When independent parties develop database schemas for the same domain, they will almost always be quite different from each other. These differences are referred to as *semantic heterogeneity*, which also appears in the presence of multiple XML documents, Web services, and ontologies—or more broadly, whenever there is more than one way to structure a body of data. The presence of semi-structured data exacerbates semantic heterogeneity, because semi-structured schemas are much more flexible to start with. For multiple data systems to cooperate with each other, they must understand each other's schemas. Without such understanding, the multitude of data sources amounts to a digital version of the Tower of Babel.

This article begins by reviewing several common scenarios in which resolving semantic heterogeneity is crucial for building data-sharing applications. We then explain why resolving semantic heterogeneity is difficult and review some recent research and commercial progress in addressing the problem. Finally, we point out the key problems and opportunities in this area.

SCENARIOS OF SEMANTIC HETEROGENEITY

Enterprise information integration.

Enterprises today are increasingly facing data management challenges that

New tools and techniques can help ease the pain of reconciling schemas.

ALON HALEVY, UNIVERSITY OF WASHINGTON

Why Your Data Won't Mix

Why Your Data Won't Mix

involve accessing and analyzing data residing in multiple sources, such as database systems, legacy systems, ERP systems, and XML files and feeds. For example, for an enterprise to obtain a "single view of customer," it must tap into multiple databases. Similarly, to present a unified external view of its data, either to cooperate with a third party or to create an external facing Web site, the enterprise must access multiple sources. As the electronic marketplace becomes more prevalent, these challenges are becoming bottlenecks in many organizations.

There are many reasons why data in enterprises resides in multiple sources in what appears to be a haphazard fashion. First, many data systems were developed independently for targeted business needs, but when the business needs changed, data needed to be shared between different parts of the organization. Second, enterprises acquire many data sources as a result of mergers and acquisitions.

Over the years, there have been multiple approaches to addressing EII (enterprise information integration) challenges. Until the late 1990s, the two leading approaches were data warehousing and building custom solutions. Data warehousing solutions had the disadvantage of accessing stale data in many cases and not being able to work across enterprise boundaries. Custom code solutions are expensive, hard to maintain, and typically not extensible.

In the late '90s, several companies offered solutions that queried multiple data sources in realtime. In fact, the term *EII* typically refers to these solutions. While the users of these systems still see a single schema (whether relational or XML), queries are translated on the fly to appropriate queries over the individual data sources, and results are combined appropriately from partial results obtained from the sources. Consequently, answers returned to the user are always based on fresh data. Interestingly, several of these companies built their products on XML platforms, because the flexibility of XML (and more generally of semi-structured data) made it more appropriate for data integration applications. A recent article surveys some

of the challenges faced by this industry.¹ More recent research has proposed peer-to-peer architectures for sharing data with rich structure and semantics.²

In any of these data-sharing architectures, reconciling semantic heterogeneity is key. No matter whether the query is issued on the fly, data is loaded into a warehouse, or data is shared through Web services or in a peer-to-peer fashion, the semantic differences between data sources need to be reconciled. Typically, these differences are reconciled by *semantic mappings*. These are expressions that specify how to translate data from one data source into another in a way that preserves the semantics of the data, or alternatively, reformulate a query posed on one source into a query on another source. Semantic mappings can be specified in a variety of mechanisms, including SQL queries, XQuery expressions, XSLT scripts, or Java code.

In practice, the key issue is the amount of effort it takes to specify a semantic mapping. In a typical data integration scenario, more than half of the effort (and sometimes up to 80 percent) is spent on creating the mappings, and the process is labor-intensive and error-prone. Today, most EII products come with some tools for specifying these mappings, but the tools are completely manual—an expert needs to specify the exact mapping between the two schemas.

Querying and indexing the deep Web. The *deep Web* refers to Web content that resides in databases and is accessible behind forms. Search engines typically do not index deep Web content because the crawlers that these engines employ cannot go past the forms. In a sense, the form can be seen as a (typically small) schema, and unless the crawler can understand the meaning of the fields in the form, it gets stuck there.

The amount and value of content on the deep Web are spectacular. By some estimates, there are one to two orders of magnitude more content on the deep Web than the surface Web. Examples of such content range from classified ads in thousands of newspapers around the world, to data in government databases, product databases, university repositories, and more.

