aims at serving as a common structure where different perspectives on the document coexist and that can be used as a basis for developing information systems supporting applications in such diverse domains as archiving, information search, library cataloguing or audiovisual production management.

The perspective in this work has been to maintain a metadata scheme that can model the various views on a multimedia document and to develop prototype database applications using the model as a reference. The intention is to demonstrate that it is possible to keep the pace between a model where the relevant concepts are maintained and applications where efficiency is required. A particular application need not implement all aspects of the model, concentrating instead on those relevant

for its purpose. The existence of the model guarantees the possibility of exchange of documents between applications preserving the metadata as far as possible.

Context and Content
The proposed metadata model has built on the concepts of context and content as applied to metadata. Traditional archive and library cataloguing metadata can be classified as contextual metadata, in the sense that it concerns mostly the description of the item in aspects such as its creation, use, relation with other items, preservation, rights, location or custody. A rather orthogonal perspective is given by content metadata, appearing in the traditional catalogs in the form of abstracts and keywords. In Web search engines, content metadata is massive, in the form of automatically collected descriptors and statistical parameters to evaluate relevance; on the other hand, the context of a document can only be accessed by searching other documents in the site containing it or by using the information in the document itself.

The proposed metadata scheme intends to explore context and content in a balanced way. For context, aspects such as creation, rights, custody and location are considered. For content, the scheme provides structured descriptions of content that make it easy to represent documents having parts in different media, and also considers high-level concepts such as entities and events about which documents are.

Text and Multimedia
A medieval parchment and a synthetic animation movie may have little in common, but the requirements for their metadata are not as different as the documents themselves. For both, it is interesting to have information on the creation, location, custody and use. For

473

both, it is useful to have description concerning the content to make them retrievable in a subject search.

The proposed metadata scheme intends to be appropriate for items of different nature and different media. It provides the constructs where media-dependent descriptors can be accommodated, and specialized descriptors for the most common types have already been established.

Structure, Volume and Granularity

Generating, manually or automatically, the metadata for an object or a collection of objects can be a time consuming task. Some objects will have metadata associated to them since their creation, while others will require an explicit task of metadata creation by the time they are included in an archive. There are objects for which metadata generation can be afforded, due to their importance, and many others that will not justify such an investment. In any case, some minimal description is required to keep track of the existence of the documents.

The provision for a hierarchical structuring of documents is an important feature when it is necessary to deal with large quantities of documents. It is possible to establish the level at which the description is made and possibly refine it at some later time, without having to reorganize what is already acquired. This can be useful when a large collection is incorporated in an archive and it is necessary to rapidly make it visible, even if not at the level of the individual item. For the systematic incorporation of new materials in a production environment, it is also important to begin the description by registering values that can be automatically extracted, and information available on sets of items, as it may be the only metadata expectable for many items in such a dynamic environment.

Standards

The metadata scheme has built on available metadata standards. The main sources are the ICA ISAD/ISAAR [3] standards for archival description and the MPEG-7 [8] (ongoing) standard for multimedia content description. Standards are important for enabling metadata interchange between applications. It is essential to the goal of maintaining a metadata scheme to keep abreast of the developments in metadata standards.

BUILDING A METADATA MODEL

The metadata model has been designed as the assembly of modules dedicated to the description of separable aspects of an object's description. Some modules are centered on contextual metadata; others on content and others include both. Parts of the model are now used to illustrate in more detail the solutions proposed along the lines presented in the previous section. The diagrams use UML syntax to present the main concepts with their attributes and relationships.

Description Units

The term Description Unit has been used to designate a metadata unit, where the description of some object is organized. The diagram for the Description Unit, presented in Figure 1, shows the descriptors currently used to describe the common features of multimedia objects.

Figure 1: List of descriptors

The set of descriptors presented results from the analysis of several examples of diverse datasets (broadcaster archives, photo museum, historic archive) and current or under development standards. However, the proposed model should not be substantially changed if minor variations in the descriptors are performed.

The object may be a single item, such as a letter, or a collection of objects that are described as a set, possibly at a higher level. A hierarchical structure exists on the description units, whereby several description units can be linked to a higher level description unit, in a many-to-one relationship. The latter is the metadata for the set of objects while the former is the metadata for each of those objects.

