Comparing and Evaluating Mapping Systems with STBenchmark

Bogdan Alexe
UC Santa Cruz
abogdan@cs.ucsc.edu

Wang-Chiew Tan
UC Santa Cruz
wctan@cs.ucsc.edu

Yannis Velegrakis
University of Trento
velgias@disi.unitn.eu

ABSTRACT

Schema mappings are fundamental building blocks in many information integration applications. Designing mappings is a time-consuming process and for that reason many mapping systems have been developed to assist in the task of designing mappings. However, to the best of our knowledge, a benchmark for comparing and evaluating these systems has not yet been developed. We demonstrate STBenchmark, a benchmark that we have developed for evaluating mapping systems. Our demonstration will showcase the different aspects of mapping systems that STBenchmark evaluates, highlight the results of our comparison and evaluation of four mapping systems, as well as make a case for the need for a standard specification input mechanism to mapping systems in order to make progress towards the development of a uniform testbed or repository for schema mappings and data exchange tasks.

1. INTRODUCTION

A fundamental problem in information integration is generating *mappings* between schemas. A mapping is a precise specification of how an instance over one schema, called the *source* schema, is to be translated into an instance over a second schema, called the *target* schema. Designing mappings is a time-consuming process due to the high heterogeneity of the typically independently developed schemas and many mapping systems have been developed to make this task easier. A mapping system is a visual programming system with the goal of making it easy for a designer to generate mappings. The interfaces of mapping systems provide a graphic representation of the source and target schema, and allow the design of mappings by associating elements between the two schemas. The association can range from very simple, such as a direct line between two elements, to more complex, such as an association of two or more elements through a graphical box that denotes a function application. The visual specification is then compiled into some executable code, referred to as the *transformation script*, that is typically expressed using languages such as XSLT, XQuery, Java or C.

Examples of mapping systems include Altova Mapforce [11], AquaLogic [5], IBM Rational Data Architect [10], Microsoft BizTalk Mapper [18], which is embedded in Microsoft Visual Stu-

dio, Stylus Studio [15], and the research prototypes HePToX [3] and Clio [8]. Visual interfaces for specifying XML-to-XML [6, 7] or relational-to-relational [21] transformations can also be considered mapping systems. Unfortunately, until today, there have been only limited efforts on evaluating the mapping [4] or the data transformation process [17] and no generic benchmark exists to compare and evaluate different mapping systems. Similar to the motivation of benchmarking relational database management systems, a benchmark for mapping systems is important for assessing the relative merits of the different systems, which in turn is important for customers in order to make the right investment decisions. In fact, a recent workshop on information integration [2] also raised the need for developing benchmarks for data exchange systems. We propose to demonstrate STBenchmark (http://www.stbentchmark.org), a benchmark we have developed for that purpose. The following section describes the challenges that were faced in the development of the benchmark while Section 3 provides an overview of the demonstration.

2. CHALLENGES AND SOLUTIONS

Expressing application scenarios. A benchmark typically consists of a set of standard application scenarios that can be tested against different systems that offer similar functionalist's (e.g., the popular TPC-H [16] and XMark [14] benchmarks). For a consequence, the application scenarios must be clearly understood and correctly interpreted by the different systems under evaluation. Unlike benchmarks for other types of systems, such as relational [16] or XML query processing engines [14], it is considerably more challenging to design a benchmark for mapping systems. A major difficulty arises from the fact that there is no standard input language or input methodology for mapping systems. Benchmarks for query engines, for instance, can exploit the standard query language, e.g., SQL or XQuery, to precisely communicate the benchmark application scenarios to different query processing engines. Mapping systems, on the other hand, do not have a standard set of visual metaphors in their graphical interfaces. Even if they do use similar metaphors, these metaphors may be interpreted differently by the various mapping systems. Consider, for example, the simple visual specification illustrated in Figure 1(a) that intuitively describes how a table is copied from one schema to another. The lines, which may be the result of a schema matching process, indicate the correspondences between the elements. This visual specification is compiled into nonequivalent XSLT scripts by different mapping systems. The XSLT script that is generated by Altova Mapforce [11] groups all names of proteins in the source, followed by all accession information and creation dates, under a single ⟨Protein⟩ tag. (ref. Figure 2(b) which is the result of applying Mapforce's XSLT script on the source instance shown in Figure 2(a).) Thus, in

