

Query Planning with Limited Source Capabilities (Extended Version)

Chen Li
Department of Computer Science
Stanford University
chenli@cs.stanford.edu

Edward Y. Chang
Electrical & Computer Engineering
University of California
echang@ece.ucsb.edu

Abstract

In information-integration systems, sources may have diverse and limited query capabilities. In this paper we show that because sources have restrictions on retrieving their information, sources not mentioned in a query can contribute to the query result, since they can provide useful bindings. In some cases we can access sources repeatedly to retrieve bindings to answer a query, and query planning thus becomes considerably more challenging. We find all the obtainable answers to a query by translating the query and source descriptions to a simple recursive Datalog program, and evaluating the program on the source relations. This program often accesses sources that are not in the query. Some of these accesses are essential, as they provide bindings that let us query sources, which we could not do otherwise. However, some of these accesses can be proven not to add anything to the query’s answer. We show in which cases these off-query accesses are useless, and prove that in these cases we can compute the complete answer to the query by using only the sources in the query. In the cases where off-query accesses are necessary, we propose an algorithm for finding all the useful sources for a query. We thus solve the optimization problem of eliminating the unnecessary source accesses, and optimize the program to answer the query.

Keywords: information-integration systems, limited source capabilities, Datalog programs.

1 Introduction

The rapid growth of the Internet is giving us access to an unprecedented number of heterogeneous information sources. Many information-integration systems (e.g., TSIMMIS [CGMH⁺94], the Information Manifold [LRO96], Garlic [RS97], Infomaster [GKD97], Disco [TRV98], Tukwila [IFF⁺99], and InfoSleuth [JBB⁺97]) have been proposed to support seamless access to these heterogeneous data sources. To perform queries on these sources, many studies [Dus97, LMSS95, Qia96, RSU95] construct answers to queries using views. These approaches are closely related to query-containment algorithms for conjunctive queries and for Datalog programs [Ul97].

In heterogeneous environments, especially in the context of World Wide Web, sources may have diverse and limited query capabilities. For example, many Web bookstores like `amazon.com` [Ama] and `barnesandnoble.com` [Bar] provide some search forms. A user fills out a form by specifying the values of some attributes, e.g., book title, author name, publisher, ISBN, so that the source can return the books satisfying the query conditions. These sources do not accept queries such as “return the information about all the books you know about.” There are many reasons for

the source restrictions, including the concerns of efficiency and security, and the limitations of the interfaces supplied by the sources.

In this paper, we consider a practical information-integration problem: *querying sources with limited capabilities*. We first show that because sources have restrictions on retrieving their information, sources not directly mentioned in a query can contribute to the query result, as shown by the following example.

EXAMPLE 1.1 Suppose that we want to compare the average prices of the books sold by `amazon.com` and `barnesandnoble.com`. Since both sources require each source query to specify at least one value of ISBN, author, or title, we cannot retrieve all their information about books. On the other hand, although we may have some known book information such as some authors, book titles, and ISBNs, the available information may not be enough to sample the two sources. However, assume that we have the access to the source `prenticehall.com`. We can retrieve from this source the authors of books published by Prentice Hall, and use these authors to query `amazon.com` and `barnesandnoble.com`. After retrieving the books, we can compare the average prices of the two sources.¹ □

Example 1.1 suggests that we can use the source `prenticehall.com` to retrieve some bindings for the author domain, and use the bindings to answer the query, although this source is not mentioned directly in the query. In some cases, as shown by the motivating example in Section 2, we can access sources repeatedly to obtain bindings to compute more results to a query. In Section 3 we propose a framework of query planning in integration systems with source restrictions. In the framework, source descriptions and a query are translated into a Datalog program, and we compute the *maximal* answer to the query by evaluating the program on the source relations. Datalog is used in the query planning since the planning process for a query can be recursive, although the query itself is not.

Being able to obtain the maximal results is desirable. However, the challenge is to return the results with the minimum cost. In other words, we do not want to involve all sources blindly during the plan-generation process. In Section 4 we show that in some cases a query does not need any bindings from off-query sources. In these cases we prove that the complete answer to the query can be computed by using only the sources in the query. In the cases where it is necessary to access off-query sources, we show in Section 5 that not all the sources that contribute bindings to the query are really necessary. We thus want to include judiciously only those sources that provide some values at a place where they are needed. We develop an algorithm for finding all the useful sources for a query. We solve the optimization problem of eliminating the unnecessary source accesses, and optimize the program to answer the query (Section 6).

We discuss in Section 7 how to explore other possibilities of obtaining bindings, e.g., by using cached data and domain knowledge. In the cases where a user may be interested in a partial answer to a query, we do not need to compute the maximal answer, which may be expensive to retrieve. We discuss how to compute a partial answer to a query, and the tradeoff between the number of results and the cost of an execution plan.

¹We use this example to illustrate the idea of accessing sources not in a query to obtain bindings. If we consider the possibility that bookstores and publishers may have deals, a better strategy would be: use the authors from `prenticehall.com` to retrieve book titles from the two bookstore sources, then use these titles to retrieve more authors. After several iterations, we average the prices of the books that are not published by Prentice Hall.

In this study we focus on a class of *connection queries*. A connection query is a natural join of distinct source views with the necessary selection and projection. (The details are described in Section 2.) Here we are taking the following universal-relation-like assumption [Ull89]: different attributes that share the same name in different views have the same meaning. However, universal-relation study did not consider restrictions of retrieving information from relations. In addition, as we will see in Section 2.2, a connection query can be generated in general cases, where our techniques are applicable.

Here is a summary of the contributions of this study:

1. We show that in information-integration systems, sources not in a query can contribute to the query result because of source restrictions. In some cases, we can obtain bindings by accessing sources repeatedly to answer a query, thus query planning in the presence of source restrictions becomes considerably more challenging.
2. We propose a query-planning framework, in which source descriptions and a query are translated into a Datalog program. We evaluate the program on the source relations to compute the maximal obtainable answer to the query.
3. We show how to decide whether accessing off-query sources is necessary. In the cases where it is not necessary, we prove that the complete answer to the query can be computed by using only the sources in the query.
4. In the cases where we need other sources to contribute bindings, we propose an algorithm for finding the useful sources and constructing an efficient Datalog program to compute the answer.

1.1 Related Work

There are two approaches to information integration [Dus97]:

1. The source-centric approach: Both user queries and source views are in terms of some global views. For each user query, the integration system needs to plan how to answer the query using source views. The Information Manifold and Infomaster follow this approach.
2. The query-centric approach: User queries are in terms of views synthesized at a mediator [Wie92] that are defined on source views. After view expansion [LYV⁺98] at the mediator, the query is translated to a logical plan that is composed of the source views. TSIMMIS follows this approach, and we follow this approach in this paper.

[Ull97] is a good survey on the differences between these two approaches. Many studies have been done by taking the source-centric approach. For example, Qian [Qia96] discussed how to use query folding to rewrite queries using views without considering limited source capabilities. Rajaraman, Sagiv, and Ullman [RSU95] proposed algorithms for answering queries using views with binding patterns. Duschka and Levy [DL97] considered source restrictions by translating source binding patterns into rules in a Datalog program, assuming that all attributes share the same domain. The paper did not discuss how to trim useless sources, thus it may generate programs that are not efficient to evaluate. Other studies [ASU79a, ASU79b, CM77, CKPS95, SY80] discussed

conjunctive-query rewriting and optimization without considering the restrictions of retrieving information from relations.

By taking the query-centric approach, [LYV⁺98] showed how to generate an executable plan of a query based on source restrictions. If the complete answer to the query cannot be retrieved, [LYV⁺98] would not answer the query, but would claim that an executable plan does not exist. In this case, our approach can still compute a partial answer. Although we take the query-centric approach in this study, our techniques for finding useful sources are also applicable to the source-centric approach, since when source views are the same as global predicates, the query-centric approach in [DL97] and our framework generate equivalent Datalog programs. Other related studies include how to optimize conjunctive queries with source restrictions [FLMS99, YLUGM99], how to describe source capabilities using a powerful language [VP97], how to compute mediator capabilities given source capabilities [YLG MU99], and how to convert data at mediators [CDSS98].

2 Preliminaries

In this section, we present our motivating example and introduce the notation used throughout the paper.

EXAMPLE 2.1 Assume that we are building a system that integrates the information from four sources of musical CDs, as shown in Table 1. Sources s_1 and s_2 have information about CDs and their songs; sources s_3 and s_4 have information about CDs, their artists, and their prices. To simplify the notation, we use attribute *Song* for song title and attribute *Cd* for CD title. The “Must Bind” column in the table indicates the attributes that must be specified at a source. For instance, every query sent to s_2 must provide a CD title. In other words, without the information about CD titles, source s_2 cannot be queried to produce answers.

Source	Contents	Must Bind
s_1	$v_1(\text{Song}, \text{Cd})$	Song
s_2	$v_2(\text{Song}, \text{Cd})$	Cd
s_3	$v_3(\text{Cd}, \text{Artist}, \text{Price})$	Cd
s_4	$v_4(\text{Cd}, \text{Artist}, \text{Price})$	Artist

Table 1: Four sources of musical CDs

Source-view schemas can be represented by a hypergraph [Ull89], in which each node is an attribute and each hyperedge is a source view. The hypergraph of the four source views is shown in Figure 1, which also shows the tuples at each source. To simplify the presentation, we use symbols t_i , c_j , and a_k to represent a song title, CD, or artist, respectively. For instance, the source view $v_1(\text{Song}, \text{Cd})$ contains two tuples: $\langle t_1, c_1 \rangle$ and $\langle t_2, c_3 \rangle$. The figure also shows the adornments of the attributes in each source view: b means that the attribute must be bound, f means that the attribute can be free.

