

Multiplex: A Formal Model for Multidatabases and Its Implementation

Amihai Motro

Department of Information and Software Engineering
George Mason University
Fairfax, VA 22030-4444
ami@gmu.edu
<http://www.ise.gmu.edu/ami>

Abstract. The integration of information from multiple databases has been an enduring subject of research for over 20 years, and many different solutions have been attempted or proposed. Missing from this research has been a uniform framework. Usually, each solution develops its own ad-hoc framework, designed to address the particular aspects of the problem that are being attacked and the particular methodology that is being used. To address this situation, in this paper we define a formal model for multidatabases, which we call Multiplex. Multiplex is a simple extension of the relational model, which may serve as a uniform abstraction for many previous ad-hoc solutions. Multiplex is based on formal assumptions of integrability, which distinguish between scheme and instance reconcilability among independent databases. Multiplex supports database heterogeneity, and it provides several degrees of freedom that allow it to model actual situations encountered in multidatabase applications. In addition, in situations in which a single answer is not obtainable (either because the global query is not answerable, or there are multiple candidate answers), Multiplex defines approximative answers. Finally, Multiplex provides a practical platform for implementation. A prototype of such an implementation is described briefly.

1 Introduction

The integration of information from multiple databases has been an enduring subject of research for over 20 years. (Surveys of this area include [5, 7, 24, 31]; collections of articles on this topic include [17, 29, 16, 13]; recent workshops include [30, 12].) Indeed, while the solutions that have been advanced tended to reflect the research approaches prevailing at their time, the overall goal has remained mostly unchanged: to provide flexible and efficient access to information residing in a collection of distributed, heterogeneous and overlapping databases.¹

A standard approach to this problem has been to integrate the independent databases by means of a comprehensive *global scheme* that models the information contained in the entire collection of databases (for example, [20, 33, 6,

¹ More generally, the problem may involve other kinds of information sources as well.

26]). This global scheme is fitted with a *mapping* that defines the elements of the global scheme in terms of elements of the schemes of the member databases. Algorithms are designed to interpret queries on the global scheme. Such *global queries* are translated (using the information captured in the mapping) to queries on the member databases; the individual answers are then combined to an answer to the global query. The global scheme and the scheme mapping constitute a *virtual database*; the main difference between a virtual database and a conventional database is that whereas a conventional database contains data, a virtual database points to other databases that contain the data. An important concern is that this query processing method be *transparent*; i.e., users need not be aware that the database they are accessing is virtual.

Much of the work in this area has been on the construction of global schemes. The main issue here is the resolution of *intensional inconsistencies* (semantic heterogeneity) among the member schemes (for example, [4, 14, 19]). Yet the complementary problem of *extensional inconsistencies* has received much less attention. This problem arises when alternative sources with overlapping information provide mutually inconsistent information, and requires methods for resolving such inconsistencies in global answers [1].

Inconsistencies result in multiple candidate answers; the dual problem also exists, in which a global query might have no answer at all. Such situations often occur when a member database becomes temporarily unavailable. In such cases, rather than reject the query altogether, it is desirable to *approximate* the global answer using whatever information that is available [9, 8].

Often missing from multidatabase research is a *uniform* framework; each solution formulates its own version of the problem, states its own assumptions, develops its own ad-hoc model, and proposes its own solutions. Often missing as well is a *formal* treatment of the subject, with precise formulation of all assumptions, definitions and algorithms.

In this paper we address these and other issues in a model and a system for multidatabases, which we call Multiplex. Several important features of the Multiplex model are elaborated below.

(1) **Extension of the relational model.** Multiplex extends the definition of relational schemes, queries and answers to an environment of multiple databases, and it provides a common ground for other definitions as they become necessary. The extension also retains the attractive *simplicity* of the relational model, with relatively few new concepts.

(2) **Formal assumptions of integrability.** Most integration methods have operated under tacit assumptions regarding the mutual consistency of the schemes and the instances of the underlying databases. Multiplex defines two kinds of inconsistency (intensional and extensional) and formulates two assumptions of integrability: the Scheme Consistency Assumption and the Instance Consistency Assumption. These explicit assumptions, which have been absent from previous work, provide an unambiguous framework, and help to classify other methods.

(3) **Full support for heterogeneity.** The simplicity and popularity of the relational model makes it an ideal integration model, and the integrated view

that Multiplex provides is indeed relational. Yet, there is no restriction on the underlying data models; the only requirement is that they communicate their results in tabular form. Consequently, the member databases in a Multiplex multidatabase may be relational, object-oriented, or, in general, stored in any software system that can respond to requests with tabulated answers.

(4) **Flexibility to model real-world situations.** The Multiplex model is distinguished by several *degrees of freedom* that allow it to model actual situations encountered in multidatabase applications. Specifically,

1. **Source unavailability.** Multiplex reflects the dynamics of multidatabase environments where member databases may become temporarily unavailable, and global queries might therefore be unanswerable in their entirety.
2. **Source inconsistency.** Multiplex accepts that requested information may be found in more than one database, and admits the possibility of inconsistency among these multiple versions.
3. **Ad-hoc integration.** Multiplex permits ad-hoc global schemes of limited scope, that cull from existing databases only the information relevant to a given application.

Intuitively, these degrees of freedom correspond to mappings (from the global scheme to the member schemes) that are *neither total, nor single-valued, nor surjective*. We note that earlier approaches to global scheme integration were based on often unrealistic assumptions that existing database schemes could be integrated completely and perfectly in a single global scheme (i.e., mappings that are total, surjective, and single-valued; sometimes even one-to-one). The complexity of existing databases quite often renders this approach unrealistic. The abovementioned degrees of freedom are therefore important, as they represent a significant departure from earlier approaches.

(5) **Approximative answers.** Because Multiplex mappings are not total, global queries may be *unanswerable*; because the mappings are not single-valued and because there is no assumption of mutual consistency among the member databases, global queries may have *several candidate answers*. In both these situations, Multiplex defines *approximative answers*. The overall concern is that when a single authoritative answer is unavailable, a multidatabase system should approximate the request as best as possible. For example, the best approximative answer to a query on the names, salaries, departments and locations of all employees, could be the names and departments of all employees, the salaries of some employees, and the locations for none of the employees. As another example, the best approximative answer to a query on the employees who earn over 50, could be the employees who earn over 40. Note that the former approximative answer was “less” than what was requested, whereas the latter was “more” than what was requested, corresponding to “below” and “above” approximations.

