



ELSEVIER

Data & Knowledge Engineering 36 (2001) 251–281

**DATA &
KNOWLEDGE
ENGINEERING**

www.elsevier.com/locate/datak

Landscaping the information space of large multi-database networks

Mike P. Papazoglou ^{a,*}, Henderik A. Proper ^{b,1}, Jian Yang ^a

^a University of Tilburg, Infolab, PO Box 90153, 5000 LE Tilburg, The Netherlands

^b Ordina Institute, Groningenweg 6, 2803 PV Gouda, The Netherlands

Received 3 August 2000; received in revised form 3 August 2000; accepted 3 August 2000

Abstract

The promises of network-accessible information are increasingly difficult to achieve. These difficulties are due to a variety of causes, such as, the rapid growth in the volume of network-available information and the increasing complexity, diversity and terminological fluctuations of the different information sources available. This paper presents a conceptual architecture for the organization *information space* across collections of component systems in a multi-database network that provides serendipity, exploration and contextualisation support so that users can achieve logical connections between concepts they are familiar with and schema terms employed in multi-database systems. Large-scale searching for multi-database schema information is guided by a combination of lexical, structural and semantic aspects of schema terms in order to reveal more meaning both about the contents of an information term and about its placement within the distributed information space. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Multi-databases; Semantic information discovery; Information retrieval; Meta-data; Database clustering

1. Introduction

The dramatic growth in global interconnectivity has placed vast amounts of data within easy reach. At the same time it has made on-demand access to widely distributed information a natural expectation for users.

A complicating factor is the difficulty in providing coherent access and correlation of information that originates from diverse and widely distributed sources. This is an involved process, not only because of the sheer volume of information available, but also because of heterogeneity in naming conventions, meanings and modes of data usage. Differences in descriptions, abstraction levels, and precise meanings of terms being used in disparate sources do not yield well at all to automation. These problems are compounded by differences in user perceptions and interpretations, and variations that may occur at autonomous sources over time. Users are

* Corresponding author. Tel.: +31 13 466 2349; fax: +31 13 466 3069.

E-mail addresses: mikep@kub.nl (M.P. Papazoglou), e.proper@acm.org (H.A. Proper), jian@kub.nl (J. Yang).

¹ Also subfaculty of Informatics, University of Nijmegen, Toernooiveld 1, 6525 ED Nijmegen, The Netherlands.

presented with the problem of gaining adequate knowledge of a potentially huge, complex and dynamic system in order to access and combine information in a coherent and logical manner. Yet multi-database systems demand from users *prior detailed knowledge* of the definition and uses of their underlying data [31]. This expectation on a user's intellectual capacities is quite unreasonable in the context of a large distributed information space.

The focus in multi-database systems is on query processing techniques and not on how to discover where the actual schema elements in the component systems reside. No particular attention is paid to how schema items are structured, what they mean and how they are related to each other across component database schemas. The user's perception of the information content in networked databases is that of a vast space of information in a large flat, disorganized set of database servers. In contrast to this, our approach to searches for widely distributed information concentrates on providing a dynamic, incremental and scalable logical organization of component database sources, and search tools that are guided by this organization.

We view user interaction with a multi-database space as comprising two major phases, the: *schema information discovery phase*, where users systematically explore the multi-database space to locate potentially useful databases, and

distributed query/transaction phase, where the requested data sets are retrieved from the candidate databases.

We consider the development of a methodical, scalable search process critical to the successful delivery of information from networked database systems. Hence, in order to provide users with tools for the logical exploration of distributed information sources a four-step process, termed *schema information discovery*, is introduced. This process encompasses the following steps:

1. *Determining* the information needs of users by means of different term suggestions.
2. *Locating* candidate database sources that address these needs.
3. *Selecting* schema items of interest from these sources.
4. *Understanding* the structure, terminology and patterns of use of these schema items which can subsequently be used for querying/transaction purposes.

The very nature of this process suggests that we should provide facilities to landscape the information available in large multi-database networks and enable the users to deal with a controlled amount of material at a time, while providing more detail as the user looks more closely.

To support the process of schema information discovery while overcoming the complexity of wide-area information delivery and management, we cannot rely on a collection of indexes which simply contain schema information exported by individual database sources. A more structured and *interactive* approach to searching is required. The precursor of such an advanced search approach assumes that we are in a position to impose some logical organization of the distributed information space in such a way that potential relationships between the component database systems in the network can be explored. In addition, to maintain scalability, this must be achieved through a decentralized mechanism which does not proceed via a one-step resolution and merging of system information into a single static monolithic structure as advocated by many conventional practices for integrating multi-database systems. These and related issues are addressed in this article.

This paper presents the concept of schema information discovery for large multi-database networks. The paper is organized as follows. Section 2 presents related work. In Section 3 a logical

organization for the semantic cross-correlation of meta-data information from component databases in a multi-database system is defined formally. This logical organization of meta-data forms the core of our conceptual architecture for information space. Section 4 presents clustering techniques that allow the information space to be populated with available database nodes. Navigation and querying mechanisms to navigate and query the resulting information space are provided in Section 5. Finally, Section 6 presents some experimentation results while Section 7 presents our conclusions and future work.

This work is an extension and elaboration of some early ideas outlined in [33,47]. In [33] we concentrated on the organization of physical data sharing in large database networks, and described how physical data sharing ties in with a precursor of the conceptual organization of the information space presented in this paper. In [47] we described techniques and algorithms used for the conceptual clustering of databases. In this paper we concentrate on an extension and formalization of these ideas and on the logical grouping of databases, according to subject and a common terminology context. In addition, we present navigation and querying techniques for understanding the semantic context of networked databases.

2. Finding information: An overview

This section starts by providing a broad discussion of the problem of finding information in vast information spaces. This is complemented by a discussion of a number of techniques from different fields for locating information.

2.1. The vastness of information space

The elementary building blocks of information space are the information assets themselves. The term *information asset* can be defined as:

any distinct information bearing entity that is accessible on a networked environment and which may be combined with other such entities connected to the same network.

A definition that truly supports the open character of the network. Examples of information assets included in this definition are:

- Web pages (including free-text, sound, images, and video fragments).
- Free-text databases, such as newsgroups, mailing list archives, etc.
- Digital libraries.
- Traditional (relational, object-oriented) databases.

This definition of information asset should give an indication of the potential vastness of information space. Although in this paper we confine ourselves to information assets in conjunction with multi-database networks, the techniques outlined herein can also be applied to other examples of information assets.

2.2. Dealing with the vastness of information space

Different techniques have been, and still are being, developed to deal with the vastness of information space.

2.2.1. Web-based searching

The growth in use of the World Wide Web (WWW) has led to the development of a variety of search engines which attempt to locate a large number of WWW documents by indexing large portions of the Web. These search engines tend to return many potentially relevant information assets. Users are still required to manually wade through large result sets in search of truly relevant assets.

Most approaches to WWW querying [7,32] concentrate only retrieval based on the contents of an information asset “as-is”. They naively assume that the user (or search engine) is explicitly aware of the structure, semantics and vocabulary differences of the information assets that are available to them. However, due to the multiplicity, complexity, and terminology fluctuation of the information available, such an assumption is not practical. Practical studies have shown that there is a critical mismatch between a user’s and the Web’s vocabulary [51]. Picking the right terms depends on how intimate searchers are with the vocabulary use in documents they wish to retrieve.

Centralized index search engines such as Lycos [35], Web Crawler [45] are manual indexing schemes that rely on techniques which “crawl” the network compiling a master index. The index can then be used as a basis for keyword searches. These systems are not scalable because they use a global indexing strategy, i.e., they attempt to build one central database that indexes everything. Such indexing schemes are rather primitive as they cannot focus their content on a specific topic (or categorize documents for that matter): as the scope of the index coverage expands, indexes succumb to problems of large retrieval sets and problems of cross-disciplinary semantic drift.

Some of the above limitations are addressed by content-based search engines such as the Content Routing System [52] and Harvest [12]. These systems generate summarized descriptions (*content labels*) of the contents of information assets. The Content Routing System creates and maintains indexes of widely distributed sites. In this distributed information retrieval system a collection of documents is described by means of a content label which in turn can be treated as a document and can be included in another collection. Content labels help users explore large information spaces. However, document collections and their labels are confined to the context of their underlying information servers. This idea has been extended in the HyPersuit system [56] by generalizing collections so that they may span documents from various servers.

The Harvest information discovery and access system [12] provides an integrated set of tools for gathering information from diverse Internet servers. It builds topic-specific content indexes (summaries from distributed information), provides efficient search mechanisms, and caches objects as they are retrieved across the Internet. Each local search engine builds a specialized directory for a certain domain of documents. Federated search engines scan those directories and form federated directories which aggregate documents according to application-specific needs.

2.2.2. Subject gateways

A subject gateway, in network-based information access, is defined as a facility that allows easier access to network-based information resources in a defined subject area [28]. Subject gateways offer a system consisting of a database and various indexes that can be searched through a Web-based interface. Each entry in the database contains information about a network-based resource, such as a Web page, Web site or document.

Advanced gateways provide facilities for enhanced searching. For example the Social Science Information Gateway (SOSIG) [53], incorporates a thesaurus containing social science terminology. This gives users the option of generating alternative terms/keywords with which to search the resource catalog. Another example of an advanced subject gateway is the Organization of Medical Networked Information (OMNI) [40] which allows users to access medical and health-related information. OMNI also facilitates searches across other databases of resources such as databases of dental resources.

The key difference between subject gateways and the popular Web search engines, e.g., Alta Vista [5], lies in the way the indexing is performed. Alta Vista indexes individual pages and not resources. For example, a large document consisting of many Web pages hyper-linked together via a table of contents would be indexed in a random fashion. In contrast this subject gateways, such as OMNI, index at the resource level, thus, describing a resource composed of many Web pages in a much more coherent fashion.

Furthermore, a subject gateway has the “luxury” of being able to focus on a specific subset of the information space. Usually an area for which some well-defined thesaurus is available.

2.2.3. Federated digital libraries

The most important problem specific to digital libraries with spatial distribution is the federation problem: making distributed collections of heterogeneous documents appear to be a single (virtually) integrated collection. Each such federation may address a specific domain area, e.g., biomedicine, computer science, social sciences and so on. In such Federated Digital Libraries (FDLs) the difficulty lies in transforming a federation of multiple semi-structured heterogeneous sources (which lack coherence) into a single logical source. Here, we are faced with at least two major technical challenges. Firstly, document handling is hard as there is a large number of documents with differing structure and differing terminology. Secondly, due to the large number and variety of documents available, unless classification schemes are employed – so that document sources can be indexed in different ways and different levels of detail – distributed searching cannot be feasible [51].

2.2.4. Multi-database systems

Multi-database (or federated) systems have as their aim the ability to access multiple autonomous databases through querying. The emphasis is on integration and sharing of distributed information and not on information discovery. A particular database may choose to export parts of its schema which are registered in a federal dictionary. A requesting database consults the federal dictionary for existing databases and then imports schema elements that it requires. While this approach might be appealing for a small number of interconnected databases it is clearly not scalable. Locating the right information in a large unstructured network of data dictionaries is extremely cumbersome, has limited potential for success and, more importantly, is error prone as it does not deal with terminology nuances.

Several research activities in the area have concentrated on the issue of creating semantically enhanced federated database dictionaries [8,10,17,38]. Construction of conceptual ontologies on the basis of domain-specific terminologies and formalisms that can be mapped to description logics are also discussed in [31]. Some of the issues relating to the identification of semantically related information can be found in [10], where the authors describe an approach that relies on an abstract global data structure to match user terms to the semantically closest available system

terms. Concepts grounded on a common dictionary are defined in a domain and schema elements from component databases are manually mapped to these concepts. A different approach is taken by Kahng and Mcleod [29], where a domain-specific classification scheme is built incrementally by considering one schema at a time and mapping its elements in a concept hierarchy. However, both these approaches tend to centralize the search within a single logical index thereby defeating scalability by introducing performance limitations for large networks.

