Conjunctive Query Containment Revisited

(Extended Abstract)

Chandra Chekuri*
Anand Rajaraman**

Department of Computer Science, Stanford University

Abstract. We consider the problems of conjunctive query containment and minimization, which are known to be NP-complete, and show that these problems can be solved in polynomial time for the class of *acyclic queries*. We then generalize the notion of acyclicity and define a parameter called *query width* that captures the "degree of cyclicity" of a query: in particular, a query is acyclic if and only if its query width is 1. We give algorithms for containment and minimization that run in time polynomial in n^k, where n is the input size and k is the query width. These algorithms naturally generalize those for acyclic queries, and are of practical significance because many queries have small query width compared to their sizes. We show that we can obtain good bounds on the query width of Q using the *treewidth* of the incidence graph of Q. Finally, we apply our containment algorithm to the practically important problem of finding equivalent rewritings of a query using a set of materialized views.

1 Introduction

Testing query containment and equivalence are fundamental problems of database theory, and are central to global query optimization in database systems. Conjunctive queries are an important class of database queries, equivalent in expressive power to SPJ queries in the relational algebra. We consider the classical problem of testing containment of conjunctive queries. The problem is well-known to be NP-complete [CM77]. In view of its practical significance, considerable attention has been devoted to finding classes of conjunctive queries that admit polynomial-time time algorithms for equivalence and minimization [ASU79a, ASU79b, JK83]. *Acyclic queries* have been extensively studied in the context of query optimization in distributed database systems, and are well-known to have desirable algorithmic properties [Yan81].

In this paper, we first present polynomial-time algorithms to test the containment of an arbitrary conjunctive query in an acyclic query and to minimize

* Email: chekuri@cs.stanford.edu. Supported by NSF Award CCR-9357849, with matching funds from IBM, Mitsubishi, Schlumberger Foundation, Shell Foundation, and Xerox Corporation.

** E-mail: anand@cs.stanford.edu. Supported by NSF grant IRI–92–23405, ARO grant DAAH04–95–1–0192, and USAF contract F33615–93–1–1339.

an acyclic query. We then introduce a new parameter of a query called the *query width,* and show that the acyclic queries are precisely the class of queries with query width 1. We generalize the query containment and minimization algorithms to arbitrary queries such that their running time is polynomial in n^k, where n is the input size and k is the query width. These results are significant not only because they naturally generalize the algorithms for acyclic queries, but also because most commonly encountered queries have small query width compared to their sizes. We relate query width to *treewidth,* an extensively studied graph-theoretic parameter, and show that we can obtain good bounds on the query width from the treewidth of the *incidence graph* of the query.

There are close connections between query containment and the problem of answering queries using materialized views [LMSS95]. This problem has recently received considerable attention because of its numerous applications, which include speeding up query evaluation, querying heterogeneous information sources, mobile computing, and maintaining physical data indepedendence ([LMSS95] provides references). For example, the Information Manifold project [LRO96] represents the contents of heterogeneous information sources as views on a common set of base relations. A query Q is "solved" by a program that uses the views to obtain information from the sources.

We consider in this paper the problem of finding an equivalent rewriting of a conjunctive query Q using a set of views $\mathcal{V}$ defined by conjunctive queries, when Q does not use repeated predicates, and show how our algorithms for query containment can be modified for this problem. A restricted variant of this problem, where neither Q nor the views in $\mathcal{V}$ use repeated predicates, is known to be NP-complete [LMSS95].

This paper is organized as follows. Section 2 presents the basic definitions and defines the problems we consider. Section 3 introduces acyclicity, query width, and treewidth, and describes how these concepts relate to one another. In Section 4 we present our algorithms for query containment, and in Section 5 we modify these algorithms to obtain algorithms for answering queries using views. Section 6 describes related work, and Section 7 concludes by describing some open problems.

2 Preliminaries

We assume a fixed set of predicates, called the *database predicates*, over which queries are posed and views are defined. All queries in this paper are *conjunctive queries*, defined in the conventional manner [Ull89]. We say query Q_1 is *contained* in query Q_2, denoted by $Q_1 \subseteq Q_2$, if for every state of the relations corresponding to the database predicates, the relation corresponding to the head of Q_1 is a subset of the relation corresponding to the head of Q_2. We say that Q_1 is *equivalent* to Q_2, denoted by $Q_1 \equiv Q_2$, if $Q_1 \subseteq Q_2$ and $Q_2 \subseteq Q_1$.

A *containment mapping* from Q_1 to Q_2 is a mapping from the variables of Q_1 to the variables and constants of Q_2 that maps each subgoal of Q_1 to a subgoal of Q_2 and also maps the head of Q_1 to the head of Q_2. For conjunctive

queries, $Q_1 \subseteq Q_2$ if and only if there is a containment mapping from Q_2 to Q_1. The problems of testing whether $Q_1 \subseteq Q_2$ and $Q_1 \equiv Q_2$ are both NP-complete [CM77].

A *view* is a conjunctive query with a unique head predicate. Let Q be a query and $\mathcal{V}$ a set of views over the same set of database predicates. Let Q' be a query over the view predicates in $\mathcal{V}$. We extend the notions of containment and equivalence in the natural manner and make statements such as $Q \subseteq Q'$ and $Q \equiv Q'$. We call Q' an *equivalent rewriting* of Q using $\mathcal{V}$ if $Q' \equiv Q$.