Suppose a user wants to find the prices of the CDs that contain a song titled t_1 . The answer can be obtained by taking the union of the following four joins: $v_1 \bowtie v_3$, $v_1 \bowtie v_4$, $v_2 \bowtie v_3$, and $v_2 \bowtie v_4$, and performing a selection $\text{Song} = t_1$ and then a projection onto the attribute *Price*. There are four CDs containing the song: $\langle c_1, a_1, \$15 \rangle$, $\langle c_1, a_1, \$13 \rangle$, $\langle c_5, a_5, \$11 \rangle$, and $\langle c_4, a_3, \$10 \rangle$. Therefore, without considering the source restrictions, the answer is $\{\$15, \$13, \$11, \$10\}$. However, due to the




Figure 1: The hypergraph representation

limited source capabilities, only the \$15 can be computed if we process each join in the query at one time (as in [HKWY97, LRO96, LYV⁺98]). The reason is that, $v_1 \bowtie v_3$ yields the \$15 in the answer; $v_1 \bowtie v_4$ cannot be executed by using only v_1 and v_4 , since v_4 requires that attribute *Artist* be specified, but we cannot bind this attribute using only these two source views. Similarly, neither of the other two joins can be executed. As a consequence, the user misses the cheaper source for CD c_1 and entirely misses CDs c_4 and c_5 . \square

In this study, we first propose a framework that can retrieve more results from sources with binding restrictions. Instead of considering each join individually, our framework involves other sources not in a join to produce more bindings to answer the join. For instance, when joining v_1 and v_4 , we also consider the information provided by v_2 and v_3 . As we will see in Section 3.3, our framework can find two additional CDs containing the song titled t_1 : $\langle c_1, a_1, \$13 \rangle$ and $\langle c_4, a_3, \$10 \rangle$, and return two additional tuples: \$13 and \$10. If the user wants to find the cheapest CD, our approach can save \$5 for the user!

2.1 Source Views

Now we give the notation used throughout the paper. Let an integration system have n sources, say, s_1, \dots, s_n . Assume that each source s_i provides its data in the form of a relational view v_i . If sources have other data models, we can use *wrappers* [HG⁺97] to create the simple relational view of data. In the case where one source has multiple relations, we can represent this source with multiple logical sources each of which exports only one relational view.

We assume that differences in ontologies, vocabularies, and formats used by sources have been resolved. In particular, if two sources share an attribute name, we assume that the attributes are equivalent (i.e., wrappers take care of any differences). Related research [MW97, PGMW95] suggests ways to deal with ontology and format differences. We assume that the schemas of the source views are defined on a global set of attributes. Each view schema is a list of global attributes, and different source views may share the same schema. For instance, in Example 2.1, we have four global attributes: *Song*, *Title*, *Artist*, and *Price*; source views v_1 and v_2 share the same schema (*Song*, *Cd*).

The query capability of each source is described as a template with a binding pattern [Ull89] representing the possible query forms that the source can accept. The adornments for the attributes in the binding pattern include b (the attribute must be bound) and f (the attribute can be free). For simplicity of exposition, we assume that each view has one template. We use v_i to stand for both the source view and its adorned template, and we believe the distinction should be clear in

context. Let $\mathcal{A}(v_i)$ denote the attributes in a source view v_i , and let $\mathcal{B}(v_i)$ and $\mathcal{F}(v_i)$ be the set of bound and free attributes in the adorned template of v_i , respectively. For instance, in Example 2.1, $\mathcal{B}(v_1) = \{Song\}$, $\mathcal{F}(v_1) = \{Cd\}$, and $\mathcal{A}(v_1) = \{Song, Cd\}$. Let \mathcal{V} denote the source views with their adornments, $\mathcal{A}(\mathcal{V})$ be the attributes in \mathcal{V} , and \mathcal{R} be an instance of the source relations of \mathcal{V} .

2.2 Queries

A user query is represented in the form

$$Q = \langle \mathcal{I}, O, \mathcal{C} \rangle$$

where \mathcal{I} is a list of input assignments of the form **attribute = constant**, O is a list of output attributes whose values the user is interested in, and \mathcal{C} is a list of *connections*. Each connection is a set of source views that connect the input attributes and the output attributes. As we will see shortly, we interpret a connection as the natural join of the views in the connection. The following are some possible ways in which \mathcal{C} could be generated:

1. It is generated by query expansion at a mediator, as in TSIMMIS [LYV⁺98].
2. It is generated by a minimal-connection algorithm, as in universal-relation systems [Ull89].
3. It is specified explicitly by the user.

For instance, the query in Example 2.1 can be represented as

$$Q = \langle \{Song = t_1\}, \{Price\}, \{T_1, T_2, T_3, T_4\} \rangle$$

in which the four connections are: $T_1 = \{v_1, v_3\}$, $T_2 = \{v_1, v_4\}$, $T_3 = \{v_2, v_3\}$, and $T_4 = \{v_2, v_4\}$. Note that there can be multiple input attributes and multiple output attributes in a query. Let $I(Q)$ and $O(Q)$ respectively denote the input attributes and the output attributes of query Q . $I(Q)$ and $O(Q)$ do not overlap because the user does not need to choose an attribute as an output attribute if its value is known. Let $\mathcal{A}(T)$ be all the attributes in a connection T .

2.3 The Answer to a Query

Suppose that T is a connection in query Q . For those tuples in the *natural join* of the relations in T that satisfy the input constraints in Q , their projections onto the output attributes are the *complete answer for connection T* . The union of the answers for all the connections in Q is the *complete answer to query Q* . Due to the limited source capabilities, the *obtainable answer to a query* is the *maximal* answer to the query that can be retrieved from the sources, using only the initial bindings in the query and the source relations \mathcal{R} .

The complete answer to a user query could be retrieved if the sources did not have limited query capabilities. However, we may get only a partial answer to the query due to the source restrictions. For instance, in Example 2.1, the complete answer to the query is $\{\$15, \$13, \$11, \$10\}$, while as we will see in Section 3.3, the obtainable answer is $\{\$15, \$13, \$10\}$. In some cases, we cannot retrieve any answer to a query, even though the complete answer to the query is not empty. The following example is a case in point.

Figure 2: The source views in Example 2.2

EXAMPLE 2.2 Consider the two source views $v_1(A, B)$ and $v_2(B, C)$ in Figure 2. Suppose that v_1 has one tuple $\langle a_0, b_0 \rangle$ and v_2 has one tuple $\langle b_0, c_0 \rangle$. Assume that a user query is $Q = \langle \{A = a_0\}, \{C\}, \{T\} \rangle$, in which the only connection $T = \{v_1, v_2\}$. The complete answer to Q is $\{c_0\}$, but the obtainable answer is empty, since we cannot obtain any binding of B . \square

Given source descriptions and a query, if the complete answer to the query cannot be computed, our framework collects as much information as possible to answer the query. In the rest of this paper, unless otherwise specified, the term *the answer for a connection* means the obtainable answer for the connection, and the *answer to a query* is the union of the obtainable answers for all the connections in the query.

3 Query Planning

In this section, we propose a framework of query planning in the presence of source restrictions. In the framework source descriptions and a query are translated into a Datalog program, which can be evaluated on source relations to answer the query. We also discuss the efficiency of the program.

3.1 Constructing the Program $\Pi(Q, \mathcal{V})$

Given source descriptions \mathcal{V} and a query Q , we translate them into a Datalog program, denoted $\Pi(Q, \mathcal{V})$, that incorporates the source restrictions and the query, and thus can be evaluated on the source relations. For instance, Figure 3 shows the Datalog program $\Pi(Q, \mathcal{V})$ for the query and the source views in Example 2.1. We use names beginning with lower case letters for constants and predicate names, and names beginning with upper case letters for variables. Note that this program is recursive, though query Q is not.

$r_1: \text{ans}(P) :- \widehat{v}_1(t_1, C), \widehat{v}_3(C, A, P)$ $r_2: \text{ans}(P) :- \widehat{v}_1(t_1, C), \widehat{v}_4(C, A, P)$ $r_3: \text{ans}(P) :- \widehat{v}_2(t_1, C), \widehat{v}_3(C, A, P)$ $r_4: \text{ans}(P) :- \widehat{v}_2(t_1, C), \widehat{v}_4(C, A, P)$ $r_5: \widehat{v}_1(S, C) :- \text{song}(S), v_1(S, C)$ $r_6: \text{cd}(C) :- \text{song}(S), v_1(S, C)$ $r_7: \widehat{v}_2(S, C) :- \text{cd}(C), v_2(S, C)$ $r_8: \text{song}(S) :- \text{cd}(C), v_2(S, C)$	$r_9: \widehat{v}_3(C, A, P) :- \text{cd}(C), v_3(C, A, P)$ $r_{10}: \text{price}(P) :- \text{cd}(C), v_3(C, A, P)$ $r_{11}: \text{artist}(A) :- \text{cd}(C), v_3(C, A, P)$ $r_{12}: \widehat{v}_4(C, A, P) :- \text{artist}(A), v_4(C, A, P)$ $r_{13}: \text{cd}(C) :- \text{artist}(A), v_4(C, A, P)$ $r_{14}: \text{price}(P) :- \text{artist}(A), v_4(C, A, P)$ $r_{15}: \text{song}(t_1) :-$
--	--

Figure 3: The Datalog program $\Pi(Q, \mathcal{V})$ in Example 2.1

Let us look at the details of how the program $\Pi(Q, \mathcal{V})$ is constructed. For each source view v_i , we introduce an EDB predicate ([Ull89]) v_i and an IDB predicate \widehat{v}_i (called the α -predicate of v_i). Predicate v_i represents all the tuples at source s_i , and \widehat{v}_i represents the *obtainable* tuples at s_i . Introduce a goal predicate *ans* to store the answer to the query: the arguments of *ans* correspond to the output attributes $O(Q)$ in Q .