2.2.5. DAI: Multi-agent systems

Most of the work on software agent systems has concentrated on improving information discovery methods on the WWW and adopt them for use within cooperating agent configurations. The protocols of the WWW provide purely keyword-based index services and look up in collections of documents. Most DAI approaches employ some form of knowledge representation to enable more sophisticated representation of information sources and inferencing abilities [9, 22,41]. Two of the most notable activities which are related to this work are: *information matchmaking* [27] and *information brokering using context logic* [22].

Matchmaking is an automated process whereby information providers and consumers are cooperating assisted by an intelligent facilitator utilizing a knowledge sharing infrastructure. Matchmaking depends on messaging and content languages and allows information providers and consumers to continuously issue and retract advertisements and requests, so that information does not become stale. This is particularly critical where information changes rapidly.

Brokering agents use distributed information resource discovery and retrieval facilities to assist service providers list and publish their services, and help seekers to find services and information of interest. Brokering agents in DAI normally employ disambiguation techniques that rely on ontological terms to make assertions that will assist them in accomplishing their mission. For instance, in business applications an information broker uses standard terms in the ontology to describe products in a particular domain and identifies correspondences between the standard term(s) and terms used by diverse companies to describe the same product. In this way consistent business semantics can be provided across different market domains and segments. Several projects that develop ontologies for use with the Web such as SHOE [25] and the Wide Knowledge Base (WebKB) [20] are used with agent technology. SHOE is a Web-based knowledge representation language that supports multiple versions of ontologies and ontology integration. WebKB uses ontologies and machine learning to achieve automatic classification of Web pages.

3. Logical architecture of a multi-database information space

In order to improve efficient searching and gathering of schema information in large multi-database networks, the first task is to partition the multi-database information space into distinct, domain-specific, categories that are meaningful to database users. These categories can be formed by using some form of topic/subject-based classification mechanism. Such classification mechanisms are common practices in library and information sciences, e.g., the INSPEC indexing and abstracting service covering most of the research literature in Computer Science and Electrical Engineering [51]. Using domain-specific classifications to create logical clusters of databases makes searches more directed, meaningful and efficient. In addition, a directory of topics created

as a result of domain-specific database categorization can also provide topic-specific searches and useful browsable organization of inter-component database schema information.

There are three basic principles that a system must address to allow for scalable information searching and gathering. Firstly, some organization of the underlying databases is needed to enable the discovery of data inter-relationships. Topic classification schemes are used for this purpose, as they summarize related information subspaces together. Secondly, this organizational structure must itself be scalable. In other words, both interactions with the resulting structure, as well as maintenance of the structure, must be scalable. Thirdly, users must be presented with a collection of tools (lexicographic, and user friendly graphical interfaces) which allows for easy exploration and interpretation of the information contents of the system. In the following, we address these issues in the context of a logical architecture for a multi-databases information space.

The logical architecture presented below has been among other inspired by ideas from the field of information retrieval, where documents are clustered to form a multi-layered information space [3,4,6] that can be navigated by users in search of relevant documents.

3.1. Basic building blocks for schema information discovery

Our approach to information searching and gathering in large database networks relies on logically partitioning a collection of networked databases into distinct topic-based categories that are meaningful to users. This occurs by creating abstract representations of these topics as logical objects (henceforth referred to as topics). Database-content clustering algorithms are employed to automatically compute sets of related component databases – via their exported meta-data terms – and associate them with an appropriate topic, see Fig. 1. The abstract objects representing these topics essentially represent centroids around which databases cluster, and are engineered to describe a particular domain. It is expected that the databases in a multi-database network target-specific narrow domains, such as Geophysics, Biomedicine, Economics, Chemical Engineering and so on.

To put the organization of a topic-based multi-database information space into perspective, we use a comprehensive example from an *Education & Training* multi-database network. This network connects educational and training service providers, publication providers, accreditation, and government agency database servers. This situation is shown in Fig. 1 which provides a conceptually holistic view and cross-correlates information from the multiple database servers (referred to as assets). In the following, we will describe this process in two broad steps.

Firstly, we employ *meta-data* to describe the structure and contents of each individual exported asset schema. Each exported schema may also contain additional meta-data elements such as those described in Section 3.4. Subsequently, distinct sets of meta-data terms are logically aggregated to describe a particular subtopic. For example, aggregation (summarization) of meta-data which abstract assets containing information about Courses, Committees, Accreditation-processes and so on, may represent a subtopic such as Accreditation. This subtopic is a composite object which contains meta-data objects describing assets and their attributes. For instance, the Course object and attributes pertaining to it such as Course-name, Credit-points, Duration, Semester, etc. would be contained in the Accreditation subtopic. The term contents of subtopics are organized in the form of graphs, which are called *context graphs*, see Section 4. Each subtopic is represented as a descriptor object in a *descriptor graph*, e.g., descriptors such as d_1, d_2, d_3, \dots in

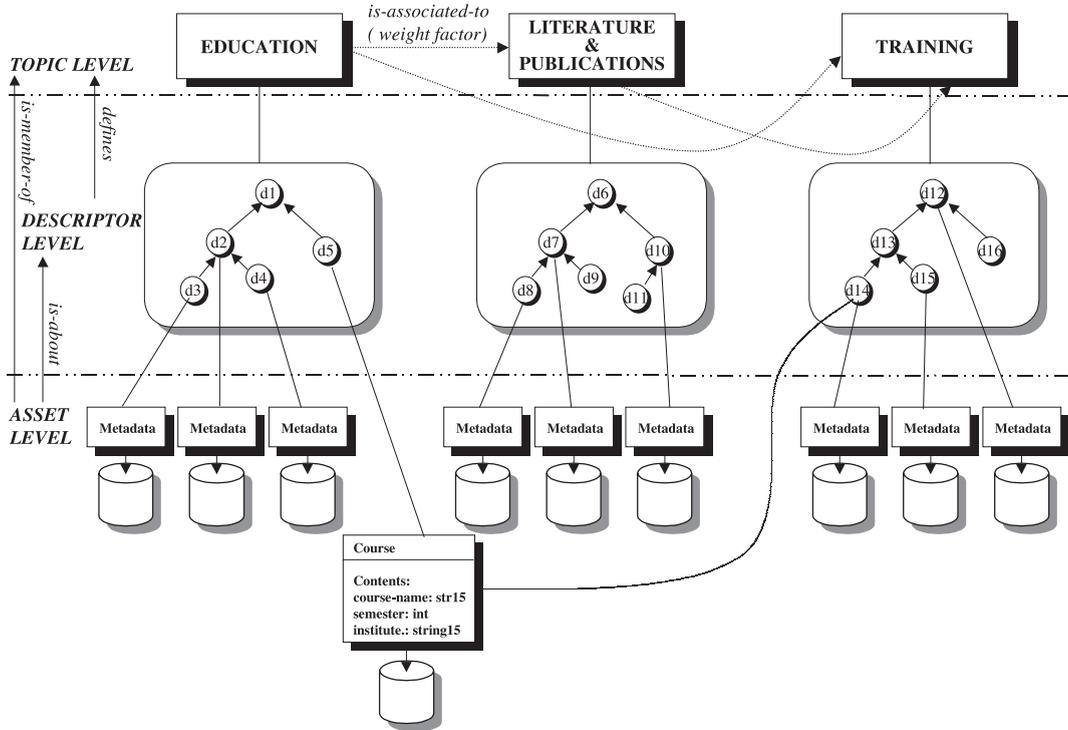


Fig. 1. Three-level organization of the information space.

Fig. 1. The aggregation of semantically related descriptors may span a specific topic. For example, semantically related subtopics such as for example Accreditation (descriptor d4) and Enrollment-Program (descriptor d5) are connected into a higher-level construct which we call a *topic*, e.g., Education, see Fig. 1. Topics thus represent semantically related database clusters (via their respective meta-data schemas) and form topically coherent groups that unfold descriptive textual summaries and an extended vocabulary of terms for their underlying assets. The area of interest of a topic is defined in terms of a *descriptor graph*. Each descriptor graph in turn is defined in terms of related context graphs. An example can be found in Fig. 2, where the meta-data for Course are represented synoptically and abbreviated context graphs for descriptors d5 and d12 are illustrated. This figure also illustrates that a meta-data object can be associated with (can contribute to) different subtopics.

Secondly, to circumvent terminology mismatches and semantic drifts between disparate schema terms, topical synoptic knowledge and a standard vocabulary for term suggestions is supported by each topic. A topic materializes a class hierarchy depicting all terms within the topic sampled by the topic, e.g., Education. Each topic is characterized by its name and the context of its terms (term hierarchy and term descriptions) for each specific subject. Terms within a topic are shown to have a distinct meaning (sense) and context. The terminology context (description graph) of a topic is based on a standard domain ontology, e.g., Education. An ontology can be defined as a linguistic representation of a conceptualization of some domain of knowledge [24]. This ontology consists of abstract descriptions of classes of objects in a vertical domain, relationships between

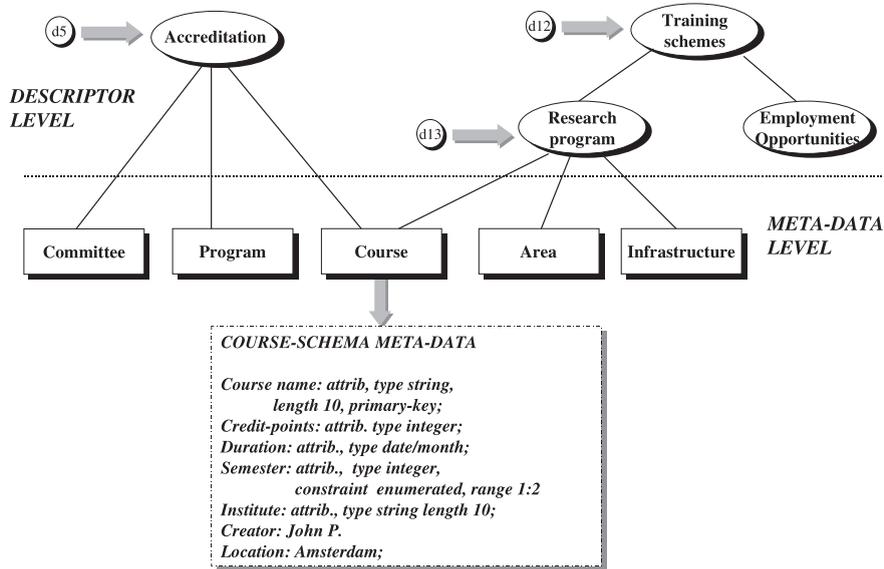


Fig. 2. Meta-data and descriptors for database schemas.

these classes, terminology descriptions and other domain-specific information and establishes a common vocabulary for interacting with between the various information sources underlying a topic. Currently, WordNet [37] is used to derive such a standard vocabulary for each of the topics. Hence, the topic structure is akin to an associative thesaurus and on-line lexicon (created automatically for each topic category). Thesaurus-assisted explanations created for each topic-based information subspace serve as a means of disambiguating term meanings, and addressing terminology and semantic problems. A topic is thus a form of a logical object (a kind of a *contextualized abstract view* over the content of large semantically related database collections) whose purpose is to cross-correlate, collate, and summarize the meta-data descriptions of semantically related network-accessible data,² and thus it is grounded on a common standard ontology.

Fig. 3 shows an excerpt from WordNet showing eight different senses of Course. By identifying the sense which a topic about Courses is about, terminology use in the context of that topic may be standardized³.

The topic-based multi-database configuration provides an appropriate frame of reference for both component database schema term indexing and user instigated searches. Fig. 1 in particular illustrates that the *Education & Training* multi-database network comprises a set of topics such as Education, Training, Literature & Publications, Employment, and so on. The topic-areas, described by each topic are interconnected by weighted links to make the searches more directed. For example, Fig. 1 illustrates that the topic Education is connected to other topics such as Literature & Publications, and Training by means of weighted links. Weighted links represent the degree of

² Topics were termed “Generic concepts” or “Global concepts” in previous work [33], [47].