The hierarchical organization of descriptions is an established practice in archival description, as is patent in recent archival standards [3]. The hierarchy is a natural structure to support

the kind of analysis that archival professionals perform on the documents they handle: the hierarchy of an archive typically reflects the hierarchy of the organization that produced the documents.

Level Flexibility and Control

A rigid hierarchic structure may impose undesirable limitations, in case the levels of the hierarchy are not adapted to the application domain. The adopted structure may be too complex or too simple for the actual collection. Also, the adopted designations of the levels will have to be quite generic and will lack significance for the specific situations. The opposite approach of leaving the organization of the collection completely open prevents the software from helping on validating the records. A compromise between both extremes is to allow the definition of a hierarchy suited to each concrete situation, but requiring the actual records to respect it once defined. This way, the proposed model combines flexibility and discipline.

The concept of Scheme is introduced to allow for the choice of an appropriate structure of levels in the organization of a collection of objects. A Scheme names a hierarchy, and has several associated Description Levels. Each Description Unit belongs to a unique Description Level. The goal has been to provide the constructs for specifying the number of levels, their designations and the admissible links between them.
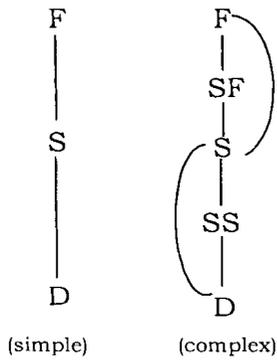
F   F
| |
| SF
| |
S   S
| |
| SS
| |
D   D
(simple)   (complex)

Figure 2: Example Schemes

In a simple case, there could be just three levels, Fonds, Series and Document for instance. Objects in the collection must be one of these kinds and documents cannot link directly to the fonds. A more complex collection might require more levels, such as SubFonds and SubSeries, with more diverse link possibilities. A document may be part of a subseries or directly belong to a series. In both cases the actual links in a collection form a tree.

It is up to the archive manager to design the most appropriate scheme for each archive. Actual descriptions will have to conform to the adopted scheme. The hierarchy of the Description Units is therefore constrained by that of the Description Levels. This provides the support for enforcing, at the application level, the correct structuring of description units. Figure 3 shows the fragment of the metadata model dedicated to the control of the hierarchy.
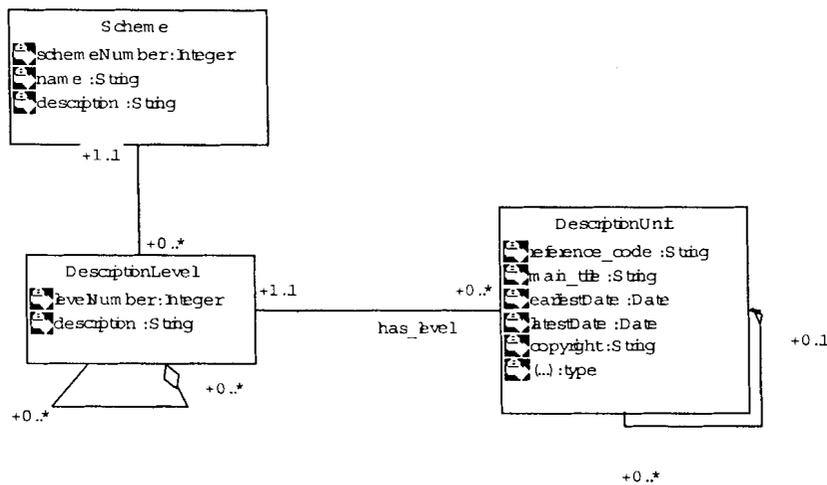
Scheme
- schemeNumber:Integer
- name :String
- description :String

+1 .1

+0 .*

DescriptionLevel
- levelNumber:Integer
- description :String

+1 .1   +0 .*

has_level

DescriptionUnit
- reference_code :String
- main_title :String
- earliestDate :Date
- latestDate :Date
- copyright :String
- (...) :type

+0 .1

+0 .*

+0 .*

Figure 3: Schemes and Description Levels

Content Description

Content-based search on multimedia objects requires content metadata. The model proposes the concept of Segment where description regarding the content is organized. A segment is always linked to some description unit providing its context. Several segments may be

476

attached to the same description unit, accounting for different parts of its content. Segments are also hierarchically organized, so that some segment may refine the description available at its parent segment.

Figure 5 shows the links between description units and segments. Segments are further specialized for text, image and video, so that specific descriptors are offered for each media.