Here, too, the challenge stems from the very wide variety in the way Web site designers model aspects of a given domain. Therefore, it is impossible for designers of Web crawlers to assume certain standard form field names and structures as they crawl. Even in a simple domain such as searching for used cars, the heterogeneity in forms is amazing. Of course, the main challenge comes from the scale of the problem. For example, the Web site at www.everyclassified.com, the first site to aggregate content from thousands of form-based sources, includes more

than 5,000 semantic mappings of Web forms in the common categories of classified ads. Later in the article, we will describe the ideas that made this Web site possible. It is important to emphasize that accessing the deep Web is even more of a challenge for the content providers than it is for the search engines. The content providers thrive on getting users' attention. In the early days of the WWW, any good database would be immediately known (e.g., IMDb for movies). The number of such databases today, however, is vast (estimated in the hundreds of thousands), and people do not know about them. Instead, people's searches start from the search box of their favorite engine, and these engines do a poor job of indexing deep Web content. Hence, if I create an excellent database of Middle Eastern recipes and put it on the Web behind a form, it may remain invisible. Ironically, I'm better off creating a set of Web pages with my recipe contents than creating an easily searchable database. Finally, it should be noted that enterprise search faces a somewhat similar problem: many of the interesting data sources within an enterprise are in databases, and even providing simple keyword search over this content is challenging.

Merchant catalog mapping. An example of semantic heterogeneity occurs in aggregating product catalogs. Consider an online retailer such as Amazon.com. Such a retailer accepts feeds of products from thousands of merchants, each trying to sell its goods online. To aggregate the vast number of feeds, online retailers prescribe a *schema*: a hierarchy of products and their associated properties. To sell their products online, the merchants need to send a feed that adheres to the prescribed schema. On the back end, however, the data at the merchant is stored in the local schema, which is likely quite different from the one prescribed by the retailer (and typically covers a small fragment of that schema). Hence, the problem is creating mappings between thousands of merchants and a growing number of recognized online retailers (roughly 10 of them in the United States at this time).

An interesting point to note about this scenario is that there is not necessarily a *single* correct semantic mapping from the merchant's schema to that of the retailer. Instead, because there are subtle differences between product categories, and products can often be mapped to several categories, there are multiple mappings that may make sense—and the *best* one is the one that ultimately sells more products.

Schema versus data heterogeneity. Heterogeneity occurs not only in the schema, but also in the actual data values themselves. For example, there may be multiple ways of referring to the same product. Hence, even

though you are told that a particular field in a merchant's data maps to `ProductName`, that may not be enough to resolve multiple references to a single product. As other common examples, there are often multiple ways of referring to companies (e.g., IBM vs. International Business Machines), people's names (which are often incomplete), and addresses. To fully integrate data from multiple sources, one needs to handle both semantic-level and data-level heterogeneity. Typically, different products have addressed these two parts of the problem in isolation. As one example, several of the products for "global spend analysis" have focused on data-level heterogeneity. This article focuses mostly on schema heterogeneity.

Schema heterogeneity and semi-structured data. The problem of semantic heterogeneity is exacerbated when we deal with semi-structured data, for several reasons. First, the applications involving semi-structured data are typically ones that involve sharing data among multiple parties; hence semantic heterogeneity is part of the problem from the start. Second, schemas for semi-structured data are much more flexible, so we are more likely to see variations in the schema. Finally, the main advantage of semi-structured data is that attributes can be added to the data at will (or even simply derived from inspecting the data itself), and once that flexibility is in place, the number of additional attributes we see is significant, and understanding their exact meaning becomes crucial. On the flip side, in many applications involving semi-structured data, it is enough to reconcile only a specific set of attributes, while we can still manipulate and display any other attribute. Specifically, we need only to reconcile those attributes that are going to be used for equating data across multiple sources.

WHY IS IT SO HARD?

The problem of reconciling schema heterogeneity has been a subject of research for decades, but solutions are few. The fundamental reason that makes semantic heterogeneity so hard is that the data sets were developed independently, and therefore varying structures were used to represent the same or overlapping concepts. In many cases, we are trying to integrate data systems that were developed for slightly (or vastly) different business needs. Hence, even if they model overlapping domains, they will model them in different ways. Differing structures are a byproduct of human nature—people think differently from one another even when faced with the same modeling goal. As a simple illustration, one of the assignments I give in my senior-level database course is to design an inventory schema based on a one-page English descrip-

Why Your Data Won't Mix

tion of what it should cover. Invariably, the schemas I get from my students are vastly different.³

From a practical perspective, one of the reasons that schema heterogeneity is difficult and time-consuming is that it requires both domain and technical expertise: you need people who understand the business meaning of each of the schemas being reconciled and people skilled in writing transformations (e.g., SQL or XQuery experts).

