# Using Semi-Joins to Solve Relational Queries 

PHILIP A. BERNSTEIN AND DAH-MING W. CHIU

Harvard University; Cambridge, Massachusetts


#### Abstract

The semi-join is a relational algebraic operation that selects a set of tuples in one relation that match one or more tuples of another relation on the joining domains. Semi-joins have been used as a basic ingredient in query processing strategies for a number of hardware and software database systems. However, not all queries can be solved entirely using semi-joins. In this paper the exact class of relational queries that can be solved using semi-joins is shown. It is also shown that queries outside of this class may not even be partially solvable using "short" semi-join programs. In addition, a linear-time membership test for this class is presented.


KEY words and phrases: semi-join, relational database, relational query processing
CR CATEGORIES: $\quad 3.74,4.33,5.25$

## 1. Introduction

The utility of very high level database query languages based on the relational database model, such as QUEL [8] and SEQUEL [4], is predicated on the ability to implement these languages efficiently. Many schemes for interpreting relational query languages in centralized systems, distributed systems, and on special parallel hardware have been devised (e.g., [9, 10, 12-14]). These schemes are essentially strategies for executing a combination of projection, restriction, and join operators from relational algebra to answer the given query [7]. Two of these methods, the distributed retrieval algorithm for SDD-1 [14] and the RAP database machine [9], use a special operator that is a composed join and projection. The properties of this operator, which we call the semi-join and denote by $\propto$, are the subject of this paper.
The semi-join operator takes the join of two relations, $R$ and $S$, and then projects back out on the domains of relation $R .{ }^{1}$ That is, it retrieves those tuples in $R$ that join with some tuple in $S$. Alternatively, one can think of semi-join as a generalization of restriction; it restricts $R$ by values that appear in $S$ 's join domain.
The value of the semi-join operator is that it can reduce the amount of effort required to do a later expensive join, while the semi-join itself is often quite cheap.

[^0]Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear. and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.
This work was supported in part by the National Science Foundation under Grants MCS-77-05314 and ENG-76-11824, by the Joint Services Program under Contract N00014-75-C-0648, and by the Defense Advanced Research Projects Agency of the Department of Defense, monitored by the Naval Electronic Systems Command under Contract N0039-77-C-0074.
Authors' addresses: P. A. Bernstein, Aiken Computation Laboratory, Harvard University, Cambridge, MA 02138; D. W. Chiu, Pierce Hall, Harvard University, Cambridge, MA 02138.
(1) 1981 ACM 0004-5411/81/0100-0025 \$00.75

Journal of the Association for Computing Machinery. Vol. 28. No. I, January 1981. pp. 25-40.

For example, suppose a query requires that the join of $R$ and $S$ be constructed. Instead of computing the join directly, one can first reduce the size of $R$ and $S$ by semi-joins (i.e., $R \propto S$ and $S \ltimes R$ ) and then construct the join of the resulting relations. No information is lost by the preliminary semi-joins. For the semi-joins to be performed, only the projection of the joining columns need be sent. If the size of these projections is small relative to the amount by which $R$ and $S$ are reduced, then the preliminary semi-joins will be profitable.

To reduce the cost of processing joins, semi-joins play a pivotal role in the query processing algorithm of SDD-1, a prototype distributed database system. The principal problem in evaluating a relational query on a distributed database is that two relations that must be joined may reside at different sites. For the join to be performed, one relation must be shipped to the site of the other. Communication is the dominant cost in a distributed database system, so minimizing the amount of data to be shipped is of prime importance [11]. Since the join must eventually be performed, one can only hope to limit communication, not eliminate it. One tactic for limiting communication is to perform first as much local processing of the query as possible; for example, all restrictions should be applied early. Another tactic is to use semi-joins. To join $R$ and $S$ at different sites, one can ship the projection of $R$ on its joining column to $S$ 's site and use a semi-join to reduce $S$ by $R$ before shipping $S$ to $R$ 's site. This will be a profitable tactic whenever the projection of $R$ on its joining column is smaller than the amount by which $S$ is reduced by the semi-join. The SDD-1 algorithm uses a heuristic hill-climbing strategy that tries to apply as many such semi-joins as are profitable.

Semi-joins are also important for the RAP machine, a hardware device tailored for relational query processing. RAP is not capable of performing joins itself and must therefore ship data to a conventional CPU to perform them. To reduce the amount of data shipped, the RAP designers have provided a hardware semi-join instruction that they recommend using to partially evaluate queries before performing joins on an external processor.

Since the database is of finite size, there are limits to how much the database can be reduced relative to a given query using semi-joins. Knowing which sequences of semi-joins can fully reduce the size of the database is important in selecting optimal retrieval strategies for RAP, SDD-1, and other systems as well. In this paper we examine the question of which semi-join sequences are efficacious. We show that relational calculus queries are naturally partitioned into two broad classes: One class can be fully evaluated using a small number of semi-joins; the other class can be neither fully evaluated using semi-joins alone nor easily reduced using a small number of semi-joins. For the one class, semi-join solution strategies are well behaved; for the other they can be dismally poor. We proceed formally by developing some basic properties of semi-joins and then showing the behavior of semi-joins with respect to these two classes of queries.

