

# ArHeX: Flexible Composition of Indexes and Similarity Measures for XML

Ismael Sanz, Rafael Berlanga  
Universitat Jaume I, Spain  
{isanz,berlanga}@uji.es

Marco Mesiti  
Università di Milano, Italy  
mesiti@dico.unimi.it

Giovanna Guerrini  
Università di Genova, Italy  
guerrini@disi.unige.it

## Abstract

*This work-in-progress paper describes the features of the ArHeX similarity-oriented XML processing toolkit [12]. ArHeX is designed to assist in the engineering of XML similarity-oriented applications, supporting the design and evaluation of suitable similarity measures and their associated indexes for each specific application.*

## 1. Introduction

There are many XML database applications that require some notions of similarity. For example, in the integration and merging of highly heterogeneous XML databases, in which there is no common schema, exact approaches are impractical due to the great (and unpredictable) structural variations of the diverse sources. Therefore, similarity measures are required to identify similar information modeled using different structures or terminologies. Moreover, in systems handling objects with complex structures (e.g., protein data, music retrieval systems, or shape databases) there is a fundamental need of identifying similar objects according to a similarity function. These applications usually require carefully tailored similarity functions that consider both structural variations and vocabulary discrepancies occurring in the collection of documents.

A crucial design issue in the development of these kinds of similarity-based systems is that a single notion of similarity that “works best” in any situation does not exist. Different users in different contexts may require different similarity functions; for instance, a biologist may wish to retrieve proteins based on a comparison with a given amino acid sequence, while another one may issue a query asking for “malaria antigen” in the associated textual description, and a third user may combine both kinds of queries. This leads to the notion of *multi-similarity* systems [1], which are designed to support multiple notions of similarity simultaneously. These systems require the possibility to combine different similarity measures depending on the characteristics of the data that need to be handled. This aspect is par-

ticular relevant when handling heterogeneous collections of XML documents because heterogeneity can occur at many different levels – from changes in the vocabulary used in tag names to complex document-wide structural variations. Moreover, retrieval should be performed efficiently. Thus, similarity measures should be coupled with specifically tailored indexing structures for their efficient computation.

Starting from these requirements, we are developing the ArHeX similarity-oriented XML processing toolkit. Key features of ArHeX are: (i) its ability to support collections which are heterogeneous at multiple levels of granularity, (ii) its flexible pattern-based query model, and (iii) its component-based architecture. These features allow ArHeX to support multiple user-defined similarity measures on top of efficient indexes. The following sections describe the features of ArHeX and the techniques it employs.

## 2 Measure Definition and Composition

A wide variety of similarity measures, both general purpose and specifically tailored, has been proposed for XML [7]. Each measure produces good results when the collections present specific characteristics and are hardly reusable in other collections. General purpose measures include metric functions such as the Manhattan or Euclidean distances; IR-like matching coefficients such as the cosine with  $tf \times idf$  weighting; entropy-based measures such as the Kullback-Leibler divergence; and structure-oriented techniques such as variants of the Tree Edit Distance algorithm [11].

These similarity measures are usually obtained through the composition of several “atomic” measures at a given *granularity level* of the XML hierarchy. For instance, a measure for complete XML documents is defined in terms of paths similarity, which in turn requires some criterion to compare the elements in the path. The following levels can be devised: the whole XML document, subtrees (i.e., regions of documents), paths, elements, links, attributes and textual content (of attributes and data content elements).

This indicates that it is possible to build frameworks for the implementation of complex XML similarity measures,

based on a library of basic component functions (implementing the atomic measures). This does not exclude the employment of ad-hoc measures, if necessary. For example, the designer of a similarity-based application in the domain of genetics may need to combine a generic text-oriented function that matches protein names with a highly specialized function that matches amino acid sequences.

We have thus defined a formal framework for the definition and composition of similarity measures relying on the granularity levels of documents and then specified software components implementing such functions that can be combined to obtain new measures specific for a given context.

**A Formal Specification of Similarity Measures.** Let  $\mathcal{DOC}$  be a set of XML documents, and  $G = \{\text{DOC}, \text{ELEMENT}, \text{ATTRIBUTE}, \text{PATH}, \text{REGION}, \text{CONTENT}, \text{LINK}\}$  the granularity levels at which documents in  $\mathcal{DOC}$  can be compared. Given a granularity level  $\gamma \in G$ , a mapping function  $m_\gamma$  can be defined for extracting from a collection  $S \subseteq \mathcal{DOC}$  the portions of documents at that granularity level. For example, the mapping functions  $m_{\text{ELEMENT}}$  and  $m_{\text{PATH}}$  applied on a collection  $S \subseteq \mathcal{DOC}$  return the set of elements (denoted  $m_{\text{ELEMENT}}(S)$ ) and paths (denoted  $m_{\text{PATH}}(S)$ ) occurring in  $S$ , respectively. A partial order relation  $\prec_G$  between granularity levels of  $G$  can be defined, representing the containment relationship between granularity levels; it can be read as “is lower-level than”. For instance,  $\text{ELEMENT} \prec_G \text{PATH}$  because the paths in  $m_{\text{PATH}}(S)$  are defined in terms of elements in  $m_{\text{ELEMENT}}(S)$ .