³ The WordNet lexicographic tool is presently used only for experimental purposes and will be replaced by an appropriate subject gateway in the near future.

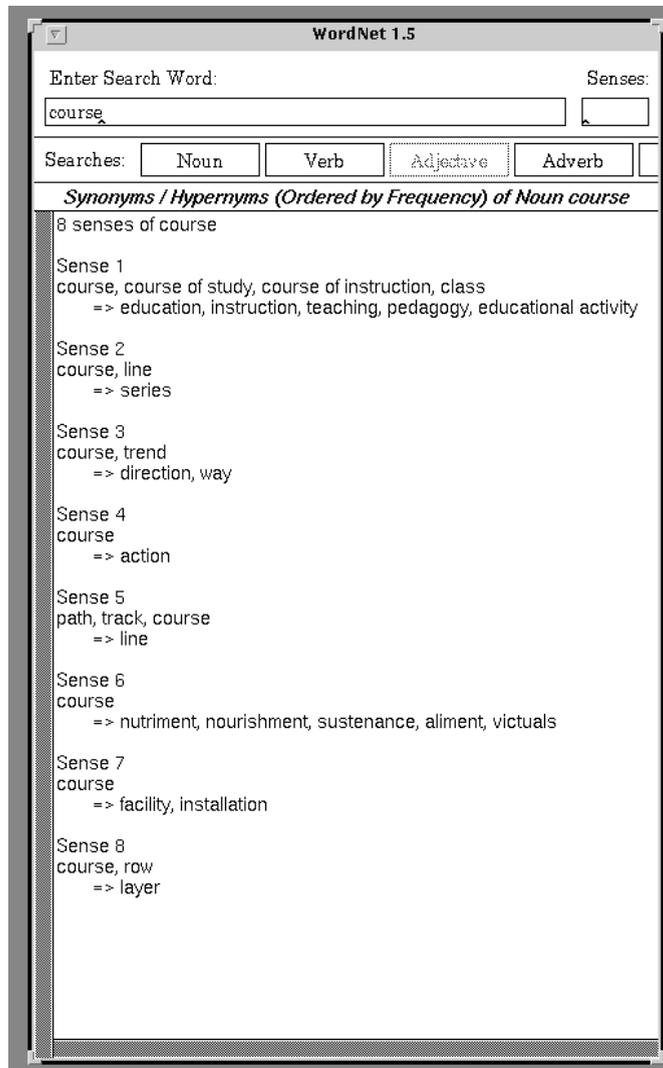


Fig. 3. Different senses of the term course.

semantic relatedness between resources and are created based on comparison of their descriptor terminology. Automatic comparison of descriptor terminology is conducted based on a set of similarity matrices. A detailed study of a similarity matrix-based algorithm for comparing descriptor terms and deriving weights can be found in [34].

A topically organized multi-database information space can be viewed as a Web-space, or a hyper-text, that encompasses collections of exported meta-data [3,4]. Such a multi-database information space partitions component databases into topically coherent groups, and presents descriptive term summaries and an extended vocabulary of terms for searching and querying the vastly distributed information space of the component databases that underly it. Thus, when dealing with a specific subtopic such as Accreditation we are not only able to source appropriate

information from remote assets based on the same topic but also to provide *matching* information about enrollment-programs, training schemes, research activities and publication data.

Assets in this network may become a member of more than one topic if they relate to their thematical foci, see Figs. 1 and 2. Individual topics are useful for browsing and searching large database collections because they organize information space in more comprehensible subspaces. For example, the Education topic provides a common terminology basis upon which database assets dealing with Enrollments, Courses, Accreditation, etc. (see Fig. 1) achieve knowledge of each others information content.

Although topics provide synoptic information about their underlying database clusters, *they do not require integration of the data sources*. This approach comes in strong contrast with approaches to semantic inter-operability based on explicit integration of conceptual schemas on the basis of semantic lexica [10,17]. The advantage of forming conceptual database clusters is that searches are goal-driven and the number of potential inter-database interactions is restricted substantially as it facilitates the distribution and balancing of resources via appropriate allocation to the various database partitions.

Formally, the basic building blocks of information space can be represented by three sets:

- let \mathcal{TO} be the set of topics,
- let \mathcal{DE} be a set of descriptors,
- let \mathcal{AS} be the set of assets (databases) in information space.

Collectively, we refer to these elements as information objects:

$$\mathcal{IO} \triangleq \mathcal{AS} \cup \mathcal{DE} \cup \mathcal{TO}.$$

As a first rule, the three base sets should not overlap, so the following axiom should apply:

[IS1] \mathcal{AS} , \mathcal{DE} , and \mathcal{TO} are mutually disjoint sets.

Overall a multi-database network can be viewed in terms of three layers, see Fig. 1. In this figure the \mathcal{AS} , \mathcal{DE} , and \mathcal{TO} information objects each define a distinct level in the multi-database information space:

- *Topics layer*. A layer consisting of the topics that have been used to classify the assets (the underlying databases).
- *Descriptors layer*. A layer consisting of all descriptors used to provide a thematical description for the topics and define what the underlying assets (the databases) are about. Each descriptor is represented in terms of a context graph which organizes the terms contained in an asset. A formal description of descriptors can be found in Section 3.3, while the process of creating and aggregating context graphs is described in Section 4.
- *Asset layer*. A layer consisting of the actual databases; being the assets of information space.

3.2. Structure of information space

As shown in Fig. 1, information space is characterized by a number of relationships between the constituting information objects in each layer. Let \mathcal{RL} define the set of relationships in information space. For relationships, the following functions are presumed to be defined:

- $\text{Src, Dst} : \mathcal{RL} \rightarrow \mathcal{IO}$, defining the source and destination of a relationship.

- Weight : $\mathcal{RL} \rightarrow \mathcal{WG}$, defining the certainty of a relationship. For weights, the existence of a combination operation \otimes and a total order $<$ is presumed. For example, when using probabilities as weights, the combination operation may be computed as:

$$w \otimes v \triangleq 1 - (1 - w) \times (1 - v).$$

To combine the weights associated to links in information space, i.e., when more precisely defining the \otimes operator, we can actually make use of the same probabilistic reasoning as has been studied in the context of probabilistic information retrieval. See, for instance, [23,49,57].

The relationships in information space should be irreflexive:

$$[\text{IS2}] \quad \forall_{r \in \mathcal{RL}} [[\text{Src}(r) \neq \text{Dst}(r)]].$$

As an abbreviation we shall use

$$x \xrightarrow[r]{w} y \triangleq \text{Src}(r) = x \wedge \text{Dst}(r) = y \wedge \text{Weight}(r) = w.$$

When we are not interested in the specific weights associated to links, we will usually omit the weight w , and simply write $x \xrightarrow{r} y$.

Two key classes of relationships between information objects can be distinguished:

- *intra-layer relationships* ($\mathcal{IA} \subseteq \mathcal{RL}$), linking information objects of the same layer, viz., horizontal dimension, and
- *inter-layer relationships* ($\mathcal{IR} \subseteq \mathcal{RL}$), linking information objects of different layers, viz., vertical dimension.

These classes of relationships are presumed to form a partition of \mathcal{RL} :

$$[\text{IS3}] \quad \mathcal{IA} \text{ and } \mathcal{IR} \text{ are a partition of } \mathcal{RL}.$$

The first class of relationships between information objects we focus on are the inter-layer relationships, i.e., those relationships that bridge information objects in the three layers. The following three classes of relationships exist between the information objects of the different layers:

- $\mathcal{AR} \subseteq \mathcal{IR}$ a set of aboutness relationships defining what the underlying assets are about,
- $\mathcal{DR} \subseteq \mathcal{IR}$ a set of defining relationships expressing the thematical focus of a topic in terms of descriptors, and
- $\mathcal{MR} \subseteq \mathcal{IR}$ a set of membership relationships identifying to which topics a given asset belongs.

The aboutness relationships in \mathcal{AR} define what the underlying databases are really about in terms of the descriptors. In other words, the aboutness relationship defines the thematical scope of the databases.

These classes of relationships are presumed to form a partition of \mathcal{IR} :

$$[\text{IS4}] \quad \mathcal{MR}, \mathcal{DR} \text{ and } \mathcal{AR} \text{ are a partition of } \mathcal{IR}.$$

For the three classes of inter-layer relationships, the following predicates may be defined:

$$x \text{ IsMemberOf}_w y \triangleq \exists_{r \in \mathcal{MR}} [x \xrightarrow[r]{w} y],$$

$$x \text{ IsAbout}_w y \triangleq \exists_{r \in \mathcal{AR}} [x \xrightarrow[r]{w} y],$$

$$x \text{ Defines}_w y \triangleq \exists_{r \in \mathcal{DR}} [x \xrightarrow[r]{w} y].$$

These relationships should indeed bridge the appropriate layers:

$$[\text{IS5}] \quad x \text{ IsMemberOf } y \Rightarrow x \in \mathcal{AS} \wedge y \in \mathcal{TO}.$$

$$[\text{IS6}] \quad x \text{ IsAbout } y \Rightarrow x \in \mathcal{AS} \wedge y \in \mathcal{DE}.$$

$$[\text{IS7}] \quad x \text{ Defines } y \Rightarrow x \in \mathcal{DE} \wedge y \in \mathcal{TO}.$$

The second class of relationships between information objects are concerned with intra-layer relationships. Many different types of relationships between information objects within a single layer may exist. For example:

- different types of associations,
- different types of part–whole relationships.

Instead of introducing a whole plethora of possible relationship types, we only focus on the following two general classes of relationships which are fairly representative:

- $\mathcal{AR} \subseteq \mathcal{IA}$ a set of associations,
- $\mathcal{PR} \subseteq \mathcal{IA}$ a set of part–whole relationships.

The membership relationship between assets and topics is fully derivable from the IsAbout and Defines relationships:

$$[\text{IS8}] \quad a \text{ IsAbout}_{w_1} d \wedge d \text{ Defines}_{w_2} t \iff a \text{ IsMemberOf}_{w_1 \otimes w_2} d.$$

Each asset must be about some descriptor with the maximum weight, e.g., 10/10, associated:

$$[\text{IS9}] \quad \forall a \in \mathcal{AS} \exists d [a \text{ IsAbout}_{1_{\mathcal{WG}}} d],$$

where $1_{\mathcal{WG}}$ denotes the maximum weight from \mathcal{WG} based on the total order. Each topic must be defined using a descriptor with the maximum weight associated:

$$[\text{IS10}] \quad \forall t \in \mathcal{TO} \exists d [d \text{ Defines}_{1_{\mathcal{WG}}} t].$$

These classes of relationships are presumed to form a partition of \mathcal{IA} :

$$[\text{IS11}] \quad \mathcal{AR} \text{ and } \mathcal{PR} \text{ are a partition of } \mathcal{IA}.$$

For the two general classes of intra-layer relationships, the following predicates may be defined:

$$x \text{ IsAssocTo}_w y \stackrel{\Delta}{=} \exists r \in \mathcal{AR} [x \xrightarrow{r}_w y],$$

$$x \text{ IsPartOf}_w y \stackrel{\Delta}{=} \exists r \in \mathcal{PR} [x \xrightarrow{r}_w y].$$

Intra-layer relationships should indeed be *intra* layer:

$$[\text{IS12}] \quad \text{For each } S \in \{\mathcal{AS}, \mathcal{DE}, \mathcal{TO}\}$$

$$x \in S \wedge x \text{ IsAssocTo } y \Rightarrow y \in S.$$

$$[\text{IS13}] \quad \text{For each } S \in \{\mathcal{AS}, \mathcal{DE}, \mathcal{TO}\}$$

$$x \in S \wedge x \text{ IsPartOf } y \Rightarrow y \in S.$$

Part-of relationships cannot be weighted. In other words, they should always be of the maximum weight:

$$[\text{IS14}] \quad x \text{ IsPartOf}_w y \Rightarrow w = 1_{\mathcal{W} \setminus \mathcal{G}}.$$

In the context of multi-database systems it does not make sense to allow for subset relationships between topics as topics are autonomous and disjoint from each other. In other words, we have:

$$[\text{IS15}] \quad x \text{ IsPartOf}_w y \Rightarrow x, y \notin \mathcal{TC}.$$

The part-of relationship is transitive and irreflexive:

$$[\text{IS16}] \quad x \text{ IsPartOf } y \text{ IsPartOf } z \Rightarrow x \text{ IsPartOf } z.$$

$$[\text{IS17}] \quad \neg(x \text{ IsPartOf } x).$$

In the context of multi-database systems, it furthermore does not make sense to introduce intra-layer relationships between the underlying databases due to the unnecessary complexity and also because all databases are strongly related to their encompassing topic.