Example 1. Suppose we have relations $part(Pname, Type)$, $supp(Sname, Saddr)$, $sales(Part, Supplier, Customer)$, and $cust(Cname, Caddr)$. Query Q asks for the types of parts bought by customers who have the same address as some supplier. Query Q' asks for types of parts sold by suppliers such that a customer at the same address buys parts of the same type.

$$\begin{aligned} Q: \ & q(T) :- sales(P,S,C) \ \& \ part(P,T) \ \& cust(C,A) \ \& \ supp(S',A) \\ Q': \ & q(T) :- sales(P,S,C) \ \& \ part(P,T) \ \& \ part(P',T) \ \& \\ & \quad sales(P',S',C') \ \& \ cust(C,A) \ \& \ supp(S',A) \end{aligned}$$

In Section 4, we present a polynomial-time algorithm to verify that $Q' \subseteq Q$.

Suppose we have the materialized views V_1, V_2, and V_3 shown below. View V_1 relates customers with the types of parts they buy, V_2 gives the address of each customer who buys some part, and view V_3 gives the address of each supplier.

$$\begin{aligned} V_1: v_1(C_1,T_1) \ &:- sales(P_1,S_1,C_1) \ \& \ cust(C_1,A_1) \ \& \ part(P_1,T_1) \\ V_2: v_2(C_2,A_2) &:- sales(P_2,S_2,C_2) \ \& \ cust(C_2,A_2) \\ V_3: v_3(S_3,A_3) \ &:- supp(S_3,A_3) \end{aligned}$$

The query Q can be equivalently rewritten using the views as follows:

$$C: q(T) :- v_1(C,T) \ \& \ v_2(C,A) \ \& \ v_3(S',A)$$

We can save a join by using the materialized views to answer the query. □

To test whether C is an equivalent rewriting of Q, we construct the *expansion* E of C by replacing each view predicate in the body of C by its definition, using different local variables for the expansion of each view predicate. It is easily seen that $C' \equiv E$. Since Q and E are defined over the same sets of database relations, we can use containment mappings between them to test their equivalence and containment, and hence the equivalence and containment of Q and C. For example, the expansion of C in Example 1 is

$$\begin{aligned} E: q(T) :- & \ sales(P_1,S_1,C) \ \& \ cust(C,A_1) \ \& \ part(P_1,T) \ \& \\ & \ sales(P_2,S_2,C) \ \& \ cust(C,A) \ \& \ supp(S',A) \end{aligned}$$

It can easily be verified that $E \equiv Q$, and so $C \equiv Q$.

The problem of finding an equivalent rewriting of a query Q using a set of views $\mathcal{V}$ was shown to be NP-complete by Levy et al. [LMSS95]. They show that the problem remains NP-complete even when the query and the views contain no

repeated predicates. In this paper we present a polynomial-time algorithm for the equivalent rewriting problem, provided the query satisfies certain conditions to be defined in Section 3.

We now define some terminology. We use *argument* to mean either a variable or a constant that appears in a query, and *query term* to mean either an argument or a query subgoal. A *variable mapping* is a function that maps a set of variables to a set of arguments. A set of variable mappings $\phi_1, \ldots, \phi_n$, whose domains are different but perhaps overlapping, are said to be *consistent* if there do not exist variable A and integers i and j such that $\phi_i(A) \neq \phi_j(A)$. If $\phi_1, \ldots, \phi_n$ are consistent, we define their *union mapping,* ϕ, to be the variable mapping whose domain is the union of the domains of $\phi_1, \ldots, \phi_n$, and $\phi(A) = B$ if there exists some i, $1 \leq i \leq n$ such that $\phi_i(A) = B$. Variable mappings are defined to be the identity mapping on predicate symbols and constants, and so we can apply a variable mapping to any query term with the obvious meaning. If $\bar{A}$ is a tuple of arguments in the domain of a variable mapping ϕ, then $\phi(\bar{A})$ is a tuple over the set of attributes $\bar{A}$, where the value of each attribute is its image under ϕ.

A *partial mapping* ϕ from query Q to Q' is a variable mapping that maps some subset of the variables of Q to variables of Q', such that if X_i is the ith head argument of Q and $\phi(X_i) = Y_i$, then Y_i is the ith head argument of Q'. A *containment mapping* from Q to Q' is therefore a partial mapping whose domain is the set of all variables of Q.

3 Acyclicity, Query Width, and Treewidth

3.1 Acyclic Queries

It is often profitable to represent a conjunctive query as a hypergraph. The nodes of the hypergraph are the constants and variables in the query. There is one hyperedge corresponding to each query subgoal, that includes the variables and constants occuring as arguments in that subgoal. Figure 1 shows the hypergraphs corresponding to queries Q and Q' of Example 1.

In the rest of this paper, we restrict ourselves for simplicity of exposition to queries whose hypergraphs are connected. However, our results generalize in a straightforward manner to queries with disconnected hypergraphs.