A similarity measure at a given granularity level  $\gamma$  can be defined as a function  $f_\gamma : m_\gamma(S) \times m_\gamma(S) \rightarrow [0, 1]$ . Moreover, it can be expressed in terms of other lower level similarity functions. For example, if we denote elements by  $e_n$  and paths by  $p_m$ , then an instance of similarity function for paths expressed in terms of a similarity function for elements is

$$f_{\text{PATH}}(p_1, p_2) = \frac{\sum_{i,j} f_{\text{ELEMENT}}(e_i, e_j)}{|p_1||p_2|}$$

where in the simplest case  $f_{\text{ELEMENT}}(e_i, e_j) = 1$  if  $e_i = e_j$  and 0 otherwise, and  $|p|$  denotes the length of path  $p$ .

**Measures as Components.** In ArHeX, a *measure component* is an implementation of a similarity function at a given granularity level. According to the definitions above, a component can depend on one or more lower-level components, but it is irrelevant *which* concrete lower-level component is used, as long as it belongs to the right granularity level. This capability of creating complex measures by combining simpler parts which are measures themselves is analogous to the *Composite* pattern commonly used in Software Engineering [6]; the general model for measure com-

position can be illustrated by the UML diagram in Fig. 1. In addition, every component can be parametrized; for instance, a component that computes similarity at the textual level may allow the user to choose whether common words (“stop words”) must be considered.

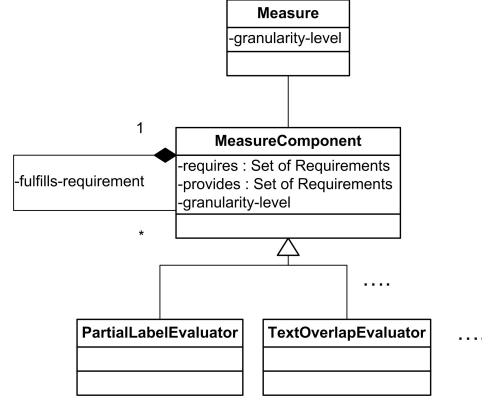


Fig. 1. Simplified UML model describing the application of the *Composite* pattern to the creation of complex measures: A *MeasureComponent*’s requirements are fulfilled by other *MeasureComponents*, creating the provides/requires hierarchy.

Using this framework, a large number of functions with different requirements can be defined for each granularity level. In order to characterize them, the partial ordering of levels  $\prec_G$  is extended into a *provides/requires hierarchy* typically used in software component engineering [10]. Each measure component is thus tagged with two extra properties, whose values are chosen from a predefined set of *features*: the *provides* property indicates the granularity level at which the function operates (e.g., all node-level functions provide the feature *nodeMatch*), while the *requires* property indicates the features that must be provided by the lower-level functions on which the function relies.

For instance, let  $C_{\text{PATH}}$  be a component that implements the  $f_{\text{PATH}}$  function defined above. Then,  $\text{provides}(C_{\text{PATH}}) = \{\text{pathMatch}\}$  and  $\text{requires}(C_{\text{PATH}}) = \{\text{elementMatch}\}$ . For the component to be usable, another component that provides *elementMatch* must be available.

The provides/requires hierarchy separates the representation of the components from the actual similarity function implemented. This allows us to implement components for generic operations (usually called “tie components” [10]) by computing an aggregated value out of the results of other components. For instance, consider the “weighted sum” tie component  $C_{\text{WSUM}}$ , that can be defined for a set of  $n$  components  $\{C_1, \dots, C_n\}$  at the same granularity level  $\gamma$  and a set of  $n$  weights  $\{w_1 \dots w_n | w_i \in \mathbb{R}\}$ :

$$\begin{aligned}
parameters(C_{WSUM}) &= \{\{C_1 \dots C_n\}, \{w_1 \dots w_n\}\} \\
provides(C_{WSUM}) &= \bigcap_i provides(C_i) \\
requires(C_{WSUM}) &= \bigcup_i requires(C_i) \\
f_{C_{WSUM}} &: G \times G \rightarrow [0, 1] \\
f_{C_{WSUM}}(o_1, o_2) &= \sum_i w_i \times f_{C_i}(o_1, o_2)
\end{aligned}$$

Fig. 2 shows an instantiation of a multilevel similarity measure that uses *WeightedSum*, the ArHeX implementation of  $C_{WSUM}$ . It combines components at the node, label, and text granularity levels. Note how the components at the node similarity level are computed using a weighted sum. Edge labels represent the requires/provides hierarchy.