Let $T = \{v_1, \dots, v_k\}$ be a connection in \mathcal{Q} . The following rule is the *connection rule* of T :

$$\mathbf{ans}(O(\mathcal{Q})) \text{ :- } \widehat{\mathbf{v}}_1(\mathcal{A}(\mathbf{v}_1)), \dots, \widehat{\mathbf{v}}_k(\mathcal{A}(\mathbf{v}_k))$$

where the arguments in predicate \mathbf{ans} are the corresponding attributes in $O(\mathcal{Q})$. For each argument in $\widehat{\mathbf{v}}_i$, if the corresponding attribute in view v_i is an input attribute of \mathcal{Q} , this argument is replaced by the initial value of the attribute in \mathcal{Q} . Otherwise, a variable corresponding to the attribute name is used as an argument in predicate $\widehat{\mathbf{v}}_i$. For instance, in Figure 3, rules r_1, r_2, r_3 , and r_4 are the connection rules of the connections T_1, T_2, T_3 , and T_4 , respectively.

Decide the domains of all the attributes in the source views, i.e., group the attributes into sets while the attributes in each set share the same domain. Introduce a unary *domain predicate* for each domain to represent all its possible values that can be deduced.² In Figure 3, we use the predicates $\mathit{song}, \mathit{cd}, \mathit{artist}$, and price to represent the domains of song titles, CD titles, artists, and prices, respectively.

Suppose that source view v_i has m attributes, say A_1, \dots, A_m . Assume the adornment of v_i says that the arguments in positions $1, \dots, p$ need to be bound, and the arguments in positions $p+1, \dots, m$ can be free. The following rule is the α -rule of v_i :

$$\widehat{\mathbf{v}}_i(\mathbf{A}_1, \dots, \mathbf{A}_m) \text{ :- } \mathit{domA}_1(\mathbf{A}_1), \dots, \mathit{domA}_p(\mathbf{A}_p), \mathbf{v}_i(\mathbf{A}_1, \dots, \mathbf{A}_m)$$

in which each domA_j ($j = 1, \dots, p$) is the domain predicate for attribute A_j . For $k = p+1, \dots, m$, the following rule is a *domain rule* of v_i :

$$\mathit{domA}_k(\mathbf{A}_k) \text{ :- } \mathit{domA}_1(\mathbf{A}_1), \dots, \mathit{domA}_p(\mathbf{A}_p), \mathbf{v}_i(\mathbf{A}_1, \dots, \mathbf{A}_m)$$

For instance, rule r_9 in Figure 3 is the α -rule of v_3 ; rules r_{10} and r_{11} are the domain rules of v_3 .

Assume that $A_i = a_i$ is in the assignment list \mathcal{I} of \mathcal{Q} , the following rule is the *fact rule* of attribute A_i :

$$\mathit{domA}_i(\mathbf{a}_i) \text{ :-}$$

For instance, rule r_{15} in Figure 3 is a fact rule of attribute Song , since we know from the query that t_1 is a song title.

The program $\Pi(\mathcal{Q}, \mathcal{V})$ is constructed in three steps:

1. Write the connection rule for each connection in \mathcal{Q} .
2. Write the α -rule and the domain rules for each source view in \mathcal{V} .
3. Write the fact rule for each input attribute in \mathcal{Q} .

In Figure 3, rules r_1, r_2, r_3 , and r_4 are the connection rules of T_1, T_2, T_3 , and T_4 , respectively. Rule r_5 is the α -rule of v_1 , and r_6 is the domain rule of v_1 . Similarly, rules r_7 to r_{14} are the α -rules and the domain rules of the other three source views. Finally, r_{15} is the fact rule of the input attribute Song . Recall that the source views in each connection link the input attributes and the output attributes in the query. Based on how program $\Pi(\mathcal{Q}, \mathcal{V})$ is constructed, we have the following proposition:

²The idea of domain predicates is borrowed from [DL97]. However, in our framework, different domains have different domain predicates, while in [DL97] only one domain predicate was used for all attributes.

Proposition 3.1 *Given source descriptions \mathcal{V} and a query \mathcal{Q} , the program $\Pi(\mathcal{Q}, \mathcal{V})$ is a safe ([Ull89]) Datalog program. \square*

3.2 Binding Assumptions

During the construction of the program $\Pi(\mathcal{Q}, \mathcal{V})$, we make the following important assumptions:

1. Each binding for an attribute must be from the domain of this attribute;
2. If a source view requires a value, say, a string, as its first argument, we will not allow the strategy of trying all the possible strings to “test” the source;
3. Rather we assume that any binding is either obtained from the user query, or from a tuple returned by another source query.

We use Example 2.1 to explain these assumptions in details. The first assumption says that we would not use an artist name as a binding for attribute *Song*. Notice that if two attributes have the same *type*, they can still be from two different domains. For example, the attributes *Song* and *Cd* share the same *string* type, but they have two different domains; the domain of *Song* is all song titles, and the domain of *Cd* is all CD titles.

Source view $v_3(Cd, Artist, Price)$ requires that a source query to source s_3 give a CD title. The second assumption says that we would *not* allow the following naive “strategy”: generate all possible strings to test whether s_3 has CDs with these strings as titles. This approach is too inefficient to be practical, since there will be an infinite number of strings that need to be generated to test s_3 . The third assumption says that each bound value of an attribute A must either be derived from the user query, or be a value of A in a tuple returned by another source query. For instance, if c_1 is a CD title returned from source s_1 , and $Cd = c_2$ is an initial binding in a query, then we know that c_1 and c_2 are two CD titles, and they can be used to query source s_3 . In Section 7 we will discuss other possibilities of obtaining bindings.

3.3 Evaluating the Program $\Pi(\mathcal{Q}, \mathcal{V})$

We evaluate the Datalog program $\Pi(\mathcal{Q}, \mathcal{V})$ on the source relations \mathcal{R} to compute the facts for predicate *ans*. Note that the v_i ’s are the only EDB predicates in $\Pi(\mathcal{Q}, \mathcal{V})$. However, in an integration system, we do not know the tuples at each source before sending source queries. Now we show how to evaluate $\Pi(\mathcal{Q}, \mathcal{V})$ to answer the query.

To evaluate the domain rules and the α -rule of a source view v_i , predicate v_i is “populated” by source queries to s_i . Suppose that the right-hand side of its domain rules and its α -rule is:

$$\text{dom}A_1(A_1), \dots, \text{dom}A_p(A_p), v_i(A_1, \dots, A_m)$$

Once we know that (a_1, \dots, a_p) are the values of the $\text{dom}A_j$ ’s ($j = 1, \dots, p$), respectively, we can send a query $v_i(a_1, \dots, a_p, A_{p+1}, \dots, A_m)$ to source s_i . This source query is guaranteed to be executable, since it satisfies the binding requirements of v_i . The results of this source query add more tuples to the predicate \hat{v}_i (for the α -rule) and the predicates $\text{dom}A_j$ ’s (for the domain rules).

After the evaluation of the program terminates, the facts for the domain predicates include *all* the obtainable values of the corresponding domains. Similarly, the α -predicate facts are *all* the obtainable tuples at the sources. Since $\Pi(Q, \mathcal{V})$ includes the connection rules for the connections in query Q , the facts for the goal predicate *ans* form the maximum answer to Q . Thus, we have proved the following theorem:

Theorem 3.1 *Given source descriptions \mathcal{V} and a query Q , for any source relations \mathcal{R} of \mathcal{V} , if we evaluate $\Pi(Q, \mathcal{V})$ on \mathcal{R} , the set of facts for the predicate *ans* is the obtainable answer to Q . \square*

Order	Source Query	Returned Tuple(s)	New Bindings(s)
1	$v_1(t_1, C)$	$\langle t_1, c_1 \rangle$	$Cd = c_1$
2	$v_3(c_1, A, P)$	$\langle c_1, a_1, \$15 \rangle$	$Artist = a_1$
3	$v_4(C, a_1, P)$	$\langle c_1, a_1, \$13 \rangle, \langle c_2, a_1, \$12 \rangle$	$Cd = c_2$
4	$v_2(S, c_2)$	$\langle t_2, c_2 \rangle$	$Song = t_2$
5	$v_1(t_2, C)$	$\langle t_2, c_3 \rangle$	$Cd = c_3$
6	$v_3(c_3, A, P)$	$\langle c_3, a_3, \$14 \rangle$	$Artist = a_3$
7	$v_4(C, a_3, P)$	$\langle c_4, a_3, \$10 \rangle$	$Cd = c_4$
8	$v_2(S, c_4)$	$\langle t_1, c_4 \rangle$	

Table 2: Evaluating the program in Figure 3

IDBs	Results	IDBs	Results
\hat{v}_1	$\langle t_1, c_1 \rangle \langle t_2, c_3 \rangle$	<i>song</i>	t_1, t_2
\hat{v}_2	$\langle t_1, c_4 \rangle \langle t_2, c_2 \rangle$	<i>cd</i>	c_1, c_2, c_3, c_4
\hat{v}_3	$\langle c_1, a_1, \$15 \rangle \langle c_3, a_3, \$14 \rangle$	<i>artist</i>	a_1, a_3
\hat{v}_4	$\langle c_1, a_1, \$13 \rangle \langle c_2, a_1, \$12 \rangle \langle c_4, a_3, \$10 \rangle$	<i>price</i>	$\$15, \$14, \$13, \$12, \$10$
<i>ans</i>	$\$15, \$13, \$10$		

Table 3: Results

Table 2 shows how to evaluate the program in Figure 3 to compute the answer to the query in Example 2.1, and Table 3 shows the results. As expected, the program computes all the possible values of song titles, CD titles, artists, and prices that can be retrieved from the four sources and the query, as well as all the obtainable tuples at the sources. The set of *ans* facts is the answer to the query. Therefore, our approach returns two more tuples, \$13 and \$10, than the approach in Section 2. Note that we cannot retrieve the tuple $\langle t_1, c_5 \rangle$ of v_2 and the tuple $\langle c_5, a_5, \$11 \rangle$ of v_4 , since we cannot retrieve the value a_5 for attribute *Artist*, no matter what legal source queries we execute.