Let E and F be hyperedges of hypergraph $\mathcal{G}$ such that the nodes in $E - F$ are *unique to* E; that is, they appear in no other hyperedge of $\mathcal{G}$. Then we call E an *ear*, the removal of E from $\mathcal{G}$ *ear removal,* and say that "E is removed in favor of F." The *GYO-reduction* of a hypergraph [Gra79, YO79] is obtained by removing ears until no further ear removals are possible. A hypergraph is *acyclic* if its GYO-reduction is the empty hypergraph; otherwise it is *cyclic.*

A query is cyclic (acyclic) if its hypergraph is cyclic (acyclic). If Q is an acyclic query, an *elimination tree* of Q is a rooted tree constructed as follows. Choose some sequence of ear removals for the hypergraph of Q. The tree has a node for each subgoal of Q, and E is a child of F in the tree whenever the hyperedge

corresponding to E is eliminated in favor of the hyperedge corresponding to F in the chosen ear removal sequence.

Example 2. The hypergraph of query Q (Figure 1(a)) is acyclic. Figure 2 shows one possible elimination tree for this hypergraph. The hypergraph of Q' (Figure 1(b)) is cyclic, because it contains no ear and is its own GYO-reduction.

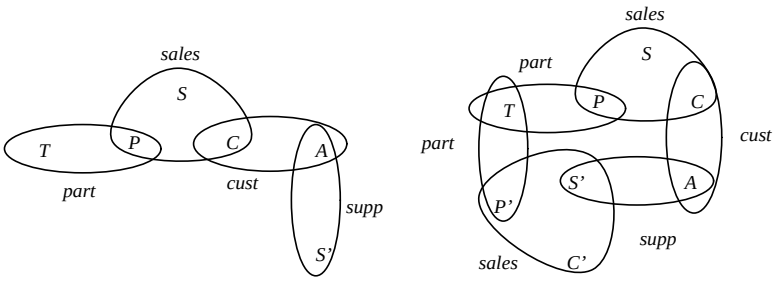

(a) Hypergraph for Q (b) Hypergraph for Q'

Fig. 1. Hypergraphs for the queries in Example 1.

Suppose Q is an acyclic query, and T is an elimination tree for Q. An important observation that follows from the definition of the elimination tree is that for any argument X of Q, the subgoals that mention X form a connected subtree of T. It is this "connectedness property" of acyclic queries that enables a polynomial-time query containment algorithm for such queries (Section 4).

Our algorithms for acyclic queries assume the elimination tree as an input. Tarjan and Yannakakis [TY84] present a simple linear-time algorithm that tests whether Q is acyclic and if so, constructs an elimination tree for it.

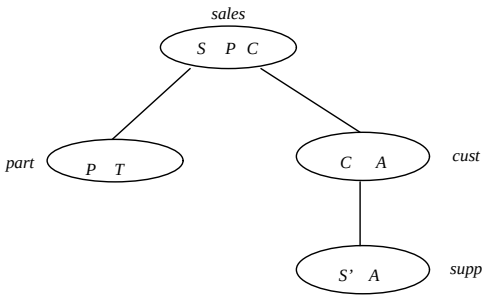

Fig. 2. Elimination tree for hypergraph in Figure 1(a).

3.2 The Width of a Query

It is natural to ask whether we can generalize the notion of query acyclicity in some way. Ideally, we would like to classify queries according to some parameter k that measures their "degree of cyclicity," such that we can design query containment algorithms whose complexity increases with k. In this section we present such a measure, which we call the *query width*.

A *query decomposition* of Q is a tree $T = (I, F)$, with a set $X(i)$ of subgoals and arguments associated with each vertex $i \in I$, such that the following conditions are satisfied:

- For each subgoal s of Q, there is an $i \in I$ such that $s \in X(i)$.
- For each subgoal s of Q, the set $\{i \in I \mid s \in X(i)\}$ induces a (connected) subtree of T.
- For each argument A of Q, the set

  $$\{i \in I \mid A \in X(i)\} \cup \{i \in I \mid A \text{ appears in a subgoal } s \text{ such that } s \in X(i)\}$$

  induces a (connected) subtree of T.

The *width* of the query decomposition is $\max_{i \in I} |X(i)|$. The *query width* of Q is the minimum width over all its query decompositions.

Example 3. Suppose *red* and *blue* are relations that represent the set of red and blue arcs in a directed graph G. Queries Q_2 and Q_3 below ask for blue arcs in subgraphs of G that satisfy certain properties.

$$Q_2 : q_2(A, B) :- blue(A, B) \ \&\ red(B, C) \ \&red(C, D) \ \&\ red(D, B)$$
$$Q_3 : q_3(X, Y) :- blue(X, Y) \ \&\ red(Y, Z) \ \&red(Z, Y) \ \&\ red(Z, Z)$$

It can be verified that Q_2 is cyclic. Figure 3(a) shows a query decomposition (of width 2) of Q_2; it can be shown that the query width of Q_2 is in fact 2. Section 4 presents an efficient algorithm to test whether $Q_3 \subseteq Q_2$. □

The elimination tree of an acyclic query (with one subgoal at each node) is a query decomposition of width 1, and acyclic queries have width 1. Moreover, the query decomposition with the smallest number of nodes for a query of width 1 is also an elimination tree. The following proposition summarizes these observations.

