An Extensible Framework for Personal Cross-Media Information Management

Adriana Ispas, Beat Signer, and Moira C. Norrie

Institute for Information Systems, ETH Zurich
CH-8092 Zurich, Switzerland
{ispas,signer,norrie}@inf.ethz.ch

Abstract. Nowadays a user's personal information space is fragmented into multiple repositories on their local machine as well as on remote servers. In order to enable later access to resources managed within such a cross-media information space, information has to be organised in a format that can be processed by automatic retrieval processes. We propose a general framework for personal information management based on extending a cross-media link server with supplemental metadata functionality. In addition to user generated information, our solution automatically derives metadata for classifying and associating resources based on direct interaction with the information space. Resources and metadata can be integrated by referencing external resources or information may be managed directly by the framework. The presented cross-media information management solution is not limited to a fixed set of predefined resources and can be extended based on a resource plug-in mechanism.

1 Introduction and Motivation

The pervasiveness of information that we face today leads to the "*ineptitude of getting at the record*" as already predicted in 1945 by Vannevar Bush [1]. A user's ability to gather collections of resources that are relevant for a specific task is often limited to partial or incomplete results. Where automatic retrieval processes fail, we are forced to remember the location or means of obtaining these resources.

The distribution of resources and metadata across various repositories, such as sites for sharing and managing multimedia, email clients or native storage systems, hinders an efficient information retrieval process. The duplicity of resources and the poor semantic interconnections between information items are consequences of information fragmentation. Another problem in the retrieval process is the human tendency to express searches in vague natural language [2] which results in the so-called *semantic gap* [3] between the actual information need and the formal expressions required by the retrieval services.

An effective solution to satisfy a user's information needs depends on providing the appropriate tools for information organisation, management and retrieval based on an integrated view of different information repositories [4]. While such an integration may be tackled by various technological approaches, the semantic

CEUR-WS.org/Vol-437/paper3.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

gap can be reduced by attaching metadata to information resources. However, the numerous metadata standards lead to a so-called "*alphabet soup*" [5] where standards are seldom interoperable and the handling of a request over multiple repositories results in an information lossy transformation process [6].

In an attempt to provide solutions for personal information management, we have to address at least two major issues. First, important metadata as well as information about a resource's original context gets lost when resources are copied between repositories. For example, once an email attachment has been copied to the file system, it loses metadata such as information about the sender, the date the document was received or the email's topic and content. A second issue deals with the capture of indirect metadata sources that can be derived from an individual's working behaviour and interaction with the information space. For instance, an event in the news that happened the same day the user created the first version of a document is a cue which might be used as an entry point when trying to recall a specific piece of information.

We propose a solution for cross-media information management based on capturing a user's semantic traces while interacting with different resources and storing this metadata within their personal cross-media information space based on the dimensions shown in Fig. 1. *Classification* results from placing resources into various semantic collections, such as named directories, the collection of resources bearing the same tag or resources used within the same website. In this way, we can gather metadata by observing explicit user-driven tagging or classifying actions within existing applications as well as classifying resources automatically based on their original contextual information. The first case is depicted in Fig. 1 as *manual classification* and the latter by *automatic classification*. Note that, in our approach, resources may be placed in several semantic collections at the same time (multiple classification). The *associations* category denotes the correlations resulting from resource usage patterns and scenarios as well as from recording annotations and metadata attached to the resources within their original context. While the *manual associations* denote explicit user annotations, *automatic associations* are system generated by gathering and structuring the relevant content.

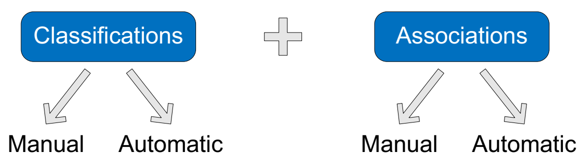

Fig. 1. Contextual correlations stored in a personal information space

We postulate that the recall of information can be improved by providing a "*selection by association*" based on a flexible and extensible model for managing both data and metadata within a personal information space. A unified view of a user's entire collection of resources is obtained by storing resources in the local repository or referencing external resources based on a resource proxy approach.