(6) **Quick adaptation to evolving environments.** Present data environments may be highly dynamic; for example, newly discovered data sources may need to be incorporated, the structure of existing data sources may change, or existing data sources may need to be deleted altogether. In Multiplex such changes are easy to effect. As we shall see, the integration consists of providing pairs of

equivalent views: a view on the global database (“the information needed”), and a view on a member database (“how it is materialized”). The complexity of these views can vary greatly: they could range from a complex calculation, to a statement that simply denotes the equivalence of two attribute names.

(7) **A practical platform for implementation.** Finally and most importantly, the Multiplex model is a practical platform for implementation. Hence, Multiplex is not only a formal model, but a practical system as well.

With respect to limitations, we note that Multiplex queries and mappings are based on conjunctive views with aggregate functions. Although a language based on conjunctive views and aggregate functions does not have the full power of the relational algebra or calculus, it is a powerful language nonetheless.

The recent explosion of on-line information sources on the Internet has increased the interest in this area significantly, with several systems that are roughly in the same class as Multiplex. These include SIMS [3], TSIMMIS [15], UniSQL [18], and the Information Manifold [2]. These systems are discussed later in this paper.

Section 2 defines the database concepts that will be used later. Section 3 discusses integrability, and defines multidatabases and multidatabase queries and answers. Section 4 extends the model to approximative answers. Section 5 describes the Multiplex prototype, Section 6 compares Multiplex to several other systems that share similar goals, and Section 7 concludes this paper with a brief summary and a discussion of further research issues.

2 Relational Databases

In this section we define the database concepts that will be used throughout this work. Our formalization of relational databases is mostly conventional.

2.1 Schemes and Instances, Views and Queries

Assume a finite set of attributes T , and for each attribute $A \in T$ assume a finite domain $dom(A)$, and assume a special value called *null* and denoted $-$, which is not in any of the domains. A *relation scheme* R is a sequence of attributes from T . A *tuple* t of a relation scheme $R = (A_1, \dots, A_m)$ is an element of $dom(A_1) \cup \{-\} \times \dots \times dom(A_m) \cup \{-\}$. A *relation instance* (or, simply, a *relation*) r of a relation scheme R is a finite set of tuples of R . A *database scheme* D is a set of relation schemes $\{R_1, \dots, R_n\}$. A *database instance* d of the database scheme D is a set of relations $\{r_1, \dots, r_n\}$, where r_i is a relation on the relation scheme R_i ($i = 1, \dots, n$). Finally, a *database* (D, d) is a combination of a database scheme D and a database instance d of the scheme D .

Let D be a database scheme. A *view* of D is an expression that defines (1) a new relation scheme V , and (2) for each instance d of D an instance v of V . v is called the *extension* of the view scheme V in the database instance d .

We shall assume that all views are of the family known as *conjunctive* views [34]. Although conjunctive views are a strict subset of the relational tuple

calculus, they are a powerful subset, corresponding to the set of relational algebra expressions with the operations *Cartesian product*, *selection* and *projection* (where the selection predicates are conjunctive).

A *query* Q on a database scheme D is a view of D . The extension of Q in a database instance d of scheme D is called the *answer* to Q in the database instance d .

As an example, given the relation schemes $Emp = (Name, Salary, Dname)$ and $Dept = (Dname, Supervisor)$, a view Emp_sup may be defined by

$$project_{Name, Supervisor} select_{Emp.Dname=Dept.Dname} Emp \times Dept$$

The scheme of Emp_sup is $(Name, Supervisor)$.

2.2 View Operations

Given views V' and V of a database scheme D , V' is a *subview* of V , denoted $V' \sqsubseteq V$, if there exists a selection-projection view W of V , such that the schemes of V' and W are identical, and in every database instance d of D , the extensions of V' and W are identical. V is then a *superview* of V' .

Note the difference between this definition of subview and the common definition of *contained view*, denoted $V' \subseteq V$, which is based on the containment of two sets of tuples. By restricting W to selection only, the concept of subview is reduced to contained view.

Assume that V' is a subview of V . The *enlargement* of V' to V is a view whose scheme is V , and for every database instance d of D , its extension is obtained by adding null values to the extension of V' in d in the attributes that are in V but not in V' .

Assume that V_1 and V_2 are subviews of V . The *subview union* of V_1 and V_2 over V , denoted $V_1 \sqcup V_2$, is the union of their enlargements to V . Assume that V_1 and V_2 are superviews of V . The *superview intersection* of V_1 and V_2 over V , denoted $V_1 \sqcap V_2$, is the intersection of their projections on V .²

When the definition of V may be assumed from the context, we shall call these operations simply the *subview union* and *superview intersection* of V_1 and V_2 . Note that these operations generalize the union and intersection operations on views, which are commonly defined only for views that have the same scheme, because when V_1 , V_2 and V all have the same scheme, these new operations are reduced to the common view operations. As we shall see, the subview union and the superview intersection will be used to provide *lower and upper bounds* of the view V .

Finally, views V_1 and V_2 of a database scheme D are *overlapping*, if there exists a view V of D , such that $V \sqsubseteq V_1$ and $V \sqsubseteq V_2$, and there exists some instance d of D in which the extension of V is non-empty.

² Note that it may be possible to *identify* separate tuples by using additional information that may be available, such as functional dependencies [25].

3 Multidatabases

We begin by defining derivative databases; i.e., databases derived from other databases. This notion is necessary to define equivalence between views of different databases, which in turn provides a method for expressing the commonality of two database schemes using scheme mappings. Scheme mappings are the basis for our definition of multidatabases. We then formulate the Scheme Consistency Assumption and the Instance Consistency Assumption. These assumptions lead to fundamental observations concerning the integrability of independent databases. We complete the description of the Multiplex multidatabase model by defining multidatabase queries.

3.1 Derivative Databases, View Equivalence and Scheme Mapping

Consider a database (D, d) . Let D' be a database scheme whose relation schemes are views of the relation schemes of D . The database scheme D' is said to be *derived* from the database scheme D . Let d' be the database instance of D' which is the extension of the views D' in the database instance d . The database instance d' is said to be *derived* from the database instance d . Altogether, a database (D', d') is a *derivative* of a database (D, d) , if its scheme D' is derived from the scheme D , and its instance d' is derived from the instance d .

In this paper we are not concerned with an effective procedure for determining whether one database is a derivative of another, a question that depends on the language for expressing views. For our purpose here, it is sufficient to note that a database may or may not be a derivative of another database.

Let (D_1, d_1) and (D_2, d_2) be two derivatives of a database (D, d) . A view V_1 of D_1 and a view V_2 of D_2 are *equivalent*, if for every instance d of D the extension of V_1 in d_1 and the extension of V_2 in d_2 are identical. Intuitively, view equivalence allows us to substitute the answer to one query for an answer to another query, although these are different queries on different schemes.