While schema heterogeneity is challenging for humans, it is drastically more challenging for programs. A program is given only the two schemas to reconcile, but those schemas are merely symbols. They do not capture the entire meaning or intent of the schemas—those are only in the minds of the designers.

Figure 1 illustrates some of the challenges in resolving semantic heterogeneity. The figure shows a typical manual schema-matching tool in which the designer needs to draw lines between the matching attributes of the two schemas. As can be seen in the example, there are several kinds of semantic discrepancies between schemas:

1. The same schema element in the two schemas is given different names (e.g., IssueDate and OrderIssueDate);
2. Attributes in the schemas are grouped into table structures (or XML nesting) in different ways (e.g., consider the subtrees of the BuyerParty element in the two schemas); and
3. One schema may cover aspects of the domain that are not covered by the other (e.g., the left schema does not have anything like OrderSummary in the right schema).

When reconciling heterogeneity from thousands

of Web forms, there are additional sources of heterogeneity. Some forms are already specialized for a particular domain (used cars, jobs), whereas in others the user needs to select a category before entering additional attributes. In some cases the location is already implicit in the form (e.g., using some hidden fields), while in others the user needs to select a city and state or zip.

Some argue that the way to resolve semantic heterogeneity is through standard schemas. Experience has shown, however, that standards have limited success and only in domains where the incentives to agree on standards are very strong. Even then, as the online retailer example illustrated, while data providers may share their data using a standard, their own data systems still employ their original schemas (and the cost of changing those systems is prohibitive). Hence, semantic heterogeneity needs to be resolved at the step where the data provider exposes its data to its counterparts.

As one thinks about resolving schema heterogeneity, it is important to note several common instances of the problem, which may shed light on the specific problem at hand:

- One schema may be a new version of the other.
- The two schemas may be evolutions of the same original schema.
- We may know that we have many sources modeling the same aspects of the underlying domain (horizontal integration).
- We may have a set of sources that cover different domains but overlap at the seams (vertical integration).

**Reconciling
Two
Disparate
Schemas**

THE STATE OF THE ART

Resolving schema heterogeneity is inherently a heuristic, human-assisted process. Unless there are very strong constraints on how the two schemas you are reconciling are different from each other, one should not hope for a completely automated solution. The goal is to reduce the time it takes human experts to create a mapping between a pair of schemas, and enable them to focus on the hardest and most ambiguous parts of the mapping. For example, the tools that enabled building the Web site at www.everyclassified.com required that we be able to map the fields of the Web form to our own schema in one minute, on average.

As would be expected, people have tried building semi-automated schema-matching systems by employing a variety of heuristics.⁴ The process of reconciling semantic heterogeneity typically involves two steps. In the first, called schema matching, we find correspondences between pairs (or larger sets) of elements of the two schemas that refer to the same concepts or objects in the real world. In the second step, we build on these correspondences to create the actual schema mapping expressions. The Clio Project at IBM Almaden is a prime example of work on building the mapping expressions.⁵

The following classes of heuristics have been used for schema matching:

- **Schema element names.** Element names (e.g., table and attribute names) carry some information about their intended semantics. Hence, by looking at the names, possibly stemming the words first, we can obtain clues for the schema matcher. The challenges involved in using names are that the use of synonyms is very common, as is the use of hypernyms (words that are specializations or generalizations). Furthermore, we often see that same word being used with different meanings (homonyms). We also see abbreviations and concatenations of words appearing in element names.
- **Data types.** Schema elements that map to each other are likely to have compatible data types, but this is certainly not a rule. In many schemas, however, the data types are underspecified (e.g., CDATA for XML). In practice, considering data types is a useful heuristic for ruling out certain match candidates.
- **Data instances.** Elements from two schemas that match each other often have similar data values. Similarities can arise in several ways: 1. values drawn from the same small domain (e.g., makes of cars or names of countries); 2. significant occurrences of the same values (e.g., superlatives describing houses for sale); or 3. patterns of values (e.g., phone numbers, price ranges). Data

instances are extremely useful when available, but one cannot rely on their availability.