3.4 Efficiency of the Program

The program $\Pi(Q, \mathcal{V})$ is constructed in a brute-force way, and it needs to be optimized. In particular, for each connection T in the query, the program may access views that are not in T . However, some of these off-connection accesses do not add anything to the query's answer. In the next two sections we discuss how to decide whether accessing off-connection views is necessary, and if necessary, what source views should be accessed.

4 Accessing Off-connection Views

In this section we discuss how to decide whether accessing off-connection views is necessary to compute the answer for a connection. In the case where it is necessary, we prove that the complete answer for the connection can be computed by using only the sources in the connection.

4.1 An Example

The following example shows that accessing all the views not mentioned in a connection is not always necessary to compute the maximal obtainable answer for the connection.




Figure 4: Source views in Example 4.1

EXAMPLE 4.1 Consider the five source views in Figure 4. Suppose that a user submits a query $\mathcal{Q} = \langle \{A = a_0\}, \{D\}, \{T_1, T_2\} \rangle$, which has two connections $T_1 = \{v_1, v_3\}$, $T_2 = \{v_2, v_3\}$. That is, the user knows that the value of A is a_0 , and wants to get the associated D values using $v_1 \bowtie v_3$ and $v_2 \bowtie v_3$. Assume that different attributes have different domains. The corresponding Datalog program $\Pi(\mathcal{Q}, \mathcal{V})$ is shown in Figure 5. For instance, rule r_1 is the connection rule of connection T_1 ; rules r_3 and r_4 are the α -rule and the domain rule of source view v_1 , respectively.

$r_1: \text{ans}(D) \quad :- \quad \widehat{v}_1(a_0, C), \widehat{v}_3(C, D)$	$r_9: \text{domD}(D) \quad :- \quad \text{domC}(C), v_3(C, D)$
$r_2: \text{ans}(D) \quad :- \quad \widehat{v}_2(a_0, B, C), \widehat{v}_3(C, D)$	$r_{10}: \widehat{v}_4(C, E) \quad :- \quad v_4(C, E)$
$r_3: \widehat{v}_1(A, C) \quad :- \quad \text{domA}(A), v_1(A, C)$	$r_{11}: \text{domC}(C) \quad :- \quad v_4(C, E)$
$r_4: \text{domC}(C) \quad :- \quad \text{domA}(A), v_1(A, C)$	$r_{12}: \text{domE}(E) \quad :- \quad v_4(C, E)$
$r_5: \widehat{v}_2(A, B, C) \quad :- \quad \text{domC}(C), v_2(A, B, C)$	$r_{13}: \widehat{v}_5(E, F) \quad :- \quad \text{domE}(E), v_5(E, F)$
$r_6: \text{domA}(A) \quad :- \quad \text{domC}(C), v_2(A, B, C)$	$r_{14}: \text{domF}(F) \quad :- \quad \text{domE}(E), v_5(E, F)$
$r_7: \text{domB}(B) \quad :- \quad \text{domC}(C), v_2(A, B, C)$	$r_{15}: \text{domA}(a_0) \quad :-$
$r_8: \widehat{v}_3(C, D) \quad :- \quad \text{domC}(C), v_3(C, D)$	

Figure 5: The Datalog program $\Pi(\mathcal{Q}, \mathcal{V})$ in Example 4.1

Consider connection T_1 . The program $\Pi(\mathcal{Q}, \mathcal{V})$ accesses the three views that are not in T_1 during the evaluation of the program. However, these off-connection accesses do not contribute to T_1 's results. The reason is that, suppose $t = \langle d \rangle$ is a tuple in the complete answer for T_1 , and t comes from tuple $t_1 = \langle a_0, c \rangle$ of v_1 and tuple $t_3 = \langle c, d \rangle$ of v_3 . By sending a query $v_1(a_0, C)$ to s_1 we can retrieve the tuple t_1 . With the new binding $C = c$ we can send a query $v_3(c, D)$ to s_3 , and retrieve tuple t_3 . Therefore, by using only the views in connection T_1 we can compute its complete answer.

Consider connection T_2 . Since we cannot get any binding for the attribute C by using only the two views in T_2 , we need v_2 and v_4 to contribute more C bindings. Thus these two off-connection views are useful to T_2 . On the other hand, $v_5(E, F)$ does not contribute to T_2 's results, because though source queries to s_5 bind more E and F values, these new bindings do not help obtain more answers for T_2 . \square

In general, given a connection T in a query \mathcal{Q} , we need to decide whether accessing the views outside T is necessary. Before giving the solution, we first introduce some notation.

4.2 Forward-closure

Definition 4.1 (forward-closure) Given a set of source views $\mathcal{W} \subseteq \mathcal{V}$ and a set of attributes $X \subseteq \mathcal{A}(\mathcal{V})$, the *forward-closure of X given \mathcal{W}* , denoted $f\text{-closure}(X, \mathcal{W})$, is a set of the source views in \mathcal{W} such that, starting from the attributes in X as the initial bindings, the binding requirements of these source views are satisfied by using only the source views in \mathcal{W} . \square

In other words, $f\text{-closure}(X, \mathcal{W})$ can be computed as follows: At the beginning, only the attributes in X are bound, and $f\text{-closure}(X, \mathcal{W})$ is empty. At each step, for each source view $v \in \mathcal{W} - f\text{-closure}(X, \mathcal{W})$, check whether $\mathcal{B}(v)$, the bound attributes of v , is a subset of the bound attributes so far. If so, add v to $f\text{-closure}(X, \mathcal{W})$, and each attribute in $\mathcal{F}(v)$, the free attributes of v , becomes bound. Repeat this process until no more source views can be added to $f\text{-closure}(X, \mathcal{W})$. Let $\mathcal{A}(f\text{-closure}(X, \mathcal{W}))$ denote all the attributes of the source views in $f\text{-closure}(X, \mathcal{W})$. Therefore, $\mathcal{A}(f\text{-closure}(X, \mathcal{W}))$ includes all the attributes that can be bound eventually by using the source views in \mathcal{W} starting from the initial bindings in X .

EXAMPLE 4.2 In Example 4.1, $f\text{-closure}(\{A\}, \{v_1, v_2, v_3\}) = \{v_1, v_2, v_3\}$, because we can use the bound attribute A to get tuples of v_1 and then bind C , which is the only bound attribute of v_2 and v_3 . In Example 2.1, $f\text{-closure}(\{Song\}, \{v_1, v_4\}) = \{v_1\}$, and $f\text{-closure}(\{Song\}, \{v_1, v_3\}) = \{v_1, v_3\}$. \square

4.3 Independent Connections

A connection T in query \mathcal{Q} is *independent* if

$$f\text{-closure}(I(\mathcal{Q}), T) = T$$

That is, the binding requirements of the source views in the connection can be satisfied by using only these source views starting from the initial bindings in $I(\mathcal{Q})$. In other words, if connection $T = \{w_1, \dots, w_k\}$ is independent, then there exists an *executable sequence* of all the source views in connection T : w_{i_1}, \dots, w_{i_k} , such that $\mathcal{B}(w_{i_1}) \subseteq I(\mathcal{Q})$, and for $j = 2, \dots, k$, $\mathcal{B}(w_{i_j}) \subseteq I(\mathcal{Q}) \cup \mathcal{A}(w_{i_1}) \cup \dots \cup \mathcal{A}(w_{i_{j-1}})$; i.e., the binding requirements of each source view in the sequence can be satisfied by the initial bindings in \mathcal{Q} and the source views before it. For instance, the connection $T_1 = \{v_1, v_3\}$ in Example 4.1 is independent since it has an executable sequence: v_1, v_3 . The following theorem shows that an independent connection does not require bindings from views outside the connection.

Theorem 4.1 *If connection T is independent, then for any source relations \mathcal{R} , we can compute the complete answer for T by using only the source views in T .* \square

Proof: Suppose that connection T has k source views, and it has an executable sequence v_1, \dots, v_k . Consider each tuple t in the complete answer for T . Assume that tuple t comes from tuples t_1, \dots, t_k of source views v_1, \dots, v_k , respectively. Since v_1, \dots, v_k is an executable sequence, the binding requirements of v_1 are satisfied by $I(\mathcal{Q})$, i.e., $\mathcal{B}(v_1) \subseteq I(\mathcal{Q})$. Thus, we can send a source query to s_1 by binding the attributes in $\mathcal{B}(v_1)$ to their initial values in \mathcal{Q} , and retrieve the tuple t_1 from v_1 . We then use the bound values of $I(\mathcal{Q}) \cup \mathcal{F}(v_1)$ to send s_2 a source query to get tuple t_2 . Repeat this process following the executable sequence, until we retrieve all the t_i 's. Therefore, by using only the views in T , we can retrieve the tuple t in the complete answer for the connection. ■

Theorem 4.2 *If connection T is not independent, then there exists an instance of the source relations \mathcal{R} , such that some tuples in the complete answer for T cannot be obtained.* □

Proof: If connection T is not independent, i.e., $f\text{-closure}(I(\mathcal{Q}), T) \neq T$, we construct an instance of source relations \mathcal{R} , such that a tuple in the complete answer for the connection cannot be obtained. Let $\mathcal{A}(T) = \{A_1, \dots, A_n\}$ be the set of attributes in T . Let tuple $t = (a_1, \dots, a_n)$, where a_i is a distinct value for attribute A_i . If A_i is in $I(\mathcal{Q})$, then its value in t , a_i , is its initial value in \mathcal{Q} . Each view v_i in T has only one tuple t_i , which is the projection of t onto the attributes $\mathcal{A}(v_i)$. Other sources are empty. Then the projection of t onto $O(\mathcal{Q})$ is in the complete answer for T . However, since $f\text{-closure}(I(\mathcal{Q}), T) \neq T$ and all other sources are empty, we cannot get the necessary bindings to retrieve the tuples t_i 's, and we cannot compute any answer for connection T . ■

Theorem 4.1 suggests that if a connection T is independent, we can compute the complete answer for T by using only the source views in T . In addition, Theorem 4.2 suggests that if connection T is not independent, there can be an instance of the source relations, such that some tuples in the complete answer for the connection cannot be retrieved. Many related studies (e.g., [FLMS99, LYV⁺98]) consider the case where a connection in a query is independent. If the connection is not independent, their algorithms give up attempting to answer the connection. However, our framework can still compute a partial answer for the connection by accessing off-connection views. In the next section, we will discuss how to decide what source views need to be accessed.

