Multimedia Information Disclosure in a Distributed Environment

Daan Velthausz
Telematics Research Centre
P.O. Box 589, 7500 AN Enschede
The Netherlands
D.Velthausz@trc.nl

Introduction

As the ubiquitousness of multimedia information has been encouraging the diversification of the user's need for information, each user may have his own individual requirements regarding both content issues and non-content issues. Examples of the first can be requirements with respect to accuracy, completeness and up-to-dateness. Examples of the latter are requirements regarding data format, processing time and cost. Not only is the information environment heterogeneous and distributed, there are also price tags attached to the different parts of the information chain. The price may depend on many different aspects, such as information quality, connect time, processing time, network usage, etc. Most research thus far has focused on content-based issues. The importance of non-content issues is a direct consequence of the use of multimedia data, and is only slowly being acknowledged. We advocate that an advanced multimedia information retrieval service should offer users the possibility to express both content and non-content criteria to guide the disclosure process.

An information retrieval strategy (the total set of decisions and actions taken throughout the conduct of a search that affect the outcome of the search) should be capable of dealing with multiple information repositories, multiple media types and different types of (individually tailorable) user needs. Furthermore, the retrieval strategy should include platform and network characteristics in the information retrieval process, since they put constraints on the amount and form of the retrievable information. As retrieval of information in a distributed and dynamic environment also requires knowledge of the characteristics of the information repositories, such as coverage, origin, access mechanism, required user effort, performance in terms of costs, access speed and usage statistics (average recall, precision) [RaRM93], a common understanding of the properties of information repositories is preferred in order to select the information repositories to be searched.

Within the ADMIRE (ADvanced MUltimedia INformation REtrieval) project we are currently developing a retrieval strategy that is capable of retrieving multimedia information in a distributed environment. This retrieval strategy is based on existing techniques, such as a depth-first approach, but also takes the non-content issues, such as time and cost aspects, into account. Another strength is its combination with the ADMIRE information model [VeBE94]. Because of this model's capability, the querying and retrieval of information can be achieved in a more standardised and user-friendly manner.
Before discussing the strategy in more detail we will briefly outline the characteristics of the ADMIRE information model.

Information model

Electronic information provision is taking place in a distributed, heterogeneous, networked environment. This is an open environment in which geographically separated information 'producers' and 'consumers' interact via a distributed system. In order to disclose information in an efficient and effective manner, a model of the available information is required. The information model should allow search and manipulation of stored information in a distributed environment without having to deal with the actual data. Conventional models are not suitable for an efficient and effective disclosure of multimedia information.

In the ADMIRE project we have developed a comprehensive and extensible multimedia information model that offers support for a wide variety of information needs (e.g. exact, fuzzy) and disclosure methods (ranging from query formulation to navigation). The ADMIRE model can be considered as a generalisation of existing multimedia information models. It provides a framework for all media types, regardless of type, size or abstraction level. The ADMIRE model structures the information in a uniform manner via information objects and relationships. An information object (a finite, independent, piece of information) is characterised by the following property types:

- *raw data*: a sequence of elementary data units;
- *format*: denotes the representation of the raw data, e.g. 'JPEG image';
- *attribute*: characterises the raw data and its format in relation to the external world, e.g. 'creation date';
- *feature*: a domain-independent representation of the raw data and the format, examples are 'colour histogram' for images, 'word occurrence' within a textual document, 'camera and object motion' of a video shot;
- *concept*: a domain-dependent and format-independent semantic interpretation of the other property types.

Note that both format and attributes add new information that cannot be inferred from the raw data; i.e. they are 'content-independent'. By using the data of the information object and (domain) knowledge, it is possible to infer the features and concepts that belong to the information object, via feature extraction and interpretation functions, respectively.

Information can often be retrieved by users at multiple levels of granularity. The minimum level is typically determined by the information provider, and can vary from an individual frame to a whole movie. In the ADMIRE model these different levels of granularity are reflected by distinguishing two kinds of information objects, viz. *basic* and *composite* information objects. A basic information object is the smallest piece of information, determined by the provider, that can be manipulated and retrieved as a whole. A composite information object is composed of one or more basic information objects and/or one or more composite information objects. By using composite information objects the data can be organised into logical components that are structured through a particular composition relation. We distinguish different types of relationships between information objects, such as temporal, spatial and semantic relations. This results often in a hierarchical structure with basic information objects as leaves. For example, a video can be subsequently decomposed into a sequence of scenes, shots and frames.

In our model, even complete information repositories can be modelled as (high-level) composite information objects via the five described property types. The properties of a composite information object can be inferred by using the properties of the constituent information objects. For example, in a repository containing two books characterised by respectively the concepts 'chimpanzee' and 'donkey', the concept 'mammal' might be inferred to characterise the coverage of repository.