This also includes the capture of information cues derived from interacting with the information space. Motivated by the argument that it is easier to remember information cues related to their data rather than structured schemas or formalisms [7], we support the information retrieval process by automatically storing cues that can later be used as entry points into the cross-media information space and linking these cues and other metadata with content elements.

We start in Sect. 2 by describing our data model for cross-media information spaces. The corresponding cross-media information management framework is introduced in Sect. 3. Section 4 provides a comparison of our approach and existing solutions for personal information management. Concluding remarks are given in Sect. 5.

2 A Data Model for Cross-Media Information Spaces

Our model for personal cross-media information spaces is defined using the semantic, object-oriented data model OM [8]. OM supports information modelling through a two-level structure of *classification* and *typing*. Typing refers to the representation of entities by means of objects with attributes and methods, whereas classification deals with their semantic roles. A particular classification is represented by a named *collection* of objects with a specified membertype. Just as types can be specialised through subtypes, the classification of objects can be specialised through subcollections. The OM model was used to define the resource-selector-link (RSL) model [9], a general link metamodel that generalises concepts found in hypermedia systems and is used to implement the iServer cross-media information platform [10]. The model presented here is an extension of RSL. Note that it is beyond the scope of this paper to describe the RSL model in detail but further information can be found in [9].

The schema of our extended model is shown in Fig. 2. The shaded rectangles denote collections of objects (classification) with the name of the collection in the unshaded part and the name of the associated type in the shaded part. The shapes with a solid outline are part of the original RSL model, whereas shapes with a dashed outline form part of the model extension. In the following we describe how the RSL concepts have been used to realise the ideas presented in the previous section and present the required extensions.

The central concept within the RSL model is the generic notion of an `entity`, the instances of which are classified by the collection `Entities`. An `entity` is an abstraction for all possible information items which might be handled in a personal information space. The abstract `resource` type is used to represent any kind of information resource such as different types of documents, images, videos or even physical resources, but also metadata structures and information cues that are semantically correlated with the information resources. Specific extensions of the `resource` type have to be implemented to represent particular types of resources. In Fig. 2, we show the extension for information cues with the type `cue` and its subtypes `visualCue` and `temporalCue`.