- **Schema structure.** Matching elements in a schema are typically related to other related schema elements. For example, in an object-oriented hierarchy, if two classes match each other, then the children of these classes often will also (at least partially) match. In XML, proximity of attributes in a DTD is significant (e.g., a `Phone` field next to `Agent` may suggest that it is the phone of the agent). Relying on such a heuristic can be brittle, however, and one of the main challenges is to find an initial match that drives the similarity of its neighbors.
- **Integrity constraints.** Considering integrity constraints on single attributes or across attributes can be useful for generating matches. For example, if two attributes are known to be keys in their respective schemas, then that provides additional evidence of their similarity. While each of these heuristics is useful, experience has shown that taking any of them in isolation leads to a brittle schema-matching solution. Hence, research has focused on building systems that combine multiple heuristics.^{6,7,8} Despite these ideas, commercial products rely on completely manual specification of semantic mappings. They help by offering visual interfaces that enable designers to draw the lines between elements of disparate schemas, while the nitty-gritty details of the mapping can often be generated on the back end. These tools already save a significant amount of time, but they do not suggest mappings to the designer.

AN EMERGING SOLUTION: LEVERAGING PAST EXPERIENCE

One of the fundamental reasons that the schema-matching solutions described here are brittle is that they exploit *only* evidence that is present in the two schemas being matched, ignoring past experience. These schemas often lack sufficient evidence to discover matches. Looking more closely at schema-matching tasks, however, it is evident that these tasks are often repetitive. Specifically, we often find that we repeatedly map schemas in the same domain into a common mediated schema. For example, creating the engine at www.everyclassified.com involved mapping thousands of Web forms in the same domain into a common schema, the one exposed to the users by the engine itself. Human experts, after seeing many schemas in a particular domain, are able to map schemas much faster because they have seen many variations on how concepts in the domain are represented in schemas.

The challenge, therefore, is to endow the schema matcher with the same capabilities: leverage past expe-

Why Your Data Won't Mix

rience. For example, once the system has been given several mappings in the domain of used cars, it should be able to predict mappings for schemas it has not seen before. As it sees more schemas in a particular domain, its predictions should become more accurate, and it should be more robust in the presence of variations.

This idea was explored for the past few years in several academic research settings^{9,10,11,12,13} and has recently been applied commercially for the first time by Transformic Inc., the creators of www.everyclassified.com. The research projects considered the use of machine learning as a mechanism for enabling a schema matcher to leverage previous experience. In machine learning, the system is provided a set of *training examples* and uses them to learn models of the domain of interest. In this context, the training examples are schema mappings that are manually constructed by domain experts and given to the system. The models of the domain enable the system to look at a new schema and predict a schema map. For example, the system can learn that the attribute concerning house descriptions typically involves a long text and includes frequent occurrences of superlatives. Furthermore, the system can learn variations on the ways people name this field in practice.

Another application of this idea is searching for Web services, namely locating Web services (or operations within them) that are relevant to a particular need. Simple keyword search does not suffice because keywords (or parameter names) do not capture the underlying semantics of the Web service. The Woogles Search Engine¹⁴ (available at www.cs.washington.edu/woogle) is based on analyzing a large collection of Web services and clustering parameter names into semantically meaningful concepts. These concepts are used to predict when two Web service operations have similar functionality.

What can you learn from the past? The paradigm of learning from past experience of performing schema-matching tasks is only in its infancy. It is interesting to take a step back and consider what one can learn from the past in this context.

We assume the *past* is given to us as a collection of schemas in a particular domain, mappings between pairs of schemas in that collection, and to the extent possible, data instances. The schemas can come from anywhere and can involve closely related domains, not necessarily modeling the same data. In many cases, such schemas can be obtained from the Web or resources such as xml.org. In others, they may be available throughout an enterprise. Such a collection of schemas is often referred to as a *corpus*, in analogy to the use of *corpora* of documents underlying information retrieval (IR) and Web-search engines. Of course, while the corpora in IR involve collections of words, here we are managing semantically richer elements, such as schemas and their instances.

The goal of analyzing a corpus of schemas and mappings is to provide hints about deeper domain concepts and at a finer granularity. Looking a bit closer at the approach, the following are examples of what we can learn from a corpus.

Domain concepts and their representational variations. As a first step, we can analyze a corpus to identify the main concepts in the domain. For example, in a corpus of book inventory schemas, we may identify the concepts of book and warehouse and a cluster of price-related elements. Even more importantly, we will discover *variations* on how these concepts are represented. The variations may differ on naming of schema elements, grouping attributes into tables, or the granularity of modeling a particular concept. Knowledge of these variations will be leveraged when we match two schemas in the domain.