5 Finding the Relevant Source Views of a Connection

If a connection T is not independent, we may get more bindings by accessing views not in T . Some of the off-connection accesses are actually essential, as they provide bindings that let us ask queries of the views in the connection, which we could not do otherwise. However, some of these off-connection accesses can be proven not to add anything to the query's results. In this section, we discuss how to eliminate the unnecessary view accesses. We first give the formal definition of relevant source views of a connection, and then propose an algorithm for finding all the relevant source views of a connection. To simplify the presentation, in the rest of the paper we assume that different attributes are from different domains.

5.1 Relevant Source Views of a Connection

Given source descriptions \mathcal{V} , a query \mathcal{Q} , and a connection T in \mathcal{Q} , a source view $v \in \mathcal{V}$ is a *relevant source view* of connection T if for some source relations \mathcal{R} , removing v from \mathcal{V} can change the

obtainable answer for connection T ; otherwise, v is *irrelevant* to connection T . In other words, a source view v is relevant to a connection T if we can miss some answers for T if we do not use v . Note that whether a source view is relevant to a connection does not depend on other connections in the query. In Example 4.1, the relevant source views of connection T_1 are the two views in T_1 , while the relevant source views of connection T_2 are v_1, v_2, v_3 , and v_4 .

Given source descriptions \mathcal{V} , a query \mathcal{Q} , and a connection T in \mathcal{Q} , we need to solve the following problem: how to find all the relevant source views of T ? The following example shows the challenge of this problem: not all the views that contribute bindings to a connection are relevant to the connection.




Figure 6: The source views in Example 5.1

EXAMPLE 5.1 Consider the five source views in Figure 6. Suppose that a user submits a query $\mathcal{Q} = \langle \{A = a\}, \{F, G\}, \{T\} \rangle$, which has one connection $T = \{v_1, v_2, v_3\}$. That is, the user knows the value of A is a , and wants to get the associated D and E values using $v_1 \bowtie v_2 \bowtie v_3$. The connection T is not independent, because we cannot bind the attributes D and E by using only the views in T starting from the initial binding in \mathcal{Q} . We need other source views to bind D and E , so that we can query s_2 and s_3 to retrieve tuples. Thus, v_4 and v_5 may be useful since they provide bindings of D and E .

However, though view v_5 can bind attribute E , it is *not relevant* to connection T . To illustrate the reason, we prove that using only v_1, v_2, v_3 , and v_4 , we can compute all the obtainable answer for T . Suppose tuple $t = \langle f, g \rangle$ is in the obtainable answer, and t comes from tuple $t_1 = \langle a, b, c \rangle$ of v_1 , tuple $t_2 = \langle b, d, e, f \rangle$ of v_2 , and tuple $t_3 = \langle c, d, e, g \rangle$ of v_3 . Since the initial value of A in the query is a , we can send a source query $v_1(a, B, C)$ to retrieve tuple t_1 from v_1 . Because attribute D is not in $I(\mathcal{Q})$, and only v_4 (with the binding pattern ff) takes D as a free attribute, then the value d of D must be derived from the result of a source query to s_4 , which includes a tuple whose D value is d . With $C = c$ and $D = d$, we can retrieve tuple t_3 from v_3 by sending a source query $v_3(c, d, E, G)$, and then retrieve tuple t_2 from v_2 by sending a source query $v_2(b, d, e, F)$. Therefore, without using v_5 , we can get the tuple t in the obtainable answer for connection T . The proof also shows that without using v_4 , we cannot get any answer for T . \square

As there may be many source views with different schemas and binding patterns, it becomes challenging to decide what source views can really contribute to the results of a connection. Before giving the algorithm for finding all the relevant source views of a connection, we require a series of definitions.

5.2 Queryable Source Views

A source view v is *queryable* if it is in $f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$. All the queryable source views are those that we may eventually query, starting from the initial bindings in $I(\mathcal{Q})$, and perhaps using several preliminary queries to other sources in order to get the bindings we need for these source views. Let \mathcal{V}_q denote all the queryable source views in \mathcal{V} , and $\mathcal{A}(\mathcal{V}_q)$ be all the attributes in \mathcal{V}_q .

Lemma 5.1 *If attribute A is in $\mathcal{A}(\mathcal{V}_q)$, then either A is in $I(\mathcal{Q})$, or there exists an executable sequence of source views, such that A is a free attribute of the tail in the sequence.* \square

Proof: If attribute A is not in $I(\mathcal{Q})$, since A is in $\mathcal{A}(\mathcal{V}_q)$, there must exist a source view v in \mathcal{V}_q , such that A is in $\mathcal{F}(v)$. Based on how $\mathcal{V}_q = f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$ is computed, there must exist an executable sequence of source views in which v is the tail. This sequence can be obtained by backtracking from v following the reverse order in which these source views are added into $f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$. \blacksquare

We cannot get any tuples from a nonqueryable source view, no matter what the source relations are. If a connection contains a nonqueryable source view, we cannot get any answer for this connection. Thus we need to consider only the *queryable connections* in \mathcal{Q} , i.e., the connections that do not have any nonqueryable source view. Clearly an independent connection is also a queryable connection, but not vice versa. For instance, in Example 4.1, connection T_2 is queryable, since both v_2 and v_3 are queryable source views, but T_2 is not independent.

5.3 Kernel, BF-chain, and Backward-closure

Definition 5.1 (kernel) Assume T is a queryable connection in query \mathcal{Q} . A set of attributes $\mathcal{K} \subseteq \mathcal{A}(T)$ is a *kernel* of T if

$$f\text{-closure}(\mathcal{K} \cup I(\mathcal{Q}), T) = T$$

and by removing any attribute A from \mathcal{K} , $f\text{-closure}((\mathcal{K} - \{A\}) \cup I(\mathcal{Q}), T) \neq T$. \square

Intuitively, a kernel \mathcal{K} of connection T is a minimal set of attributes in $\mathcal{A}(T)$ such that, if the attributes in \mathcal{K} have been bound, together with the initial bindings in $I(\mathcal{Q})$, we can bind all the attributes $\mathcal{A}(T)$ by using only the source views in T . In Example 4.1, $\{C\}$ is a kernel of connection T_2 , because $f\text{-closure}(\{C\} \cup I(\mathcal{Q}), T_2) = f\text{-closure}(\{C, A\}, T_2) = T_2$. In Example 5.1, $\{D\}$ is a kernel of the connection T , while $\{D, E\}$ is not. Since a kernel of a connection must be minimal, it cannot share any attribute with $I(\mathcal{Q})$.

We can get a kernel of connection T by shrinking the set of attributes $X = \mathcal{A}(T) - I(\mathcal{Q})$ as much as we can while X satisfies: $f\text{-closure}(X \cup I(\mathcal{Q}), T) = T$. When X cannot be smaller, it will be a kernel of T . An independent connection has only one kernel: the empty set ϕ . A nonindependent connection must have only nonempty kernels. It may have multiple kernels, as shown by the following example.

EXAMPLE 5.2 Figure 7 shows a hypergraph of four source views. The binding patterns for $v_1(A, B, C)$, $v_2(C, D, E)$, and $v_3(E, F, A)$ are all *bff*, and the binding pattern for $v_4(E, G)$ is *ff*. Assume a user query is $\mathcal{Q} = \langle \{B = b_0\}, \{A, C, E\}, \{T\} \rangle$, in which the only connection is $T =$




Figure 7: Multiple kernels of a connection

$\{v_1, v_2, v_3\}$. T has three kernels: $\{A\}$, $\{C\}$, and $\{E\}$. For instance, $\{A\}$ is a kernel because $f\text{-closure}(\{A\} \cup I(\mathcal{Q}), T) = f\text{-closure}(\{A, B\}, \{v_1, v_2, v_3\}) = \{v_1, v_2, v_3\} = T$. \square

Definition 5.2 (BF-chain) A sequence of queryable source views w_1, \dots, w_k (i.e., each $w_i \in \mathcal{V}_q$) forms a *BF-chain* (bound-free chain) if for $i = 1, \dots, k - 1$, $\mathcal{F}(w_i) \cap \mathcal{B}(w_{i+1}) \neq \phi$. The source views w_1 and w_k are the *head* and the *tail* of the BF-chain, respectively. \square

In other words, for every two adjacent source views in a BF-chain, the free attributes of the first one overlap with the bound attributes of the second, and thus the first source view contributes bindings to the second one. In Example 4.1, (v_4, v_2, v_1, v_3) is a BF-chain, in which v_4 is the head and v_3 is the tail.

Definition 5.3 (backward-closure) Suppose that A is an attribute in $\mathcal{A}(\mathcal{V}_q)$. The *backward-closure of A* , denoted $b\text{-closure}(A)$, is the set of queryable source views that can be “backtracked” from A by following some BF-chain in a reverse order, in which A is a free attribute of the tail in each BF-chain. \square

More precisely, $b\text{-closure}(A)$ can be computed as follows: Start by setting $b\text{-closure}(A)$ to those source views in \mathcal{V}_q that take A as a free attribute. For each view $v \in \mathcal{V}_q - b\text{-closure}(A)$, if there is a view $w \in b\text{-closure}(A)$ such that $\mathcal{F}(v) \cap \mathcal{B}(w) \neq \phi$, then add v to $b\text{-closure}(A)$. Repeat this process until no more queryable source views can be added to $b\text{-closure}(A)$. In Example 4.1, the backward-closure of attribute C is $\{v_1, v_2, v_4\}$. The backward-closure of a set of attributes $X \subseteq \mathcal{A}(\mathcal{V}_q)$, denoted $b\text{-closure}(X)$, is the union of all the backward-closures of the attributes in X , i.e., $b\text{-closure}(X) = \bigcup_{A \in X} b\text{-closure}(A)$.