## 2. Relations and Semi-Joins

A relation is a subset of the Cartesian product of its domains. We distinguish between relation names, denoted by $\left\{R_{1}, R_{2}, \ldots\right\}$, and their corresponding relations (i.e., sets of tuples), denoted by $\left\{S_{1}, S_{2}, \ldots\right\}$. A database $D$ is defined to be an ordered set of relation names $\mathbf{R}$. A database state is, in turn, a pair $\langle\mathbf{S}$, env〉, where $\mathbf{S}$ is an ordered set of relations and $e n v: \mathbf{R} \rightarrow \mathbf{S}$. For notational convenience we will assume that $\mathbf{R}=\left\{R_{1}, \ldots, R_{n}\right\}, \mathbf{S}=\left\{S_{1}, \ldots, S_{n}\right\}$, and $\operatorname{env}\left(R_{i}\right)=S_{i}$ for $i=1, \ldots, n$.

Note that we use the standard mathematical definition of a relation as a fixed
subset of the Cartesian product of its domains. This differs slightly from some uses in the database literature where relation means a "time-varying subset of the Cartesian product of its domains" that changes as operations are applied (e.g., [6]). Thus, in our model an operation maps one database state into another. In practice, of course, only a single copy of the database is maintained. Update operations modify the copy; retrieval operations produce a temporary database that is the subset of the real database that the query requested. We adopt our model for mathematical simplicity only and do not suggest that entire databases be created and destroyed as a consequence of each operation.

Given two database states $\mathbf{S}^{\prime}$ and $\mathbf{S}^{\prime \prime}$, we define $\mathbf{S}^{\prime} \cup \mathbf{S}^{\prime \prime}=\left\{S_{i}^{\prime} \cup S_{i}^{\prime \prime} \mid i=1, \ldots, n\right\}$ to be their componentwise union. Also, we define a partial order $\leq$ on database states. We write $\mathbf{S}^{\prime} \leq \mathbf{S}^{\prime \prime}$ if $S_{i}^{\prime} \subseteq S_{i}^{\prime \prime}, i=1, \ldots, n$. Similarly, $\mathbf{S}^{\prime}<\mathbf{S}^{\prime \prime}$ if $S_{i}^{\prime} \subsetneq S_{i}^{\prime \prime}$, $i=1, \ldots, n$.

We denote the semi-join operation by the symbol $\propto$. The semi-join of $R_{i}$ on domain $A$ with $R_{j}$ on domain $B$ is defined as a function from database states into database states: $\left\langle R_{i}{ }_{A} \propto_{B} R_{j}\right\rangle(\mathbf{S})=\mathbf{S}^{\prime}$, where $S_{k}^{\prime}=S_{k}$ for $k \neq i$ and $S_{i}^{\prime}=$ $\left\{t_{1} \in S_{i} \mid \exists t_{j} \in S_{j}\right.$ such that $\left.t_{i} . A=t_{j} . B\right\}$. The relational calculus notation $t . X$ means "the value of domain $X$ of tuple $t$." When there is only one semi-join on $R_{i}$ by $R_{j}$, we drop domain references and simply write $R_{i} \propto R_{j}$.

We limit ourselves to semi-joins involving single domains. Our justification is practicality. To perform a semi-join of, say, $R_{i}$ on domains $A$ and $B$ with $R_{j}$ on domains $C$ and $D$ requires projecting $R_{j}$ on a pair of domains $C$ and $D$. The size of such a projection will generally be quite large and therefore the benefit of the semijoin is unlikely to exceed its cost. For this reason, multidomain semi-joins are ignored in SDD-1 [14] and RAP [9]. Still, we do allow multidomain semi-joins in a limited context. If a pair of domains in a relation are always treated as a single composite domain (i.e., neither domain ever participates in a semi-join by itself), then the composite domain can be considered to be atomic and all of our results follow correctly.

Lemma 1. If $S_{1}=S_{1}^{\prime} \cup S_{1}^{\prime \prime}$, then
(a) $\left\langle R_{2} \ltimes R_{1}\right\rangle\left(\left\{S_{1}, S_{2}\right\}\right)=\left\langle R_{2} \ltimes R_{1}\right\rangle\left(\left\{S_{1}^{\prime}, S_{2}\right\}\right) \cup\left\langle R_{2} \ltimes R_{1}\right\rangle\left(\left\{S_{1}^{\prime \prime}, S_{2}\right\}\right) ;$
(b) $\left\langle R_{1} \ltimes R_{2}\right\rangle\left(\left\{S_{1}, S_{2}\right\}\right)=\left\langle R_{1} \ltimes R_{2}\right\rangle\left(\left\{S_{1}^{\prime}, S_{2}\right\}\right) \cup\left\langle R_{1} \ltimes R_{2}\right\rangle\left(\left\{S_{1}^{\prime \prime}, S_{2}\right\}\right)$.

Proof. Follows directly from the definition of $\propto$.
For evaluating a relational database query, in general many semi-joins are required. Hence we define a semi-join program to be a sequence of semi-joins, denoted $\sigma=\left\langle R_{i_{1}} \ltimes R_{i_{2}}, \ldots, R_{i_{2 m-1}} \ltimes R_{i_{2 m}}\right\rangle$, whose meaning is simply the composition of the semi-joins, $\sigma=\left\langle R_{i_{1}} \propto R_{i_{2}}\right\rangle \cdot \cdots \cdot\left\langle R_{i_{2 m-1}} \propto R_{i_{2 m}}\right\rangle$. The following lemma shows that semi-join programs obey a simple monotonicity property.