Figure 1: Mapping Scenarios: (a) Copy (b) Denormalization

Figure 2: (a) Source Instance (b,c,d) Target Instances

the case when the source instance consists of more than one protein, the target instance generated by the XSLT script does not even conform to the target schema. In fact, to specify a transformation that copies the source instance, Mapforce requires an additional line between the Protein elements in Figure 1(a). The XSLT script that is generated by Stylus Studio [15] creates a single ⟨Protein⟩ tag within which there are single ⟨Name⟩, ⟨Accession⟩ and ⟨Created⟩ tags. Only the name, accession information and creation date of the first protein in the source are listed under the ⟨Name⟩, ⟨Accession⟩ and ⟨Creation⟩ tags respectively. (ref. Figure 2(c).) Microsoft's BizTalk Mapper [18], IBM Rational Data Architect [10] and Clio [12] generate XSLT scripts that return a copy of the source instance. (ref. Figure 2(d).)

The above example is an indication of the fact that the benchmark application scenarios cannot be specified as visual specifications. To overcome this problem, the benchmark application scenarios were designed as *mapping scenarios*. A *mapping scenario* is a triple $(S, T, \mathcal{P})$, where S and T are the source and target schema, respectively, and $\mathcal{P}$ is a precise description of how an instance of S is to be translated into an instance of T. Since mapping systems do not take mapping scenarios as input, the benchmark user is required to express each mapping scenario as a visual specification through the graphical interface of the mapping system.

STBenchmark consists of three different components, each one used to evaluate a different aspect of a mapping systems:

(i) Basic Mapping Scenarios and Source Instances. STBenchmark offers a set of basic mapping scenarios that we believe represents a minimum set of transformations that should be readily supported by any mapping system. This means that a mapping designer should be able to obtain the desired executable code through the visual interface of the mapping system without having to modify the executable code. The set consists of 11 basic mapping scenarios derived by a careful analysis of constructs commonly needed across different information integration applications, such as data exchange, data warehouses, XML publishing, schema evolution, real-world mapping specifications as well as the scientific literature in these topics and the experience of the authors in these areas. Some of the basic mapping scenarios include copying (as depicted in Figure 1(a)), horizontal and vertical partitioning, object fusion, identifier generation, normalization and denormalization (see Figure 1(b)), flattening and nesting, and constant value assignment. While we believe that these scenarios capture the majority of data transformations that occur in practice and have wide industry relevance, they are not intended to represent all possible mapping scenarios that may occur in practice.

For each mapping scenario, STBenchmark also provides a source instance which consists of data extracted from real-world instances. Hence, the executable code generated by the mapping systems can be evaluated using these "real" instances.
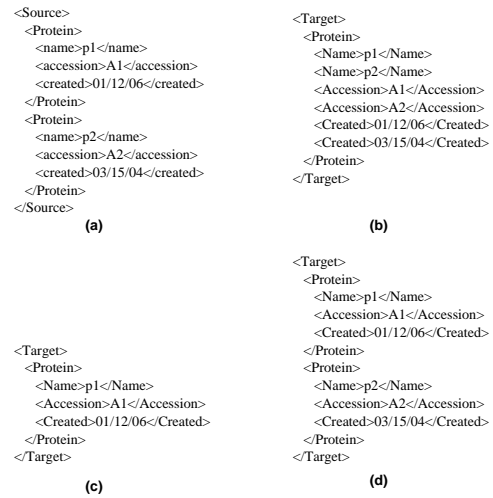
(ii) Generating Complex Mapping Scenarios and Instances. STBenchmark also provides a mapping scenario generator (SGen) and an instance generator (IGen) that can be used to test a mapping system with mapping scenarios (along with source instances) of varying complexity and size.

SGen takes as input a set of configuration parameters and returns as output a mapping scenario. SGen is capable of generating complex mapping scenarios of arbitrary sizes by combining and intermixing basic mapping scenarios in different ways based on the configuration parameters. Hence, the mapping scenarios generated by SGen can simulate real world mappings where many different transformations usually occur simultaneously from one schema to another. For example, the configuration parameters can be used to specify schema properties such as the level of nesting in the schemas, the number of subelements of each schema element, the number of elements involved in a join between two elements, the length of the join paths formed in the schemas, the kind of joins (star or chain) etc. By sampling values from Gaussian distributions based on the configuration parameters, SGen is able to generate schemas that look natural.