An alternative to this solution could be just to extend the hierarchy of description units downwards until an arbitrary level of detail. However, when dealing with items at the sub-document level, many of the attributes loose

meaning and applicability. For instance, the creator is not relevant when describing a fragment of a text, part of a picture or a sequence in a movie. That information is probably already stored in the description unit corresponding to the complete document. On the other hand, the diversity of attributes that may be of significance for all the possible contents advises a richer and more flexible support. A segment usually contains information that describes (part of) the contents of the document and for which little contextual information is required. Segments are specialized according to the type of material and some of them have an intrinsic time dimension, as is the case with video.
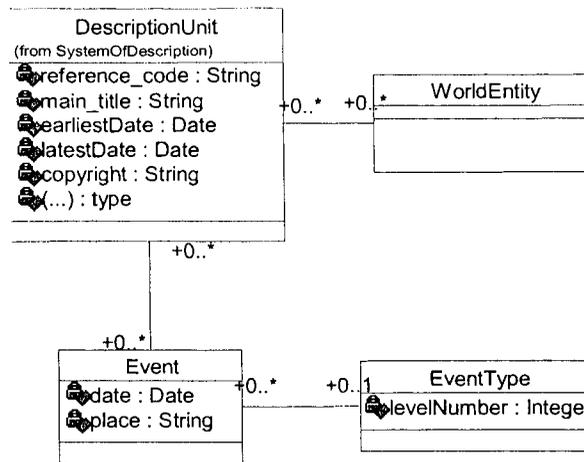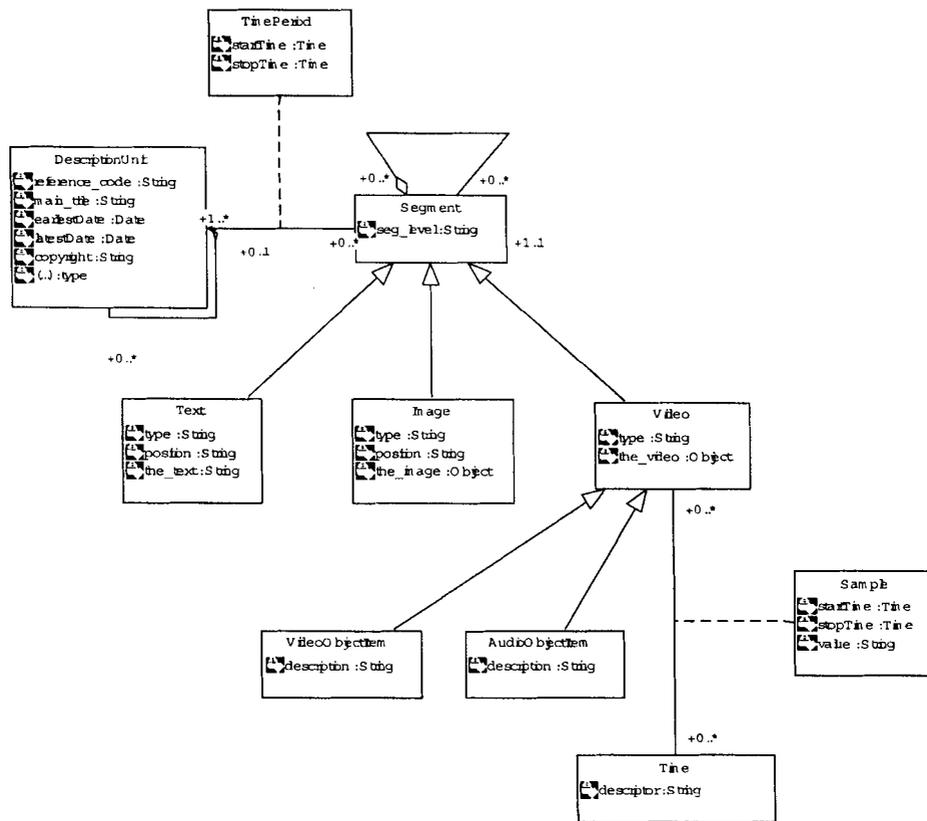
Figure 4: About

Figure 5: Segment

Many documents may refer to the same event or world entity. In order to support authority files for these two categories and thus improve normalization on its representation, they are taken out of the description unit as another device to facilitate contents based search or handle content at a higher semantic level, Segments in contrast are devoted to more syntactic or physical characteristics like color histograms, actual text or frequency contents of a speech.