Lemma 2. If $\mathbf{S}^{\prime} \leq \mathbf{S}^{\prime \prime}$, then $\sigma\left(\mathbf{S}^{\prime}\right) \leq \sigma\left(\mathbf{S}^{\prime \prime}\right)$.
Proof. We proceed by induction on the length of $\sigma$. The basis step goes as follows: Suppose $\sigma=\left\langle R_{2} \ltimes R_{1}\right\rangle$. Then,

$$
\begin{array}{rlrl}
\sigma\left(\left\{S_{1}^{\prime \prime}, S_{2}^{\prime \prime}\right\}\right)= & \sigma\left(\left\{S_{1}^{\prime}, S_{2}^{\prime \prime}\right\}\right) \cup \sigma\left(\left\{S_{1}^{\prime \prime}-S_{1}^{\prime}, S_{2}^{\prime \prime}\right\}\right) & & \text { (by Lemma 1(a)) } \\
= & \sigma\left(\left\{S_{1}^{\prime}, S_{2}^{\prime}\right\}\right) \cup \sigma\left(\left\{S_{1}^{\prime}, S_{2}^{\prime \prime}-S_{2}^{\prime}\right\}\right) & & \\
& \cup \sigma\left(\left\{S_{1}^{\prime \prime}-S_{1}^{\prime}, S_{2}^{\prime \prime}\right\}\right) & \text { (by Lemma 1(b)) }
\end{array}
$$

Thus $\sigma\left(\left\{S_{1}^{\prime}, S_{2}^{\prime}\right\}\right) \leq \sigma\left(\left\{S_{1}^{\prime \prime}, S_{2}^{\prime \prime}\right\}\right)$. Since all relations other than $S_{2}$ remain unchanged by $\sigma, \sigma\left(\mathbf{S}^{\prime}\right) \leq \sigma\left(\mathbf{S}^{\prime \prime}\right)$.

The induction step follows similarly.

## 3. Queries

We are interested in applying semi-joins to evaluate relational database queries. We will express our queries in a notation similar to that of relational calculus [7]. Formally, an equi-join qualification $q$ is a conjunction of clauses of the form ( $R_{i} . X=R_{j} . Y$ ), where $X$ and $Y$ are domains of $R_{i}$ and $R_{j}$, respectively. Each equi-join qualification $q$ defines an equi-join query $Q$ that produces a relation from each database state as follows:

$$
Q(\mathbf{S})=\left\{\left\langle t_{1}, \ldots, t_{n}\right\rangle \in S_{1} \times \ldots \times S_{n} \mid q\left(t_{1}, \ldots, t_{n}\right)\right\}
$$

(Since each $t_{i}$ is a tuple, $\left\langle t_{1}, \ldots, t_{n}\right\rangle$ is a tuple of tuples.) In words, $Q(\mathbf{S})$ is the subset of $S_{1} \times \cdots \times S_{n}$ that satisfies $q$. We say that two queries are equivalent, denoted $Q_{1} \equiv Q_{2}$, if $Q_{1}(\mathbf{S})=Q_{2}(\mathbf{S})$ for all S (this is called "strong equivalence" in [2]).

Notice that equi-join queries do not include one-relation clauses of the form ( $R_{i} . X=$ constant), nor do they include target lists. One-relation clauses were excluded because they are generally evaluated using special techniques before semi-joins are applied. For example, in SDD-1 one-relation clauses correspond to local operations that do not require semi-joins [14]; in RAP one-relation clauses are evaluated in one rotational delay using the MARK operation [9]. In any case, by creating a relation containing a single constant tuple, these one-relation clauses can be considered in equi-join queries.

Target lists were excluded because the semi-join does not have the power to "project out" certain domains. A single projection operation applied to the result of an equi-join query is always sufficient to obtain the full power of target lists. ${ }^{2}$

Our definition of query does allow multidomain joins between two relations by conjoining two clauses. For example, to join $R_{i}$ on domains $W$ and $X$ with $R_{j}$ on domains $Y$ and $Z$, the clauses $\left(\left(R_{i}, W=R_{j} . Y\right) \wedge\left(R_{i} . X=R_{j} . Z\right)\right)$ are used.

We also allow one-relation joins of the form ( $R_{i} . A=R_{i} . B$ ). Like simple onevariable clauses, these one-relation joins are also generally evaluated using special techniques before semi-joins are applied. However, syntactically eliminating them from our class of queries is futile, since a one-relation join can be in the closure of $q$ even if it is not in $q$ itself. For example, $\left(R_{i} . A=R_{j} . C\right)$ and $\left(R_{j} . C=R_{i} . B\right)$ implies ( $R_{i} . A=R_{i} . B$ ). So, although we will want to think of these one-relation joins as executing "for free," we still need them in our definition of query. More will be said on this issue in Section 4.

It will occasionally be convenient to examine the components of $Q(\mathbf{S})$ that come from each of $S_{1}, \ldots, S_{n}$. We introduce projection for this purpose. The notation $Q(\mathbf{S})\left[R_{i}\right]$ denotes $\left\{t_{i} \in S_{i} \mid\left\langle t_{1}, \ldots, t_{i}, \ldots, t_{n}\right\rangle \in Q(\mathbf{S})\right\}$. Similarly, $\mathbf{S}\left[R_{i}\right]$ denotes $S_{i}$.

## 4. Solving Queries by Semi-Join Reduction

4.1 Reductions. The main purpose of semi-joins is to reduce the number of tuples involved in the evaluation of a query. Let $Q$ be a query and $S$ be a database state. We define REDUCTIONS $(Q, \mathbf{S})=\left\{\mathbf{S}^{\prime} \mid \mathbf{S}^{\prime} \leq \mathbf{S}\right.$ and $\left.Q\left(\mathbf{S}^{\prime}\right)=Q(\mathbf{S})\right\}$. A full reduction $\mathbf{S}^{*}$ of $\mathbf{S}$ with respect to $Q$ is an element of REDUCTIONS $(Q, \mathbf{S})$ such that there is no $\mathbf{S}^{\prime}$ in REDUCTIONS $(Q, S)$ with $\mathbf{S}^{\prime}<\mathbf{S}^{*}$. We denote $\mathbf{S}^{*}\left[R_{i}\right]$ by $S_{i}^{*}$.