Given two different database schemes, which are both derivatives of the same scheme (the “reference scheme”), we express their commonality by means of scheme mappings.

Assume two database schemes D_1 and D_2 , which are both derivatives of a database scheme D . A *scheme mapping* (D_1, D_2) is a collection of view pairs $(V_{i,1}, V_{i,2})$ ($i = 1, \dots, m$), where each $V_{i,1}$ is a view of D_1 , each $V_{i,2}$ is a view of D_2 , and $V_{i,1}$ is equivalent to $V_{i,2}$.

As an example, the equivalence of attribute *Salary* of relation scheme *Emp* in database scheme D_1 and attribute *Sal* of relation scheme *Employee* in database scheme D_2 is indicated by the view pair

$$(\text{project}_{\text{Salary}} \text{Emp}, \text{project}_{\text{Sal}} \text{Employee})$$

As another example, given the relation schemes $\text{Emp} = (\text{Name}, \text{Title}, \text{Salary}, \text{Supervisor})$ in database scheme D_1 , and $\text{Manager} = (\text{Ename}, \text{Level}, \text{Sal}, \text{Sup})$

in database scheme D_2 , the retrieval of the salaries of managers is performed differently in each database, as indicated by the view pair

$$(\text{project}_{Name,Salary} \text{select}_{Title=manager} Emp, \text{project}_{Ename,Sal} Manager)$$

From a practical point of view, note that the number of views in scheme mappings may be controlled, by mapping the “largest” views possible. The equivalence of their corresponding subviews may then be inferred.

3.2 Multidatabase

A *multidatabase* is

1. A scheme D .
2. A collection $(D_1, d_1), \dots, (D_n, d_n)$ of databases.
3. A collection $(D, D_1), \dots, (D, D_n)$ of scheme mappings.

The first item defines the scheme of a multidatabase, the second item defines the member databases in the multidatabase environment, and the third item defines a mapping from the multidatabase scheme to the schemes of the member databases.

The “instance” of a multidatabase consists of a collection of global view extensions that are available from the member databases. Specifically, the views in the first position of the scheme mappings specify the “contributed information” at the global level, and the views in the second position describe how these contributions are materialized.

As defined in Section 3.1, scheme mappings allow to substitute certain views in one database with equivalent views in another database. In a multidatabase, the former database is the global database, and the latter is a member database.

This definition may be considered a formalization of *virtual databases* defined in [26]. Scheme mapping may be considered an abstraction of different solutions that have been advanced to the task of relating global schemes to schemes of member databases (e.g., [20, 33, 6, 26]).

3.3 Integrability Assumptions

The purpose of multidatabases is to integrate information from several, independent databases. Of course, the problem of integration is trivial, unless the information sources are *inconsistent*: i.e., a portion of the real world is described differently by more than one source. Such inconsistencies fall into two categories: (1) *intensional inconsistencies*, and (2) *extensional inconsistencies*.

Intensional inconsistencies, often referred to as *semantic heterogeneity*, are defined as differences in *modeling*. For example, differences in relation schemes, or in the semantics of individual attributes (e.g., measurement units). Extensional inconsistencies surface only *after* all intensional inconsistencies have been resolved, at a point where the systems participating in a specific transaction may

be assumed to have identical intensional representation for all overlapping information. At that point it is possible that two information sources would provide different answers to the same query.

We shall assume that there exists a single (hypothetical) database that represents the real world. This ideal database includes the usual components of scheme and instance. Its scheme is the perfect model, and its instance are the perfect data. We now formulate two assumptions. These assumptions are similar to the Universal Scheme Assumption and the Universal Instance Assumption [25], although their purpose here is quite different. These two assumptions are statements of the *integrability* of the given databases. They use the definition of derived databases in Section 3.1.

The Scheme Consistency Assumption (SCA). All database schemes are *derivatives* of the real world scheme. That is, every relation scheme is a view of the real world scheme. The meaning of this assumption is that the different ways in which reality is modeled are all correct; i.e., there are no *modeling errors*, only *modeling differences*. To put it in yet a different way, all intensional inconsistencies among the independent database models are reconcilable.

The Instance Consistency Assumption (ICA). All database instances are *derivatives* of the real world instance. That is, in each database instance, every relation instance is derived from the real world instance. The meaning of this assumption is that the information stored in databases is always correct; i.e., there are no factual *errors*, only different *representations* of the facts. In other words, all extensional inconsistencies among the independent database instances are reconcilable.

Although these assumptions have not been articulated before in the context of database integration, tacit assumptions have often been made. Most previous work on scheme integration has tacitly subscribed to the Scheme Consistency Assumption, and the different approaches to scheme integration are therefore implementations of specific techniques for reconciling modeling inconsistencies. With few exceptions in the areas of logic databases [32] and data fusion [1], most previous work on database integration has tacitly subscribed to the Instance Consistency Assumption as well, thus avoiding any possibility of data inconsistency.

The Multiplex model assumes that the Scheme Consistency Assumption *holds*, meaning that all differences among database schemes are reconcilable, and that the Instance Consistency Assumption *does not hold*, allowing the possibility of irreconcilable differences among database instances.

In other words, the member databases are all assumed to have schemes that are derivatives of a hypothetical real world database scheme; these schemes are related through the multidatabase scheme, which is yet another derivative of this perfect database scheme. But the member database instances are not assumed to be derivatives of the real world instance.

Clearly, without subscribing to the SCA, it is not possible to integrate a given set of databases. On the other hand, subscribing to the ICA would not reflect the reality of independently maintained databases.

3.4 Discussion

Our definition of multidatabases provides four important “degrees of freedom”, which reflect the realities of multidatabase environments.

First, the mapping from D to the member schemes is not necessarily *total*; i.e., not all views of D are expressible in one of the member databases (and even if they are expressible, there is no guarantee that they are mapped). This models the dynamic situation of a multidatabase system, where some member databases might become temporarily unavailable. In such cases the corresponding mappings are “suspended”, and some global queries might not be answerable in their entirety. Similarly, if an authorization mechanism is enforced, a user may not have permission to some views.

Second, the mapping is not necessarily *surjective*; i.e., the member databases may include views that are not expressible in D (and even if they are expressible, there is no guarantee that they are mapped). This models the pragmatism of multidatabases, which usually cull from existing databases only the information which is relevant to a specific set of applications. For example, a large database may share only one or two views with the multidatabase.

Third, the mapping is not necessarily *single-valued*; i.e., a view of D may be found in several member databases. This models the realistic situation, in which information is found in several overlapping databases, and provides a formal framework for dealing with multidatabase inconsistency. Recall that if we do not assume that the Instance Consistency Assumption holds, then we do not assume that the member instances are all derived from a single instance. Thus, the inclusion of view pairs (V, V_1) and (V, V_2) in two scheme mappings of a multidatabase does not imply that the extensions of V in the member databases are identical. Rather, it implies that they *should be* identical.