In other words, in this context we have:

$$[\text{IS18}] \quad x \text{ IsAssocTo } y \Rightarrow x, y \notin \mathcal{AS}.$$

$$[\text{IS19}] \quad x \text{ IsPartOf } y \Rightarrow x, y \notin \mathcal{AS}.$$

In future research we will apply the results as presented in this paper to more general forms of assets, such as Web-pages, documents, etc. In these latter cases it is indeed sensible to cater for associative and part-of relationships between the assets. For example, documents referring to each other imply an associative relationship while a chapter is-part-of a book.

The set of descriptors that is associated to a topic by the Defines relationship is essentially the thematic scope of a topic:

$$\text{Thema}(t) \triangleq \{d \mid d \text{ Defines } t\}.$$

The theme of a topic should have a unique top element based on the IsPartOf relationship on descriptors:

$$[\text{IS20}] \quad \forall t \in \mathcal{TC} \exists! x \in \text{Thema}(t) [\neg \exists y \in \text{Thema}(t) [x \text{ IsPartOf } y]].$$

If two topics are associated to each other with some weight, then this association should somehow be reflected by the associations between descriptors associated to the respective topics:

$$[\text{IS21}] \quad t_1 \text{ IsAssocTo}_w t_2 \Rightarrow w = \bigotimes_{r \in B(t_1, t_2)} \text{Weight}(r), \text{ where}$$

$$B(t_1, t_2) \triangleq \{r \in \mathcal{AR} \mid \exists d_1 \in \text{Thema}(t_1), d_2 \in \text{Thema}(t_2) [d_1 \xrightarrow{r} d_2]\}.$$

The set of descriptors used to express what an asset is about, is referred to as the *aboutness* of that asset:

$$\text{Aboutness}(a) \triangleq \{d \mid a \text{ IsAbout } d\}.$$

The aboutness of an asset should have a unique top element as well:

$$[\text{IS22}] \quad \forall a \in \mathcal{AS} \exists! x \in \text{Aboutness}(a) [\neg \exists y \in \text{Aboutness}(a) [x \text{ IsPartOf } y]].$$

The aboutness of assets is a rather complex relationship. In [11,43] detailed studies have been made of what aboutness really is. In these publications several rules may be found governing the aboutness relationship. An example of such a rule would be:

$$[IS23] \quad x \text{ IsPartOf}_{w_1} y \wedge x \text{ IsAbout}_{w_2} c \Rightarrow y \text{ IsAbout}_{w_1 \otimes w_2} c.$$

For more detailed rules on aboutness, refer to, e.g. [11,43].

3.3. Descriptors

The descriptors defining the theme of a topic play a very important role in information space. It provides users, as well as the system, with a description of a context in terms of some suitable language. Topics in information space may be described using such mechanisms as: keywords, index expressions [18], noun phrases, or different forms of conceptual graphs [54]. Most readers will be familiar with keywords to express what a text database or document is about. The language of index expressions was introduced as a description language that would allow keywords to be put in relation to each other. For example sentences of the form:

Training scheme contains research programs

rather than simple keywords:

{training, scheme, research, program}.

The effectiveness of index expressions in an information retrieval context has been studied and reported in [14]. The language of index expressions is actually a subset of the language of noun phrases. For example, the language used in this paper employs terms such as part-of, pertains-to, contains, is-a, is-related-to and conceptual graph descriptions. The use of conceptual graph like languages for retrieval purposes has been studied in, e.g. 2,39,13]. These publications point out that it is useful to separate concepts and their description as a specific concept may quite well be represented in different languages using a specific description for each of the languages.

In this subsection the descriptors are defined in more detail. As keywords and index expressions can both be represented as conceptual graphs, we base the formal definition of descriptors on conceptual graphs.

To define descriptors formally, we start out from two base sets:

- a set \mathcal{CO} of concepts, and
- a set \mathcal{ED} of edges.

In addition, the following functions are needed:

- $\text{Src, Dst} : \mathcal{ED} \rightarrow \mathcal{CO}$, defining the source and destination of an edge.
- $\text{Name} : (\mathcal{ED} \cup \mathcal{CO}) \rightarrow \text{String}$, the name of a relationship or a concept. The names of the relationships could, for example, be based on the ones as pre-defined by WordNet, such as: hypernyms, hyponyms, part-of, and pertains-to.
- $\text{Explain} : \mathcal{CO} \rightarrow \text{String}$, some additional explanation of the concept.

Together \mathcal{CO} , \mathcal{ED} , Src , Dst , Name , and Explain define a *concept space*.

Using this concept space, the set of (possible) descriptors can be defined as follows:

$$\mathcal{DE} \triangleq \left\{ E \subseteq \mathcal{ED} \left| \begin{array}{l} \text{The set of edges } E \text{ spans a connected subgraph and} \\ \forall c, d \in \text{Concepts}(E) [\text{Name}(c) = \text{Name}(d) \Rightarrow c = d] \end{array} \right. \right\},$$

where $\text{Concepts}(E) \triangleq \bigcup_{e \in E} \{\text{Src}(e), \text{Dst}(e)\}$. Note that within one descriptor, the names of concepts must be unique. For descriptors $d, e \in \mathcal{DE}$, the *IsPartOf* relationship can be defined formally as:

$$d \text{ IsPartOf}_{1_{\mathcal{W}\mathcal{G}}} e \triangleq d \subseteq e.$$

The *IsPartOf* relationship for the descriptors d and e can simply be defined formally as:

$$d \text{ IsPartOf } e \triangleq d \subseteq e.$$

By using ontologies more contextual information may be added to the descriptions of concepts. For instance, in Fig. 3, a WordNet [37] style context for the Course concept is shown. There it is shown what particular meaning (sense) and context of the word Course is the focus of the concept. As mentioned before, two different concepts (in different topics) may indeed have the same name, while they are concerned with different senses of the same name/term.

3.4. Additional meta-data

To each of the information object, additional meta-data attributes may be associated [55] to provide semantic richness. Such meta-data attributes may for example be concerned with:

authorship, dateofcreation, medium, fileformat, pricing, quality, and location.

Meanwhile different emerging meta-data standards have come into being. For example, Dublin Core [19] or PICS [44]. Meta-data attributes may be formalized as follows:

- Name : $\mathcal{IO} \rightarrow \text{Names}$ providing the name of the information object.
- Location : $\mathcal{AS} \rightarrow \text{URI}$ yielding the physical location/address of an asset (database). For example, in terms of a Universal Resource Indicator (URI, such as the well-known URL from the World Wide Web).
- Layer : $\mathcal{IO} \rightarrow \{\text{Topics, Descriptors, Assets}\}$ returning the name of the layer an information object belongs to.

For Layer we obviously have:

$$[\text{IS24}] \quad \forall t \in \mathcal{TO} [\text{Layer}(t) = \text{Topics}] \wedge \forall c \in \mathcal{DE} [\text{Layer}(c) = \text{Concepts}] \wedge \forall a \in \mathcal{AS} [\text{Layer}(a) = \text{Assets}].$$

4. Clustering of assets

In the following we describe a general methodology that aids in clustering databases and creating their corresponding topic nodes in information space. Key criteria that have guided this methodology are: scalability, design simplicity and easy to use structuring mechanisms.

4.1. Describing databases and subtopics

In order to initially cluster component databases, a high level description of the database contents in the form of a descriptor must first be derived.

To demonstrate this consider a set of assets (databases) that contribute towards the Accreditation subtopic and deal with academic institutions and accreditation processes. These assets contain entities such as Courses, Committees, Accreditation, Processes, etc. In order to become

part of a topic-based multi-database network the schema terms of these databases are represented in the form of a *context graph*. This graph is essentially a form of a *conceptual graph* and inter-connects terms (concepts) on basis of their semantic relatedness. To achieve this we use a variant of an information retrieval technique called, the *star technique*. With this technique, a concept is selected and then all concepts related to it are placed in a class [30]. Concepts not yet in a class are selected as new seeds until all concepts are assigned to a class. Classes in the star technique are not related to each other. The variant of the star technique that we are using starts with a concept, then an additional concept that is related to an already selected concept is represented as a another class and is explicitly connected to the selected concept. The new concept is then selected as a pivot and the process is repeated until no new concepts can be added. For example, the context graph for the Accreditation descriptor, is depicted in Fig. 4 and is shown to contains nodes which contain the terms Committee, Institutions, Courses, etc. and their related terminology.

The context graph edges depict inter-connections (association, generalization, specialization or containment) between the concepts within a particular asset. Concept inter-relations are determined on the basis of a reference lexicographic substrate that underlies the concepts in information space. For this purpose, as already explained, we use the lexicographic system WordNet [37] that supports semantic concept matching through the use of an extensive network of word meanings connected by a variety of textual and semantic relations.

If d is the context graph (descriptor) that is derived for a database (asset) a using the algorithm as sketched above, then we know: $a \text{ lsAbout } d$. Furthermore, using the lsPartOf relationship for descriptors, we can also infer that for any context graph e , if $e \text{ lsPartOf } d$ we have $a \text{ lsAbout } e$.

To facilitate clustering and discovery of information, we require that a subtopic (descriptor) can be totally described in terms of three sections which contain a synoptic description of the meta-data content of the asset; associations between meta-data terms in the form of a semantic-net; and finally, links from these descriptions to other related assets in the network. This information can be viewed by users of the system once they have chosen a component database that potentially matches their interests (see Section 5).

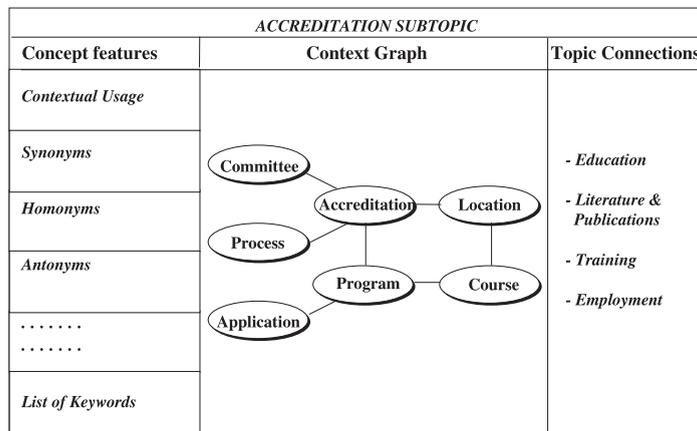


Fig. 4. Describing a subtopic.

Fig. 4 illustrates that each database node contains the following sections:

- *concept features*,
- *context graph*, and
- *topic connections* section.

The *concept features* section contains additional information concerning the concepts used in the context graph, essentially providing the Explain function for concepts in the context graph. These descriptions include abstract descriptions of terms in the domain such as their sense (unique), relationships between these terms, composition of terms, terminology descriptions, hypernym, hyponym, antonyms-of, part-of, member-of, pertains-to relations, contextual usage (narrative descriptions), a list of keywords, and other domain-specific information, that apply to the entire collection of members of a topic. Moreover, it may include other useful details such as: geographical location of databases, access authorization and usage roles, explanations regarding corporate term usage and definitions, domains of applicability, charge costs, and so on. The feature descriptions entries are partially generated on the basis of WordNet and contain information in the form represented in Figs. 3 and 5.