Lemma 3. For each $Q$ and $\mathbf{S}$ there exists a unique $\mathbf{S}^{*}$ and $\mathbf{S}^{*}=\left\{Q(\mathbf{S})\left[R_{1}\right], \ldots\right.$, $\left.Q(\mathbf{S})\left[R_{n}\right]\right\}$.

Proof. Let $\mathbf{S}_{\text {red }}=\left\{Q(\mathbf{S})\left[R_{1}\right], \ldots, Q(\mathbf{S})\left[R_{n}\right]\right\}$ and let $\mathbf{S}^{*}$ be a full reduction. It

[^1]$\mathbf{S}=\left\{\boldsymbol{S}_{1}, S_{2}, S_{3}\right\}$.

| $\underline{S_{1}(A B)}$ | $S_{2}(C$ D) | $\underline{S_{3}(E F)}$ |
| :---: | :---: | :---: |
| 01 | 12 | 23 |
| 34 | 45 | 5 |

Let $Q$ be a query with qualification $q$.
$q=\left(R_{1} \cdot B=R_{2} \cdot C\right) \wedge\left(R_{2} . D=R_{3} \cdot E\right) \wedge\left(R_{3} \cdot F=R_{1} . A\right)$.
$\mathbf{S}^{*}=\{\varnothing, \varnothing, \varnothing\}$.
Yet for each join clause, the corresponding semi-join does not reduce at all.
Fig. I. An example where $S^{*}$ is not obtainable by semi-joins.


Fig. 2. (a) Query graph $G_{Q_{1}}$ for query $Q_{1}$ defined by qualification $q_{1}=\left(R_{1}, A=R_{2}, B\right) \wedge$ ( $R_{2} \cdot B=R_{3} . C$ ). (b) Query graph $G_{Q_{2}}$ for query $Q_{2}$ defined by qualification $q_{2}=\left(R_{1}, A=R_{2}, B\right) \wedge$ $\left(R_{2} \cdot B=R_{1}, C\right) \wedge\left(R_{3} . C=R_{1} \cdot A\right)$. (c) Query graph $G_{Q_{3}}$ for query $Q_{3}$ defined by qualification $q_{3}=$ $\left(R_{1} \cdot A=R_{2} \cdot B\right) \wedge\left(R_{2} \cdot C=R_{3} \cdot D\right) \wedge\left(R_{3} \cdot E=R_{1} \cdot F\right)$.
must be that $\mathbf{S}_{\text {red }} \leq \mathbf{S}^{*}$, for otherwise $Q\left(\mathbf{S}^{*}\right) \subsetneq Q(\mathbf{S})$. Since this holds for any full reduction $\mathbf{S}^{*}$, it follows that $\mathbf{S}_{\text {red }}$ is the unique full reduction.

The principal issue of this paper is discovering when semi-join programs can completely resolve an equi-join query. This is equivalent to asking when a semi-join program can obtain $\mathbf{S}^{*}$. Unfortunately, $\mathbf{S}^{*}$ cannot always be obtained by semi-joins. For example, consider $\mathbf{S}$ and $Q$ in Figure 1. In this case, semi-joins cannot reduce $\mathbf{S}$ at all with respect to $Q$, yet $S^{*}=\{\varnothing, \varnothing, \varnothing\}$.

We now show a class of queries for which $\mathbf{S}^{*}$ can always be obtained by semijoins. Later we will show that for any query outside this class, semi-join programs cannot produce $S^{*}$ in general.
4.2 Tree Queries. Given a query $Q$ with qualification $q$ we define its corresponding query graph $G_{Q}\left(V_{Q}, E_{Q}\right)$ as (see Figure 2):
$V_{Q}=$ set of all relation names referenced by $q$
$E_{Q}=\left\{(i, j) \mid i \neq j\right.$ and some clause of $q$ references both $R_{i}$ and $\left.R_{j}\right\}$.

Since we allow more than one join between two relations, $E_{Q}$ is a multiset (i.e., may contain duplicates), so $G_{Q}$ is a multigraph.

We partition the set of all equi-join queries into two classes. We call a query a tree query either if its query graph is a tree or if it is equivalent to a query whose query graph is a tree. We denote the set of all tree queries by $T Q$ and all other equi-join queries by $C Q$ (for cyclic queries).

Consider the set of queries in Figure 2. Since $G_{Q_{1}}$ is a tree, $Q_{1} \in T Q$. Although $G_{Q_{2}}$ has a cycle, $Q_{1} \equiv Q_{2}$, so that $Q_{2} \in T Q$ also. The cycle in $G_{Q_{3}}$ is inherent; $Q_{3}$ is not equivalent to any query whose graph is a tree, so $Q_{3} \in C Q$. A simple efficient algorithm that tests whether a query is in $T Q$ appears in Section 5.