Proposition 1. *A query is acyclic if and only if its query width is 1.*

3.3 Treewidth

Computing the query width of a query Q is NP-complete. Our algorithms assume that given a query Q and some constant k, we can determine efficiently whether its query width is bounded by k and if so, construct a query decomposition of width no more than k. It is open whether there is a polynomial-time algorithm

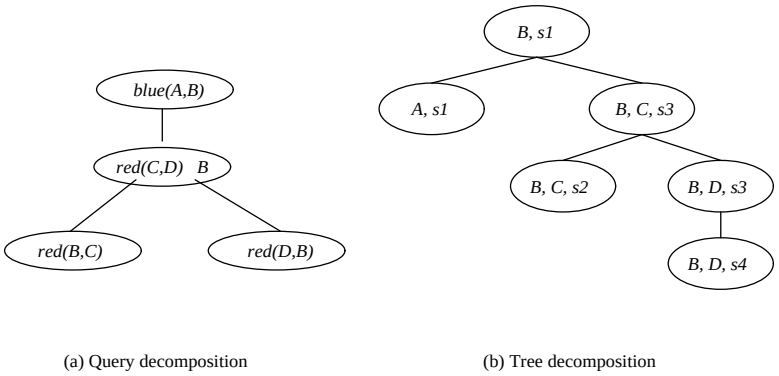

(a) Query decomposition (b) Tree decomposition

Fig. 3. A query decomposition and a tree decomposition for query Q_2 in Example 3

for the above problem. However, we can obtain an upper bound on the query width by using the closely related notion of *treewidth*.

The *incidence graph* $G_Q = (V, E)$ of query Q has a vertex for each argument and for each subgoal of Q. There is an edge between an argument X and a subgoal s whenever X occurs in s. Figure 4 shows the incidence graph of query Q_2 from Example 1, where we use s_1, s_2, s_3, s_4 to denote the subgoals of Q_2 in order from left to right.

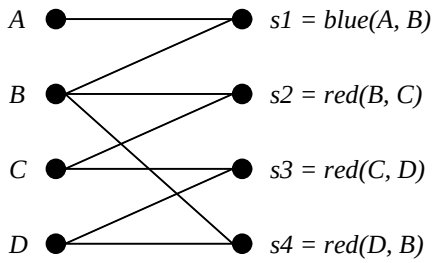

Fig. 4. The incidence graph for the query Q_2.

A *tree decomposition* of a graph $G = (V, E)$ is a tree $T = (I, F)$, with a set $X(i) \subseteq V$ associated with each vertex $i \in I$, such that the following conditions are satisfied:

- For each $v \in V$, there is an $i \in I$ with $v \in X(i)$.
- For all edges $(v, w) \in E$, there is an $i \in I$ with $v, w \in X(i)$.
- For each $v \in V$, the set $\{i \in I \mid v \in X(i)\}$ induces a (connected) subtree of T.

The *width* of the tree decomposition is $\max_{i \in I} |X(i)| - 1$. The *treewidth* of G is

the minimum width over all its tree decompositions. [3] Figure 3(b) shows a tree decomposition (width 2) of the incidence of graph in Figure 4.

The *treewidth of a query* is the treewidth of its incidence graph. We can show easily that every tree decomposition of the incidence graph of Q is also a query decomposition of Q. (For example, Figure 3(b) is another query decomposition for query Q_2.) Therefore, the query width of a query Q is certainly not more than one greater than its treewidth (due to the -1 in the definition of treewidth). In fact, we can show the following result, where $tw(Q)$ is the treewidth of Q and $qw(Q)$ is the query width of Q.

Proposition 2. *For any query Q, $tw(Q)/a \leq qw(Q) \leq tw(Q)$, where a is the maximum predicate arity in Q.*

We assume that all predicate arities are bounded by some constant in this paper, and so Proposition 2 implies that the treewidth approximates the query width to within a constant factor. Proposition 2 is useful because treewidth is an extensively studied concept in graph theory. Bodlaender [Bod93] presents an efficient algorithm to determine, for a given k, whether the treewidth of a graph is bounded by k. The running time of the algorithm is exponential in k but linear in the size of the graph.

Theorem 3 [Bod93]. *For all $k \in N$, there exists a linear time algorithm that tests whether a given graph $G = (V, E)$ has treewidth at most k, and if so, outputs a tree decomposition of G with treewidth at most k which has at most $|V| - k$ nodes.*

4 Algorithms for Query Containment

Section 4.1 presents a polyomial-time algorithm to test whether $Q' \subseteq Q$, when Q is an acyclic query (there are no restrictions on Q'). Our approach is to construct partial mappings from Q to Q' and successively merge partial mappings until we either can no longer merge mappings or have found a containment mapping from Q to Q'. The connectedness property allows us to merge partial mappings in polynomial time. In Section 4.2 we extend the algorithm to work with query decompositions of arbitrary width for Q. Given a decomposition of width k, the algorithm runs in time polynomial in n^k where n is the sum of the sizes of Q and Q'.

4.1 Containment Algorithm for Acyclic Queries

Let $T = (I, F)$ be an elimination tree for Q, and let s_i be the subgoal of Q corresponding to node $i \in I$. The algorithm maintains a relation $Map_i(\bar{A}_i)$ at each node i. The attributes $\bar{A}_i$ of the relation are the arguments of s_i (repeated arguments appear only once in $\bar{A}_i$). We use S_i to denote the set of subgoals in

[3] The -1 in the definition of width ensures that trees have treewidth 1.

the subtree rooted at i. In the algorithm below, $\ltimes$ denotes the natural semijoin operator.