Fourth, while the definition assumes that the member databases adhere to the relational model defined here, they need not be relational, or even of the same data model. Recall that the only purpose of the views in the second position of the scheme mappings is to describe how the views in the first position are materialized. Therefore, the member databases need not be relational, and the views in the second position need not be relational expressions. The only requirement is that they compute tabular answers.

3.5 Multidatabase Queries

An essential part of the definition of a database model is its query language. The definition of a language must provide for syntax, as well as semantics; that is, one must define not only how queries are written, but also their extension in any database instance. In this section we consider multidatabase queries.

Syntactically, a multidatabase query is simply a query Q of the scheme D . Intuitively, the answer to a multidatabase query Q should be obtained by transforming it to an equivalent query of the views in the first position of the scheme mappings (the available information). These views would then be materialized (using the view definitions in the second position of the scheme mappings), and

the translated query would be processed on these materialized views. Formally, the required transformation of Q is stated as follows.

Let $D = \{R_1, \dots, R_n\}$ denote a database scheme, and let $M = \{V_1, \dots, V_m\}$ denote a set of views of D . Translate a given query Q_D of the database scheme to an equivalent query Q_M of the view schemes.

However, a solution to this translation problem may not exist, or there could be multiple solutions. To observe that multiple solutions may exist, consider a database with a relation $R = (A, B, C)$ and views $V_1 = \text{project}_{A,B}R$ and $V_2 = \text{project}_{A,C}R$, and consider the query $Q = \text{project}_AR$. Q can be answered from both V_1 or V_2 . To observe that a solution may not exist, consider a database with two relations $R = (A, B)$ and $S = (B, C)$, and one view $V = R \bowtie S$, and consider the query $Q = \text{select}_{A=a}R$. Clearly, Q cannot be answered from the view V , because the join would not necessarily include all of R 's tuples.

This translation problem (for conjunctive queries and views) has been addressed by Larson and Yang [21,22], by Levy et al. [23] and by Brodsky and Motro [8]. We shall assume that a translation algorithm exists which is sound and complete; i.e., it computes all the correct translations that exist.

3.6 Multidatabase Answers

Assume a multidatabase with scheme D and mapped views M . The *answer* to a query Q on this multidatabase is the set of answers produced by a sound and complete translation algorithm. There are two possible cases:

1. When the translation algorithm produces more than one solution, these solutions may evaluate to different answers. Each such answer is a *candidate answer*. The answer to Q is the set of all candidate answers.³
2. When a solution to the translation problem does not exist, the answer to Q is the empty set of answers. This empty set of answers should be interpreted as *answer unavailable*.⁴

Of course, if we were to subscribe to the Instance Consistency Assumption, then all the solutions generated in the first case would be guaranteed to evaluate to the same answer, and the answer to Q would be this unique answer. In theory, this answer may be evaluated from an arbitrary solution to the translation problem. In practice, however, some realistic model of *cost* should be adopted, and the *cheapest* solution should be chosen. Note that under the Instance Consistency Assumption, it is possible that the translation algorithm will have no solutions. Hence, the possibilities under the ICA are one answer or no answer.

4 Approximative Answers

From a user perspective, each legitimate database query should evaluate to a single answer. The multidatabase answers defined in the previous section deviate from this ideal in two cases: when no answer is available, and when several

³ Note that possibly none of these answers is consistent with the real world answer.

⁴ Note the difference between an empty set of answers and an empty answer.

different answers are available. In either case, it is clear that a single *perfect* answer (i.e., an answer identical to the real world answer) cannot be determined from the multidatabase environment. At best, the system can provide an *approximation* of this elusive perfect answer. In this section, we extend the Multiplex model to handle these situations.

Intuitively, a global query cannot be translated to an equivalent query of the available views, because the mapping of the global scheme to the member schemes is not total; i.e., some information “promised” in the global scheme cannot be “delivered”. The most common reason for this is that some member databases are not responding. But a similar situation would occur if some of the requested information cannot be delivered due to insufficient permissions, or due to some resource having been exhausted before the entire answer could be obtained (e.g., time, or some other cost measure). In situations where a query Q cannot be rewritten as an equivalent query of the available views, an issue of great importance is how well can Q be *approximated* using the available views; i.e., what is the best approximation of Q that can be evaluated from the views?

Intuitively, a global query is translatable to different equivalent queries over the available views, because the mapping of the global scheme to the member schemes is not single-valued; i.e., there exists a view of the global scheme that can be materialized in more than one way. This happens when two view pairs of the mapping have the *same* view in their first position. Obviously, for every translation that uses one view, there is an equivalent translation that uses the other view. More generally, it happens when two pairs have *overlapping* views in their first position, as this implies that the intersection view can be mapped in two different ways. The most common reason for such multivalued mappings is that the information resources have overlapping information. Unless the ICA holds, these different translations could evaluate to different answers. In situations where there are different ways to rewrite Q as an equivalent query of the available views, and they evaluate to different answers, an issue of great importance is whether any one specific answer should be *preferred* over the others, or how the answers could be *combined* into a single answer.

Our discussion of approximative answers is divided into two. First we assume that the ICA holds. Next we assume that the ICA does not hold. We begin with concepts that will be used in both situations.

4.1 Sound and Complete Answers

Recall our assumption of a (hypothetical) database that represents the real world perfectly. Under the Scheme Consistency Assumption (which is adopted in Multiplex) all database schemes are derivatives of this database, and without loss of generality we assume now that the scheme of the available database is *identical* to the real world model. Let (D, d_0) denote the real world database, and let (D, d) denote the actual database. Therefore, the database instance d is an *estimate* of the real database instance d_0 .

Consider a query Q on the database scheme D . Let q denote its answer in the database instance d , and let q_0 denote its answer in the database instance

d_0 ; i.e., q_0 is the perfect answer, and q is its estimate from the actual database instance. Following [27], we say that q is a *sound* answer if $q \subseteq q_0$, and q is a *complete* answer if $q \supseteq q_0$. If q is both sound and complete then q has *integrity*.

Clearly, the perfect answer q_0 lies “between” any sound answer q_s and any complete answer q_c : $q_s \subseteq q_0 \subseteq q_c$. When the perfect answer cannot be computed from the available database, sound and complete answers serve as “below” and “above” *approximations*. Clearly, it is desirable to obtain the “largest” sound answer and the “smallest” complete answer. Together, these provide the tightest approximation for the perfect answer.