Relationships between concepts. Given a set of concepts, we can discover relationships between them, and the ways in which these relationships are manifested in the representation. For example, we can find that the Books table typically includes an ISBN column and a foreign key into an Availability table, but that ISBN never appears in a Warehouse table. These relationships are useful to prune candidate schema matches that appear less likely. They can also be used to build a system that provides advice in *designing* new schemas.

Domain constraints. We can leverage a corpus to find integrity constraints on the domain and its representations. For example, we can observe that ISBN is a foreign key into multiple tables involving books, and hence possibly an identifier for books, or discover likely data types for certain fields (e.g., address, price). Constraints may have to do with *ordering* of attributes. For example, in a corpus of Web forms about cars for sale, we may discover that the make attribute is always placed before the model

and price attributes, but occurs after the new/used attribute. Typically, constraints we discover in this way are *soft constraints*, in the sense that they are sometimes violated, but can still be taken as rules of thumb about the domain. Therefore, they are extremely useful in resolving ambiguous situations, such as selecting among several candidate schema matches.

LOOKING FORWARD

The need for flexible data-sharing systems, within and across enterprises, is only in its infancy. The tools we have today lag far behind customer needs. The problem is only exacerbated by the fact that much more of the data we need to manage is semi-structured and is often the result of trying to *extract* structure from unstructured data. Hence, we need to manage data where the values, attributes names, and semantics are often uncertain.

Going forward, there are two major challenge areas: dealing with drastically larger schemas and dealing with vastly more complex data-sharing environments. In both of these areas, we may have to change the way we think.

Larger schemas and schema search. The techniques described in this article deal with small to medium-size schemas (including up to hundreds of elements). To their credit, these techniques gracefully handle large numbers of such schemas. It is well known that many real-world schemas have thousands of schema elements (tables and attributes), the SAP schemas being a prime example. The challenge of creating schema mappings is considerably harder here: you cannot even view the entire schema on a screen or multiple screens.

Two principles need to guide the work on mapping larger schemas. The first is that the schema-matching tools need to incorporate advanced information visualization methods.¹⁵ Much of the challenge in designing schema mappings for large-scale schemas is ensuring that the attention of the designers is constantly directed to the right place, that they can view hypothetical mappings and remove them easily, and that they can see effectively how one fragment of the mapping may affect other fragments. The system should also be able to explain why certain match predictions are being made.

The second principle requires changing the way we think of schema matching. Specifically, I am proposing a *schema search engine*. The engine contains a set of indexes on the elements of a particular schema (or set of schemas). The engine takes as input schema elements (e.g., table name, attribute, XML tag), schema fragments, or combinations of schema fragments and data instances. The engine returns a ranked list of schema elements in

the indexed schema that are candidate matches. The interface should be as simple as we see today in search engines. The justification for such a tool is that much of the work in schema mapping is simply finding where in a huge schema there are relevant fragments to the part of the schema that is currently under consideration, or alternatively, finding a relevant schema in a large collection of schemas. Hence, instead of focusing on tools that solve the entire problem but are inherently *brittle*, build *robust* tools that bring the users closer to their needs. The Woogole Web-service search engine previously described is an example of such an engine, searching over Web service operations rather than schemas and their fragments.

Managing dataspace. A much greater challenge facing the data management community is to raise the abstraction level at which data is managed. Today, we have powerful systems for managing data at the level of a single database system (whether relational, XML, or in some other model). The data management challenges we face, however, are at a much higher level: we need to manage a *dataspace*, rather than a database.

A dataspace consists of a set of *participants* and a set of *relationships*. Participants are individual data sources: relational databases, XML repositories, text databases, Web services, data-stream systems, sensor deployments, or any other element that stores or delivers data. Some participants may be *transparent*, with a full language for posing queries; a prime example is a traditional relational DBMS. Other participants may be *opaque*—offering limited interfaces for posing queries (usually supported by specific programs); examples are Web services, stored procedures, and other software packages. In addition, some participants may have no structure to their data (e.g., text) or only some structure (e.g., code collections). Examples of dataspace include: an enterprise, the desktop, a library, large scientific projects, a smart home, or a battlefield.

A dataspace should be able to model any kind of relationship between two (or more) participants. In the extreme case, a relationship is a full schema mapping that enables arbitrary data exchange and query reformulation among participants. In other cases, the relationship can express simple dependencies, where the details are not known precisely (e.g., one participant is an evolved version of another). The relationships can have a temporal aspect (e.g., how frequently data is exchanged), or have one that is a mirror or backup of the other.