In the sequel we will only consider equi-join queries whose query graphs are connected. A query whose query graph is disconnected produces a result that is the

Let $Q$ be a query with qualification $q$.

$$
\begin{aligned}
q= & \left(R_{1} \cdot A=R_{2} \cdot B\right) \wedge\left(R_{1} \cdot C=R_{3} \cdot D\right) \wedge\left(R_{2} \cdot E=R_{4} \cdot F\right) \\
& \wedge\left(R_{3} \cdot G=R_{5} \cdot H\right) \wedge\left(R_{3} . J=R_{6} . K\right)
\end{aligned}
$$


$\sigma_{Q, 1}=\left\langle R_{2} \ltimes R_{4}, R_{3} \propto R_{5}, R_{3} \ltimes R_{6}, R_{1} \ltimes R_{2}, R_{1} \ltimes R_{3}\right\rangle$.
Fig. 3. An example reducing semi-join program.
Cartesian product of database substates produced by each connected component. That is, since these connected components are not joined in any way, there can be no interaction between them with respect to semi-joins. Hence, there is no loss of generality in treating the components separately.

We can now state our main result.
Theorem 1. For any $Q \in T Q$, there is a semi-join program $\sigma_{Q}$ such that for all $\mathbf{S}$, $\sigma_{Q}(\mathbf{S})=\mathbf{S}^{*}$ and $\left|\sigma_{Q}\right|=2 n-2$, where $n$ is the number of relations referenced by $Q$. Furthermore, for each $i, 1 \leq i \leq n, S_{i}^{*}$ can be obtained from $\mathbf{S}$ in $n-1$ semi-joins.

Theorem 1 says that there is always a short semi-join program that produces $\mathbf{S}^{*}$ for $Q \in T Q$. To prove Theorem 1 , we construct a semi-join program that obtains $S_{i}^{*}$ from $S$.

We begin by briefly reconsidering one-relation clauses of the form ( $R_{i} . A=R_{i} . B$ ). These clauses, we will assume, are applied to the database state before any semi-joins are executed. They are assumed to be free and do not appear in the semi-join programs that we construct to solve $Q$.

Let $Q \in T Q$ be a query whose graph is a tree, and let $i \in V_{Q}$ be a relation referenced by $Q .{ }^{3}$ We use $G_{Q}$ as a control structure to guide our semi-join program that evaluates $Q$. In $G_{Q}$, choose $i$ as the root of the tree and suppose $i$ has $m-1$ children, $m \geq 1$. For simplicity, we rename the relations so that the root of the tree is 1 (formerly $i$ ) and its children are $2, \ldots, m$. Each of these children is the root of a query subtree, and therefore they define the set of queries $Q_{2}, \ldots, Q_{m}$. We now recursively define the reducing semi-join program for 1 in $Q$, denoted $\sigma_{Q, 1}$, as follows:
(i) if $m=1$, then $\sigma_{Q, 1}=\langle \rangle$, the empty program;
(ii) else $\sigma_{Q, 1}=\left\langle\sigma_{Q_{2}, 2}, \ldots, \sigma_{Q_{m}, m}, R_{1} \ltimes R_{2}, R_{1} \ltimes R_{3}, \ldots, R_{1} \ltimes R_{m}\right\rangle$.

To produce $\sigma_{Q, i}$, rename the relations with their original indices.
In words, $\sigma_{Q, 1}$ executes one semi-join per edge in $G_{Q}$, in a breadth-first leaf-to-root order. Equivalently, reducing semi-join programs can be produced mechanically from queries using the detachment algorithm of Wong and Youssefi [13]. An example reducing semi-join program appears in Figure 3.

The following lemma shows that $\sigma_{Q, i}$ is the program we are looking for to prove Theorem 1.

Lemma 4. Let $Q$ be a query with qualification $q$ whose query graph is a tree. Choose $i \in V_{Q}$ as the root of the tree. Then $\sigma_{Q, i}(\mathbf{S})\left[R_{i}\right]=S_{i}^{*}=Q(\mathbf{S})\left[R_{i}\right]$. That is, $\sigma_{Q, i}$ fully reduces $S_{i}$ with respect to $Q$.

[^2]Proof. The proof is by induction on the height of the tree $G_{Q}$ (the height of a tree is the length of the longest path from the root to a leaf). Choose $i \in V_{Q}$ as the root of the tree.

Basis Step. Suppose the height of $G_{Q}$ is 0 . Since no join clauses on one variable are possible, $q$ is empty. Hence $\mathbf{S}=\mathbf{S}^{*}$. Since $\sigma_{Q, i}$ is empty, $\sigma_{Q, i}(\mathbf{S})=\mathbf{S}^{*}$.

Induction Step. Suppose the lemma is true for all queries whose query graph is a tree of height $<p$. We show it to be true if the height of $G_{Q}$ equals $p$. As before, let us rename the relations so that $i$ becomes 1 and $i$ 's children are $2, \ldots, m, m>1$. Let $Q_{2}, \ldots, Q_{m}$ be the queries, with qualifications $q_{2}, \ldots, q_{m}$, defined by the subtrees (of $G_{Q}$ ) rooted by each child. Finally, let $c_{j}, 2 \leq j \leq m$, be the join clause linking relation $j$ to relation 1, and let $q^{\prime}=c_{2} \wedge \cdots \wedge c_{m}$.