IGen, the instance generator, is built on top of ToXGene [1] and takes as input a schema and a set of configuration parameters and returns as output an instance that conforms to the given schema. The configuration parameters are used to specify the characteristics of the instance to be generated, such as the number of complex elements in the instance, the maximum length of generated string values and the allowed ranges for numeric values.

SGen and IGen have been implemented in such a way that they are able to reproduce the same output when given the same input. Hence, SGen and IGen produce repeatable results independently of factors such as time, hardware platforms, or operating systems. Towards this goal, SGen and IGen are implemented in Java and make use of the fact that the random number generators produce identical streams of pseudo-random numbers when given the same input seed.

It is also worth mentioning that by using SGen and ignoring T and $\mathcal{P}$ in the output of SGen, SGen is in effect a schema generator.

(iii) A Usability Model for Evaluating Mapping System Interfaces. Since all mapping systems that we have encountered provide a graphical interface to help reduce the effort required by a
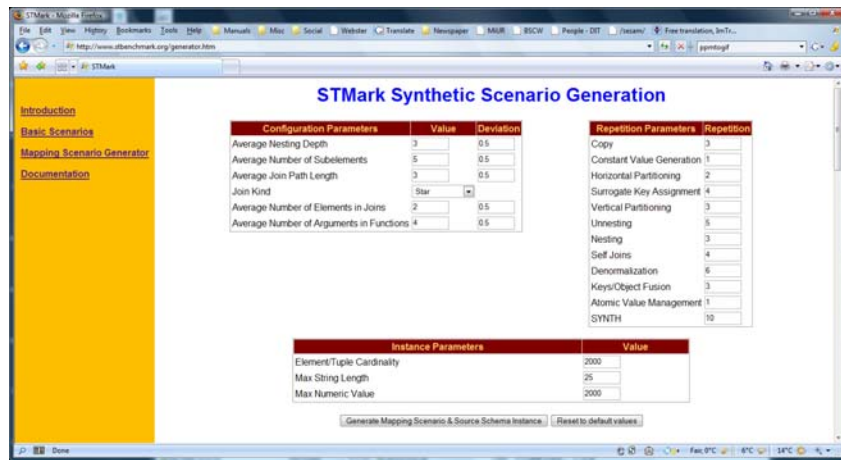
Figure 3: STBenchmark's graphical user interface for the configuration parameters of SGen and IGen.

mapping designer to generate mappings, STBenchmark also comes with a simple usability (SU) model that is able to provide a first-cut measure on the amount of effort required of a mapping designer to generate mappings through a graphical interface. In this model, effort is simply quantified as the number of mouse actions and the number of keystrokes used for text input. Three different types of mouse actions are captured by the model, namely, dragging actions, single and double mouse clicks. A cost model that quantifies dragging actions and text inputs with higher costs than single or double mouse clicks is also provided. While the SU model can provide a quick measure on the effort required to implement the same mapping scenario across different mapping systems, we emphasize that it is not meant to replace a much needed comprehensive human-computer interaction study (which is not the subject of this paper) on the usability of mapping systems.

It becomes clear from the above three components that STBenchmark is not intended to benchmark the schema matching process which is concerned with obtaining a set of mapping elements, where each mapping element indicates how elements of one schema relate to elements of the other schema. Some mapping systems are equipped with a matching module that (semi-)automatically derives these mapping elements, while in many others, the mapping elements are manually specified through the visual interface. Thus, considering the matching process in the benchmark may not have been fair for the latter category of systems. Furthermore, a number of proposals already exist for schema matching benchmarks [20, 19]. These could be easily integrated in STBenchmark if needed.

3. DEMONSTRATION OVERVIEW

Our demonstration will illustrate how STBenchmark is used to evaluate mapping systems. In the course of our demonstration, we will also describe some interesting findings about the mapping systems that we have evaluated based on STBenchmark.