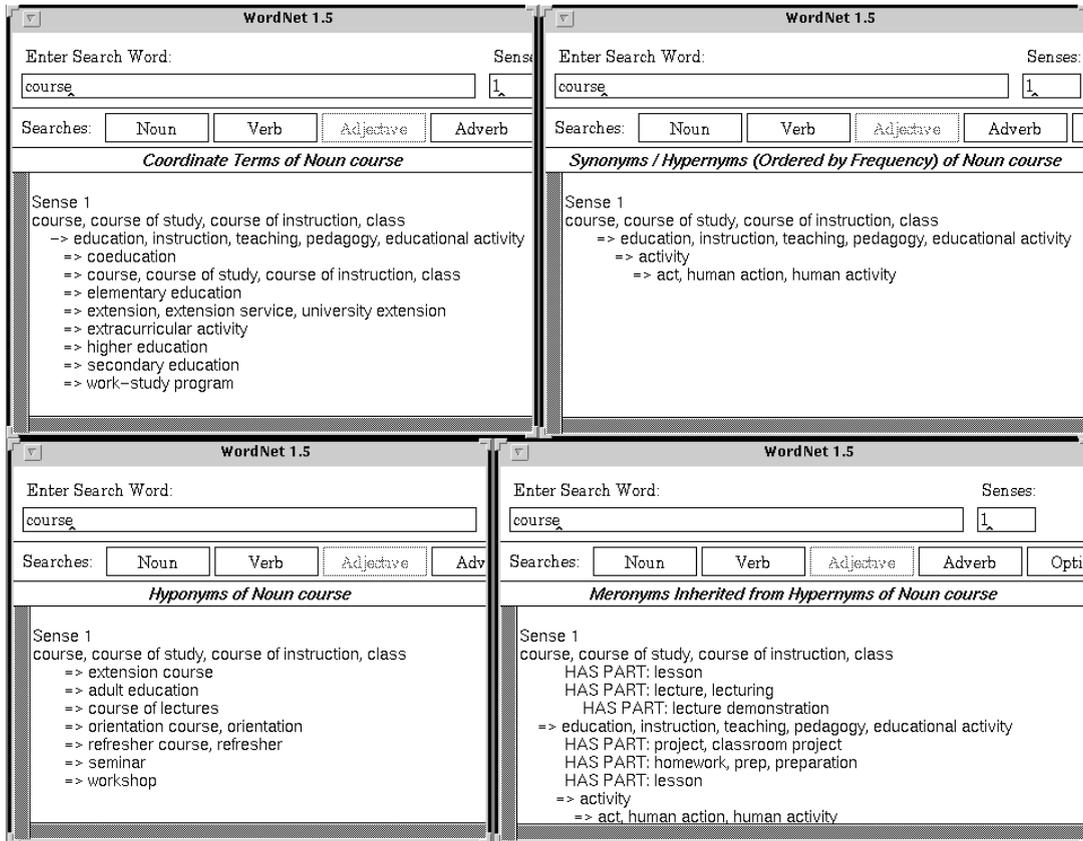


Fig. 5. Further WordNet details for Course.

The *context graph* section contains a non-directed graph which connects the concepts as described by the *concept features* section. Except for viewing purposes when navigating information space, the context graph used in the clustering of databases to form topics. Each of the concept nodes defines (in conjunction with its respective entry in the feature descriptions window) a common structured vocabulary of terms and term relationships relevant to that concept. Finally, the topics connection section shows to which topics the Accreditation descriptor is related to in the network.

4.2. Similarity-based clustering of databases

Similarity-based clustering of database schemas organizes databases into related groups based on the concepts that are referred to by the *concept features* section of their database description (see Fig. 4) they contain and the link structure of their context graphs.

Our clustering algorithm determines the similarity between two graphs (representing two different database schema meta-data) based on both concept similarity and link similarity factors. This is accomplished in two steps. Firstly, a pairwise-similarity of nodes in two context graphs is computed. From this an initial “pairing” of the nodes is determined. In the second step a comparison of the link structure of two context graphs is made based on the inter-node pairings and a semantic distance value is calculated. We chose this concept/link similarity-based algorithm because it is relatively easy to implement and avoids generating very large clusters.

Concept-based similarity. This is calculated using cluster analysis techniques [21] to identify co-occurrence probabilities – representing the degree of similarity – between two discrete concepts.

Our similarity metric is based on the meaning of the collection of terms representing the terminological context (viz., semantic-levels) of a particular concept, e.g., Course, see Fig. 3. The comparison is based on: a conversion of each context graph node (Committee, Process, Subject, Course, etc. (see Fig. 4) to a corresponding matrix of noun terms (containing the entire terminological context of a concept); and a subsequent comparison of terms within these matrices.

A matrix a for concept c with m columns, and n rows, should be organized such that:

1. Each column contains synonyms describing the same concept

$$\forall_{1 \leq i < j \leq m; q \leq k \leq n} [a_{i,q} \text{ described the same concept as } a_{j,k}].$$

For example, Course, Course-of-study, Course-of-lectures, etc.

2. There must be a unique column in which the first element corresponds to the name of the concept

$$\exists!_i [\text{Name}(c) = a_{i,1}].$$

3. The columns in the matrix should be ordered from general to more specific. In other words, terms to the left should (pairwise) not be more specific than terms to the right

$$\forall_{1 \leq i < j \leq m; 1 \leq k \leq n} [a_{i,k} \text{ is not a more specific term than } a_{j,k}].$$

For example, Education, Educational-activity, are more general terms than Course, while Computer-science-course is a more specific term than Course.

Similarity analysis is mainly based on statistical co-occurrences of terms based on techniques which have been successfully used for automatic thesaurus generation of textual databases [21,

50]. In fact, we base our term-based similarity on the improved *cosine* formula [50] which is used to calculate the semantic distance between the vector for an item in a hierarchical thesaurus and the vector for a query item. To provide the right ontological context for semantic term matching, we use again the massive semantic net WordNet [37].

Comparison of the conceptual structure of two context graphs. To determine the structural and semantic similarity between two graphs, we based our algorithms regarding conceptual similarity between terms on heuristics-guided spreading activation algorithms, and on work in the information retrieval area presented in [48]. These approaches take advantage of the semantics in a hierarchical thesaurus representing relationships between index terms. The algorithms calculate the conceptual closeness between two index terms, interpreting the conceptual distance between two terms as the topological distance of the two terms in the hierarchical thesaurus. During this process similarity between nodes is established by considering the edges separating the nodes in the context graph as well as the actual graph structure. Some initial results regarding the comparison and clustering process can be found in [36].

Once similarity between the context graphs describing what the underlying databases are about has been established, the context graphs can be aggregated to create topics. The aggregation of the context graphs from various component databases, results in the clustering of inter-related database schema terms, and forms a descriptor graph at the top level (see Fig. 6). This aggregation is performed on the basis of the reference lexicographic substrate (WordNet). The aggregation algorithm employed does not integrate the aggregated databases, as is the usual case with other approaches [31], but rather links concepts at the topics level with semantically equivalent concepts in its underlying cluster of database context graphs.

For each database cluster, a topic is created to represent the area of interest that the group embodies, e.g., a Education topic for the Accreditation, Tertiary-Education and Enrollment-Program databases, see Fig. 1. The aggregated context graph is used as a base to define the thematic scope of the particular topic.

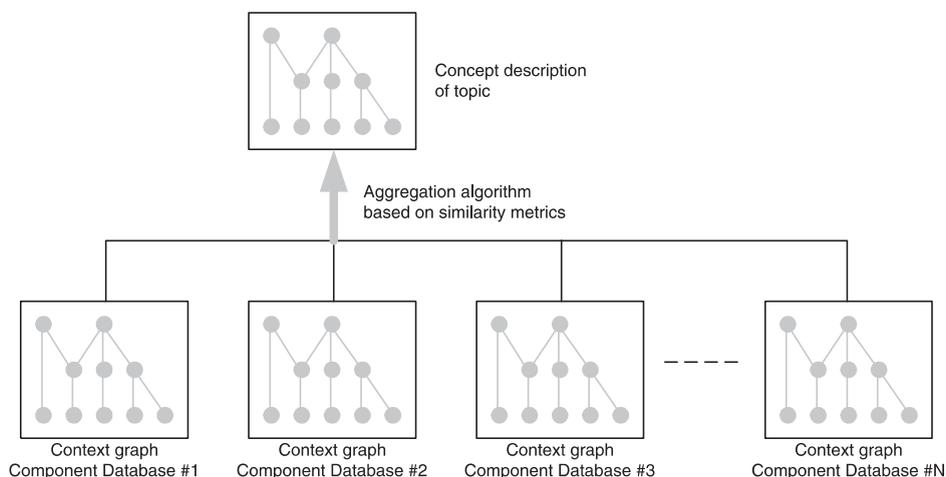


Fig. 6. Clustering interrelated component schema terms.

When clustering the databases, some cut-off value needs to be employed to express how similar two databases should be for them to be placed in the same cluster. This cut-off value then essentially expresses the internal *cohesion* of the resulting cluster (and topic) [34].

4.3. Deriving subtopics

When some topic t contains a large number of databases, it may be useful to split t into several smaller topics. This can be done by re-clustering the underlying set of databases using a higher cut-off value. Each of the resulting new topics has an internal cohesion that is higher than the original topic. If T is the set of resulting subtopics, then these may be linker to the original (large) topic using the *IsPartOf* relationship:

$$s \in T \Rightarrow s \text{ IsPartOf}_{1_{w^*}} t,$$

where 1_{w^*} represents the maximal value from *Weight*.

4.4. Inter-relating topics

Once the topics have been created, it is also possible to use the similarity algorithm to compute the similarity between the (aggregated) context graphs associated to topics. This similarity can then be interpreted as the similarity between topics. Using some cut-off value, the similarity relationship between topics can be used to fill the *IsAssocTo* relationship between topics. The similarity weight can be used as a weight on these relationships. In other words, if c is the cut-off value, and $\text{Sim}(t_1, t_2)$ expresses the similarity between two topics, we have

$$\text{Sim}(t_1, t_2) = w \wedge w \geq c \Rightarrow t_1 \text{ IsAssocTo}_w t_2.$$

5. Navigation and querying

Schema information discovery spans a spectrum of activities ranging from a search for a specific data-item(s) (contained in possibly several component databases) to a non-specific desire to understand what information is available in these databases and the nature of this information. This section is concerned with navigating and querying information space as techniques for schema information discovery.

The architecture discussed so far in this paper has been implemented in [34]. This architecture needs appropriate extensions for querying and navigation. This is on-going research based on several ideas that have been tested and implemented elsewhere in the literature. In particular, [1,15,3] presents techniques that have been applied in the context of documents, while in the work reported in [26] these techniques are translated to query formulation on databases. The latter publication is in line with the application of stratified hyper-media as a way of navigating a multi-database information space.

5.1. Stratified architecture

As mentioned before, information space can be grouped into an assets, descriptors, and a topics layer. These three layers allow us to organize information space as a three-level stratified hyper-media as reported in, e.g. [1,15,3]. Organizing a multi-database information space as a stratified hyper-media enables users to navigate information space in a natural way.

In a stratified hyper-media, each layer consists of a set of nodes and links. For a set X of information objects in a particular layer, \mathcal{AS} , \mathcal{DE} , and \mathcal{TO} , the corresponding layer is defined by: $\langle X, \text{IsAssocTo}^X, \text{IsPartOf}^X \rangle$, where:

$$x \text{ IsAssocTo}^X y \stackrel{\Delta}{=} x, y \in X \wedge x \text{ IsAssocTo } y,$$

$$x \text{ IsPartOf}^X y \stackrel{\Delta}{=} x, y \in X \wedge x \text{ IsPartOf } y.$$

The three layers of information space can then be defined as follows:

1. Assets layer: $\langle \mathcal{AS}, \text{IsAssocTo}^{\mathcal{AS}}, \text{IsPartOf}^{\mathcal{AS}} \rangle$.
2. Descriptors layer: $\langle \mathcal{DE}, \text{IsAssocTo}^{\mathcal{DE}}, \text{IsPartOf}^{\mathcal{DE}} \rangle$.
3. Topics layer: $\langle \mathcal{TO}, \text{IsAssocTo}^{\mathcal{TO}}, \text{IsPartOf}^{\mathcal{TO}} \rangle$.