From the definition of query, we have

$$
\left\langle t_{1}, \ldots, t_{n}\right\rangle \in Q(\mathbf{S}) \quad \text { iff } \quad t_{j} \in S_{j}(1 \leq j \leq n) \quad \text { and } \quad q\left(t_{1}, \ldots, t_{n}\right)
$$

It immediately follows that $t_{1} \in Q(S)\left[R_{1}\right]$ iff

$$
\begin{equation*}
\exists t_{j} \in S_{j}(2 \leq j \leq n) \quad \text { such that } \quad q\left(t_{1}, \ldots, t_{n}\right) \tag{1}
\end{equation*}
$$

Expanding $q$ into its component clauses, (1) holds iff
$\exists t_{j} \in S_{j}(2 \leq j \leq m) \quad$ such that

$$
\begin{align*}
& \left(q^{\prime}\left(t_{1}, \ldots, t_{m}\right)\right. \text { and } \\
& \left.\exists t_{k} \in S_{k}(m<k \leq n) \text { such that }\left(q_{2}\left(t_{1}, \ldots, t_{n}\right) \wedge \ldots \wedge q_{m}\left(t_{1}, \ldots, t_{n}\right)\right)\right) \tag{2}
\end{align*}
$$

Each $q_{j}, 2 \leq j \leq m$, ranges over a set of relations that is disjoint from every other $q_{k}$, $k \neq j$. Thus, since each $q_{j}$ is evaluated by $Q_{j}$,

$$
\begin{align*}
& \exists t_{j} \in S_{j}(2 \leq j \leq m) \quad \exists t_{k} \in S_{k}(m<k \leq n) \quad \text { such that } \\
& \quad\left(q_{2}\left(t_{1}, \ldots, t_{n}\right) \wedge \cdots \wedge q_{m}\left(t_{1}, \ldots, t_{n}\right)\right) \text { iff } \exists t_{j} \in S_{j}(2 \leq j \leq m) t_{j} \in Q_{j}(\mathbf{S})\left[R_{j}\right] . \tag{3}
\end{align*}
$$

Since each $G_{Q \rho} 2 \leq j \leq m$, is a tree of height $<p$, by the induction hypothesis we have that $Q_{j}(\mathbf{S})\left[R_{j}\right]=\sigma_{Q_{j} j}(\mathbf{S})\left[R_{j}\right]$. Substituting this into (3) and combining (3) with (2), we find that (1) holds iff

$$
\begin{equation*}
\exists t_{j} \in S_{j}(2 \leq j \leq m) \quad \text { such that } \quad\left(q^{\prime}\left(t_{1}, \ldots, t_{m}\right) \text { and } t_{j} \in \sigma_{Q_{j}, j}(\mathbf{S})\left[R_{j}\right]\right) \tag{4}
\end{equation*}
$$

Let $S_{j}^{+}=\sigma_{Q_{j}} j(\mathbf{S})\left[R_{j}\right]$. Each clause of $q^{\prime}$ can be evaluated by a semi-join, so (4) holds iff

$$
\begin{align*}
& \exists t_{j} \in S_{j}(2 \leq j \leq m) \quad \text { such that } \\
& \quad\left\langle t_{1}, \ldots, t_{m}\right\rangle \in\left\langle R_{1} \propto R_{2}, \ldots, R_{1} \ltimes R_{m}\right\rangle\left(\left\{S_{1}, S_{2}^{+}, \ldots, S_{m}^{+}\right\}\right) . \tag{5}
\end{align*}
$$

But $\sigma_{Q .1}(\mathbf{S})\left[R_{1}, \ldots, R_{m}\right]$ exactly calculates $\left\langle\mathrm{R}_{1} \propto R_{2}, \ldots, R_{1} \propto R_{m}\right\rangle\left(\left\{S_{1}, S_{2}^{+}, \ldots\right.\right.$, $\left.S_{m}^{+}\right\}$), so $t_{1} \in Q(\mathbf{S})\left[R_{1}\right]$ iff $\exists t_{j} \in S_{j}, 2 \leq j \leq m$, such that $\left\langle t_{1}, \ldots, t_{m}\right\rangle \in$ $\sigma_{Q, 1}(\mathbf{S})\left[R_{1}, \ldots, R_{m}\right]$. It immediately follows that $Q(\mathbf{S})\left[R_{1}\right]=\sigma_{Q, 1}(\mathbf{S})\left[R_{1}\right]$, as desired.

Note that $\sigma_{Q, i}$ produces $S_{i}^{*}$ but not $S_{j}^{*}, j \neq i$. That is, only relation $S_{i}$ is fully reduced by $\sigma_{Q, i}$. Intuitively, the reason is that each relation (except for $S_{i}$ ) is only reduced by a proper subset of the clauses in the query. More specifically, each relation $S_{j}$ is reduced by clauses in the subtree of $Q$ rooted at $j$ (considering $i$ as the root of $Q$ 's tree). To produce $\mathbf{S}^{*}$, we first perform $\sigma_{Q, i}$ and then apply semi-joins from the root $i$ down the tree toward the leaves. The end effect is that each relation is reduced by every clause. Proceeding formally, the program we want is $\sigma_{Q}=$ $\left\langle\sigma_{Q, .}, \sigma_{Q, i}^{-1}\right\rangle$, where $\sigma_{Q . i}^{-1}$ is defined (with the usual renaming of relations) as follows:
(i) if $m=1$ then $\sigma_{Q, 1}^{-1}=$ ( $\rangle$, the empty program;
(ii) else $\sigma_{Q_{.1}}^{-1}=\left\langle R_{2} \ltimes R_{1}, \ldots, R_{m} \ltimes R_{1}, \sigma_{Q_{2} .2}^{-1}, \ldots, \sigma_{Q_{m}, m}^{-1}\right\rangle$,
where $2, \ldots, m$ are the children of 1 as before. The principal property of $\sigma_{Q}$ is that for each $j \in V_{Q}, \sigma_{Q, j}$ is embedded in $\sigma_{Q}$. That is, for each $j, \sigma_{Q, j}$ can be obtained from $\sigma_{Q}$ simply by excising certain semi-joins from $\sigma_{Q}$. Clearly for all $\mathbf{S}, \sigma_{Q}(\mathbf{S}) \leq \sigma_{Q, j}(\mathbf{S})$. The proof of Theorem 1 now follows directly.
Proof of Theorem 1. Let $Q \in T Q$ and assume $G_{Q}$ is a tree. Let $\mathbf{S}^{\prime}=\sigma_{Q}(\mathbf{S})$. Since for each $j \in V_{Q}, \sigma_{Q, j}$ is embedded in $\sigma_{Q}$, it follows from Lemma 4 that $\mathbf{S}^{\prime}=\mathbf{S}^{*}$. $\sigma_{Q, j}$ contains $n-1$ semi-joins and $\sigma_{Q}$ contains $2 n-2$ semi-joins.