Consider a *subanswer* of Q (a subview *enlarged* with nulls to the scheme of Q). The elements of the subanswer are elements of q_0 ; hence the subanswer is a sound answer.⁵ Consider a *superanswer* of Q (a superview *projected* on the scheme of Q). The elements of q_0 are elements of the superanswer; hence the superanswer is a complete answer. Intuitively, combining subanswers and superanswers (through union and intersection) will provide us with the aforementioned “largest” sound answer and “smallest” complete answer.

However, when the ICA does not hold, inconsistencies among subanswers and superanswers are possible. For example, we may encounter a subanswer and a superanswer where the former is not contained in the latter (i.e., some elements of the subanswer are not in the superanswer).

It is possible to define models where such inconsistencies are resolved by preferring certain answers over others (based on known external properties of the answers). This version of Multiplex assumes that no such additional information is available, in effect accepting all answers, subanswers and superanswers as “equally good”. Inconsistencies are then resolved on the basis of *voting*.

Sound answers establish that certain data are *included* in q_0 ; complete answers establish that certain data are *excluded* from q_0 . Therefore, the soundness or completeness of an answer may be interpreted as a claim (a *vote*) on each element of the domain of the answer: a sound answer is a *yes* vote for its members, and a *maybe* vote for all its non-members. A complete answer is a *maybe* vote for its members and a *no* vote for all its non-members.

The assumption that all information is “equally good” is interpreted that each subanswer *claims* to be sound, each superanswer *claims* to be complete, and each answer *claims* to be both sound and complete. Our assumption that no additional information is available implies that all claims have the same likelihood of being correct; i.e., their individual votes have equal weights.

We now propose this three-valued operation to combine conflicting votes:

	yes	no	maybe
yes	yes	maybe	maybe
no	maybe	no	maybe
maybe	maybe	maybe	maybe

⁵ We consider a tuple with null values to be sound if its non-null values match those of a tuple of q_0 .

Briefly, according to this operation the final verdict is definite if and only if all votes are consistent. It is easy to verify that this operation is associative, and thus the order in which the candidate answers are considered is immaterial.

Consider an example with two answers q_1 and q_2 . There are four subsets of the answer space that are treated homogeneously: $q_1 \cap q_2$, $q_1 - q_2$, $q_2 - q_1$, and $\overline{q_1 \cup q_2}$. The subsets for which the vote is *yes* suggest a sound answer; the subsets for which the vote is not *no* suggest a complete answer. It is easy to verify that the elements in $q_1 \cap q_2$ have a *yes* vote, elements in $q_1 - q_2$ and in $q_2 - q_1$ have a *maybe* vote, and all other elements have a *no* vote. Hence, $q_1 \cap q_2$ is voted as a sound answer, and $q_1 \cup q_2$ is voted as a complete answer.

4.2 Approximation under the ICA

If we assume that the Instance Consistency Assumption holds, then the only possible anomaly are queries for which “full answers” (answers that respond to the entire query) are unavailable. We define “partial answers” (which the ICA would guarantee to be mutually consistent), and we show how they should be combined into *approximative answers*, consisting of the largest sound and the smallest complete answers.

Because the ICA holds, then every *subview* of Q that can be expressed with the available views would evaluate to a sound approximation of the answer to Q , and any *superview* of Q that can be expressed with the available views would evaluate to a complete approximation of the answer to Q .

As an example, assume the global scheme $Emp = (Ename, Salary, Department, Location)$ and the views

$$\begin{aligned} V_1 &= project_{Ename, Salary} Emp \\ V_2 &= project_{Ename, Salary} select_{Department=design} Emp \\ V_3 &= project_{Ename, Department} Emp \\ V_4 &= project_{Ename, Department} select_{Salary>40} Emp \\ V_5 &= project_{Department, Location} Emp \end{aligned}$$

and consider these two queries:

$$\begin{aligned} Q_1 &= project_{Ename, Salary, Department, Location} Emp \\ Q_2 &= project_{Ename, Department} select_{Salary>50 \wedge Location=midtown} Emp \end{aligned}$$

Consider first Q_1 . Normally, it would be answered by joining V_1 , V_3 and V_5 . Suppose that V_1 and V_5 are not available. By substituting V_2 for V_1 , the system can provide the names and departments of *all* the employees, and the salaries of *some* of the employees, but locations for *none* of the employees. In effect, this information is the *union* of two *subviews* of Q_1 .

Consider now Q_2 . Normally, it would be answered by joining V_1 , V_3 and V_5 and selecting $Location = midtown$ and $Salary > 50$. Suppose that V_1 is not available. By joining V_3 and V_5 and selecting $Location = midtown$ the system can provide the names and departments of the midtown employees. By itself, V_4 provides the names and departments of the employees who earn over 40. Each

of these answers *contains* the requested answer. The *intersection* of these two *superviews* of Q_2 provides the requested information for the midtown employees who earn over 40, a set which is “close” to the answer.

Let $D = \{R_1, \dots, R_n\}$ denote a database scheme, let $M = \{V_1, \dots, V_m\}$ denote a set of views of D , and let Q be a query of D . A view V of M is a *maximal sound approximation* of Q using M , if $V \sqsubseteq Q$, and for every view V' of M : $V' \sqsubseteq Q \Rightarrow V' \sqsubseteq V$. A view V of M is a *minimal complete approximation* of Q using M , if $V \supseteq Q$, and for every view V' of M : $V' \supseteq Q \Rightarrow V' \supseteq V$.

The following results follow from mathematical set theory. The maximal sound approximation of Q using M is the *union* of all the subviews of Q that are expressible with views of M : $\bigcup\{W_M \mid W_M \sqsubseteq Q\}$. The minimal complete approximation of Q using M is the *intersection* of all the superviews of Q that are expressible with views of M : $\bigcap\{W_M \mid W_M \supseteq Q\}$.

Assume a multidatabase with scheme D and mapped views M . When a query Q on this multidatabase has no answer, the *approximate answer* to Q is the pair of maximal sound approximation and minimal complete approximation:

$$\langle \bigcup\{W_M \mid W_M \sqsubseteq Q\}, \bigcap\{W_M \mid W_M \supseteq Q\} \rangle \quad (1)$$

Observe that, because the ICA holds, when Q is answerable in its entirety, this approximation converges to a single answer, as described in Section 3.6. Note that a minimal complete approximation is not guaranteed, because a query might have no superviews that can be expressed with the available views. Assume a database scheme with two relations R_1 and R_2 and one view $V_1 = R_1$, and consider the query $Q = R_2$. Clearly, Q cannot be expressed with the available views; moreover, there is no view of the available views that is a superview of Q . On the other hand, a maximal sound approximation is guaranteed because a subview (possibly empty) of a query always exists.

