METOKIS -TOWARDS A SEAMLESS CONTENT AND KNOWLEDGE EXCHANGE INFRASTRUCTURE

Wernher Behrendt, Sunil Goyal, Rupert Westenthaler

Salzburg Research Forschungsgesellschaft, Austria

{wernher.behrendt | sunil.goyal | rupert.westenthaler}@salzburgresearch.at

ABSTRACT

The building of interoperable knowledge based multimedia information systems requires a dialog between those research communities that make advances in content classification through e.g. image analysis, those working on standardisation of multimedia meta data, e.g. the MPEG community and those working on advanced knowledge-based end user applications that can feed and use the emerging Semantic Web. We present the working hypothesis of a new European research project which aims at developing a seamless exchange infrastructure between different data, knowledge and content management architectures and applications. The ultimate goal of the research is to demonstrate ways in which distributed information systems can manipulate complex information objects at a level of abstraction which is close to a conceptual model as it would be used by humans. We illustrate the motivation for the research and present the building blocks of the overall system whose status is work in progress.

1. INTRODUCTION

The research community for Image Analysis and Multimedia Interactive Services is making steady progress towards better handling of image data and of automating the generation of metadata from low-level image data. Researchers in the adjacent field of knowledge and content engineering are at the same time, postulating new kinds of information systems which are meant to put into practice, the vision of "Ambient Intelligence" [1]. Using this Ambient Intelligence Space people and organisations can exchange complex information objects that combine multimedia, hypertextual navigation and statements about their content in some knowledge representation language. The seamless mediation across system borders will be enabled by Web Services or software agents. In this paper, we explore the design concepts behind such systems from the viewpoint of METOKIS, a European research project in the 6th Framework Programme. The aim of METOKIS is to define a common structure for Knowledge-and-Content objects, based on previous research done by partners, and to define a complementary infrastructure which these objects use as their "habitat" and which provides operators to manipulate the knowledge objects in more flexible ways than is currently possible.

2. EXCHANGE OF MEDIA IS NOT SEAMLESS AT PRESENT

We shall illustrate some of the current discontinuities that hinder the implementation of an Ambient Intelligence Space for knowledge based multimedia at present by illustrating a scenario from the field of cultural studies.

Experts in Arts and Literature explore a field called *intertextual studies*. They are interested in the relationships that exist between different works of art, irrespective of the mode of artistic expression. For example, the famous book "Don Quixote" by Cervantes is a parody on a romantic epic called "Orlando Furioso" by Ariost. In a recent project, a multimedia authoring tool was built that allowed the cultural experts to create hypermedia links between different multimedia resources (e.g. an image of the paintings, the corresponding text in the novel, and the equivalent scene in the musical). At the same time, the experts were able to use an ontology which they had previously defined, in order to ascribe a formal type to each of these relationships. This way they were able to express that "Don Quixote" is_Parody_of "Orlando Furioso". Similarly, a specific scene of the video may be ontologically annotated as "The concept of fighting windmills" and the same is true for the corresponding part of the text in the original novel. But in today's working environment, as soon as these knowledge-and-content object leaves its specialist presentation environment, it loses its semantic richness! For example, if we display it as HTML we have to translate the semantic relationships into navigational links thus losing their conceptual specificity: "Parody" and "Medial transformation" both become just "GoTo"!

The challenges posed to both research communities by the above application domain, is as follows: Assuming some cultural expert has created an intertextual thread about idealism and the shared concept of "fighting windmills", can the image analysis community through large scale image search extend such intertextual threads by adding more examples?

Conversely, will there be a way for the knowledge based systems community to associate partially recognised media ("there is a person holding a spear, and there is an unrecognised other object in the image) with existing ontologies? The argument put forward here is that even if both research communities could in principle, solve the research problems of recognition and classification, we would still need a common information structure to place the partial truths that our recognition/classification components have found.