Summarising, besides the support for flexible, ad-hoc, content-based queries (via concepts, features and attributes), the model is capable of dealing with logical and layout-based queries. This approach facilitates a uniform, flexible and media-independent disclosure of information. [VeBE94] provides a detailed description of the ADMIRE model and shows how the model can be employed effectively to disclose multimedia information and to deal with various kinds of information needs.

Retrieval strategy

In a heterogeneously distributed environment, it is likely that different matching methods (which determine the degree of similarity between the query and information objects) are used, even for the same property type, e.g. due to different representations within different repositories. Hence, the retrieval strategy needs to handle and integrate the outcome (i.e. relevance estimations) of different matching methods.

To be able to compare the similarity scores of different information objects with each other, the different scores need to be combined into one common similarity score. If these scores are determined via the same method, this single similarity score can be computed straightforward, based on the operators and assigned weights (specified in the query). If the similarity scores are determined via different matching methods, they are likely to be expressed in different units. If it is not possible to convert these scores into a common unit, one might use the ranking information of the information objects from the matching method. A more sophisticated approach is to normalise the scores, e.g. via score/(max.-score) per matching method. We are currently investigating which method to include in our strategy.

The dynamic nature of the distributed information environment requires that the retrieval strategy includes an information discovery mechanism, since it is likely that not all properties of each information object are directly retrievable at the beginning of a search. The 'discovery' of information objects is performed through the logical structure of the starting composite information object, i.e. a repository describing other repositories.

At the repository aggregation level not all information objects, stored in the repository, are directly visible. They are only available via the properties of the repository (composite) information object. If, for example, a user wants to find titles of books containing images of blue butterflies while no selection for an information repository has been made or specified in the query, it is likely that one searches an index of repositories, obtains a ranking and then searches the top few repositories of interesting (composite) information objects. The ranking of potential interesting repository (composite) information objects is determined based on their properties. This requires, however, that the properties in the query can be 'transformed' into properties of the repository aggregation level. For example, a query indicating the concept 'chimpanzee' might be transformed into the concept 'mammal' at the repository level to select the most appropriate repository.

One may argue that, given a universal relevance estimation for the information objects, the whole information space can be searched for potentially relevant information objects by simply selecting the appropriate number, say n, most relevant ones. Since the search takes place in a distributed environment, the location of where these relevance estimations are computed, becomes important, especially since we take cost and timing aspects (e.g. transport time) into account. If the computation takes place at the user (home/client) side, this means that if k (one level higher) composite information objects need to be searched, $k * n$ information objects need to be retrieved, while $(k - 1) * n$ of these information objects are discarded without being seen by the user. In this situation it is better to adopt a more sophisticated merging strategy. The algorithm to determine how many information objects should be retrieved, embedded in our strategy, is based on [CaLC95]. Instead of using ranking scores we use normalised estimated relevance scores. If the user requests at most *n* results, the estimated relevance for the information objects is taken into account when determining the amount of objects to be retrieved from these information objects. This method favours for information objects from repositories with high scores but also enables relevant objects from a poor repository to be ranked highly.

Furthermore, since we take cost and time aspects into account, the whole information space cannot be investigated completely. The retrieval strategy must make decisions as to which composite information object to investigate. We currently use a simple depth-first approach, where the estimated relevance score determines the next appropriate information object to be investigated in more detail. Time and cost can simply be taken into account by checking if there is time and money left to search the next most promising branch. A more sophisticated approach takes the timing and money aspects into consideration in the relevance estimation score. This new measure should be the basis for selecting the appropriate information objects to be investigated in more detail. Since time and money can be 'spent' during a search session, the estimated relevance scores cannot be considered constant. Hence, a recalculation of the estimated relevance scores should be performed. The question remains as to when this recalculation should be performed. If, for instance, one branch search is completed, the leafs are reached and matched, and the relevance estimation is recomputed for all next accessible branches.

Current status

Within the ADMIRE project we are currently designing and implementing a prototype of the retrieval strategy, in which the information is modelled according to the ADMIRE information model.

References

[CaLC95] J. P. Callan, Z. Lu, W.B. Croft, *Searching Distributed Collections With Inference Networks.* In 18th ACM SIGIR 1995.

[RaRM93] Rao, R., D.M. Russell, J.D. Mackinlay, *System Components for Embedded Information Retrieval from Multiple Disparate Information Sources*. In Proceedings of the ACM Symposium on User Interface Software and Technology, Atlanta, November 1993, p. 23-33.

[VeBE96] Velthausz, D.D., Bal, C.M.R., Eertink, E.H., *A Multimedia Information Object Model for Information Disclosure*. To appear in proceedings of the MMM'96, Third International Conference on MultiMedia Modelling, Toulouse, France, 12-15 November 1996