First, we will show that visual specifications cannot be used to specify benchmark application scenarios (or test cases) to mapping systems. We show this by demonstrating that different mapping systems interpret the same visual specification differently. For example, we will implement scenarios similar to the one presented in Figure 1(a) on four different mapping systems and show that, as in the case presented in Section 2, mapping systems may generate nonequivalent XSLT transformation scripts. To make our demon-stration complete, we will illustrate that these XSLT scripts are nonequivalent by executing them against a simple source instance and show that we obtain different outputs under the same visual specification.

Second, we will showcase our 11 basic mapping scenarios and demonstrate how they can be implemented in different mapping systems. In the course of implementing a mapping scenario with different mapping systems, we will also demonstrate the differences in the level of support from different mapping systems in specifying mappings. For example, we will demonstrate that mapping scenarios, such as the one in Figure 1(b), can be easily implemented through the graphical interface of some mapping systems. We will also demonstrate how the same mapping scenario may be difficult or impossible to implement through the visual interface of some other mapping systems. In addition, we demonstrate that the degree of difficulty in implementing these mapping scenarios can be captured through our simple usability model.

Third, we will demonstrate how SGen can be used to generate complex mapping scenarios by (1) combining extended versions of basic mapping scenarios and (2) intermixing extended versions of basic mapping scenarios. Type (1) mapping scenarios model real-world mappings where different types of transformations occur in parallel in different parts of a schema while type (2) mapping scenarios model real-world mappings where different types of transformations occur simultaneously in the same part of a schema. We will also demonstrate that SGen can combine scenarios of types (1) and (2) and generate mapping scenarios of arbitrary sizes. In addition, we will also demonstrate how our instance generator IGen can be used for creating instances that conform to a given schema. We will also show how SGen and IGen can be used together to generate mapping scenarios and instances that conform to the source schemas of the mapping scenarios, as well as how these instances can be used for the data transformations specified in the mapping scenarios. Finally, we will also demonstrate how SGen and IGen can be used as general purpose schema and instance generators. Fig. 3 shows our graphical user interface for SGen and IGen.

Fourth and last, our demonstration will also illustrate some of the interesting observations we gathered while evaluating four mapping systems with STBenchmark. For example, we will demonstrate some peculiarities in the XSLT scripts generated by some mapping systems and show, through the use of SGen and IGen, how they affect the performance of the generated XSLT scripts in

general. We will also show how some mapping systems drastically reduce the amount of effort required to implement a mapping scenario in general through a graceful use of a schema matching module that goes hand-in-hand with the semantics of visual metaphors used in the graphical interface. Finally, we will demonstrate how the automatic handling of source and target schema constraints can help reduce the amount of effort required to implement a mapping scenario.

4. CONCLUDING REMARKS

Benchmarks are typically evaluating three main aspects of systems: expressiveness, correctness, and performance. Depending on the nature of the systems that are to be evaluated, a benchmark may focus more on one of these aspects as opposed to the others. For instance, query engine benchmarks, such as TPC-H [16], evaluate the expressiveness and correctness of the engines by providing a series of queries that need to be answered correctly, but they emphasize mostly on how well the engines scale. Towards this goal, they provide instance generators in order to test the engines with instances of different sizes. On the other hand, information integration system benchmarks, like THALIA [9], focus more on the expressiveness and correctness aspect. They provide a series of test scenarios along with data instances, all drawn from real life applications. In this work, we are proposing to demonstrate STBenchmark, a benchmark that we have developed for comparing and evaluating mapping systems. To the best of our knowledge, this is the first benchmark developed for such systems. One of the advantages of STBenchmark is that it equally focuses on the evaluation of all the different aspects of a mapping system. Through the set of basic mapping scenarios, it evaluates its expressive power. Through the real life source instances accompanying every basic mapping scenario, it evaluates its correctness. Through the use of the SGen and IGen modules it evaluates how well the systems scale not only in terms of instance size but also in terms of schema size and of transformation complexity. Finally, since a mapping system is a graphical tool, STBenchmark, through its usability model, evaluates the effort required by the mapping designer in the mapping generation process.