By the definitions of kernel, BF-chain, and backward-closure, we have the following lemmas and theorem.

Lemma 5.2 *If \mathcal{K} is a kernel of a queryable connection T and A is an attribute in \mathcal{K} , then A is not in $\mathcal{A}(f\text{-closure}((\mathcal{K} - \{A\}) \cup I(\mathcal{Q}), T))$. That is, starting from the attributes of $(\mathcal{K} - \{A\}) \cup I(\mathcal{Q})$ as the initial bindings, we cannot bind attribute A by using only the source views in T .* \square

Proof: Suppose that attribute A is in \mathcal{K} , and A is also in $\mathcal{A}(f\text{-closure}((\mathcal{K} - \{A\}) \cup I(\mathcal{Q}), T))$. Then starting from the attributes of $(\mathcal{K} - \{A\}) \cup I(\mathcal{Q})$ as the initial bindings, we can bind attribute A by using only the source views in T . Therefore, we can bind all the attributes in \mathcal{K} , and then bind all

the attributes in $\mathcal{A}(T)$ (since \mathcal{K} is a kernel of T). Thus, $\mathcal{K} - \{A\}$ would be a kernel of connection T . Then \mathcal{K} could not be a kernel since it is not minimal. ■

Lemma 5.3 *If A_1 and A_2 are two attributes, and there is a BF-chain such that A_1 is a bound attribute of the head and A_2 is a free attribute of the tail, then $b\text{-closure}(A_1) \subseteq b\text{-closure}(A_2)$.* □

Proof: Since A_2 is a free attribute of the BF-chain tail, we can backtrack from A_2 along the BF-chain until we reach A_1 in the head. Thus all the views on the BF-chain are in $b\text{-closure}(A_2)$. Based on how $b\text{-closure}(A_2)$ is computed, all the views in $b\text{-closure}(A_1)$ are also added into $b\text{-closure}(A_2)$ during the computation of $b\text{-closure}(A_2)$. Therefore, $b\text{-closure}(A_1) \subseteq b\text{-closure}(A_2)$. ■

Lemma 5.4 *If connection T has two different kernels $\mathcal{K}_1, \mathcal{K}_2$, then $b\text{-closure}(\mathcal{K}_1) = b\text{-closure}(\mathcal{K}_2)$.* □




Figure 8: Proof of Lemma 5.4

Proof: The main idea of the proof is shown in Figure 8. Since connection T has two different kernels, T cannot be independent, and both \mathcal{K}_1 and \mathcal{K}_2 are not empty. Since \mathcal{K}_1 is a kernel of T , for each attribute in \mathcal{K}_1 , say A_1 , by Lemma 5.2, we have $A_1 \notin \mathcal{A}(f\text{-closure}((\mathcal{K}_1 - \{A_1\}) \cup I(\mathcal{Q}), T))$. We also have $f\text{-closure}(\mathcal{K}_1 \cup I(\mathcal{Q}), T) = T$, while $f\text{-closure}((\mathcal{K}_1 - \{A_1\}) \cup I(\mathcal{Q}), T) \neq T$.

$$\mathcal{A}(f\text{-closure}((\mathcal{K}_1 - \{A_1\}) \cup I(\mathcal{Q}), T))$$

cannot be a subset of \mathcal{K}_2 , because otherwise starting from the attributes of $(\mathcal{K}_1 - \{A_1\}) \cup I(\mathcal{Q})$ as initial bindings and using only the source views in T , we could bind all the attributes in \mathcal{K}_2 , and then bind all the attributes in T (since \mathcal{K}_2 is a kernel of T), and $\mathcal{K}_1 - \{A_1\}$ would be a kernel.

Let A_2 be an attribute in \mathcal{K}_2 that is not in $\mathcal{A}(f\text{-closure}(\mathcal{K}_1 - \{A_1\}) \cup I(\mathcal{Q}), T)$. As shown in Figure 8, since A_2 must be in $\mathcal{A}(f\text{-closure}(\mathcal{K}_1) \cup I(\mathcal{Q}), T)$, then either $A_2 = A_1$, or there must exist a BF-chain, such that all the source views on the BF-chain are in T , and the head of the BF-chain takes A_1 as a bound attribute, and the tail takes A_2 as a free attribute. Note that in Figure 8, the attributes in $I(\mathcal{Q})$ may overlap with the attributes on the BF-chain.

If $A_2 = A_1$, then $b\text{-closure}(A_1) = b\text{-closure}(A_2) \subseteq b\text{-closure}(\mathcal{K}_2)$. If $A_2 \neq A_1$, then the above BF-chain exists. By Lemma 5.3, $b\text{-closure}(A_1) \subseteq b\text{-closure}(A_2) \subseteq b\text{-closure}(\mathcal{K}_2)$. Then we have $b\text{-closure}(\mathcal{K}_1) \subseteq b\text{-closure}(\mathcal{K}_2)$, since $b\text{-closure}(\mathcal{K}_1) = \bigcup_{A \in \mathcal{K}_1} (b\text{-closure}(A))$. Similarly, we can prove $b\text{-closure}(\mathcal{K}_2) \subseteq b\text{-closure}(\mathcal{K}_1)$. Therefore, $b\text{-closure}(\mathcal{K}_1) = b\text{-closure}(\mathcal{K}_2)$. ■

Lemma 5.4 shows that if a connection has multiple kernels, then the backward-closures of all these kernels are the same. For instance, in Example 5.2, the connection $T = \{v_1, v_2, v_3\}$ has three kernels: $\{A\}$, $\{C\}$, and $\{E\}$, and they have the same backward-closure: $\{v_1, v_2, v_3, v_4\}$.

5.4 Finding Relevant Source Views of a Connection

Now we show how to find all the relevant source views of a connection by giving the following theorem:

Theorem 5.1 *If \mathcal{K} is a kernel of a queryable connection T , then $b\text{-closure}(\mathcal{K}) \cup T$ are all the relevant source views of connection T .* \square




Figure 9: The backward-closure of a kernel of a queryable connection

We will give the proof shortly. Using Theorem 5.1, for a queryable connection T , we can find all its relevant source views by computing $b\text{-closure}(\mathcal{K}) \cup T$. If T is independent, then it has only one kernel, the empty set ϕ , whose backward-closure is empty. Thus only the source views in T are relevant to T , and this claim is consistent with Theorem 4.1. If the connection is not independent, as shown in Figure 9, we find a kernel \mathcal{K} of T (the B_i 's in the figure), and compute the backward-closure $b\text{-closure}(\mathcal{K})$. Then we find the relevant source views of T by taking the union of $b\text{-closure}(\mathcal{K})$ and T . Note that the backward-closures of different attributes in the kernel may overlap, and they may also overlap with the source views in T .

The Proof of Theorem 5.1

Proof: We need to prove that, for a kernel \mathcal{K} of a queryable connection T :

1. All the source views that are not in $b\text{-closure}(\mathcal{K}) \cup T$ are irrelevant to T .
2. Every source view on T is relevant to T .
3. Every source view in $b\text{-closure}(\mathcal{K})$ is relevant to T .

Proof of (1) We need to show that, by considering only the source views in $b\text{-closure}(\mathcal{K}) \cup T$, we can get the obtainable answer for connection T . Suppose $T = \{v_1, \dots, v_k\}$. For each tuple t in the obtainable answer for T , assume that t comes from the tuples t_1, \dots, t_k of source views

v_1, \dots, v_k , respectively. The values of the attributes $I(\mathcal{Q})$ in these t_i 's must be their initial values in \mathcal{Q} . Now consider the attributes in \mathcal{K} and how their values in the t_i 's are derived. These values must be obtained by some source queries using a set of source views, say \mathcal{V}_1 , whose free attributes include \mathcal{K} . By the definition of $b\text{-closure}(\mathcal{K})$, the source views in \mathcal{V}_1 must be in $b\text{-closure}(\mathcal{K})$. In order to send source queries to the sources of \mathcal{V}_1 , the binding requirements of the views in \mathcal{V}_1 must be satisfied. Let \mathcal{A} be the attributes of \mathcal{V}_1 such that the values of \mathcal{A} are not the initial values in \mathcal{Q} . Then the values of \mathcal{A} must be obtained using a set of source views, say \mathcal{V}_2 , whose free attributes include \mathcal{A} . Clearly the source views in \mathcal{V}_2 are also in $b\text{-closure}(\mathcal{K})$. We can see that all the source views used to derive the values of the attributes \mathcal{K} in the t_i 's are in $b\text{-closure}(\mathcal{K})$. After we get these values, following the same idea as in the proof of Theorem 4.1, we can get tuple t in the obtainable answer for T by using only the source views in T . Therefore, by using the source views in $b\text{-closure}(\mathcal{K}) \cup T$, we can get all the tuples in the obtainable answer for T .

Proof of (2) By Lemma 5.1, we can construct an instance of source relations \mathcal{R} such that the obtainable answer for T is not empty. But if we remove any source view in T , the obtainable answer for T becomes empty.

Proof of (3) For every source view v_i in $b\text{-closure}(\mathcal{K})$, we prove that v_i is relevant to connection T by constructing an instance of source relations \mathcal{R} , such that without using v_i , we cannot get a tuple in the obtainable answer for T . By the definition of $b\text{-closure}(\mathcal{K})$, there must be a BF-chain, as shown in Figure 9, such that the head of the BF-chain is v_i , and the tail (denoted w_1 in the figure) takes an attribute in \mathcal{K} (denoted B_1) as a free attribute.

\mathcal{R} is constructed in four steps. Start by setting all the source relations to empty. In step 1, populate the relations in connection T following the same idea as in the proof of Theorem 4.2. Let b_1, b_2, \dots be the values of attributes B_1, B_2, \dots in Figure 9, and t be the projection of the join of these tuples onto the attributes $O(\mathcal{Q})$. After we finish the construction of \mathcal{R} , tuple t will be in the obtainable answer for T .