If $G_{Q}$ is not a tree, then since $Q \in T Q$, there is an equivalent query $Q^{\prime}$ whose graph is a tree. The proof goes as before, using $Q^{\prime}$ and the fact that $Q^{\prime}(\mathbf{S})=Q(\mathbf{S})$, by definition of equivalence.
4.3 Cyclic Queries. Queries in $C Q$ are very badly behaved with respect to semi-joins compared to those in $T Q$. Not only are semi-joins incapable of obtaining $\mathbf{S}^{*}$ in general, but even the best possible reduction of $\mathbf{S}$ may only be obtainable by a very long semi-join program.
Let $Q$ be a query with qualification $q$. We define $\operatorname{PROGRAMS}(Q)$ to be the set of all semi-join programs that only perform semi-joins that correspond to clauses of $q$. Let $\mathbf{S}$ be a database state. We define $\propto$-REDUCTIONS $(Q, \mathbf{S})=\left\{\mathbf{S}^{\prime} \mid \mathbf{S}^{\prime} \leq \mathbf{S}\right.$, $Q\left(\mathbf{S}^{\prime}\right)=Q(\mathbf{S})$, and $\exists \sigma \in \operatorname{PROGRAMS}(Q)$ such that $\left.\sigma(\mathbf{S})=\mathbf{S}^{\prime}\right\}$. In words, $\propto$ REDUCTIONS $(Q, S)$ is the set of all database states that can be obtained from $\mathbf{S}$ by semi-join programs that correspond to $Q$. From the definition it immediately follows that $\propto-\operatorname{REDUCTIONS}(Q, \mathbf{S}) \subseteq$ REDUCTIONS $(Q, \mathbf{S})$. A full $\ltimes$-reduction $\mathbf{S}^{\times}$of $\mathbf{S}$ with respect to $Q$ is an element of $\propto-\operatorname{REDUCTIONS}(Q, S)$ such that there is no $\mathbf{S}^{\prime}$ in $\ltimes$-REDUCTIONS $(Q, \mathbf{S})$ with $\mathbf{S}^{\prime}<\mathbf{S}^{\star}$.

Lemma 5. For each $Q$ and $\mathbf{S}$ there is a unique full $\propto$-reduction, $\mathbf{S}^{\star}$.
Proof. If $\mid \propto$-REDUCTIONS $(Q, \mathbf{S}) \mid=1$, then $\ltimes$-REDUCTIONS $(Q, \mathbf{S})=\{\mathbf{S}\}$ and we are done. Otherwise, let $\mathbf{S}^{1}$ and $\mathbf{S}^{2}$ be two full $\propto$-reductions of $\mathbf{S}$ with respect to $Q$ that were produced, respectively, by $\sigma_{1}$ and $\sigma_{2}$. Since $\mathbf{S}^{1}$ is a full $\propto$-reduction, $\sigma_{2}\left(\mathbf{S}^{1}\right)=\mathbf{S}^{1}$. Since $\mathbf{S}^{1} \leq \mathbf{S}$, by Lemma $2 \sigma_{2}\left(\mathbf{S}^{1}\right) \leq \sigma_{2}(\mathbf{S})=\mathbf{S}^{2}$. So $\mathbf{S}^{1} \leq \mathbf{S}^{2}$. The symmetric argument shows that $\mathbf{S}^{2}=\sigma_{1}\left(\mathbf{S}^{2}\right) \leq \sigma_{1}(\mathbf{S})=\mathbf{S}^{1}$. Hence $\mathbf{S}^{1}=\mathbf{S}^{2}$.

Theorem 1 says that $\mathbf{S}^{*}=\mathbf{S}^{\star}$ for tree queries. However, for cyclic queries there are states $\mathbf{S}$ for which $\mathbf{S}^{*}$ is not in $\propto-\operatorname{REDUCTIONS}(Q, \mathbf{S})$. To show that $\mathbf{S}^{*}$ and $\mathbf{S}^{\star}$ are not always identical for cyclic queries, we will look at a special subclass of $C Q$ that has minimal cycles.