METAMEDIA DESCRIPTION SCHEME

Besides the fragments described in the previous section, which characterize the metadata model and reflect its main conceptual aspects, the scheme also includes other parts of the description that further describe context, such as the modules for Creator, Owner and Storage.

Media and Storage

An object may have several physical instances that may be stored in different places using various supports and

478

formats. The Media and Storage part of the scheme specifies the relations between a description unit and its storage in some support. There are specialized sets of descriptors for analog and for digital supports.
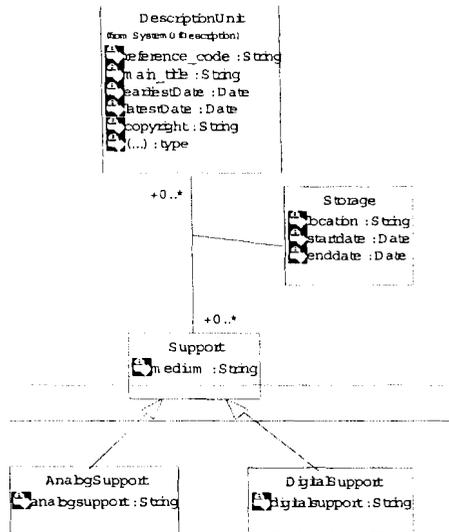

Figure 6: Support and Storage

Creator

Creators and owners of description units are also captured in the model. Figure 7 shows the part of the model that is rich enough to contain the history of the creator of the document, a person or an institution that may have suffered several reorganizations or changed the name along the time. A preferred main identification is supported and relevant locations and periods may be associated.
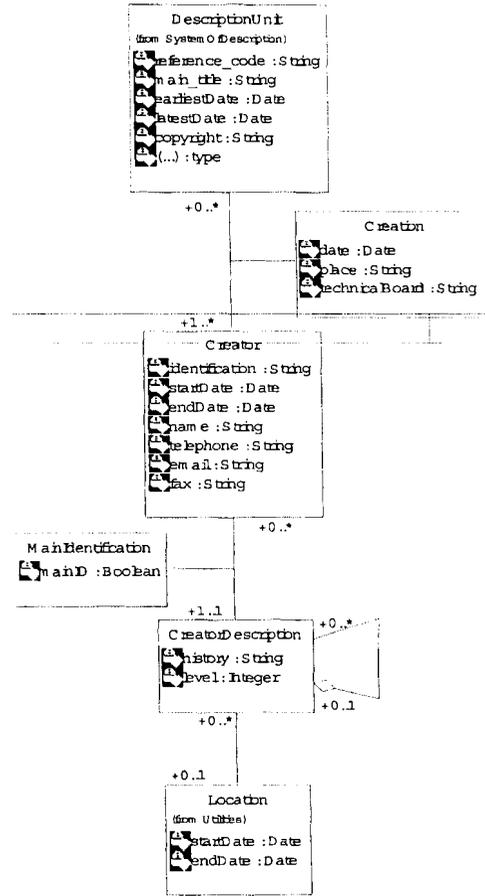

Figure 7: Creator

Owner

The owner is an archive presently in charge of the document or at some point in its history. Besides contact information, this module includes provision for the localization of copies of the document, and for other material (description units) that is logically related.

Package Assembly

The package diagram shown in Figure 9 brings together the packages where the

479

parts of the metadata model are organized. The description unit is common to all packages, as it identifies the object being described and assembles the most generic descriptors. The SystemOfDescription package includes the hierarchy level control, the description unit and the segment. This is the central package, dealing with contextual metadata (in DescriptionUnit) with content (in Segment) and with the structure of the description (in Scheme and DescriptionLevel). The Creator, Owner, Storage and About Packages coincide with the fragments of the metadata scheme already described under the same titles.
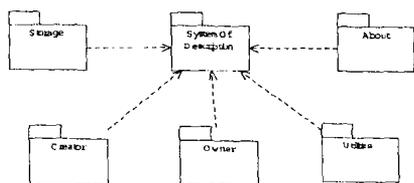

Figure 8: Package diagram

RELATION WITH METADATA STANDARDS

Producing metadata is recognized as an expensive and time-consuming task. To explore its various applications it is essential to be aware of its semantics so that, when documents with associated metadata are exchanged, the receiver can interpret and use the metadata. Metadata standards can serve this purpose. If metadata is collected according to schemes that conform to standards then it is possible for one application to export documents and associated metadata and for another application, knowing the standard that is being used, to import the metadata into its own format, using the semantics established by the standard.