Algorithm AcyclicContainment

1. Initialize the relations as follows. For each partial mapping ϕ from Q to Q' that maps s_i to some subgoal of Q', the tuple $\phi(\bar{A}_i)$ is in $Map_i(\bar{A}_i)$.
2. Process tree nodes bottom-up as follows. Suppose i is a node of T all of whose children have been processed. For each child j of i:

$$Map_i(\bar{A}_i) := Map_i(\bar{A}_i) \ltimes Map_j(\bar{A}_j)$$

3. $Q' \subseteq Q$ if and only if the relation at the root of T is nonempty.

Example 4. Figure 5 shows the relations created by the algorithm when the elimination tree in Figure 2 is used to determine whether $Q' \subseteq Q$ in Example 1. The attribute names at each node give the relation schema for that node. Step 1 creates the relations in dashed boxes and Step 2 results in the relations in the solid boxes. The numbers at the nodes show the order in which they are processed in Step 2. We conclude that $Q' \subseteq Q$ since the relation at the root is nonempty after Step 2. □

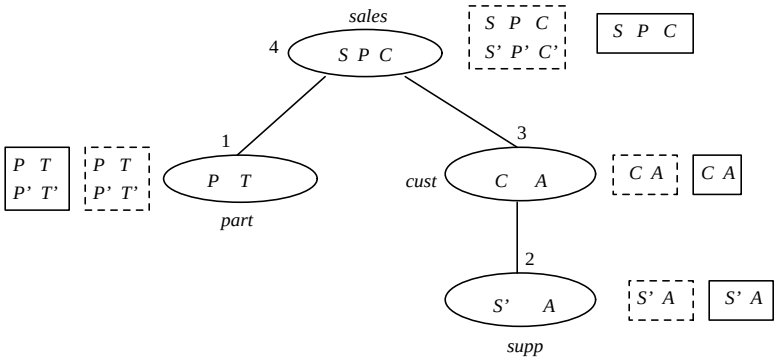

Fig. 5. Running Algorithm AcyclicContainment on Example 1 using the tree in Figure 2.

Lemma 4. *Algorithm AcyclicContainment correctly determines whether* $Q' \subseteq Q$.

Proof: We use induction on the number of nodes processed, with the following induction hypothesis: After node i is processed, tuple $t \in Map_i(\bar{A}_i)$ if and only if there is a partial mapping ϕ from Q to Q' whose domain is the set of subgoals S_i, such that $\phi(\bar{A}_i) = t$. Thus, when the root of T has been processed, the relation

at the root is nonempty if and only if there is a containment mapping from Q to Q'.

The induction hypothesis holds for the leaves, because of the way the relations are initialized in Step 1. For the induction, assume that we have processed internal node i with children $j_1, \ldots, j_r$, and let $t \in Map_i$. Step 1 ensures that there is a partial mapping ψ from Q to Q' with domain s_i such that $\psi(\bar{A}_i) = t$. Step 2 assures us that for each j_k, there is a tuple $t_k \in Map_{j_k}$ that agrees with t on the attributes in $\bar{A}_i \cap \bar{A}_{j_k}$. By the induction hypothesis, there is a partial mapping ϕ_k from Q to Q' with domain S_{j_k} such that $\phi_k(\bar{A}_{j_k}) = t_k$.

The mappings ψ and ϕ_k are consistent. To see this, suppose X is a variable in the domains of both ϕ_k and ψ. By the connectedness property, $X \in \bar{A}_{j_k}$ and $X \in \bar{A}_i$. Since t and t_k agree on their common attributes, ψ and ϕ_k agree on X. Moreover, the mappings $\psi, \phi_1, \ldots, \phi_r$ are also consistent. To see this, suppose X is a variable in the domains of ϕ_k and ϕ_l. By the connectedness property, $X \in \bar{A}_{j_k}$, $X \in \bar{A}_{j_l}$, and $X \in \bar{A}_i$. Therefore, ψ, ϕ_k and ϕ_l agree on X. Let ϕ be the partial mapping from Q to Q' that is the union of $\psi, \phi_1, \ldots, \phi_r$. Then ϕ satisfies the conditions of the induction hypothesis.

Conversely, let ϕ be a partial mapping with domain S_i. Let $\phi_1, \ldots, \phi_r$ be the projection of ϕ on the sets of variables in $S_{j_1}, \ldots, S_{j_r}$, and let ψ be the projection of ϕ on the variables in s_i. By the induction hypothesis, there is a tuple $t_k \in Map_{j_k}$, $k = 1, \ldots, r$, such that $\phi_k(\bar{A}_{j_k}) = t_k$. After Step 1, there is a tuple $t \in Map_i$ such that $\psi(\bar{A}_i) = t$. Since t agrees with the tuples $t_1, \ldots, t_r$ on all common attributes, it will remain in Map_i after the sequence of joins in Step 2. □

Theorem 5. *Algorithm AcyclicContainment determines whether $Q' \subseteq Q$ in time $O(N_Q N_{Q'} \log N_{Q'})$ using space $O(N_Q N_{Q'})$, where N_Q and $N_{Q'}$ are the sizes of Q and Q', respectively.*

Proof (Sketch). Correctness follows from Lemma 4. The space and time complexities follow from the observation that the cardinality of each relation Map_i is bounded by the number of subgoals in Q', and the number of such relations is exactly the number of subgoals in Q. The semijoins in Step 2 have to implemented as sort-merge semijoins to achieve the time complexity in the theorem. □

Corollary 6. *Given an acyclic query Q, there is an algorithm to minimize Q in time $O(N_Q^3 \log N_Q)$ using space $O(N_Q^2)$, where N_Q is the size of Q.*

4.2 Generalizing the Algorithm

We now generalize Algorithm AcyclicContainment to test whether $Q' \subseteq Q$, where we are given a query decomposition $T = (I, F)$ of width k for Q. Let $X(i)$ be the set of terms associated with node i of T. As before, we associate a relation $Map_i(\bar{A}_i)$ with node i, where $\bar{A}_i$ is constructed as follows:

1. For each argument $A \in X(i)$, $A \in \bar{A}_i$.
2. For each subgoal $s \in X(i)$, $s \in \bar{A}_i$.
3. For each subgoal $s \in X(i)$, and each argument A that occurs in s, $A \in \bar{A}_i$.