3. SELF DESCRIBING CONTENT

Despite all the work that is being done in developing e.g. OWL, the Web Ontology Language there will simply be no Semantic Web, in our view, unless we start working on an environment that is able to associate operational semantics with ontological mark-up. The reason for this is that all semantic descriptions remain vacuous until used in an operational context (see Hendler for a similar assessment [2]). The description languages offered so far have not addressed how they will connect up with the user-level tasks that will rely on system operations to be carried out on the basis of semantic annotations. An operational semantics needs to define operators and operands. Likewise, a seamless digital media (i.e. content) and knowledge environment needs to have a common view on these operators and operands in order to define flows of information through this environment.

This section looks at four aspects that we believe are central to self-description of content, and can be summarised in the following diagram (Fig. 1). They are an attempt to shed light on the confusion surrounding the ad-hoc distinctions that abound, between data, meta data, and knowledge about content.

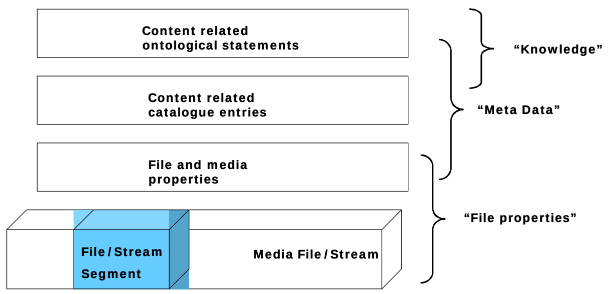

Fig. 1 Functional layering for content self description

At the bottom of the layered content object is the actual media data as it will be used by the media rendering software. Next come the file related properties, which identify content (fragments, shapes or the whole document) and the associated properties such as ownership, access restrictions, file length, encoding used etc. The third level provides the first real world description of the objects beneath, by describing them through either existing content standards which essentially create a view on both, lower level properties and "knowledge" level properties e.g. by allowing keywords or even a simplified description of the actual media content or using a richer ontology (WordNet [3]). At the fourth level, we expect to find statements that tell us the existence of different real world relationships that hold amongst objects in the content.

Each layer provides a functional view starting from identifying content at a system level; identifying objects at a Content Management System level to finally identifying real world relationships. The need for richer semantic is required at each of these three levels. The existence of a standard with richer semantics (WordNet) or poorer semantics (e.g. Dublin Core [4]) at the middle layer implies that the description of the objects is accordingly richer or poorer respectively, yet with the layer fulfilling its core responsibility of describing objects. What the research communities need to agree on is a conceptual model, which makes it clear "where to put information depending on its characteristics". For example, having detected a moving object according to an MPEG-4 definition means putting the shape information of that object into the second layer of the object. Having detected that the moving object is a person riding a horse might be a catalogue entry for a content management system. Having identified the riding person as Don Quixote attacking the Windmills and being thrown off his horse would be a narrative that should be described at the fourth level of the model. What is important is to have specific functionality associated with each layer and to have clearly defined interfaces between the layers.

The following three sections summarise the main characteristics of prototype research systems that have already attempted to clarify in which way multimedia content and knowledge about the content could be usefully organised.

3.1. Enhanced Multimedia Meta Objects (EMMOs)

The original motivation for EMMOs [5] came from the realisation that a body of knowledge is typically not just attached to singular concepts, but arises from an understanding of how one concept (or aspect of a concept) relates to another. Chen's Entity Relationship Model [6] captures this intuition for the design of relational databases and extensions of the model even hold for object-oriented database modelling. In addition to the basic relationship modelling, EMMOs are defined as a higher level data structure with predefined generic kinds of interfaces. This way, the EMMO affords a standardised and simplified manipulation language at a higher level of conceptual modelling than is possible with a direct mapping from the (E)E-R model to a relational or object-oriented DBMS.

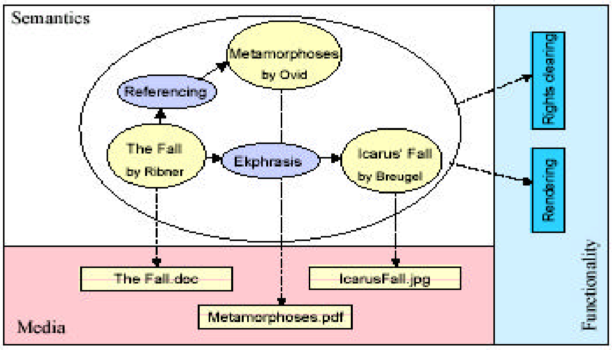

Fig. 2 Capturing functionality, semantics and media related information through EMMOs [7]

For example, one way of accessing an EMMO is through its media level interface. This interface is able to present to an application, an MPEG-7 compatible view of the object. Likewise, the EMMO defines a knowledge level interface through which one can access the ontology according to which the content of the EMMO has been annotated. Finally, the actual instantiations of concepts and relationships in a specific knowledge thread can be accessed and rendered via a transformer, into navigable hypermedia. Future work includes the extension of the transformers to create hypermedia threads "on the fly" as responses to queries using OWL-QL [8].