The above definitions are general and can be used for all sorts of information assets, e.g., documents. However, in the case of multi-database systems for \mathcal{AS} both relationships $\text{IsAssocTo}^{\mathcal{AS}}$, $\text{IsPartOf}^{\mathcal{AS}}$ are empty whereas for the \mathcal{TO} layer the $\text{IsPartOf}^{\mathcal{TO}}$ is empty, see the corresponding axioms in Section 3.

As \mathcal{AS} , \mathcal{DE} , and \mathcal{TO} are disjoint sets, these layers are disjoint as well. When navigating a stratified hyper-media architecture, users may not only want to navigate within a single layer, they will also want to navigate between layers.

The *JumpsTo* relationship is a generalization of *IsMemberOf*, *Defines* and *IsAbout* and is used for inter-layer navigation. This generalization is defined as follows:

$$x \text{ JumpsTo } y \stackrel{\Delta}{=} \exists_{r \in \mathcal{R}} [\{x, y\} = \{\text{Src}(r), \text{Dst}(r)\}].$$

This three-tier architecture is the key ingredient to finding information in distributed, scalable systems. It generates a semantic hierarchy for database schema terms in layers of increasing semantic detail. Most searches will initially target the richest semantic level, viz., the topics layer, and percolate via the descriptors layer to the assets layer in order to provide access to the contents of an asset. This type of content-based clustering of the searchable information space provides convenient abstraction demarcators for both the users and the system to make their searches more targeted, scalable and effective. This methodology results in a simplification of the way that information pertaining to a large number of inter-related database schemas can be viewed and more importantly it achieves a form of global visibility [42].

5.2. Node presentation

When navigating the stratified hyper-media as derived from an information space, users travel from information object to information object via intra-layer or inter-layer links. Each information object is presented to a user by providing three types of information to the user:

- the current location in information space,
- the possible steps to continue the journey in information space,
- more information about the current location.

Let f be the information object in information space where the user is currently at, then the above three types of information can be further specialized as:

- The description of current location consists of $\text{Name}(f)$, $\text{Layer}(f)$, and $\text{Size}(f)$.
- The possible continuations consists of six sets:

1. Refinements: $\{x|x \text{ IsPartOf } f\}$.
2. Enlargements: $\{x|f \text{ IsPartOf } x\}$.
3. Associations: $\{x|f \text{ IsAssocTo } x\}$.
4. Jumps to the topics layer: $\{x \in \mathcal{TO}|f \text{ JumpsTo } x\}$.
5. Jumps to the descriptors layer: $\{x \in \mathcal{DE}|f \text{ JumpsTo } x\}$.
6. Jumps to the assets layer: $\{x \in \mathcal{AS}|f \text{ JumpsTo } x\}$.

In describing the possible continuations to users, the system may choose to use $\text{Name}(x)$. In the case of descriptors, it may be useful to also include some synoptic version of the context graph.

- Other information about the current focus f may consist of any other meta-data available on f .

In presenting this information to a user, the system may opt to use multiple screens. For example, Fig. 7 illustrates how a user moves from the topic to the descriptor and the asset layer when looking for information relating to Education. Here we assume that the user started from the Education topic layer and that the concepts within the descriptor layer are organized in ascending order of specificity, e.g., from Tertiary-Education to Enrollment-Program, Accreditation and so on. Once the user arrives at a desired descriptor term, e.g., Tertiary-Education, then s/he can move to its asset level, e.g., Course descriptions. For instance, the user can move from the term activity, which is a highly abstract term, to more concrete terms such as Course-of-study, Course-of-lectures, and so on. We can move between these concepts by using the refinement field (operation). In Fig. 7 we assume that the user was interested in finding assets dealing with Accreditation.

The order in which the continuation options are listed in the nodes should ideally be determined by the relevance of these options to the user. For example, we consider using the mechanisms presented in [16] where the relevance of these options can be derived based on a user's search behavior up to a certain point in time. This may be combined with the relative weights that can be associated to the links in the information space.

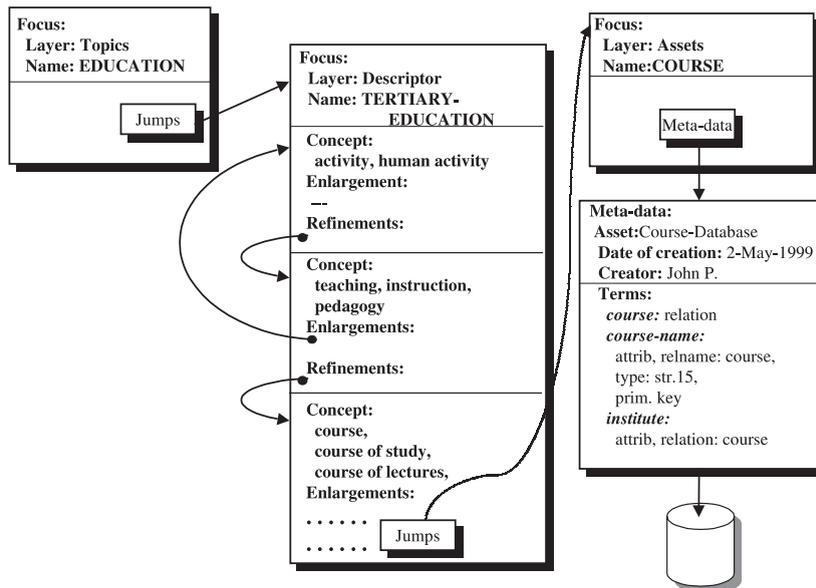


Fig. 7. Screens representing information objects.

5.3. Navigation techniques

There are two basic modes in which searching of the system may be organized. These search modes depend upon the nature of the information a user is attempting to access, and how this information relates to the database that the user is operating from. Serendipity, exploration and contextualization are supported by navigating the stratified hyper-media spanned by information space. In such cases the user is interested in finding out about a particular topic rather than a specific information (schema) item. We call this former form of exploration *index-driven*.

Alternatively, if a user is seeking data which is closely related or allied to her/his local database, then searching may be organized around the links of this database to other topics in information space. We refer to this form of exploration as *concept-driven*. Concept-driven querying is the subject of a previous publication [46]. In this paper, we are mainly concerned with index-driven searches.

Index-driven navigation allows the users to deal with a controlled amount of material at a time, while providing more detail as the user looks more closely. This form of searching is related to the dynamic indexing schemes and incremental discovery of information requirements for information discovery. An index-driven search will usually start out from some name $n \in \text{Names}$ of a concept. From this point, the search will gradually percolate down to the required level of specificity. This process starts by treating n as a descriptor and matching this to the set of available concepts. This results in a set ($D(n)$ say) of potentially interesting descriptors. This set may become rather large. It may therefore be sensible to limit this set to the elements that are lowest (i.e., least specific) in the *IsPartOf* hierarchy:

$$\{d \in D(n) \mid \neg \exists e \in D(n) [e \text{ IsPartOf } d]\}.$$

By presenting the resulting set of concepts, users can consequently home in on the intended interpretation of n , and consequently refine their focus of interest by navigating through the descriptors layer.

If a user is already focussed on a specific topic (t say), the set $D(n)$ can be limited in advance to the set $\text{Thema}(t) \cap D(n)$.

5.4. Querying of domain meta-data

When the user needs to further explore the search target, *intensional*, or schema queries [42] – which return meta-data terms from selected schema terms – can be posed to further restrict the information space and clarify the meaning of the information items under exploration. Such domain-specific queries should not be confused with queries which target the data content of the assets (to which we refer to as *distributed* queries/transactions). Intensional queries are particularly useful for assisting users who are unfamiliar with the vocabulary of terms that can be used in connection with distributed queries/transactions or with the range of information that is available for responding to distributed queries. Sample intensional queries related to the descriptions as depicted in Fig. 3 may include the following:

Query-1. Find the set of common super-terms of course.

Query-2. Find all terms more specific than course and all their parts under sense education.

Query-3. Find the smallest common super-term of course of lectures and workshop.

Query-4. Find all parts of the term course.

Query-5. Which are the common properties of refresher course and seminar?

Query-6. Find all terms which contain the properties lesson and classroom project.

Query-7. What is the definition of the term refresher course?

All of the above queries – except for the last one – are rather intuitive. The last query returns a narrative description of the requested term in English (if available).

Finally, when users feel sufficiently informed about the contents and structure of asset schema terms they have explored, they can pose meaningful distributed database requests which target the data content of the relevant component databases.

6. Experimentation and performance evaluation

The framework that we described in this paper was implemented on Sun Sparc stations under Solaris 2 using GNU C++ and CGI scripts. In order to evaluate automated clustering a test platform based on the clustering of about hundred networked relational databases has been created. There are two basic areas of experimentation being pursued. Firstly, there is the question of how well the initial automatic clustering of databases based on each asset's meta-data description can be performed. That is, the question of finding appropriate initial information and relationships in the presence of large numbers of database sources. The types of experiments performed here are somewhat allied with the field of information retrieval and clustering. The second set of experiments deals with the processing and communications necessary to support the underlying distributed structure by which the generic concepts and their inter-relationships are implemented, queried and updated. This second group of experiments thus has its roots in the fields of distributed/parallel processing and communications performance. Details of both sets of experiments can be found in [34]. In the following we summarize the first set of experiments.

The first set of experiments are based on the notion of *precision* and *recall* (as defined within IR). The precision of the retrieval is the ratio of relevant documents to retrieved documents. A precision value of 1.0 (or 100%) indicates that all the documents retrieved were relevant. Recall gives an indication of how many relevant documents were actually retrieved – and is the ratio of relevant retrieved documents to relevant documents. To achieve this, a collection of a hundred relational databases were procured from a large organization's collection of database systems. A similarity matrix was created based on the comparison of these asset descriptions with each other. Each row of the matrix represents a ranking of descriptions – assets – with respect to queries. The ranking within each row of the similarity matrix is a consequence of our asset description comparison algorithm, and this is compared against a “perfect” series of queries and their respective asset rankings. We generated average recall and precision across many queries to gain an impression of the retrieval effectiveness of the system, i.e., how close the automatically generated ranking of assets are to the perfect manual rankings. We made use of the simple Dice similarity coefficient [21] for all experiments, and a baseline was generated by using “raw” (unprocessed) database attributes as the description of an asset. During the experiments two basic factors were varied: the asset description and the comparison algorithm. The asset description was varied by adding to the raw terms linguistic aspects of the terms (the effect of sense, synonym, homonym, etc. inclusion), level of abstraction hierarchy (the effect of ontological descriptive terms), and structural information (the ordering of terms) based on WordNet. Three basic type of comparison algorithms were used during experimentation: one based on simplistic term matching, a second

identifying structural similarity based on the notion of spreading activation energy [36], and a third hierarchical approach to address the perceived difficulties with the previous two methods.

A perfect description and comparison process produces a similarity matrix which ranks databases such that 100% precision is achieved and all relevant databases are retrieved (100% recall). This situation is shown in the upper most curve of the graph in Fig. 8 – where the manually generated perfect similarity matrix has been compared to itself - returning an average precision of 100%. The lower curve of the graph shows the raw versus perfect matrix comparison – and an average precision of only 55%. The table on the left-hand side of the figure documents the recall precision data and indicates the percentage precision change between the perfect and raw similarity matrices. The points on the precision recall graph can be interpreted as follows. For the raw case at a recall level of 10 (meaning that when 10% of the relevant databases have been retrieved) 80.8% of those databases retrieved are actually relevant.