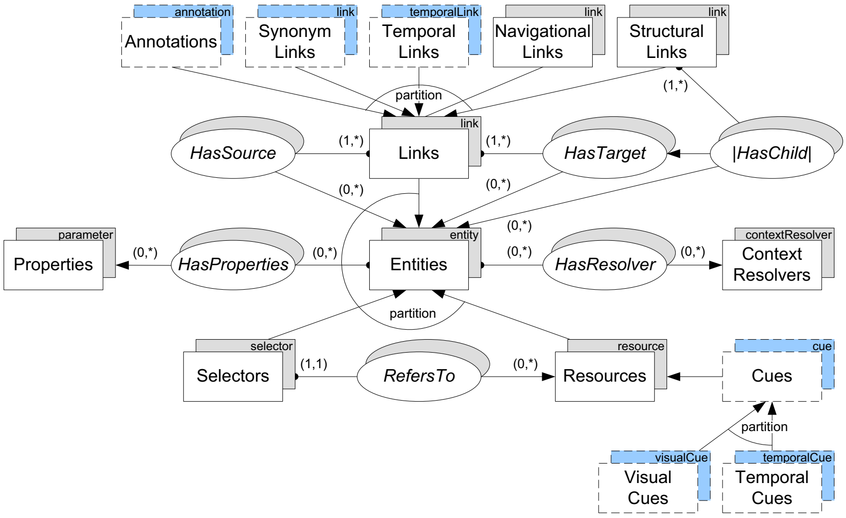

Fig. 2. Data model for cross-media information spaces

Often the possibility to address an entire resource is not flexible enough. For example, it may be just the section of a web page mentioning the contact details of a person that raised the user's attention. The `selector` entity subtype enables parts of resources to be addressed. As in the case of the resources, selectors may be specialised for concrete types of information item fragments. Note that it is not our aim to store all information resources in a centralised repository, but rather to provide unified access to resources in accordance with their relevance for specific tasks. Some resources, including metadata information, may act as proxy elements to external information units. A flexible classification of resources is supported through the creation of new collections and subcollections at runtime and the multiple classification feature offered by the OM model.

The `link` concept offered by the RSL model is used to store associations between information resources by connecting several source and target entities. This allows links to be defined between different types of resources, parts of resources or even link entities. To illustrate how information is represented in our extended model, let us have a look at the following scenario:

First thing in the morning, Adriana checks her email. An interesting talk organised by the department has been announced. She immediately searches the Web to find more details about the speaker.

We have to create three new resource objects representing the mentioned event, the speaker and the consulted web page. We then create a temporal cue object to represent the current date and system time from which we can later also derive the time of the day (e.g. "in the morning"). The semantic associations between objects are modelled through a temporal link having the temporal cue object as the source and the objects representing the event, the speaker and their personal web page as the targets.

In addition to the new temporal link type, we provide other link extensions represented by the `StructuralLinks`, `SynonymLinks` and `Annotations` collections shown in Fig. 2. The structural links may be used to represent specific resource structures. For example, we could represent the structure of an email as having a sender, one or several receivers, potential attachments, a subject and the content, which again may have a certain structure.

The synonym links are used to connect various representations of what is essentially the same information. For example, we can link together the address book entry, the contact details section of a paper and a web profile referring to the same person, thus obtaining the "synonym ring" of the resource. User annotations are recorded as links between the resource being annotated and the resource representing the content of the annotation. This means that a resource may be annotated not only with textual information but with any other resource, resource fragment or even with other links.

We mentioned earlier that we also collect metadata from a resource's original context. While it is out of the scope of this paper to describe the metadata import mechanism, note that the same principle of linking entities through specific types of links can be used. Therefore, a set of new link types is created in accordance with the specification of the underlying metadata structure standards. For example, to connect resources to their *Dublin Core*[1] descriptions we use links annotated with resources representing the Dublin Core element set.

Finally, to refine the retrieval process, the model also includes the `Properties` and the `Context Resolvers` elements. Whereas additional information such as a link's weight might be stored as a specific link property, the context resolvers control an entity's "visibility" that might be based on information derived from specified cues or other contextual information.

3 Cross-Media Information Management Framework

The main components of the proposed cross-media information management framework are shown in Fig. 3. An extensible *Resource Tracker* component observes user interactions with different *Data Providers* such as the local file system, the mail server or the Web. Specific plug-ins help to improve the granularity by, for example, tracking visited Flickr albums and pictures or sections of PDF documents. Each resource access is detected by the corresponding resource tracker extension and reported to the *Information Manager* component which checks if the resource has been accessed before and is already available in the personal information space. If necessary, a new resource object is created and classified in different semantic collections based on criteria such as its type or origin. Another process creates resource objects in correspondence to any metadata associated with the resource in its original context. In this way, external metadata is either copied into the local repository or referenced by metadata resource proxies. Finally, the *Annotation Service* is responsible for recording

[1] http://www.dublincore.org

any user-defined annotations of visited resources. Note that if the original data provider offers some annotation functionality, the framework will rely on it, otherwise specific annotation components have to be implemented. The *Information Manager* is also responsible for creating specific links between a content resource and its related metadata and annotation resources. In addition, resource access patterns are tracked and stored as a special type of links between succeeding resources. Once the user shifts their work to another resource (e.g. another web page or document), the same mechanism is triggered again.