3.2. Information Objects for Knowledge Trading

In the INKASS project an attempt is made to provide a comprehensive model of knowledge resources, which enables organisations to trade knowledge objects of any sort in B2B eCommerce market places. The model comprises at present 10 ontology "facets" [9] which specify different aspects of the information object. For example, the *community* facet acts as an interface to agents that are involved in the trading or manipulation of the information object.

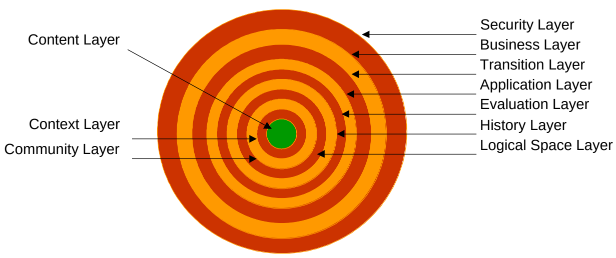

Fig. 3 Information Ontology for Knowledge Assets

The knowledge-market-driven approach taken by Information Objects complements the EMMO approach which focused its contribution on detailing the relationship between knowledge representation and segments of digital media, while acknowledging the necessity and existence of other layers of functionality, e.g for trading of rights. Future work in the METOKIS project is aimed at merging the models.

3.3. Core Ontology for Information Integration

One of the most significant bodies of work towards integration of knowledge assets with media resources has come from the area of digital libraries and cultural heritage institutions. The ABC Ontology proposed by Lagoze and Hunter [10] had the goal of harmonising a number of metadata standards in the world of digital library systems binding different metadata schemes by defining a generic "foundational" ontology [11].

Although the model is likely to guarantee migration from individual data models to a common data model, the foundational ontology (as it does not take into account operational semantics) may be biased towards static semantic equivalence that may not be maintainable once the process-specific semantics enter the scope of the model. With a functional view on building ontology-based systems, each layer can perform its operations without affecting the semantics of the overall system. The information objects therefore need to be self describing in terms of their functional characteristics covering every aspect from medial, low level multimedia properties, catalog entries to knowledge level primitives including both domain semantics and operational semantics.

4. METOKIS - BUILDING AN INFRASTRUCTURE FOR EXCHANGING KNOWLEDGE AND CONTENT OBJECTS

Experience and history show that many technological advances happen through some sort of duality: today's level of motorisation would be impossible if we still had the road system of the 19th Century. Developing an advanced authoring tool to combine knowledge and multimedia made us realise that we were facing this sort of dilemma: There is, at present no infrastructure that can make use of the fancy knowledge networks we can capture in an EMMO. Similarly, the Information Objects of the INKASS project rely at present on being exploited through a single-product content management middleware. Whenever these "intelligent objects" hit the normal world, they loose all their advanced properties. For example, exporting an EMMO into Macromedia Director is possible in principle through XML, but the ontology relationships which makes the EMMO "smart" are lost in the process. The conclusion is that it will be difficult for the new technology to demonstrate its value when it is constrained by yesterday's infrastructure.

4.1. Knowledge Content Objects

Knowledge Content Objects will bring together the lessons learned by the projects described in the previous section. Since there is now a body of world-wide experience with knowledge based multimedia document models, the next advance will be to bring those together to achieve greater momentum.