4.3 Approximation when the ICA Does Not Hold

Assume now that the Instance Consistency Assumption does not hold. In constructing an approximation for a query, we must consider both full answers and partial answers (subviews and superviews of the answer).

Let $D = \{R_1, \dots, R_n\}$ denote a database scheme, let $M = \{V_1, \dots, V_m\}$ denote a set of views of D , and let Q be a query of D . The result of the query translation algorithm is three sets of views:

1. q_1, \dots, q_k are the available answers,
2. s_1, \dots, s_m are the available subviews of the answer, and
3. c_1, \dots, c_n are the available superviews of the answer.

We assume that all these answers are *independent* of each other; that is, none of the translations are subviews of each other.

After the appropriate enlargements and projections, we have a total of $k + m + n$ answers that *vote* on the entire answer space. Every q_i votes *yes* on its members and *no* on all other elements; every s_i votes *yes* on its members, and

maybe on all other elements; and every c_i votes *maybe* on its members and *no* on all other elements. Observe, however, that tuples may now be only partially specified (i.e., with null values). The question whether a tuple is a member of an answer is indeed whether a tuple *could* be a member of an answer.

The set of tuples for which the vote was unanimously *yes* is adopted as a sound estimate. The tuples for which the vote was not unanimously *no* is adopted as a complete estimate:

$$\langle \{t \mid t \text{ has only } \textit{yes} \text{ votes}\}, \{t \mid t \text{ does not have all } \textit{no} \text{ votes}\} \rangle \quad (2)$$

Observe that these are now *estimates* of a sound lower bound and a complete upper bound: since the candidate answers and partial answers are not assumed to be sound or complete, the lower approximation is not guaranteed to be sound, and the upper approximation is not guaranteed to be complete.

Our assumption here is the most general: (1) that the ICA does not necessarily hold, and (2) that the translation algorithm delivers both full answers and partial answers (i.e., subviews and superviews). When more specific assumptions are made, this solution is simplified. First, if only full answers are considered, then these estimates are reduced, respectively, to the intersection and union of the full answers: $\langle \bigcap_{i=1}^k q_i, \bigcup_{i=1}^k q_i \rangle$. Second, if the ICA holds, but only partial answers are available, then these estimates are reduced, respectively, to the subview union and the superview intersection defined in Formula (1). Finally, if the ICA holds, and a full answer is available, then these estimates converge to the simple answer defined at the end of Section 3.6.

The above estimates are aimed at maximizing soundness (the lower estimate) and completeness (the upper estimate). It is also possible to define answers “in between” these two. For example, the set of tuples for which at most one vote was not *yes*, the set of tuples for which at most two votes were not *yes*, and so on. This creates a sequence of containing answers that increase overall completeness while reducing overall soundness. It is then possible to provide the most sound answer which meets a user requirement on *minimal answer size*, or the largest answer which meets a user requirement on *minimal soundness*, and so on.

As mentioned earlier, when additional assumptions on the candidate answers can be made, more elaborate consolidation techniques may be developed. For example, instead of assuming that the alternative answers are of equal quality, we could assume that their individual quality measures are given, and adapt the voting scheme accordingly [28].

5 Implementation

As explained in the introduction, the Multiplex model is intended to provide both a formal foundation for research in multidatabases, and a practical architecture for a multidatabase system. Of course, actual implementations must consider additional details; yet the general principles of the Multiplex model would be upheld. In this respect it is similar to the relational model itself, whose

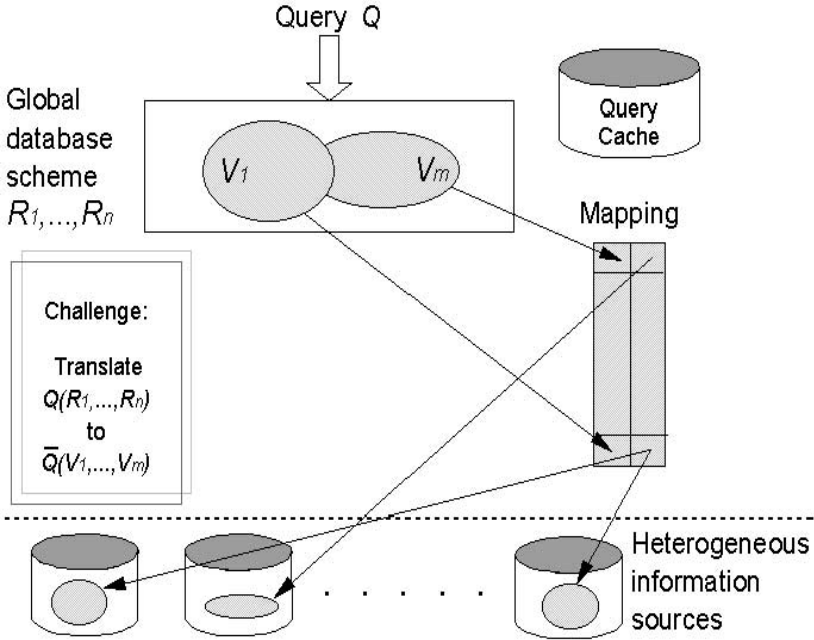


Fig. 1. A diagrammatic view of the Multiplex system.

formal definitions must be augmented with practical considerations in any implementation (e.g., optimization).

A prototype of the multidatabase model described in this paper has been implemented. A simple diagram of the system is shown in Figure 1.

The architecture and features of the Multiplex DBMS has six functional components: (1) user interface, (2) query parser, (3) query translator, (4) view retriever, (5) query optimizer, and (6) query processor, and it uses two sources of metadata: (1) a database scheme file, and (2) a database mapping file.

The user interface, the query parser, the query optimizer, the query processor, and the database scheme file are functionally similar to those of generic DBMS. There are two significant differences between a generic DBMS and Multiplex: (1) Multiplex does not have any relations; “instead” it has a scheme mapping file that matches views of the global scheme with views of the member databases. (2) Between the parsing and optimizing phases, which in a generic DBMS follow each other, the Multiplex query translator translates the global query to a query of the “available views”, and the Multiplex view retriever obtains the necessary views from the member databases. The translated query is then optimized and processed in the retrieved views in the same way as in a generic DBMS. Indeed, Multiplex retrieves views into a commercial DBMS (Oracle), and then submits the translated query to be optimized and processed by this DBMS.

For brevity, we only mention here a few notable aspects of the Multiplex system.