We call attributes of type 1 and 2 *independent attributes* and attributes of type 3 *dependent attributes* (the reason for the nomenclature will become apparent later). Let $S(i)$ denote the set of terms associated with the nodes in the subtree of T rooted at i. The algorithm for query containment is now identical to Algorithm AcyclicContainment except for the initialization step.

Algorithm QueryContainment

1. (Initialize.) For each partial mapping ϕ from Q to Q' that maps all the terms in $X(i)$, include the tuple $\phi(\bar{A}_i)$ in Map_i.
2. (Propagate bottom-up.) Process tree nodes bottom-up as follows. Suppose i is a node of T all of whose children have been processed. For each child j of i:
$$Map_i(\bar{A}_i) := Map_i(\bar{A}_i) \ltimes Map_j(\bar{A}_j)$$
3. $Q' \subseteq Q$ if and only if the relation at the root of T is nonempty.

Example 5. Figure 6 shows the relations created by the algorithm when testing whether $Q_3 \subseteq Q_2$ in Example 3. We use the query decomposition in Figure 3(a) for Q_2. (We could have used the decomposition in Figure 3(b) as well.) In the figure $s_1, \ldots, s_4$ are the subgoals of Q_2 from left to right and $c_1, \ldots, c_4$ are the subgoals of Q_3 from left to right. The relation schema is show at each node, with independent attributes followed by dependent attributes. Step 1 creates the relations in the dashed boxes, and Step 2 results in the relations in the solid boxes. We conclude the $Q_3 \subseteq Q_2$ since the relation at the root is nonempty after Step 2. □

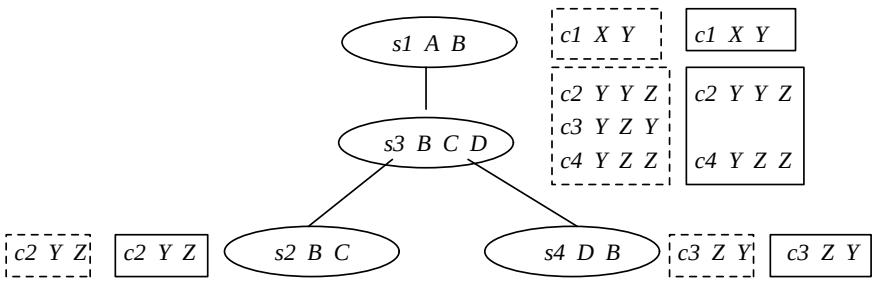

Fig. 6. Running Algorithm QueryContainment

Theorem 7. *Given a tree decomposition of width k for Q, algorithm QueryContainment determines whether $Q' \subseteq Q$ in time $O(k^2 N_Q N_{Q'}^k \log N_{Q'})$ using space $O(N_Q N_{Q'}^k)$, where N_Q and $N_{Q'}$ are the sizes of Q and Q', respectively.*

5 Algorithms for Answering Queries using Views

Let $p_1, \ldots, p_n$ be distinct predicates, and let Q be the query

$$Q : q(\bar{X}) \ :- \ p_1(\bar{X_1}) \ \& \ \ldots \ \& \ p_n(\bar{X_n})$$

We wish to determine whether there is an equivalent rewriting of Q using an arbitrary set of views $\mathcal{V}$.

Call $V \in \mathcal{V}$ an *interesting view* if there is a variable mapping ϕ that maps the subgoals of V into subgoals of Q (ϕ need not map the head of V to the head of Q). The mapping ϕ is unique for a given interesting view V: since Q has no repeated predicates, there is a unique destination for each subgoal of V. Let $V_1, \ldots, V_m$ be the interesting views in $\mathcal{V}$ with associated mappings $\phi_1, \ldots, \phi_m$, and let $v_i(\bar{Y_i})$ be the head of view V_i. Then the *canonical rewriting of Q using $\mathcal{V}$* is

$$C : q(\bar{X}) \ :- \ v_1(\phi_1(\bar{Y_1})) \ \& \ \ldots \ \& \ v_m(\phi_m(\bar{Y_m}))$$

Example 6. In Example 1, all three views V_1, V_2, and V_3 are interesting for query Q. The rewriting C is the canonical rewriting in this case. □

Lemma 8. *There is an equivalent rewriting of Q using $\mathcal{V}$ if and only if $C' \subseteq Q$, where C' is the expansion of the canonical rewriting C.*