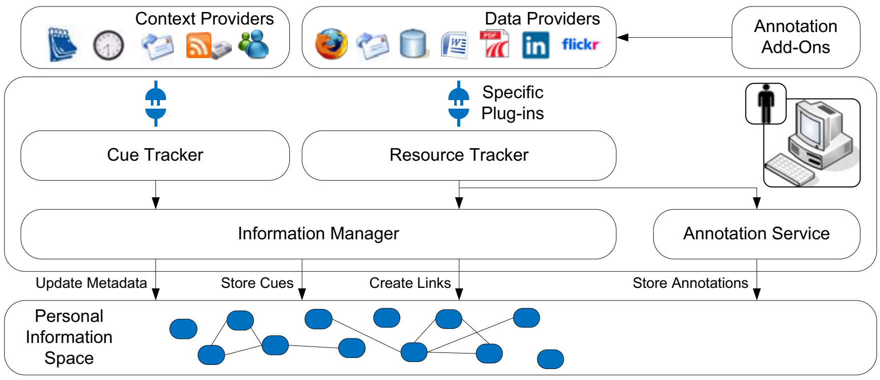

Fig. 3. Personal cross-media information management framework

The *Cue Tracker* component monitors various *Context Providers*. While Fig. 3 shows calendar applications, the date and system time information, email clients as well as other context providers, new sources of information cues may be added based on a plug-in mechanism. In accordance with devised policies, the *Cue Tracker* reacts to specific events within the tracked sources. For example, temporal cues are constructed based on time intervals marking off the duration of a conversation, the morning of a certain day or a news event possibly spanning several days. Once cues are detected, the *Cue Tracker* contacts the *Information Manager* which creates specific cue objects within the personal information space. Cues are maintained in a collection of active cues as long as their extent has not yet expired. As new content resources are detected, these are connected through specific links with all active cues. Since our system supports links with multiple sources and multiple targets, each resource is added as a new target of a link having the cue as source. Access counters attached to links monitor how often a link has been visited either in relation to a search or to connect more resources. Co-occurrences of cues within the active cues collection may entail further semantic information that we plan to investigate.

Retrieving resources may be performed by simply following links between resources. A user may specify unstructured details on specific topics, such as information cues that they remember. Once the objects corresponding to a user's query are found, trails of links may be constructed and connected resources may be returned.

4 Related Work

Several other projects have recognised the need for a uniform information space and the importance of semantic metadata for the retrieval of resources. Earlier proposals have indicated the necessity of paradigms for information space organisation. While these approaches followed just one or a few dimensions of the information space, organising resources in terms of time [11] or other attributes derived from content and context [12], we propose a framework capable of dealing with an extensible set of information cues. These cues result from classifications and associations generated while a user is interacting with their information space and they are represented as additional resources in the information space. In contrast to the *Stuff I've Seen* [13] system where various landmarks guide the user to the target of their research within a computed search result, our approach allows metadata resources to participate in both the computation and the presentation of search operations.

Related approaches tend to use relational databases [14] or RDF [15] for the realisation of the personal information space. We support multiple classification and associative linking by using the OM and RSL models, while being independent of both data presentation and representations by supporting a plug-in mechanism for various data source representations.

The *MyLifeBits* [14] system manages personal information spaces by harnessing the functionality of a database system and storing all the digital information resources, metadata and cues within a single database. Various views are constructed based on typed links between resources and by offering the possibility to store resources in multiple collections. The storage component provides a limited set of resources (e.g. email) and the database schema needs to be changed to support additional types of resources. Furthermore, only simple binary links with one source and one target resource are supported. While *MyLifeBits* is faced with the problem of providing automatic classification within their fixed set of database tables, our underlying OM model caters for flexible classification structures and the membership within collections may be easily recomputed based on contextual information and information cues. Since our solution is not based on a fixed set of resources, it can be adapted to various metadata standards by providing resource extensions. Furthermore, since we can reference external resources, it is also possible to monitor changes in a resource's original metadata.