The key distinguishing feature of dataspace management is that integrations should evolve over time and as needed, but data should be accessible in some form from the very start. This means that simple queries

Why Your Data Won't Mix

(e.g., keyword queries) should always be supported on every participant in the dataspace without any effort. As the owners of the dataspace want to create more tight integration between sources and support more complex queries across participants, they can create more detailed semantic mappings as necessary. In addition, the management of dataspaces should consider the entire lifecycle of data, including its acquisition, curation, query and update, evolution, and analysis. The initial ideas on managing dataspace¹⁶ are only starting to intrigue the research community. Practitioners have so far embraced the idea with enthusiasm. ☺

REFERENCES

1. Halevy, A. Y., Ashish, N., Bitton, D., Carey, M., Draper, D., Pollock, J., Rosenthal, A., and Sikka, V., 2005. Enterprise information integration: successes, challenges and controversies. In *Proceedings of the ACM SIGMOD Conference*.
2. Aberer, K. 2003. Peer to peer data management: introduction to a special issue. *SIGMOD Record* 32(3).
3. Halevy, A. 2003. Learning about data integration challenges from day one. *SIGMOD Record* 32(3): 16-17.
4. Rahm, E., and Bernstein, P.A. 2001. A survey of approaches to automatic schema matching. *VLDB Journal* 10(4): 334-350.
5. Yan, L. L., Miller, R. J., Haas, L. M., and Fagin, R. 2001. Data driven understanding and refinement of schema mappings. In *Proceedings of the ACM SIGMOD*.
6. Do, H.-H., and Rahm, E. 2002. COMA—a system for flexible combination of schema-matching approaches. In *Proceedings of the International Conference on Very Large Databases (VLDB)*.
7. Doan, A., Domingos, P., and Halevy, A. 2001. Reconciling schemas of disparate data sources: a machine learning approach. In *Proceedings of the ACM SIGMOD*.
8. Madhavan, J., Bernstein, P., and Rahm, E., 2001. Generic schema matching with cupid. In *Proceedings of the International Conference on VLDB*.
9. See Reference 7.

10. Halevy, A., Etzioni, O., Doan, A., Ives, Z., Madhavan, J., McDowell, L., and Tatarinov, I. 2003. Crossing the structure chasm. In *Proceedings of the First Biennial Conference on Innovative Data Systems Research (CIDR)*.
11. He, B., and Chang, K. C.-C. 2003. Statistical schema integration across the deep Web. In *Proceedings of the ACM SIGMOD*.
12. Hess, A., and Kushmerick, N. 2003. Learning to attach semantic metadata to Web services. In *Proceedings of the International Semantic Web Conference*.
13. Madhavan, J., Bernstein, P., Doan, A., and Halevy, A. 2005. Corpus-based schema matching. In *Proceedings of the International Conference on Data Engineering (ICDE)*.
14. Dong, X. L., Halevy, A. Y., Madhavan, J., Nemes, E., and Zhang, J. 2004. Similarity search for Web services. In *Proceedings of the International Conference of VLDB*.
15. Robertson, G. G., Czerwinski, M. P., and Churchill, J. E. 2005. Visualization of mappings between schemas. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
16. Franklin, M., Halevy, A., and Widom, J. 2005. Data-spaces: a new abstraction for data management.

LOVE IT, HATE IT? LET US KNOW

feedback@acmqueue.com or www.acmqueue.com/forums

ACKNOWLEDGMENTS

The ideas espoused in this paper have benefited from discussions with and hard work by my colleagues and students. In particular, I'd like to thank Phil Bernstein, Anhui Doan, Luna Dong, Dana Florescu, Zack Ives, and Jayant Madhavan. The vision of dataspaces is the brainchild of discussions with Mike Franklin, Dave Maier, and Jennifer Widom.

ALON HALEVY is a professor of computer science at the University of Washington. He received his Ph.D. in computer science from Stanford University in 1993. His research interests are in data integration, semantic heterogeneity, personal information management, management of XML data, Web-site management, peer-data management systems, and the intersection between database and AI technologies. He was a co-developer of XML-QL, which later contributed to the development of the XQuery standard for querying XML data. In 1999, Halevy co-founded Nimble Technology, one of the first companies in the enterprise information integration space. In 2004, he founded Transformic Inc., a company that creates search engines for the deep Web.

© 2005 ACM 1542-7730/05/1000 \$5.00