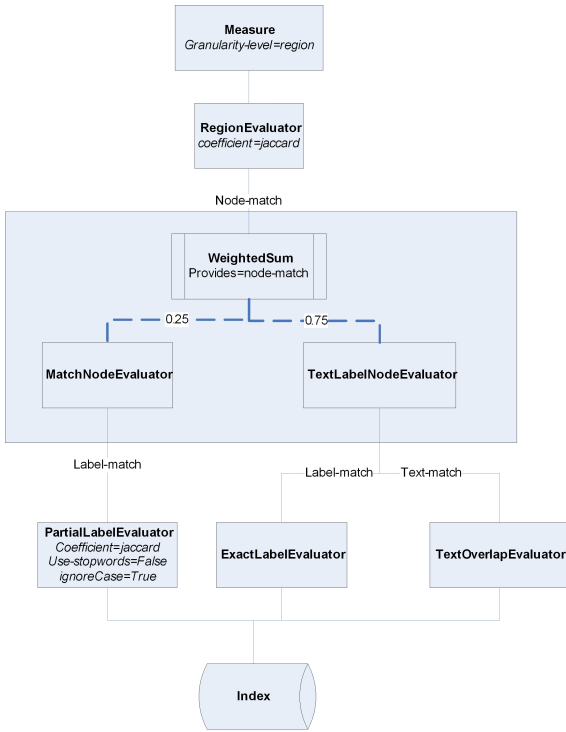


Fig. 2. Component structure of a similarity measure

**A Formal Model of Components.** We have just shown how to create components that can be glued together to form flexible similarity measures. However, the data engineer needs to answer higher-level questions. For instance: given a particular component, which other components are available to fulfill its requirements? Or, is this component sound (i.e. all of its requirements are correctly fulfilled)? This is particularly important in collaborative methodologies, in which sharing components among independently-working engineers is crucial.

Our approach is to encapsulate all of the required consistency rules in a declarative formalism, using a suitable Description Logic (DL). DLs provide a set of reasoning tasks (subsumption, instance checking, relation checking, concept consistency and knowledge base consistency) exploited in our context to automate many of the tasks that must be performed by the data engineer when designing a multi-similarity system. In addition, sufficiently expressive DLs provide “inverse functional” roles, which are exactly equivalent to candidate keys in a database. This is useful to support the semi-automatic specification of indexes as a combination of index components.

The mapping between the model previously outlined and a DL-based representation is straightforward. The rules (*A-Box* in DL parlance) can be expressed in terms of the concepts *Component* and *Feature*, and the roles *provides* and *requires*. The terminological knowledge (*T-Box*) reflects the current state of the system. An example of an instance of *RegionEvaluatorComponent* used in Fig. 2 is:

$$\begin{aligned}
&RegionEvaluator : Component \\
&(RegionEvaluator, regionMatch) : provides \\
&(RegionEvaluator, nodeMatch) : requires
\end{aligned}$$

Complex concepts and terminological rules are built on top of these instances. For example, the set of all components at the “region” granularity level is expressed as

$$\begin{aligned}
RegionEvaluatorComponent &\doteq \\
&Component \sqcap \exists provides.regionMatch
\end{aligned}$$

The DL reasoner automatically classifies the *RegionEvaluator* component as an instance of *RegionEvaluatorComponent*. Similar rules can be defined for consistency checking as outlined above.

### 3 Index Composition

Composition-based systems have already been applied in the context of schema matching (e.g. COMA++ [4]). However, a purely functional approach where components operate directly over the actual data is clearly not appropriate in our large-scale Web data context. To compute these functions efficiently, it is necessary to exploit suitable indexes on the base XML data. For instance, consider a measure component for word similarity that uses  $tf \times idf$  weights. The only way to compute such a measure efficiently in a large collection is to add the global frequency of each word to the index information. In ArHeX, this index information is called *index components*. ArHeX indexes are built as a composition of a base XML index plus a set of such components. This leads us to a second feature-based provides/requires hierarchy that associates measures to index

components. For example, the measure depicted in Fig. 2 relies on the availability of index components that provided facilities for partial label matching. The possibility of exploiting indexes on pre-computed aggregate values of composite measures is being evaluated as well.