The Dublin Core [1] has appeared as a standard comprising a minimal set of descriptors suitable for online resource discovery. It has evolved with contributions from library and archive specialists and is now taken as a component by more elaborate standards. As the DC imposes no structure on descriptors, it is easily included in a metadata scheme.

The Resource Description Framework (RDF) [7] is a W3C Recommendation concerning metadata for online resources. The RFD is not a metadata framework, but rather a proposal for interoperability between metadata schemes. Its aim is to provide a common syntax for use by diverse metadata schemes. RDF establishes a common syntax for the exchange of metadata, assuming that descriptions are available in a markup language (XML). The Metamedia model is specified in a format easily translated into an implementation in a database system, but a parallel markup format is also being developed. The markup format is essential for conversion between standards and for the interchange of information between applications. The descriptors used in the markup version of the scheme can be specified as RDF properties so that external applications can interpret Metamedia object descriptions according to their meaning. In the broader context of the Semantic Web [15] where the RDF fits, the meaning of descriptors would be established according to an ontology for the kind of objects described.

The MPEG and MHEG family of standards have focused on video coding and multimedia presentation description, respectively. MPEG-7 [13], has set a different goal and will be an ISO standard for multimedia content description. Its intention is to establish a

480

set of descriptors and description schemes for the content of audiovisual materials. The standard also includes a *Description Definition Language* to be used for defining customized description schemes. MPEG-7 has a strong connection with previous MPEG standards for video coding and keeps with their tradition in the choice of descriptors for audio and video features; for these descriptors it is expected that automatic tools will be available in the near future.

The development of the Metamedia scheme was triggered by the contribution to the initial phase of the MPEG-7 standard. The development of the scheme has taken into account both the content-directed requirements of MPEG-7 and the context-rich description of the ISAD(G). In its present state, a balance has been established between the two. Descriptors according to the Dublin Core can also be easily included in the model. Besides that, the markup version of the description scheme expressed in XML will be a convenient format for metadata interchange and is suitable for expression in RFD syntax, making it possible to integrate with other metadata specifications.

CONCLUSIONS AND ONGOING WORK

The Metamedia Description scheme has evolved from an early scheme developed as an object-oriented model for the ISAD(G) standard [10] to broader schemes proposed as Description Schemes to the MPEG-7 Committee in February 1999 to the unified scheme presented here that has been the result of the Metamedia Project [12]. In the continuation, metadata is further explored from the preservation point of view and having as targets objects with

embedded behavior[11]. The metadata model is the basis for an effort that also includes the development of markup representations for the information in the model and its use in information exchange and in electronic publishing.

Prototype
The metadata model can be translated to an operational database system without major modifications. A relational prototype has been built based on the metadata model and populated with 3 sets of data: descriptions of documents from an historic archive, photos (both description and digitized versions) and video from the MPEG-7 test sets [14]. The insertion of data and the conversions required on metadata have shown that obtaining uniformity of description requires a substantial amount of extra work that we expect will be viable to automate in part. The uniformity of description in the metadata scheme allows queries to retrieve items from the database regardless of their media.

Project Team and Opportunity
The project team includes two research groups from the Porto University, concentrated on information systems and video descriptors, and Minho University, specialized on language processing. There is also a group of three archivists integrating the team and bringing into it the tradition and perspectives of the respective community in a dialog not easy at the beginning but eventually very rich and in the basis of the results obtained. The opportunity of this research comes from the relatively early stage of mutual knowledge of the several research communities interested on multimedia databases: the archivists, the producers of multimedia contents, the information systems builders and the developers of

481

digital archives and libraries. The recent and active work on standards relevant to the subject is another proof of the actuality of the subject.