In step 2, as shown in Figure 10, populate the source relations on the BF-chain from v_i to w_1 in such a way that only following this BF-chain we can get the value b_1 of attribute B_1 . Follow the same idea as in the proof of Theorem 4.2 to populate the source relations on the BF-chain, except that the B_1 value in the tuple of w_1 is b_1 . For each source view w_j on the BF-chain, w_j may have a set of bound attributes $Z(w_j) \subseteq \mathcal{B}(w_j)$ that do not overlap with the free attributes of the previous source view on the BF-chain. For view v_i , $Z(v_i) = \mathcal{B}(v_i)$. For instance, in Figure 10, region I includes the attributes of $\mathcal{F}(w_1)$, region II includes the attributes of $\mathcal{F}(w_2)$, and region III includes the attributes $Z(w_1) = \mathcal{B}(w_1) - \mathcal{F}(w_2)$. Regarding the attributes in $Z(w_j) \cap I(\mathcal{Q})$, their values in the tuple of w_j in step 2 are their initial values in \mathcal{Q} . The values of all other attributes in $Z(w_j) - I(\mathcal{Q})$ in the tuples are some new distinct values.




Figure 10: Populating the relations on the BF-chain in Figure 9.

In step 3, for each source view w_j on the BF-chain, if $Z(w_j) - I(\mathcal{Q})$ is not empty, consider each attribute A in it. By Lemma 5.1, there exists an executable sequence of source views such that A is a free attribute of the tail. Populate the source relations of this sequence following the same idea as in the proof of Theorem 4.2, except that the value of A in the tuple of the tail is the value of A in the tuple of w_j that we chose in step 2.

In step 4, consider each attribute in $\mathcal{K} - \{B_1\}$, say B_j . By Lemma 5.1, there exists an executable sequence of source views such that B_j is a free attribute of the tail. Populate the source relations of this sequence following the same idea as in the proof of Theorem 4.2, except that the value of B_j in the tuple of the tail is b_j , which is the value that we chose for B_j in step 1.

After these four steps, it is possible that more than one tuple has been added to a source relation. Based on how the relations are populated, all the source views that are populated are in $b\text{-closure}(\mathcal{K}) \cup T$, and we can get the b_i 's for the B_i 's in the connection by using all these populated source views. Following the same idea of the proof of Theorem 4.1, we can prove that by using all these source views, we can get tuple t in the obtainable answer for T .

Note that the only shared values among these source relations are those initial values in \mathcal{Q} , those b_i 's in connection T , and the values of the attributes in some $Z(w_j) - I(\mathcal{Q})$ during step 2, in which w_j is a source view on the BF-chain from v_i to w_1 . Since \mathcal{K} is a kernel of T , B_1 cannot be bound by using only the source views in T starting from the initial bindings in $I(\mathcal{Q})$. On the other hand, without using value b_1 of B_1 , we cannot get tuple t in the answer for T . Based on how \mathcal{R} is constructed, the only way to get b_1 of B_1 is to use the source relations on the BF-chain from v_i to w_1 , plus those source views that are populated in step 3. Without using view v_i , we cannot get the value b_1 of B_1 , so we cannot get tuple t in the answer for T . Therefore, v_i is a relevant source view of connection T . ■

5.5 The Algorithm FIND_REL

Using Theorem 5.1, we give the algorithm FIND_REL that finds all the relevant source views of a queryable connection in a query. The algorithm is shown in Figure 11.

The Algorithm FIND_REL

Input: Source views \mathcal{V} , a query \mathcal{Q} , a queryable connection T in \mathcal{Q} .

Output: All the relevant source views of T .

METHOD:

1. Compute all the queryable source views $\mathcal{V}_q = f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$;
2. Compute a kernel \mathcal{K} of connection T ;
3. Compute the backward-closure $b\text{-closure}(\mathcal{K})$.

Return $b\text{-closure}(\mathcal{K}) \cup T$ as all the relevant source views of connection T .

Figure 11: The algorithm FIND_REL

Let us analyze the complexity of the algorithm. Suppose that there are n source views in \mathcal{V} . Consider a queryable connection T with m source views and k attributes. Assume it takes $O(1)$ time to check whether a set of attributes is a subset of another set of attributes. As described in Section 5.2, we can get all the queryable source views by computing $f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$. Step 1 thus

can be done in $O(n^2)$ time. Step 2 can be done by following the approach described in Section 5.3, which shrinks the attributes in $\mathcal{A}(T) - I(\mathcal{Q})$ as much as they can be shrunk. Since for each set of attributes $X \subseteq \mathcal{A}(T) - I(\mathcal{Q})$, it takes $O(m^2)$ time to compute $f\text{-closure}(X \cup I(\mathcal{Q}), T)$, step 2 can be done in $O(km^2)$ time.

In step 3, for each attribute A in a kernel \mathcal{K} of T , $b\text{-closure}(A)$ can be computed in $O(n^2)$ time because during the computation, we can keep a set of attributes \mathcal{A}_b as the union of the $\mathcal{B}(w_i)$'s for every w_i , where w_i is a source view in $b\text{-closure}(A)$ that has been computed so far. At each step, for each queryable source view v that is not in the current $b\text{-closure}(A)$, we check whether $\mathcal{F}(v) \cap \mathcal{A}_b \neq \phi$. If so, v is added to $b\text{-closure}(A)$. Thus step 3 can be done in $O(kn^2)$ time. Therefore, the total time complexity of finding the relevant source views of the connection is $O(n^2) + O(km^2) + O(kn^2) = O(k(m^2 + n^2)) = O(kn^2)$.

6 Constructing an Efficient Program

In this section we show how to use the algorithm FIND_REL to construct a more efficient program than that constructed by the algorithm in Section 3, and show how to remove some useless rules in the program.

Given source descriptions \mathcal{V} and a query \mathcal{Q} , we can find the relevant source views of all the connections in \mathcal{Q} as follows:

1. Compute all the queryable source views $\mathcal{V}_q = f\text{-closure}(I(\mathcal{Q}), \mathcal{V})$;
2. Remove the nonqueryable connections, i.e., the connections that have a nonqueryable view;
3. Compute the relevant views for each queryable connection by calling the algorithm FIND_REL;
4. Take the union of all these relevant source views.

We use only these relevant source views (denoted \mathcal{V}_r) of query \mathcal{Q} to construct a Datalog program $\Pi(\mathcal{Q}, \mathcal{V}_r)$ in the same way as $\Pi(\mathcal{Q}, \mathcal{V})$ is constructed. For instance, in Example 4.1, all the five source views are queryable. By calling the algorithm FIND_REL we find that views v_1 and v_3 are relevant to connection T_1 ; views v_1, v_2, v_3 , and v_4 are relevant to connection T_2 . Therefore, the relevant views for both connections are v_1, v_2, v_3 , and v_4 . We use these four views to construct a more efficient program, which can be obtained by dropping the rules r_{13} and r_{14} in Figure 5.

In addition, some useless rules in the program $\Pi(\mathcal{Q}, \mathcal{V}_r)$ can be removed since they do not contribute to the answer. For instance, in Example 4.1, the user is not interested in the B and E values, so rules r_7 and r_{12} in Figure 5 can be dropped. Rules r_9 and r_{10} can also be removed since the predicates in their heads are not used by other rules. Figure 12 shows the optimized program that can compute the same answer as before, and the new program is more efficient to evaluate.

In general, the useless rules in $\Pi(\mathcal{Q}, \mathcal{V}_r)$ can be found as follows: Scan through all the rules in the program $\Pi(\mathcal{Q}, \mathcal{V}_r)$, except for the connection rules. For each rule r , check whether the IDB predicate in its head is used by other rules in the program. If not, rule r is useless and can be removed from the program. Repeat this process until no useless rules can be found in the program.

$r_1: \text{ans}(D) \quad :- \widehat{v}_1(a_0, C), \widehat{v}_3(C, D)$	$r_6: \text{domA}(A) \quad :- \text{domC}(C), v_2(A, B, C)$
$r_2: \text{ans}(D) \quad :- \widehat{v}_2(a_0, B, C), \widehat{v}_3(C, D)$	$r_8: \widehat{v}_3(C, D) \quad :- \text{domC}(C), v_3(C, D)$
$r_3: \widehat{v}_1(A, C) \quad :- \text{domA}(A), v_1(A, C)$	$r_{11}: \text{domC}(C) \quad :- v_4(C, E)$
$r_4: \text{domC}(C) \quad :- \text{domA}(A), v_1(A, C)$	$r_{15}: \text{domA}(a_0) \quad :-$
$r_5: \widehat{v}_2(A, B, C) \quad :- \text{domC}(C), v_2(A, B, C)$	

Figure 12: The optimized Datalog program in Example 4.1

7 Discussions

In this section we explore other possibilities for obtaining bindings during the query planning of a user query, e.g., by using cached data and domain knowledge. We also discuss how to answer a query when the user is interested in a partial answer, not necessarily the maximal answer.

7.1 Obtaining Bindings

Theorem 4.1 suggests that accessing off-connection views is only necessary for nonindependent connections. So far, we have assumed that the bindings of a domain are either from a user query or from other source queries. If cached data are available, they can be incorporated into the program $\Pi(\mathcal{Q}, \mathcal{V})$ for a query \mathcal{Q} and source descriptions \mathcal{V} . Suppose that we have a cached tuple $t_i(a_1, \dots, a_n)$ for source view $v_i(A_1, \dots, A_n)$. The following rules are added to the program $\Pi(\mathcal{Q}, \mathcal{V})$:

$$\begin{array}{ll} \widehat{v}_1(a_1, \dots, a_n) & :- \\ \text{dom}A_i(a_i) & :- \quad (i = 1, \dots, n) \end{array}$$

The predicates $\text{dom}A_1, \dots, \text{dom}A_n$ are the domain predicates for the attributes A_1, \dots, A_n , respectively. The first rule says that tuple $t_i(a_1, \dots, a_n)$ is a obtained tuple of source view v_i . The other fact rules represent the bindings for the corresponding domains. The new rules can contribute more answers to the query. Some views that were nonqueryable when we considered only the initial bindings in \mathcal{Q} may now become queryable with the new bindings from the cached data. In general, if we have some information about a domain, we can always incorporate the information into the program $\Pi(\mathcal{Q}, \mathcal{V})$ by adding the corresponding fact rules.