## 4 Issues in Query Processing

The dual hierarchy of measure and index components leads naturally to two levels of algebraic optimization. First, the tree obtained by the composition of measure components is just a variant of an expression tree, and thus familiar expression optimization techniques adapted from compiler technology, such as common subexpression elimination, can be applied on it. Then, using the mapping of measure components to index components, the resulting expression is translated to a physical algebra, on which cost-based optimizations can be performed to obtain an execution plan.

Our flexible measure definition model seems naturally fitted to a top- $k$  query processor. There has been a lot of recent work describing approaches for optimizing top- $k$  queries in an XML context; see for example [8, 9, 13]. Many of these approaches use variants of Fagin’s threshold algorithm, which is suitable for monotonic aggregation functions [5]. In our context, however, there is no guarantee that the user-defined functions have this property; even a simple case like a weighted sum with negative coefficients (e.g., used to penalize certain features) is non-monotonic.

We are currently evaluating techniques to perform approximate top- $k$  queries in the presence of non-monotonic aggregation functions. To avoid evaluating all candidate results, we are studying sampling techniques to find promising subsets of the collection in which “good” results are more probable. Analogous techniques have been used for related problems in different contexts; see e.g. [3]. Our preliminary experiments show promising results, but many issues still remain open.

## 5 Summary

In this work-in-progress paper we have addressed the ArHeX approach to multi-similarity systems. ArHeX supports a flexible, component-oriented way to define measures, and Description Logic-based meta-data facilities that help support multiple measures simultaneously. Our goal is to build a framework which can be used to build highly tailored similarity-oriented systems, enhancing the adaptability of current approaches such as XML pattern relaxation techniques [2] and diff algorithms for change detection in XML documents [14].

In the paper, we also raised the open problem of efficiently evaluating top- $k$  queries in non-monotonic aggregation functions, of practical importance in our context.

An online version of the current prototype is available at <http://krono.act.uji.es/arhex/Demo>.

## References

- [1] S. Adali, P. Bonatti, M. L. Sapino, and V. S. Subrahmanian. A Multi-similarity Algebra. *SIGMOD*, 402–413, 1998.
- [2] S. Amer-Yahia, L. V. S. Lakshmanan, and S. Pandit. FlexPath: Flexible Structure and Full-text Querying for XML. *SIGMOD*, 83–94, 2004.
- [3] D. Berleant, L. Xie, and J. Zhang. Statool: A Tool for Distribution Envelope Determination (DENV), an Interval-based Algorithm for Arithmetic on Random Variables. *Reliable Computing*, 9(2):91–108, 2003.
- [4] H. H. Do and E. Rahm. COMA - A System for Flexible Combination of Schema Matching Approaches. *VLDB*, 610–621, 2002.
- [5] R. Fagin, A. Lotem, and M. Naor. Optimal Aggregation Algorithms for Middleware. *J. Computer and System Sciences*, 66:614–656, 2003.
- [6] E. Gamma, R. Helm, R. Johnson, J. Vlissides. Design Patterns: Elements of Reusable Object-Oriented Software. Addison Wesley, 1995.
- [7] G. Guerrini, M. Mesiti, and I. Sanz. An Overview of Similarity Measures for Clustering XML Documents. *Web Data Management Practices: Emerging Techniques and Technologies*, 56–78. Idea Group, 2006.
- [8] R. Kaushik, R. Krishnamurthy, J. Naughton, and R. Ramakrishnan. On the Integration of Structure Indexes and Inverted Lists. *SIGMOD*, 779–790, 2004.
- [9] A. Marian, S. Amer-Yahia, N. Koudas, and D. Srivastava. Adaptive Processing of Top- $k$  Queries in XML. *ICDE*, 162–173, 2005.
- [10] F. Plasil and S. Visnovsky. Behavior Protocols for Software Components. *IEEE Trans. Softw. Eng.*, 28(11):1056–1076, 2002.
- [11] S. M. Selkow. The Tree-to-tree Editing Problem. *Information Processing Letters*, 184–186, 1977.
- [12] I. Sanz, M. Mesiti, G. Guerrini, and R. Berlanga. ArHeX: An Approximate Retrieval System for Highly Heterogeneous XML Document Collections. Demo at *EDBT, LNCS(3896)*, 1186–1189, 2006.
- [13] M. Theobald, G. Weikum, and R. Schenkel. Top- $k$  Query Evaluation with Probabilistic Guarantees. *VLDB*, 648–659, 2004.
- [14] Y. Wang, D. J. DeWitt, and J.-Y. Cai. X-Diff: An Effective Change Detection Algorithm for XML Documents. *ICDE*, 519–531, 2004.