1. Multiplex uses the HTTP protocol for communicating with the member databases. In other words, it is a World Wide Web application. The user interface is accessible via WWW browsers, and client databases are invoked by so-called “cgi scripts”.
2. For its query translation, Multiplex uses the DRP software package [11], which has been modified and improved. For example, the translated query is modified further in an attempt to reduce transmission costs, using techniques known from distributed query optimization [10].
3. With respect to heterogeneity, presently, Multiplex can retrieve information from six kinds of sources: (1) relational (using Oracle), (2) object-oriented (using Ode), (3) simple files (using Unix shell scripts), and (4) Wide-area information services (WAIS, using SWISH 1.1 and WWWAIS 2.5) (5) spreadsheets (Microsoft Excell’s HTML output), and (6) menu-based (the XLibris library retrieval system).
4. Multiplex answers (refer to Formula (2)) are presented, using color, as two relations: one contains the *sound estimate*, the other contains the tuples that augment the sound estimate to a *complete estimate* (i.e., the union of both relations is the complete estimate), and users are advised that the answer to their query is estimated to contain the first relation *plus* a subset of the second relation.
5. Multiplex extends the query language of conjunctive queries with *aggregate functions*. A language based on conjunctive queries with aggregation provides a fairly powerful querying tool.

Overall, the linkage between the global database and the contributing sources is rather quick, requiring only entering pairs of equivalent queries in the mapping file. The result is that new information sources may be “plugged-in” very quickly.

6 Comparison with Other Approaches

As mentioned in the introduction, there has been considerable work in the area of multidatabases. A comprehensive discussion of every project or product is beyond the scope of this paper. In this section we compare our model and system to four different works, representing fairly different approaches.

UniSQL [18] is an example of a multidatabase system based on a comprehensive mapping of its global database scheme to the component database schemes. UniSQL provides an exhaustive framework for handling schematic heterogeneity (i.e., intensional inconsistencies) among the participating databases. Its reliance on predefined, comprehensive mappings dictates that UniSQL may not be as suitable for ad-hoc integration, in which (1) relatively small portions of the component sources are of interest (their entire schemes possibly being irrelevant, unavailable or incomprehensible), and (2) component sources change

frequently, with new sources being added and existing sources undergoing structural changes, or becoming altogether obsolete.

The TSIMMIS project [15] is an example of a system that is based on mediators and wrappers. *Mediators* [35] are software modules designed to deal with representation and abstraction problems that occur when trying to use data and knowledge resources. Mediators are understood to be active and knowledge-driven. *Wrappers* [35] are simpler software interfaces that allow a heterogeneous information sources to appear as if they conform to a uniform design or protocol. For example, a wrapper could be built to make a legacy database respond to a subset of SQL queries, as if it were a relational database. Multiplex makes fairly standard use of wrappers. With respect to mediators, a Multiplex query (a global view) may be considered a new “object”. Its translation produces an ad-hoc “mediator”, describing how the global object is to be constructed from the presently available sources. The advantage of such “dynamic mediation” are two: (1) Whereas with “static” mediators all integrated “objects” must be anticipated and predefined, in Multiplex an unlimited number of global objects may be defined spontaneously. (2) Static mediators need to be redefined whenever the available information sources change, whereas Multiplex only needs to have its mapping updated.

The approach of SIMS [3] to the integration problem is somewhat different. SIMS creates a *domain model* of the application domain, using a knowledge representation language to establish a fixed vocabulary describing objects in the domain, their attributes, and the relationships among them. Given a global query, SIMS identifies the sources of information that are required to answer the query and reformulates the query accordingly. SIMS is similar to Multiplex in that both do not rely on pre-programmed mediators, making the addition of new sources relatively simple. In both systems new sources have only to be *described* to the system. In SIMS, this description is in the knowledge representation language, using terms in the shared domain model; in Multiplex it is via pairs of equivalent views. Arguably, the SIMS descriptions are more demanding, but may allow the system to perform additional tasks. In contradistinction, Multiplex makes no claims of “intelligence”; it is a direct extension of relational model concepts, without the costs, risks, and possibly some benefits of a “knowledge-based” approach.

In many ways, the Information Manifold (IM) [2] is similar to SIMS. IM uses an object-relational model to integrate the various information sources, called *sites*. The individual sites are described and related to the global scheme, called the *world-view*, using the knowledge description language Classic. Like Multiplex, global query processing requires translation from the global set of relations to the set of available views. Like SIMS, and unlike Multiplex, the selection of relevant sites depends heavily on the quality of the site descriptions.

Finally, it should also be noted that none of these systems considers extensional inconsistencies (“too much data”) and their handling of partial answers (“too little data”) is fairly limited.

7 Conclusion

The Multiplex model that was described in this paper is both *formal* and *pragmatic*. It is a formal extension of the relational database model to multidatabases. This formalization reflects the important pragmatic issues encountered in actual multidatabase environments, and it can serve as the formal model behind many previous ad-hoc integration models that have already been designed.

The Multiplex model is also simple to implement, and when various pragmatic issues are addressed properly, it should prove to be highly practical. Towards this goal, we discuss several open issues and research directions.

Multidatabase design may be described as a mediation between the information needed (as expressed in the global database scheme) and the information available. The mediation process generates mappings that match information available (a view of some member database) with information needed (a view of the global scheme). An interesting issue that was not addressed in this paper is the *design* of the mapping. For example, a problem that might concern the designer is whether a present set of mapped views “covers” the global scheme. Formally, given a database scheme $D = \{R_1, \dots, R_n\}$, does a set of views $M = \{V_1, \dots, V_m\}$ guarantee that every query of D is expressible with M . Moreover, is it possible to characterize the queries that are not expressible, thus suggesting a view that would complement M ?

As defined in this paper, both the relational model the Multiplex model do not include integrity constraints. We are currently extending Multiplex to include global integrity constraints. Like views, these constraints would be mapped to constraints that are defined and enforced in the member databases. Global integrity constraints would provide a means for pruning the set of candidate answers.

One of the unique features of Multiplex is that its approach to cross-source inconsistencies does not require user involvement, essentially providing approximations that are based on voting. Experience accumulated so far indicates that in many instances it is highly desirable to provide user control over the process of inconsistency resolution. Our current development efforts (informally dubbed Multiplex II) would provide comprehensive and flexible control of the process, by means of *conflict resolution strategies*. These strategies would consider *properties* of the data (such as quality, currentness, or cost) as well as their actual *values*. Users would be given options such as picking the highest quality data among the alternatives, or consolidating the alternatives by their average. The strategies for resolving conflicts would be stated as part of the multidatabase design.

Acknowledgment: This work was supported in part by ARPA grant, administered by the Office of Naval Research under Grants No. N0014-92-J-4038 and N0060-96-D-3202. The author is grateful to Philipp Anokhin, who developed the prototype described in Section 5, and to Igor Rakov, who implemented an earlier version of the system. Both also provided important comments on this paper. The author is also indebted to Alex Brodsky, for his many insightful comments and suggestions.