Fig. 9 introduces incrementally aspects of the description term generation process for assets. In the first experiment term strings of the raw data have been *filtered* and *normalized* to produce basic description terms. Filtering entails eliminating all solely system related information from the databases while normalization involves conversion of strings – where possible – to corresponding meaningful terms found in WordNet. The results of this experiment are shown by the curve labeled “Raw + filter/norm” and indicate a precision improved by an average increase of 11.8% with respect to the raw descriptions. A similar situation occurs in the second curve of Fig. 9, termed “Raw + super”, where the raw description has been augmented with hierarchy (super) level terms based on the ontological organization of WordNet. Finally, the curve “Raw + filt/norm/super”

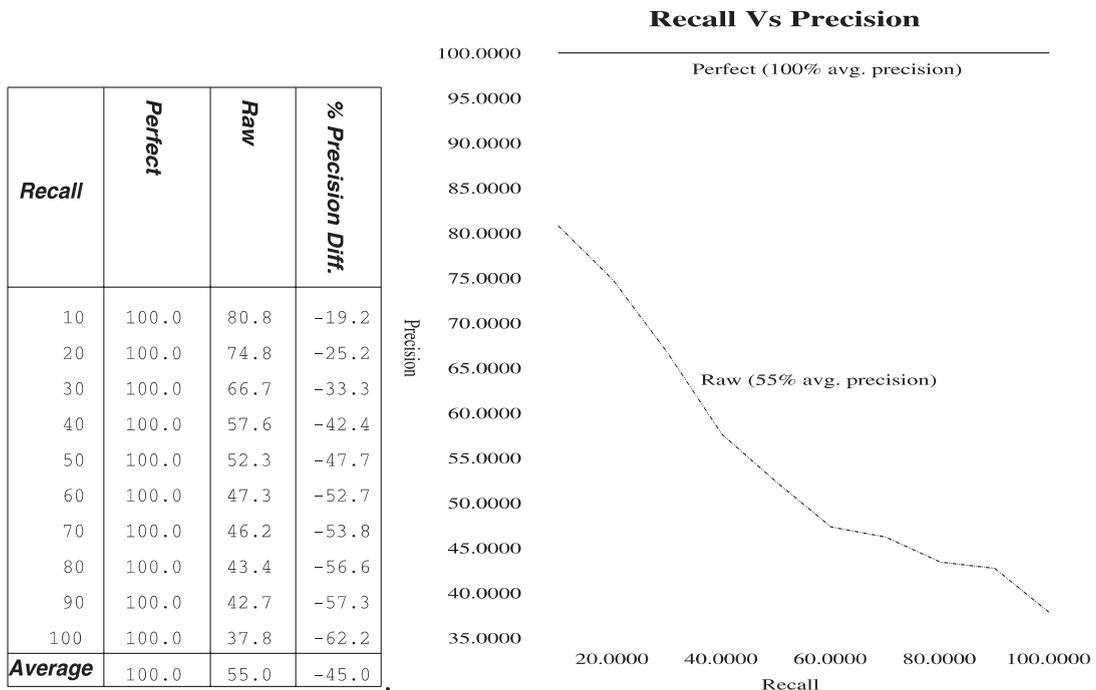


Fig. 8. Perfect and raw term (attribute) precision and recall.

Recall	Perfect	Raw	Raw + filt/norm	Raw + super	Raw + filt/norm/super
10	100.0	80.8 (-19.2)	86.6 (-13.4)	93.2 (-6.8)	92.3 (-7.7)
20	100.0	74.8 (-25.2)	81.5 (-18.5)	87.6 (-12.4)	90.0 (-10.0)
30	100.0	66.7 (-33.3)	78.2 (-21.8)	85.5 (-14.5)	88.2 (-11.8)
40	100.0	57.6 (-42.4)	74.9 (-25.1)	83.0 (-17.0)	86.5 (-13.5)
50	100.0	52.3 (-47.7)	70.6 (-29.4)	79.6 (-20.4)	84.5 (-15.5)
60	100.0	47.3 (-52.7)	66.0 (-34.0)	66.2 (-33.8)	80.2 (-19.8)
70	100.0	46.2 (-53.8)	64.8 (-35.2)	61.5 (-38.5)	81.2 (-18.8)
80	100.0	43.4 (-56.6)	61.3 (-38.7)	55.4 (-44.6)	80.1 (-19.9)
90	100.0	42.7 (-57.3)	45.1 (-54.9)	51.7 (-48.3)	70.1 (-29.9)
100	100.0	37.8 (-62.2)	38.5 (-61.5)	40.6 (-59.4)	44.6 (-55.4)
Average	100.0	55.0 (-45.0)	66.8 (-33.2)	70.0 (-30.0)	80.0 (-20.0)

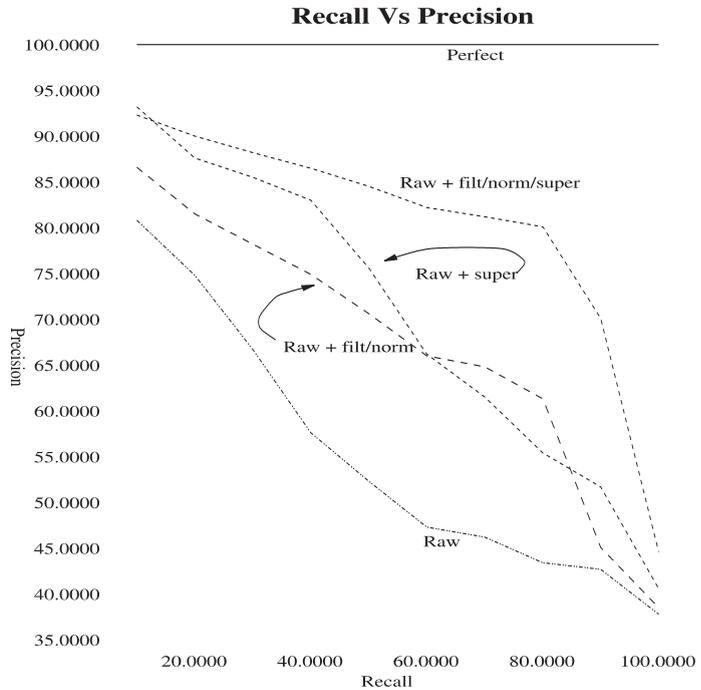


Fig. 9. Augmented raw term (attribute) descriptions.

combines the factors of filtering, normalization and higher level term descriptions. In [34] further experiments were conducted by including linguistic properties of terms, semantic, structural and branching aspects of terms and a variety of IR-based algorithms. By and large, the cumulative effects of these factors are greater than either of the previous graphs – indicating the importance of the standard ontology and term interpretation for locating schema information from multiple databases, see Fig. 10.

7. Summary and future work

The topic-based organization of a multi-database network supports semantic reconciliation of autonomous interconnected data sources as it helps the users understand what information is available through the network; helps them categorize and configure their information demands on the basis of the information available to them; and assists them to semantically disambiguate their specified terms against those provided by the database schemas in a multi-database network. This architecture enables users to gather and rearrange information from multiple networked databases in an intuitive and easily understandable manner. Large-scale searching is guided by a combination of lexical, structural and semantic aspects of schema terms in order to reveal more meaning both about the contents of a requested information item and about its placement within a given database context. Experience with this configuration suggests the clustering mechanisms used provide a valuable discovery service to end users, and that the logical organization used supports the ability of the system to scale with modest increases in topic label sizes.

Recall	Struct	Struct+super	Struct+hyp/syn/sense	Struct+hyp/syn/sense/super	Struct+syn/sense	Struct+syn/sense/super
10	90.0	92.6	89.6	93.0	90.6	92.9
20	85.3	88.4	80.2	89.9	85.1	89.1
30	79.9	86.9	75.3	87.7	80.4	87.0
40	77.0	86.3	72.5	87.2	77.1	86.3
50	74.0	85.4	69.4	86.4	74.2	85.4
60	70.7	85.1	67.7	84.7	70.8	85.1
70	68.7	84.7	65.6	83.3	68.7	84.7
80	64.3	83.9	63.0	81.5	64.6	83.8
90	56.2	70.5	57.3	66.2	56.3	70.4
100	42.2	46.1	51.9	50.4	42.1	46.1
Average	70.8	81.0	69.3	81.0	71.0	81.1

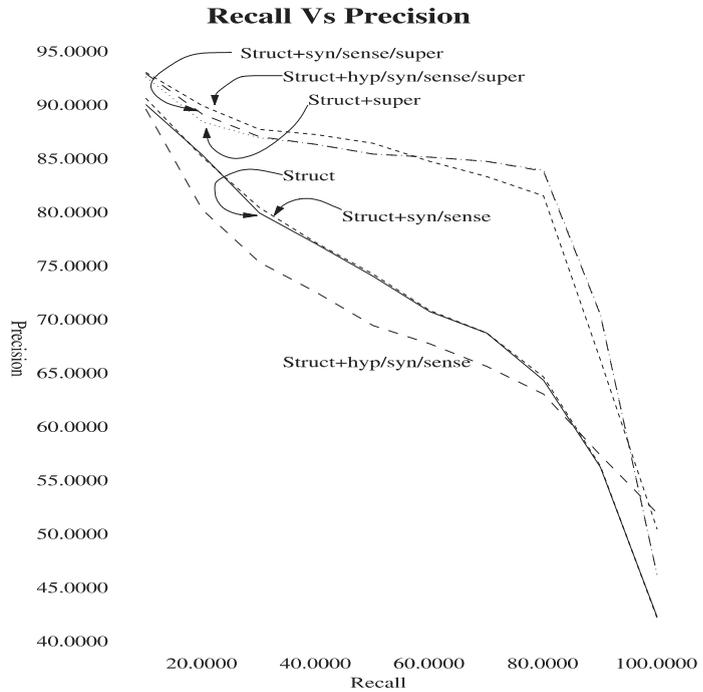


Fig. 10. Cumulative term descriptions.

In the current prototype link weights are established at a clustering phase on a tentative basis only. However, it is expected that during execution link weights to topics may need to be updated (strengthened or weakened) over time depending on interaction, new topics may be formed, and existing topics may need to merge. The next suite of experiments to be performed will deal with the characteristics of the link weight update and topic split/merge processes. From this policies will be developed (e.g., delayed/batch updating of topic information), and then evaluated.

Acknowledgements

We wish to acknowledge Steven Milliner for his help with earlier publications on the same topic and for his implementation and performance experiments in connection with the system reported herein. We also wish to acknowledge the reviewers of this paper for their constructive comments which resulted in improving the quality of the paper.

References

[1] M. Agosti, A. Archi, R. Colotti, R.M. Di Giorgi, G. Gradenigo, B. Inghirami, P. Matiello, R. Nannuci, M. Ragona, New perspectives in information retrieval techniques: a hypertext prototype in environmental law, in: Online Management 89, Proceedings of the 13th International Online Information, London, UK, 1989.
 [2] M. Agosti, F. Crestani, G. Gradenigo, Towards data modelling in information retrieval, J. Inform. Sci. 15 (6) (1989).