Proof. (If.) Suppose $C' \subseteq Q$. We will show that $Q \subseteq C'$; it follows that $Q \equiv C'$ and C is an equivalent rewriting of Q using $\mathcal{V}$. Let S_i denote the set of subgoals contributed to C' by the expansion of $v_i(\phi_i(\bar{Y_i}))$. There is a unique variable mapping ψ_i that maps the subgoals in S_i to subgoals of Q, corresponding to the unique mapping ϕ_i from V_i to Q. Moreover ψ_i is the identity on the variables in $\phi_i(\bar{Y_i})$. For any i and j, the only variables common to the subgoals in S_i and S_j occur in $\phi_i(\bar{Y_i})$ and $\phi_j(\bar{Y_j})$. Since ψ_i and ψ_j are both the identity mapping on these variables, ψ_i and ψ_j are consistent. It follows that $\psi_1, \ldots, \psi_m$ is a consistent set of mappings. Let ψ be the union mapping of $\psi_1, \ldots, \psi_m$. Then ψ is a containment mapping from C' to Q, showing that $Q \subseteq C'$.

(Only if.) Conversely, suppose the following is an equivalent rewriting of Q using $\mathcal{V}$:

$$R : q(\bar{X}) \ :- \ u_1(\bar{Z_1}) \ \& \ \ldots \ \& \ u_l(\bar{Z_l})$$

Let R' be the expansion of R. Since $Q \equiv R'$, there is a containment mapping τ from Q to R' and a containment mapping ρ from R' to Q. The projection of ρ on to the expansion of u_i maps the subgoals of the corresponding view to subgoals to Q, and so u_i is an interesting view for $i = 1, \ldots, l$. Let us assume without loss of generality that u_i is identical to v_i, $i = 1, \ldots, l$.

Consider the query T defined as follows:

$$T : q(\bar{X}) \ :- \ v_1(\rho(\bar{Z_1})) \ \& \ \ldots \ \& \ v_l(\rho(\bar{Z_l}))$$

Let T' be the expansion of T. We have constructed T from R in a such a manner as to ensure that T is an equivalent rewriting of Q and $T' \equiv Q$. Since there is a unique mapping ϕ_i from each interesting view V_i to Q, it must be the case

that $\phi_i(\bar{Y}_i) = \rho(\bar{Z}_i)$. Thus, each subgoal in T is also a subgoal of the canonical rewriting C, and therefore $C \subseteq T$ and $C' \subseteq T'$. Since $T' \equiv Q$, we have shown that $C' \subseteq Q$. □

Theorem 9. *Let Q be a query without repeated predicates and $\mathcal{V}$ a set of views. Given a query decomposition of width k for Q, there is an algorithm to test whether there is an equivalent rewriting of Q using $\mathcal{V}$ in time $O(k^2 N_Q |\mathcal{V}|^k \log |\mathcal{V}|)$ and space $O(N_Q |\mathcal{V}|^k)$, where N_Q is the size of Q and $|\mathcal{V}|$ is the size of $\mathcal{V}$.*

Proof. Follows from Lemma 8 and Theorem 7. □

The size of the canonical rewriting depends on the number of interesting views and is independent of the size of Q. We can modify the query containment algorithms of Section 4 slightly to obtain a rewriting that uses no more subgoals than the query. With each tuple t in a relation at a node of the query decomposition, we must store also the names of the views whose subgoals participate in the partial mapping corresponding to t. After Step 2 of the containment algorithm, we traverse the tree top-down and choose a set of consistent tuples, one from each node of the tree. We can then drop all but the views corresponding to the chosen tuples from the canonical rewriting C to obtain a shorter rewriting D such that $D \equiv Q$.

Theorem 10. *Let Q be a query without repeated predicates and $\mathcal{V}$ a set of views, such that Q has n subgoals. Given a query decomposition of width k for Q, there is an algorithm that tests whether there is an equivalent rewriting of Q using $\mathcal{V}$, and if so, produces a rewriting of Q using $\mathcal{V}$ that has no more than n subgoals. The algorithm runs in time $O(k^2 N_Q |\mathcal{V}|^k \log |\mathcal{V}|)$ and uses space $O(N_Q |\mathcal{V}|^k)$, where N_Q is the size of Q and $|\mathcal{V}|$ is the size of $\mathcal{V}$.*

6 Related Work

Aho et al. [ASU79a, ASU79b] gave polynomial-time minimization and equivalence algorithms for conjunctive queries corresponding to *simple tableaux.* Their results were extended by Johnson and Klug [JK83] to the class of *fanout-free queries.* The classes of fanout-free queries and queries with simple tableaux are incomparable to the classes of queries considered in this paper. The algorithms of Aho et al. and Johnson and Klug also differ from ours in that they cannot be generalized to test query containment instead of query equivalence. Recently, Qian [Qia96] showed independently that acyclic queries admit polynomial time algorithms for containment and minimization. Our work treats acyclic queries as a special case of queries with width 1, and so Qian's algorithm for query containment falls out as a special case of ours.

Yang and Larson [LY85, YL87] considered the problem of finding rewritings for SPJ queries using SPJ views. In their analysis, they considered what amounts to 1-1 mappings from the views to the query, and therefore their algorithm can miss some rewritings. The problem of finding equivalent rewritings was studied formally by Levy et al. [LMSS95], who showed the problem to be NP-complete.