Metamedia-2

The main goal of this follow-up project is to further investigate the combination of context and content description, producing a model for multimedia archival description and tools to support an environment for it. The concept of multimedia object is extended here to encompass software components. Multimedia objects are no longer just passive items but, following the trend of object-oriented programming, they often include the code that elicits their behavior. The preservation of these objects is a big challenge for archives. Taking into account the rate of production, the bulk of the description of multimedia objects must be generated at creation time. This requires standards to be supported by the production tools rather than being a concern just for an a posteriori cataloguing activity. A parallel exists with the good practices of software development, which require that documentation be generated along the way. So, the techniques that are used to automatically document software components may prove useful to generate descriptions for multimedia objects. In the opposite direction, the requirements of the archives may be explicitly handled by the software development tools. The project results will include a model (a Description Scheme in MPEG-7 terminology) for contextual and contents description of the extended multimedia archives, covering aspects relevant for software components. The implications of this extension on the description definition language will be investigated. Tools such as parsers, syntax-directed editors and output generators will be derived

and form the basis of a software environment for multimedia archives. A database model, extending the one already developed in project Metamedia, will be used to store the rich structured data they produce. The tools, the database model and a suitable user interface will all be packaged as a software framework for multimedia archives that can be instantiated to specific applications. A concrete application domain has been selected, namely audio-visual production. The AV-production archive will be MPEG-7 compliant and serve as a test bed for intelligent search strategies and mechanisms that take advantage of the multi-faceted nature of the stored metadata.

REFERENCES

1. Dublin Core Metadata Initiative. URL: http://purl.oclc.org/dc/.

2. Murtha Baca, editor. Metadata-Pathways to Digital Information. The J. Paul Getty Trust, 1998.

3. Committee on Descriptive Standards. ISAD(G): General International Standard Archival Description-2nd Edition. Technical report, International Council on Archives, 2000.

4. Committee on Descriptive Standards. ISAAR(CPF): International Standard Archival Authority Record for Corporate Bodies, Persons and Families, International Council on Archives, 1996. URL:http://www.ica.org/isaar_e.html.

5. MPEG~Requirements Group. MPEG-7 Context and Objectives. Technical report, ISO/MPEG,

1998.

6. MPEG~Requirements Group. MPEG-7 Requirements Document V.5. Technical report, ISO/MPEG, 1998.

7. Semantic Web Activity: Resource Description Framework (RDF) *URL: http://www.w3.org/RDF/.*

8. Multimedia Description Scheme Group. Text of ISO/IEC 15938-5/CD Information Technology-Multimedia Content Description Interface Part 5 Multimedia Description Schemes. Technical report, ISO/MPEG, 2000.

9. C.M. Sperberg-McQueen and Lou. Burnard. Guidelines for Electronic Text Encoding and Interchange (TEI P3). Technical report, Association for Computers and the Humanities, Association for Computational Linguistics, Association for Literary and Linguistic Computing, 1994.

10. Archivum Project Team. Archivum: System of Objects with Temporal Support for Archival Description. Technical report, INESC/FEUP, 1999.

11. Metamedia Project Team. Metamedia2: Metadata for the Preservation and Retrieval of Multimedia Components. Project Description, INESC/FEUP, 2001.

12. Metamedia~Project Team. Metamedia: Metadata for Multimedia Databases. Technical report, INESC/FEUP, 2001.

13. Moving Pictures Expert Group ISO/IEC~JTC1/SC29 WG11. The MPEG Home Page. *URL: http://www.cselt.it/mpeg/.*

14. Goethals, Karen Metadata-Oriented Multimedia Information Retrieval. M.Sc. Dissertation, FEUP, 2001.

15. W3C Semantic Web Activity, *URL: http://www.w3.org/2001/sw/.*

ABOUT THE AUTHORS

Cristina Ribeiro (born in Espinho, 1958) got a Ph.D. in Informatics, Artificial Intelligence branch, at Universidade Nova de Lisboa, 1993. She is currently Auxiliary Professor at the Department of Electrical and Computer Engineering, Faculdade de Engenharia da Universidade do Porto (FEUP), lecturing on Data Structures and Information Retrieval. She has been a Researcher at INESC since 1985. Her research interests include constraint logic programming and information retrieval.
E-mail: *mcr@fe.up.pt*

Gabriel David (born in Lisboa, 1958) got a Ph.D. in Informatics, Artificial Intelligence branch, at Universidade Nova de Lisboa, 1994. He is currently Auxiliary Professor at the Department of Electrical and Computer Engineering, Faculdade de Engenharia da Universidade do Porto (FEUP), where he is also the Director of the Master Program on Information Management. He has been a Researcher at INESC since 1985. His main research interests are in Information Systems and Information Management. He is currently the leader of the project MetaMedia-2 (Portuguese FCT) on multimedia archives.
E-mail: *gtd@fe.up.pt*