- [3] M. Agosti, R. Colotti, G. Gradenigo, A two-level hypertext retrieval model for legal data, in: A. Bookstein, Y. Chiaramella, G. Salton, V.V. Raghavan (Eds.), *Proceedings of the 14th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM Press, Chicago, Illinois, October 1991.
- [4] M. Agosti, G. Gradenigo, P.G. Marchetti, A hypertext environment for interacting with large textual databases, *Inform. Process. Manage.* 28 (3) (1992).
- [5] Digital's AltaVista search engine project, <http://www.altavista.com/>.
- [6] M. Agosti, M. Melucci, F. Crestani, Automatic authoring and construction of hypermedia for information retrieval, *ACM Multimedia Systems*, 1995.
- [7] P. Atzeni, G. Mecca, P. Merialdo, Weave the Web, in: *Proceedings of the 23rd International Conference on Very Large Data Bases*, Athens, Greece, 1997.
- [8] Y. Arens, et al., Retrieving and integrating data from multiple information sources, *Int. J. Cooperative Inform. Systems* 2 (2) (1993).
- [9] R. Burke, K.J. Hammond, Combining databases and knowledgebases for assisted browsing, in: *Proceedings of the AAAI'95 Symposium on Information Gathering from Distributed, Heterogeneous Environments*, Palo Alto, California, March 1995.
- [10] M. Bright, A. Hurson, S. Pakzad, Automated resolution of semantic heterogeneity in multidatabases, *ACM Trans. Database Systems* 19 (2) (1994).
- [11] P.D. Bruza, M. Lalmas, Logic-based Information Retrieval: Is it really worth it? in: H.J. Zimmerman (Ed.), *Proceedings of the Fourth European Congress on Intelligent Techniques and Soft Computing – EUFIT96*, Wissenschaftsverlag, Aachen, Germany, 1996.
- [12] C.M. Bowman, et al., Harvest: A Scalable, Customizable Discovery and Access System, Technical report CU-CS 732-94 (revised March 1995), University of Colorado-Boulder, Computer Science Department, Boulder, Colorado, March, 1995.
- [13] T. Brasethvik, A semantic modeling approach to metadata, *Internet Res.* 8 (5) (1998).
- [14] P.D. Bruza, Hyperindices: A Novel Aid for Searching in Hypermedia, in: A. Rizk, N. Streitz, J. Andre (Eds.), *Proceedings of the European Conference on Hypertext – ECHT 90*, Cambridge University Press, Cambridge, 1990.
- [15] P.D. Bruza, Th.P. van der Weide, Two level hypermedia – an improved architecture for hypertext, in: A.M. Tjoa, R. Wagner (Eds.), *Proceedings of the Data Base and Expert System Applications Conference (DEXA 90)*, Springer, Vienna, Austria, 1990.
- [16] F.C. Berger, Th.P. van der Weide, A feedback mechanism for query by navigation, in: R. Sacks-Davis, J. Zobel (Eds.), *Proceedings of the Sixth Australasian Database Conference, ADC'95*, Australian Computer Science Communications, Adelaide, Australia, January 1995.
- [17] S. Castano, V. De Antonellis, Semantic dictionary design for database interoperability, in: *13th International Conference on Data Engineering*, Birmingham, UK, April 1997.
- [18] T. Craven, Linked phrase indexing, *Inform. Process. Manage.* 14 (6) (1978).
- [19] Dublin Core Metadata Initiative, <http://purl.org/DC/>.
- [20] M. Craven et al., Learning to extract symbolic knowledge from the World Wide Web, in: *Proceedings of the AAAI-98 Conference on Artificial Intelligence*, 1998.
- [21] B. Everitt, Cluster analysis, Heinemann Educational Books, UK, 1981.
- [22] R. Fikes, R. Englemore, A. Farquhar, W. Pratt, Network-based information brokers, in: *Proceedings of the AAAI'95 Symposium on Information Gathering from Distributed, Heterogeneous Environments*, Palo Alto, California, March 1995.
- [23] N. Fuhr, Probabilistic models in information retrieval, *The Comput. J.* 35 (3) (1992).
- [24] T.R. Gruber, Towards principles for the design of ontologies used for knowledge sharing, in: N. Guarino, R. Poli (Eds.), *Formal Ontology in Conceptual Analysis and Knowledge Representation*, Kluwer Academic Publishers, Deventer, The Netherlands, 1993.
- [25] J. Heflin, J. Hendler, S. Luke, Coping with changing ontologies in a distributed environment, in: *AAAI-99 Workshop on Ontology Management*, Orlando, USA, July 1999.
- [26] A.H.M. ter Hofstede, H.A. Proper, Th.P. van der Weide, Query formulation as an information retrieval problem, *The Comput. J.* 39 (4) (1996).
- [27] D. Kuokka, L. Harad, Supporting information retrieval via matchmaking, in: *Proceedings of the AAAI'95 Symposium on Information Gathering from Distributed, Heterogeneous Environments*, Palo Alto, California, March 1995.
- [28] J. Kirriemuir et al. Cross-searching subject gateways, *D-Lib Magazine*, January 1998.
- [29] J. Kahng, D. McLeod, Dynamic classificational ontologies: Mediation of information sharing in cooperative federated database systems, in: M.P. Papazoglou, G. Schlageter (Eds.), *Cooperative Information Systems: Trends and Directions*, Academic Press, London, 1997.
- [30] G. Kowalski, *Information Retrieval Systems: Theory and Implementation*, Kluwer Academic Publishers, Deventer, The Netherlands, 1997.
- [31] V. Kashyap, A. Sheth, Semantic heterogeneity in global information systems: the role of metadata, context and ontologies, in: M.P. Papazoglou, G. Schlageter (Eds.), *Cooperative Information Systems: Trends and Directions*, Academic Press, London, 1997.

- [32] A. Levy, A. Rajaraman, J.J. Ordille, Querying heterogeneous information sources using source descriptions, in: Proceedings of the 23rd International Conference on Very Large Data Bases, Bombay, India, September 1996.
- [33] S. Milliner, A. Bouguettaya, M. Papazoglou, A scalable architecture for autonomous heterogeneous database interactions, in: Proceedings of the 21th VLDB Conference, Zurich, Switzerland, September 1995.
- [34] S. Milliner, Resolution of conceptual heterogeneity: achieving context via dynamic ontology, Ph.D. thesis, Queensland University of Technology, Brisbane, Australia, 2000.
- [35] M.L. Mauldin, J.R. Levitt, Web-agent related Research at the CMT, in: Proceedings of the 17th Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval, Dublin, Ireland, July 1994.
- [36] S. Milliner, M. Papazoglou, H. Weigand, Linguistic tool based information elicitation in large heterogeneous database networks, in: R.P. vandeRiet, J.F.M. Burg, A.J. vanderVos (Eds.), Proceedings of the Second Workshop on Applications of Natural Language to Databases (NLDB'96), Amsterdam, The Netherlands, June 1996.
- [37] G.A. Miller, R. Beckwith, C. Fellbaum, D. Gross, K.J. Miller, Introduction to wordnet: an on-line lexical database, *J. Lexicogr.* 3 (4) (1990).
- [38] D. McLeod, A. Si, The design and experimental evaluation of an information discovery mechanism for networks of autonomous database systems, in: 11th International Conference on Data Engineering, Taiwan, February 1995.
- [39] S.H. Myaeng, Using conceptual graphs for information retrieval: a framework for representation and flexible inferencing, in: Proceedings of the Symposium on Document Analysis and Information Retrieval, Las Vegas, Nevada, March 1992.
- [40] OMNI, Organizing Medical Networked Information, <http://omni.ac.uk/>.
- [41] T. Oates, N. Nagendra, V. Lesser, Cooperative information gathering: a distributed problem solving approach, Technical report 94-66 (version 2), Department of Computer Science, University of Massachusetts, Amherst, Massachusetts, 1994.
- [42] M.P. Papazoglou, Unraveling the semantics of conceptual schemas, *Commun. ACM* 38 (9) (1995).
- [43] H.A. Proper, P.D. Bruza, What is information discovery about, *J. Amer. Soc. Inform. Sci.* 50 (9) (1999).
- [44] Platform for Internet Content Selection (PICS), <http://www.w3.org/PICS/>.
- [45] B. Pinkerton, Finding what people want: experiences with the WebCrawler, in: Proceedings of the First International Conference on the WWW, Geneva, Switzerland, May 1994.
- [46] M.P. Papazoglou, S. Milliner, Pro-active information elicitation in wide-area information networks, Proceedings of the International Symposium on Cooperative Database Systems for Advanced Applications, World Scientific, Japan, 1996.
- [47] M.P. Papazoglou, S. Milliner, Subject-based organization of the information space in multi-database networks, Proceedings of the 10th International Conference CAISE'98 on Advanced Information Systems Engineering, Lecture Notes in Computer Science, Springer, Pisa, Italy, 1998.
- [48] R. Rada, E. Bicknell, Ranking documents based on a thesaurus, *J. Amer. Soc. Inform. Sci.* 40 (5) (1989).
- [49] C.J. van Rijsbergen, *Information Retrieval*, Butterworths, London, UK, 1990.
- [50] G.E. Salton, C. Buckley, Term-weighting approaches in automatic text retrieval, *Inform. Retrieval Manage.* 24 (5) (1988).
- [51] R.B. Schatz, et al., Interactive term suggestion for users of digital libraries, First ACM International Conference on Digital Libraries, Bethesda, Maryland, March, 1996.
- [52] M.A. Sheldon, Content routing: a scalable architecture for network-based information discovery Ph.D. thesis, Massachusetts Institute of Technology, Boston, Massachusetts, December 1995.
- [53] SOSIG: The Social Science Information Gateway, <http://www.sosig.ac.uk/>.
- [54] J.F. Sowa, *Conceptual Structures: Information Processing in Mind and Machine*, Addison-Wesley, Reading, Massachusetts, 1984.
- [55] S. Weibel, J. Godby, E. Miller, R. Danierl, Metadata Workshop Report, Dublin, Ohio, March 1995.
- [56] R. Wiess et al., HyPersuit: A hierarchical network search engine that exploits content-link hypertext clustering, in: Seventh ACM Conference on Hypertext, Washington DC, March 1996.
- [57] S.K.M. Wong, Y.Y. Yao, A probabilistic inference model for information retrieval, *Inform. Systems* 16 (3) (1990).

Michael P. Papazoglou is a full professor and director of the Infolab at the University of Tilburg in the Netherlands. His scientific interests include cooperative information systems, heterogeneous database systems, object-oriented systems and modelling, distributed computing, digital libraries and electronic commerce where he has authored approximately 100 journal articles and refereed conference papers and edited 10 books. He has chaired several prestigious conferences and serves on several scientific committees and advisory boards for international journals. His research has been funded by the European Commission, the Australian Research Council, the Japanese Society for the Promotion of Science, and Departments of Science and Technology in Europe and Australia. Papazoglou is a golden core member of the IEEE and a recipient of the prestigious IEEE certificate of appreciation for his contributions to Computer Science as distinguished Visitor of the IEEE.

H.A. (Erik) Proper is currently principal scientist at the Ordina Institute for Research and Innovation, Gouda, The Netherlands. He is also seconded as a senior research fellow to the University of Nijmegen, The Netherlands. Proper has co-authored several journal papers, conference publications and books. His main research interests include information architectures, business/IT alignment, conceptual modelling, information retrieval and information discovery. Proper received his Master's degree in May 1990 and his Ph.D. (with distinction) in April 1994 from the University of Nijmegen, The Netherlands. After receiving his Ph.D., Proper became a senior research fellow at the Computer Science Department of the University of Queensland, Brisbane, Australia. During that period he also

conducted research in the Asymetrix Research Lab at that University for Asymetrix Corp, Seattle, Washington. In 1995 he became a lecturer at the School of Information Systems from the Queensland University of Technology, Brisbane, Australia. During this period he was also seconded as a senior researcher to the Distributed Systems Technology Centre (DSTC), a Cooperative Research Centre funded by the Australian government. As of 1997, Dr. Proper works mainly in industry. First as a consultant at Origin, Amsterdam, The Netherlands, and now as a principal scientist at the Ordina Institute for Research and Innovation, Gouda, The Netherlands.

Jian Yang is an assistant professor at Infolab, Tilburg University. She received her Ph.D. in Multidatabase Systems area from The Australian National University in 1995. Her Ph.D. thesis was on “Interoperability Support in Multidatabase Systems”. Before she joined INFOLAB, she worked as a senior research scientist at the Division of Mathematical and Information Science, CSIRO in Canberra (1998-2000), and as a lecturer (assistant professor) at the Department of Computer Science, The Australian Defence Force Academy (1993–1998). Her research interests include: objected-oriented analysis and design; interoperability issues in digital libraries and e-commerce; query languages and query optimization; distributed query processing; Web-based application design; materialized view design and data warehousing.