Rajaraman et al. [RSU95] extended the results of Levy et al. to queries and views with *binding patterns*. Chaudhuri et al. [CKPS95] considered the problem of finding rewritings for SPJ queries and views, when the queries and views use bag semantics instead of the usual set semantics. They also suggest a way to extend a traditional relational query processor to choose between different rewritings based on their costs, a question we do not address in this paper. Qian [Qia96] presents a polynomial-time algorithm that, given an acyclic query Q and a set of views $\mathcal{V}$, determines whether there is a rewriting using $\mathcal{V}$ that is contained in (but not necessarily equivalent to) Q. Levy et al. [LRU96] study the problem of finding an equivalent rewriting when the set of views is possibly infinite, albeit encoded in some finite fashion. The Information Manifold system [LRO96] implements heuristics that speed up the search for a rewriting of a query by eliminating irrelevant views.

Acyclic queries and acyclic database schemes have been studied extensively because their structural properties permit efficient algorithms. Several algorithms for acyclic database schemes were given by Yannakakis [Yan81]. Our query containment algorithm for acyclic queries is closely related to the one in [Yan81] for computing projections of acyclic joins. Treewidth is extensively studied in the graph-theoretic literature, and several intractable problems admit polynomial-time algorithms on graphs of constant treewidth [Bod93].

7 Open Problems

Our work raises some interesting open problems. While we obtained bounds on the query width of Q based on the treewidth of its incidence graph, it would be useful to have an efficient algorithm that produces query decompositions of small query width, analogous to the algorithm of Bodlaender [Bod93] for decompositions of small treewidth. The connectedness property of acyclic queries leads to several efficient algorithms [Yan81]; it may be possible to generalize many of these algorithms to queries of small query width. Finally, we would like to extend our results to queries with binding patterns [RSU95] and queries with built-in predicates.

Acknowledgements

The inspiration for considering the treewidth of the incidence graph came from the work of Khanna and Motwani [KM96]. We also thank Jeff Ullman for comments on earlier versions of this paper.

References

[ASU79a] A.V. Aho, Y. Sagiv, and J.D. Ullman. Efficient optimization of a class of relational expressions. *ACM Transactions on Database Systems*, 4(4):435–454, December 1979.

[ASU79b] A.V. Aho, Y. Sagiv, and J.D. Ullman. Equivalence of relational expressions. *SIAM Journal on Computing*, 8(2):218–246, May 1979.

[Bod93] H.L. Bodlaender. A linear time algorithm for finding tree-decompositions of small treewidth. In *Proceedings of the 25th ACM Symposium on the Theory of Computing*, pages 226–234, 1993.

[CKPS95] S. Chaudhuri, R. Krishnamurthy, S. Potamianos, and K. Shim. Optimizing queries with materialized views. In *Proceedings of the Eleventh International Conference on Data Engineering*, pages 190–200, 1995.

[CM77] A.K. Chandra and P.M. Merlin. Optimal implementation of conjunctive queries in relational databases. In *Proceedings of the Ninth ACM Symposium on Theory of Computing*, pages 77–90, 1977.

[Gra79] M.H. Graham. On the universal relation. Technical report, University of Toronto, Ontario, Canada, 1979.

[JK83] D.S. Johnson and A. Klug. Optimizing conjunctive queries that contain untyped variables. *SIAM Journal on Computing*, 12(4):616–640, November 1983.

[KM96] S. Khanna and R. Motwani. Towards a syntactic characterization of PTAS. In *Proceedings of the 28th ACM Symposium on the Theory of Computing*, 1996.

[LMSS95] A.Y. Levy, A.O. Mendelzon, Y. Sagiv, and D. Srivastava. Answering queries using views. In *Proceedings of the Fourteenth ACM Symposium on Principles of Database Systems*, pages 95–104, 1995.

[LRO96] A.Y. Levy, A. Rajaraman, and J.J. Ordille. Querying heterogeneous information sources using source descriptions. In *Proceedings of the 22nd International Conference on Very Large Data Bases*, 1996.

[LRU96] A.Y. Levy, A. Rajaraman, and J.D. Ullman. Answering queries using limited external query processors. In *Proceedings of the Fifteenth ACM Symposium on Principles of Database Systems*, pages 227–237, 1996.

[LY85] P.A. Larson and H.Z. Yang. Computing queries from derived relations. In *Proceedings of the Eleventh International Conference on Very Large Data Bases*, pages 259–269, 1985.

[Qia96] X. Qian. Query folding. In *Proceedings of the Twelfth International Conference on Data Engineering*, 1996.

[RSU95] A. Rajaraman, Y. Sagiv, and J.D. Ullman. Answering queries using templates with binding patterns. In *Proceedings of the Fourteenth ACM Symposium on Principles of Database Systems*, pages 105–112, 1995.

[TY84] R.E. Tarjan and M. Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing*, 13(3):566–579, 1984.

[Ull89] J.D. Ullman. *Principles of Database and Knowledge-Base Systems, Volume II: The New Technologies*. Computer Science Press, Rockville, MD, 1989.

[Yan81] M. Yannakakis. Algorithms for acyclic database schemes. In *Proceedings of the Seventh International Conference on Very Large Data Bases*, pages 82–94, 1981.

[YL87] H.Z. Yang and P.A. Larson. Query transformation for PSJ-queries. In *Proceedings of the Thirteenth International Conference on Very Large Data Bases*, pages 245–254, 1987.

[YO79] C.T. Yu and M.Z. Ozsoyoglu. An algorithm for tree-query membership of a distributed query. In *Proceedings of IEEE COMPSAC*, pages 306–312, 1979.