We may also obtain bindings by using some known domain knowledge. For example, suppose that we have a source view $\text{student}(\text{name}, \text{dept}, \text{GPA})$ with the binding pattern bbf . That is, every query to this source must supply a name and a department of a student, so that the student's GPA can be returned. Assume we know that all the students at the source are in four departments: $\{CS, EE, Physics, Chemistry\}$. Then we can use these four departments as bindings for attribute dept to query the source, and we do not need other sources to contribute bindings for this attribute.

7.2 Computing a Partial Answer

In some cases a user may be interested in a partial answer to a query. Thus we do not need to compute the maximal answer, which may be expensive to retrieve. Theorem 4.1 suggests that if a connection is independent, its complete answer can be computed by using only the views in the query. If a connection T is not independent, we can find a kernel \mathcal{K} of T . We access some sources in $b\text{-closure}(\mathcal{K})$ to obtain bindings for the attributes in \mathcal{K} , and compute a partial answer for the

connection. Notice that we may access only a subset of the backward-closure of \mathcal{K} , since we are not interested in the maximal answer for T . In addition, we need to consider the tradeoff between the number of results and the number of source accesses. The more sources we access, the more bindings we can retrieve, and the more answers we can compute for the connection. We decide how many source queries to send based on how many results the user is interested in.

8 Conclusion

In information-integration systems, especially in the context of World Wide Web, sources may have restrictions on retrieving their information, because they have diverse and limited query capabilities. We need to consider the source restrictions while answering a user query, since we may not be able to retrieve all the data from the sources. In this paper we showed that sources not directly mentioned in a query can contribute to the query result, since they provide useful bindings to the query. In some cases we can access sources repeatedly to compute more results to the query. We proposed a framework of query planning in the presence of source restrictions. In the framework, a user query and source descriptions are translated into a Datalog program, and we evaluate the program on the source relations to compute the maximal obtainable answer to the query. Our framework supports recursive query planning because of the expressive power of Datalog. In addition, we showed that accessing off-query sources is not always necessary. In the case where these off-query accesses do not contribute to the query results, we proved that the complete answer to the query can be computed by using only the views in the query. In the case where off-query accesses are necessary, we gave an algorithm for finding all the relevant sources to the query. By using this algorithm we can trim the unnecessary view accesses and construct an efficient Datalog program to compute the answer. Currently we are investigating how to evaluate the program to compute the answer to a query efficiently.

Acknowledgments: We thank Jeff Ullman for his valuable comments on this material.

References

- [Ama] Amazon.com. <http://www.amazon.com/>.
- [ASU79a] Alfred V. Aho, Yehoshua Sagiv, and Jeffrey D. Ullman. Efficient optimization of a class of relational expressions. *ACM Transactions on Database Systems (TODS)*, 4(4):435–454, 1979.
- [ASU79b] Alfred V. Aho, Yehoshua Sagiv, and Jeffrey D. Ullman. Equivalences among relational expressions. *SIAM Journal on Computing*, 8(2):218–246, 1979.
- [Bar] Barnesandnoble.com. <http://www.barnesandnoble.com/>.
- [CDSS98] Sophie Cluet, Claude Delobel, Jerome Simeon, and Katarzyna Smaga. Your mediators need data conversion! *ACM SIGMOD Conference*, pages 177–188, 1998.
- [CGMH⁺94] Sudarshan S. Chawathe, Hector Garcia-Molina, Joachim Hammer, Kelly Ireland, Yannis Papakonstantinou, Jeffrey D. Ullman, and Jennifer Widom. The TSIMMIS project: Integration of heterogeneous information sources. *IPSJ*, pages 7–18, 1994.

- [CKPS95] Surajit Chaudhuri, Ravi Krishnamurthy, Spyros Potamianos, and Kyuseok Shim. Optimizing queries with materialized views. *International Conference on Data Engineering (ICDE)*, pages 190–200, 1995.
- [CM77] Ashok K. Chandra and Philip M. Merlin. Optimal implementation of conjunctive queries in relational data bases. *STOC*, pages 77–90, 1977.
- [DL97] Oliver M. Duschka and Alon Y. Levy. Recursive plans for information gathering. *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, IJCAI-97*, 1997.
- [Dus97] Oliver M. Duschka. Query planning and optimization in information integration. *Thesis, Computer Science Department, Stanford University*, 1997.
- [FLMS99] Daniela Florescu, Alon Levy, Ioana Manolescu, and Dan Suciu. Query optimization in the presence of limited access patterns. *ACM SIGMOD Conference*, pages 311–322, 1999.
- [GKD97] Michael R. Genesereth, Arthur M. Keller, and Oliver M. Duschka. Infomaster: An information integration system. *ACM SIGMOD Conference*, pages 539–542, 1997.
- [HGMN⁺97] Joachim Hammer, Hector Garcia-Molina, Svetlozar Nestorov, Ramana Yerneni, Markus M. Breunig, and Vasilis Vassalos. Template-based wrappers in the TSIMMIS system. *ACM SIGMOD Conference*, pages 532–535, 1997.
- [HKWY97] Laura M. Haas, Donald Kossmann, Edward L. Wimmers, and Jun Yang. Optimizing queries across diverse data sources. *Very Large Data Bases (VLDB) Conference*, pages 276–285, 1997.
- [IFF⁺99] Zachary Ives, Daniela Florescu, Marc Friedman, Alon Levy, and Dan Weld. An adaptive query execution engine for data integration. *ACM SIGMOD Conference*, pages 299–310, 1999.
- [JBB⁺97] Roberto J. Bayardo Jr., Bill Bohrer, Richard S. Brice, Andrzej Cichocki, Jerry Fowler, Abdelsalam Helal, Vipul Kashyap, Tomasz Ksiezyk, Gale Martin, Marian H. Nodine, Mosfeq Rashid, Marek Rusinkiewicz, Ray Shea, C. Unnikrishnan, Amy Unruh, and Darrell Woelk. Infosleuth: Semantic integration of information in open and dynamic environments (experience paper). *ACM SIGMOD Conference*, pages 195–206, 1997.
- [LMSS95] Alon Y. Levy, Alberto O. Mendelzon, Yehoshua Sagiv, and Divesh Srivastava. Answering queries using views. *ACM Symposium on Principles of Database Systems (PODS)*, pages 95–104, 1995.
- [LRO96] Alon Y. Levy, Anand Rajaraman, and Joann J. Ordille. Querying heterogeneous information sources using source descriptions. *Very Large Data Bases (VLDB) Conference*, pages 251–262, 1996.
- [LYV⁺98] Chen Li, Ramana Yerneni, Vasilis Vassalos, Hector Garcia-Molina, Yannis Papakonstantinou, Jeffrey D. Ullman, and Murty Valivet. Capability based mediation in TSIMMIS. *ACM SIGMOD Conference*, pages 564–566, 1998.

- [MW97] David A. Maluf and Gio Wiederhold. Abstraction of representation for interoperation. *International Symposium on Methodologies for Intelligent Systems (ISMIS)*, pages 441–455, 1997.
- [PGMW95] Yannis Papakonstantinou, Hector Garcia-Molina, and Jennifer Widom. Object exchange across heterogeneous information sources. *International Conference on Data Engineering (ICDE)*, pages 251–260, 1995.
- [Qia96] Xiaolei Qian. Query folding. *International Conference on Data Engineering (ICDE)*, pages 48–55, 1996.
- [RS97] Mary Tork Roth and Peter M. Schwarz. Don’t scrap it, wrap it! A wrapper architecture for legacy data sources. *Very Large Data Bases (VLDB) Conference*, pages 266–275, 1997.
- [RSU95] Anand Rajaraman, Yehoshua Sagiv, and Jeffrey D. Ullman. Answering queries using templates with binding patterns. *ACM Symposium on Principles of Database Systems (PODS)*, pages 105–112, 1995.
- [SY80] Yehoshua Sagiv and Mihalis Yannakakis. Equivalences among relational expressions with the union and difference operators. *Journal of the ACM*, 27(4):633–655, 1980.
- [TRV98] Anthony Tomasic, Louiqa Raschid, and Patrick Valduriez. Scaling access to heterogeneous data sources with DISCO. *IEEE Transactions on Knowledge and Data Engineering*, 10(5):808–823, 1998.
- [Ull89] Jeffrey D. Ullman. *Principles of Database and Knowledge-base Systems, Volumes II: The New Technologies*. Computer Science Press, New York, 1989.
- [Ull97] Jeffrey D. Ullman. Information integration using logical views. *International Conference on Database Theory (ICDT)*, pages 19–40, 1997.
- [VP97] Vasilis Vassalos and Yannis Papakonstantinou. Describing and using query capabilities of heterogeneous sources. *Very Large Data Bases (VLDB) Conference*, pages 256–265, 1997.
- [Wie92] Gio Wiederhold. Mediators in the architecture of future information systems. *IEEE Computer*, 25(3):38–49, 1992.
- [YLG MU99] Ramana Yerneni, Chen Li, Hector Garcia-Molina, and Jeffrey D. Ullman. Computing capabilities of mediators. *ACM SIGMOD Conference*, pages 443–454, 1999.
- [YLUG M99] Ramana Yerneni, Chen Li, Jeffrey D. Ullman, and Hector Garcia-Molina. Optimizing large join queries in mediation systems. *International Conference on Database Theory (ICDT)*, pages 348–364, 1999.