In addition to what we plan to demonstrate as described earlier, we also hope to share some of our experience with mapping systems through this demonstration: For example, our experience indicates that mapping systems share a lot of commonality in terms of the visual metaphors and constructs used for designing mappings. However (and unfortunately), there is currently a lack of standard in interpreting these visual metaphors and constructs. Hence, even though mapping systems take visual specifications as input, the transformation functions in our basic mapping scenarios are not specified as visual specifications and we leave to the benchmark user the task of creating the appropriate visual specification of each mapping scenario on each mapping system. From our study, we believe that it is crucial to develop a standard (either by standardizing the interpretation of visual metaphors and constructs and extending them to achieve more expressiveness, e.g. [13], or by developing a standard mapping specification language) for specifying inputs to mapping systems. Such a standard will not only serve to standardize the specifications of basic mapping scenarios in STBenchmark and the output of SGen, but also serve as an important step towards the development of a uniform testbed and repository for schema mappings and data exchange tasks [2].

References

[1] D. Barbosa, A. O. Mendelzon, J. Keenleyside, and K. A. Lyons. Toxgene: An extensible template-based data generator for xml. In *WebDB*, pages 49–54, 2002.

[2] Bertinoro Workshop on Information Integration. http://www.dis.uniroma1.it/~lenzerin/INFINT2007/.

[3] A. Bonifati, E. Q. Chang, T. Ho, and L. V. S. Lakshmanan. HepToX: Heterogeneous Peer to Peer XML Databases, 2005.

[4] A. Bonifati, G. Mecca, A. Pappalardo, S. Raunich, and G. Summa. Schema mapping verification: the spicy way. In *EDBT*, pages 85–96, 2008.

[5] V. R. Borkar, M. J. Carey, D. Engovatov, D. Lychagin, T. Westmann, and W. Wong. XQSE: An XQuery Scripting Extension for the AquaLogic Data Services Platform. In *ICDE*, pages 1229–1238, 2008.

[6] D. Braga, A. Campi, and S. Ceri. *QBE* (*q*uery y xample): A visual interface to the standard xml query language. *ACM TODS*, 30(2):398–443, 2005.

[7] M. Erwig. Xing: a visual xml query language. *J. Vis. Lang. Comput.*, 14(1):5–45, 2003.

[8] L. M. Haas, M. A. Hernández, H. Ho, L. Popa, and M. Roth. Clio Grows Up: From Research Prototype to Industrial Tool. In *SIGMOD*, pages 805–810, 2005.

[9] J. Hammer, M. Stonebraker, and O. Topsakal. THALIA: Test Harness for the Assessment of Legacy Information Integration Approaches. In *ICDE*, pages 485–486, 2005.

[10] IBM Rational Data Architect. http://www.ibm.com/software/data/integration/rda.

[11] Altova MapForce, Version 2007 rel.3 sp1. http://www.altova.com.

[12] L. Popa, Y. Velegrakis, R. J. Miller, M. A. Hernández, and R. Fagin. Translating Web Data. In *VLDB*, pages 598–609, 2002.

[13] A. Raffio, D. Braga, S. Ceri, P. Papotti, and M. A. Hernández. Clip: a Visual Language for Explicit Schema Mappings. In *ICDE*, 2008.

[14] A. R. Schmidt, F. Waas, M. L. Kersten, M. J. Carey, I. Manolescu, and R. Busse. XMark: A Benchmark for XML Data Management. In *VLDB*, pages 974–985, 2002.

[15] Stylus Studio, XML Enterprise Suite, Release 2. http://www.stylusstudio.com.

[16] TPC Transaction Processing Performance Council. http://tpc.org.

[17] P. Vassiliadis, A. Karagiannis, V. Tziovara, and A. Simitsis. Towards a Benchmark for ETL Workflows. In *QDB*, pages 49–60, 2007.

[18] Microsoft Visual Studio 2005, Version 8.0.50727.42. http://msdn2.microsoft.com/en-us/ie/bb188238.aspx.

[19] B. Yao, T. Ozsu, and N. Khandelwal. XBench benchmark and performance testing of XML DBMSs. In *ICDE*, pages 621–633, 2004.

[20] Y.Lee, M. Sayyadian, A. Doan, and A. Rosenthal. eTuner: Tuning Schema Matching Software using Synthetic Scenarios. *VLDB Journal*, 16(1):97–122, 2007.

[21] M. Zloof. Query-By-Example: A Data Base Language. *IBM Sys. Journal*, 16(4):324–343, 1977.