4.2. Knowledge Content Carrier Architecture

The novel aspect of METOKIS will lie in its approach to defining the infrastructure by which knowledge content objects can be created, manipulated, and exchanged. Our aim is to make knowledge workers and their tasks the focus of attention. So the working hypothesis is that the functionality of the surrounding systems should be defined in terms of the support they give to these user tasks. Moreover, we view content as preconditions and post-conditions of the user tasks. This way, there will be some similarity between workflow systems and the way we model content flows which will depend on our definitions of user task models. There will be three different end user applications: one is the definition of treatment protocols for clinical trials in which we expect a close relationship with existing workflow models. A second application concerns a publishing workflow for educational products, which is supported by semi-automatic content aggregation through a content management system. The third application is an information supply environment which supports senior executives in their decision making, by combining various kinds of news feeds into a personalised information environment. Each of the applications has a different technical focus: the clinical trials software will investigate the relationship between knowledge content objects and user interfaces on the one hand, and relational databases as storage backends on the other. The content aggregation application will investigate to what extent knowledge content objects can automate the process of building new educational products from existing resources. The senior executives information system will extend an existing electronic news filtering system to a personalised, targeted filtering and aggregation platform via knowledge level modelling.

The system will have to also show that not only can each of the applications be built and run in isolation, but it also needs to allow for cross-application content exchange. For example, doctors who have defined a clinical trial for some drug should be able to receive a personalised news environment in which information relevant to their trial plays a prominent role.

4.3. Knowledge Content Carrier Protocol

One of the discussions currently under way is to what extent the middleware proposed by METOKIS will actually lead to a system architecture in the traditional sense. If one thinks in terms of groups of functionality that should be provided by the middleware then it is also possible to define a set of operators within a protocol between the components of existing web based, knowledge and content management systems.

5. CONCLUSIONS

We have presented the working hypotheses for an advanced information systems architecture based on open content standards whose objective it is to enable producers of content to create units of added value which can be traded and manipulated in a controlled fashion. Furthermore, by combining object level, system level and protocol level into a coherent view that acknowledges advanced features of current systems (mobile code, rule bases, ontologies, object-based multimedia standards) we hope to contribute to the development of a next generation world-wide web that embraces the new possibilities, but keeps the feature that made the original WWW the unforeseen success that we have all witnessed: this feature was simply: content, system and protocol interoperability!

6. REFERENCES

[1] ISTAG 2001, Scenarios for Ambient Intelligence in 2010. Final Report, February 2001, ftp://ftp.cordis.lu/pub/ist/docs/istagscenarios2010.pdf

[2] Hendler J., “Agents and the Semantic Web,” IEEE Intelligent Systems, vol. 16, no. 2, Mar./Apr. 2001, pp. 30–37.

[3] Miller, G. A., Beckwith, Fellbaum, Gross, and Miller, K. “Introduction to WordNet: An On-line Lexical Database.” 1993.

[4] Dublin Core Metadata Element Set, Version 1.1, 1999. http://www.dublincore.org/documents/dces/

[5] S. Reich, Behrendt, Eichinger, and Schaller, Document Models for Navigating Digital Libraries. Proc. of the International Conference on Digital Libraries, Kyoto, 2000.

[6] Chen P., The Entity-Relationship Model - Toward a Unified View of Data. ACM Trans. Database Syst. 1(1): 9-36 (1976)

[7] Schellner K., Westermann U., Zillner S., Klas W. - CULTOS: Towards a World-Wide Digital Collection of Exchangeable Units of Multimedia Content for Intertextual Studies TR-2003302, Dept. of Computer Science and Business Informatics, University of Vienna, June, 2003

[8] Fikes, R.; Hayes, P.; & Horrocks, I. OWL-QL - A Language for Deductive Query Answering on the Semantic Web. Knowledge Systems Laboratory, KSL-03-14, Stanford University, Stanford, CA, 2003

[9] Abecker A., et al. Towards an Information Ontology for Knowledge Asset Trading, 2003 http://imu.iccs.ntua.gr/Papers/C59-ICE_2003 INKASS_Final.pdf

[10] Lagoze and Hunter, The ABC Ontology and Model, Proceedings of the International Conference on Dublin Core and Metadata Applications 2001, pp160--176, National Institute of Informatics, Tokyo, Japan,

[11] Hunter, J., “Enhancing the semantic interoperability of multimedia through a core ontology”. IEEE Transactions on Circuits and Systems for Video Technology, 13, 2003.