References

1. M.A. Abidi and R.C. Gonzalez, editors. *Data Fusion in Robotics and Machine Intelligence*. Academic Press, 1992.
2. A. Levy and D. Srivastava and T. Kirk. Data model and query evaluation in global information systems. *Journal of Intelligent Information Systems*, 5(2):121–143, Sep. 1995.
3. Y. Arens, C. A. Knoblock, and W.-M. Shen. Query reformulation for dynamic information integration. *Journal of Intelligent Information Systems*, 6(2/3):99–130, June 1996.
4. C. Batini, M. Lenzerini, and S. B. Navathe. A comparative analysis of methodologies for database schema integration. *Computing Surveys*, 18(4):323–364, Dec. 1986.
5. Y. Breitbart. Multidatabase interoperability. *SIGMOD Record*, 19(3):53–60, Sep. 1990.
6. Y. Breitbart, P. L. Olson, and G. R. Thompson. Database integration in a distributed heterogeneous database system. In *Proceedings of the IEEE Computer Society Second International Conference on Data Engineering*, pp. 301–310, 1986.
7. M. W. Bright, A. R. Hurson, and S. H. Pakzad. A taxonomy and current issues in multidatabase systems. *Computer*, 25(3):50–60, March 1992.
8. A. Brodsky and A. Motro. The problem of optimal approximations of queries using views and its applications. Technical Report ISSE-TR-95-104, Department of Information and Software Engineering, George Mason University, May 1995.
9. O. P. Buneman, S. Davidson, and A. Watters. Federated approximations for heterogeneous databases. *Data Engineering*, 3(2):27–34, Aug. 1989.
10. S. Ceri and G. Pelagatti. *Distributed Databases: Principles and Systems*. McGraw-Hill, 1984.
11. N. Coburn. *Derived Relation Prototype: User Guide*. Department of Computer Science, University of Waterloo, 1988.
12. P. Drew, R. King, D. McLeod, M. Rusinkiewicz, and A. Silberschatz. Report of the workshop on semantic heterogeneity and interoperation in multidatabase systems. *SIGMOD Record*, 22(3):47–56, Sep. 1993.
13. G. Wiederhold, editor. Special issue: Intelligent integration of information. *Journal of Intelligent Information Systems*, 6(2/3), June 1996.
14. D. Fang, J. Hammer, and D. McLeod. The identification and resolution of semantic heterogeneity in multidatabase systems. In *Proceedings of the First International Workshop on Interoperability in Multidatabase Systems*, pp. 136–143, 1991.
15. H. Garcia-Molina, Y. Papakonstantinou, D. Quass, A. Rajaraman, Y. Sagiv, J. Ullman, and J. Widom. The TSIMMIS approach to mediation: Data models and languages. In *Proceedings of the Second International Workshop on Next Generation Information Technologies and Systems*, pp. 185–193, 1995.
16. A. R. Hurson, M. W. Bright, and S. H. Pakzad, editors. *Multidatabase Systems: An Advanced Solution for Global Information Sharing*. IEEE Computer Society Press, 1994.
17. Y. Kambayashi, M. Rusinkiewicz, and A. Sheth, editors. *Proceedings of the First International Workshop on Research Issues on Data Engineering: Interoperability in Multidatabase Systems*, 1991.
18. W. Kim and J. Seo. Classifying schematic and data heterogeneity in multidatabase systems. *IEEE Computer*, 24(12):12–18, 1991.

19. R. Krishnamurthy, W. Litwin, and W. Kent. Interoperability of heterogeneous databases with semantic discrepancies. In *Proceedings of the First International Workshop on Interoperability in Multidatabase Systems*, pp. 144–151, 1991.
20. T. A. Landers and R. L. Rosenberg. An overview of Multibase. In H.J. Schneider, editor, *Distributed Databases*, North-Holland, 1982.
21. P.-A. Larson and H. Z. Yang. Computing queries from derived relations. In *Proceedings of the Eleventh International Conference on Very Large Data Bases*, pp. 259–269, 1985.
22. P.-A. Larson and H. Z. Yang. Computing queries from derived relations: Theoretical foundations. Technical Report CS-87-35, Department of Computer Science, University of Waterloo, Aug. 1987.
23. A. L. Levy, A. O. Mendelzon, Y. Sagiv, and D. Srivastava. Answering queries from views. In *Proceedings of the 14th Symposium on Principles of Database Systems*, pp. 95–104, 1995.
24. W. Litwin, L. Mark, and N. Roussopoulos. Interoperability of multiple autonomous databases. *Computing Surveys*, 22(3):267–293, Sep. 1990.
25. D. Maier. *The Theory of Relational Databases*. Computer Science Press, 1983.
26. A. Motro. Superviews: Virtual integration of multiple databases. *IEEE Transactions on Software Engineering*, SE-13(7):785–798, July 1987.
27. A. Motro. Integrity = validity + completeness. *ACM Transactions on Database Systems*, 14(4):480–502, Dec. 1989.
28. A. Motro and I. Rakov. Not all answers are equally good: Estimating the quality of database answers. In *Flexible Query-Answering Systems*, pp. 1–21. Kluwer, 1997.
29. H.-J. Schek, A. Sheth, and B.D. Czejdo, editors. *Proceedings of the Third International Workshop on Research Issues on Data Engineering: Interoperability in Multidatabase Systems*, 1993.
30. P. Scheuermann, C. Yu, A. Elmagarmid, H. Garcia-Molina, F. Manola, D. McLeod, A. Rosenthal, and M. Templeton. Report on the workshop on heterogeneous database systems. *SIGMOD Record*, 19(4):23–31, Dec. 1990.
31. A. P. Sheth and J. A. Larson. Federated database systems for managing distributed, heterogeneous and autonomous databases. *Computing Surveys*, 22(3):183–236, Sep. 1990.
32. V. S. Subrahmanian. Amalgamating knowledge bases. *ACM Transactions on Database Systems*, 19(2):291–331, June 1994.
33. M. Templeton, D. Brill, S. K. Dao, E. Lund, P. Ward, A. L. P. Chen, and R. McGregor. Mermaid — a front-end to distributed heterogeneous databases. In *Proceedings of IEEE*, volume 75, number 5, pp. 695–708, May 1987.
34. J. D. Ullman. *Principles of Database Systems*. Computer Science Press, 1982.
35. G. Widerhold. Glossary: Intelligent integration of information. *Journal of Intelligent Information Systems*, 6(2/3):281–291, June 1996.