The *Social Semantic Desktop* [15] is an initiative for creating a user's personal Semantic Web on their desktop. The approach heavily relies on RDF and ontologies for managing personal information spaces. Data items from various software applications are captured, converted to RDF, indexed and stored in local or remote RDF repositories. One of their main concerns is mapping services to translate RDF graphs from one ontology to another and the definition of a context model ontology is still a challenge to be solved. A limitation of the Social Semantic Desktop approach lies in representation of relations between entities as binary properties[2].

[2] http://www.w3.org/TR/swbp-n-aryRelations/

5 Conclusions

We have presented a cross-media information management framework for managing a user's personal information. An essential feature of our proposed solution is the flexibility and extensibility in terms of both resource representations as well as the capture of metadata and information cues. Metadata and information cues are constituent parts of the information space together with the information resources, supporting the information retrieval process through a "selection by association" paradigm. Last but not least, the presented approach is flexible in storing resources either locally or integrating external resources based on a resource proxy mechanism.

References

1. Bush, V.: As We May Think. Atlantic Monthly **176**(1) (Jul 1945)
2. Jain, R. In: Infoscopes, Multimedia Information Systems. Borko Furht (1996)
3. Smeulders, A.W., Worring, M., Santini, S., Gupta, A., Jain, R.: Content-Based Image Retrieval at the End of the Early Years. IEEE Transactions on Pattern Analysis and Machine Intelligence **22**(12) (Dec 2000)
4. Franklin, M., Halevy, A., Maier, D.: From Databases to Dataspaces: A New Abstraction for Information Management. SIGMOD Record **34**(4) (Dec 2005)
5. Smith, J.R., Schirling, P.: Metadata Standards Roundup. IEEE Multimedia **13**(2) (Apr-Jun 2006)
6. Foulonneau, M., Riley, J.: Metadata for Digital Resources: Implementation, Systems Design and Interoperability. Chandos Publishing (Oxford) (2008)
7. Alvarado, C., Teevan, J., Ackerman, M.S., Karger, D.: Surviving the Information Explosion: How People Find Their Electronic Information. Technical Report AIM-2003-006, MIT, Cambridge, USA (Apr 2003)
8. Norrie, M.C.: An Extended Entity-Relationship Approach to Data Management in Object-Oriented Systems. In: Proc. ER '93, Arlington, USA (Dec 1993)
9. Signer, B., Norrie, M.C.: As We May Link: A General Metamodel for Hypermedia Systems. In: Proc. ER 2007, Auckland, New Zealand (Nov 2007)
10. Signer, B.: Fundamental Concepts for Interactive Paper and Cross-Media Information Spaces. PhD thesis, ETH Zurich (May 2006) Dissertation ETH No. 16218.
11. Fertig, S., Freeman, E., Gelernter, D.: Lifestreams: An Alternative to the Desktop Metaphor. In: Proc. CHI '96, New York, USA, ACM (1996)
12. Dourish, P., Edwards, W.K., LaMarca, A., Salisbury, M.: Presto: An Experimental Architecture for Fluid Interactive Document Spaces. ACM Transactions on Computer-Human Interaction **6**(2) (1999)
13. Ringel, M., Cutrell, E., Dumais, S., Horvitz, E.: Milestones in Time: The Value of Landmarks in Retrieving Information from Personal Stores. In: Proc. Interact 2003, Zurich, Switzerland (Sep 2003)
14. Gemmell, J., Bell, G., Lueder, R.: MyLifeBits: A Personal Database for Everything. Technical Report MSR-TR-2006-23, MS Research, San Francisco, USA (Feb 2006)
15. Sauermann, L., Bernardi, A., Dengel, A.: Overview and Outlook on the Semantic Desktop. In: Proc. 1st Workshop on The Semantic Desktop at ISWC 2005, Galway, Ireland (Nov 2005)