Let $Q$ be a query with qualification $q$. The closure of $Q$, denoted $Q^{+}$, is a query whose qualification, denoted $q^{+}$, includes $q$ and all clauses implied by $q$ under transitivity. (E.g., if ( $R_{1} . A_{1}=R_{2} . A_{2}$ ) and ( $R_{2} \cdot A_{2}=R_{3} \cdot A_{3}$ ) are in $q$, then ( $R_{1} \cdot A_{1}=$ $R_{3} . A_{3}$ ) is in $q^{+}$.) Clearly $Q^{+} \equiv Q$. Also, if $Q_{1} \equiv Q_{2}$, then $Q_{1}^{+}=Q_{2}^{+}$. For if $Q_{1}^{+} \neq Q_{2}^{+}$, then there is a clause in $q_{1}^{+}-q_{2}^{+}$, say; by selecting a state $\mathbf{S}$ such that some tuple in the Cartesian product of the relations in $\mathbf{S}$ satisfies $q_{2}^{+}$but does not satisfy $q_{1}^{+}-q_{2}^{+}$, we have $Q_{1}(\mathbf{S}) \neq \mathrm{Q}_{2}(\mathbf{S})$, or $Q_{1} \neq Q_{2} .{ }^{4}$

Lemma 6. $Q_{1} \equiv Q_{2}$ iff $Q_{1}^{+}=Q_{2}^{+}$.
Proof. Follows from the above argument.
A qualification $q$ is called proper cyclic if it is of the form $\Lambda_{i=1}^{p}\left(R_{i}, A_{i, 2}=\right.$ $\left.R_{i+1} \cdot A_{i+1,1}\right)$, where $p>2$, 1 is the successor of $p$, and all $A_{i, j}$ are distinct. A query is proper cyclic if its qualification is proper cyclic.
${ }^{4}$ This result does not hold when tuple variables are introduced; see [5].


Fig. 4. The database constructed for the proof of Lemma 7.
Intuitively, a proper cyclic qualification corresponds to a nonreducible cycle. If $Q_{1}$ is proper cyclic, then $Q_{1}=Q_{1}^{+}$. Also, if $Q_{1} \equiv Q_{2}$, then (by Lemma 6) $Q_{1}=Q_{1}^{+}=Q_{2}^{+}$. But no clause of $Q_{1}$ is implied by the others, so $Q_{1}=Q_{2}^{+}$implies $Q_{1}=Q_{2}$. Hence there is no other query equivalent to $Q_{1}$.

Proper cyclic queries are badly behaved with respect to semi-joins as indicated by the following lemma.

Lemma 7. For any proper cyclic query $Q$, there exists a database state $\mathbf{S}$ such that (i) $\mathbf{S}^{\star} \neq \mathbf{S}^{*}$, and (ii) the fewest number of semi-joins required to obtain $\mathbf{S}^{\star}$ is $O(m)$, where $m$ is the number of tuples of some relation in $\mathbf{S}$.

Proof. Let $Q$ have qualification $q=\Lambda_{i=1}^{p}\left(R_{i} \cdot A_{i, 2}=R_{i+1} \cdot A_{i+1,1}\right)$, where $p \geq 2$. For convenience we assume that all domains of all relations are integer valued. Choose some integer $m>1$ and a set of distinct integers $\left\{a_{1}, \ldots, a_{m}\right\}$, where $a_{j} \geq 2 p$ for $1 \leq j \leq m$. In $\mathbf{S}$ let all of the domains not referenced by $q$ have arbitrary values, so that in constructing $\mathbf{S}$, we consider all relations to be binary. In $\mathbf{S}$ let $S_{1}=$ $\left\{(0,1\rangle,\langle p, p+1\rangle,\left\langle a_{m}, 1\right\rangle\right\} \cup\left\{\left\langle a_{j}, a_{j+1}\right\rangle: 1 \leq j<m\right\}$, let $S_{p}=\{\langle p-1, p\rangle$, $(2 p-1,0\rangle\} \cup\left\{\left\langle a_{j}, a_{j}\right\rangle: 1 \leq j \leq m\right\}$, and for each $k, 1<k<p$, let $S_{k}=\{\langle k-1$, $k\rangle,\langle p+k-1, p+k\rangle\} \cup\left\{\left(a_{j}, a_{j}\right\rangle: 1 \leq j \leq m\right\}$ (see Figure 4).
To show that $\mathbf{S}^{*} \neq \mathbf{S}^{\star}$, we observe that the tuples $\{\langle k-1, k\rangle,\langle p+k-1$, $p+k\rangle\}$ are in each $\mathbf{S}_{k}^{\kappa}, 1 \leq k<p$, and $\{\langle p-1, p\rangle,\langle 2 p-1,0\rangle\}$ are in $S_{p}$. Yet $\mathbf{S}^{*}=\{\varnothing, \ldots, \varnothing\}$, so $\mathbf{S}^{*} \neq \mathbf{S}^{\times}$. In addition, $\mathbf{S}$ is constructed so that only one possible semi-join can reduce the size of any relation in $\mathbf{S}$, and furthermore every database produced by a sequence of reducing semi-joins has this property. In fact, the only (nonredundant) sequence of semi-joins that can produce $\mathbf{S}^{\star}$ is

```
do \(m\) times;
    do \(k=1\) to \(p\) by \(1 ;\)
        \(R_{k+1} \propto R_{k} ;\)
    end;
end;
```

where $I$ is the successor of $p$. This requires $p m$ semi-joins.
The proof of Lemma 7 actually supports a stronger statement about semi-join programs for proper cyclic queries. The constructed database (Figure 4) is designed so that each semi-join only reduces the database by one tuple. Even if one were willing to be satisfied with a reduced database larger than $\mathbf{S}^{\star}$, obtaining such a reduction would still be slow. To obtain $\mathbf{S}^{\prime} \geq \mathbf{S}^{\star}$, the number of semi-joins required equals the number of tuples in $\mathbf{S}$ minus the number in $\mathbf{S}^{\prime}$.
Of course, not every cyclic query is proper cyclic. However, we can strengthen Lemma 7 to cover all cyclic queries. We begin by showing that a certain type of proper cyclic qualification is embedded in every cyclic query.

A qualification $q^{\prime}$ is a proper cyclic subqualification of $q$ if
(1) $q^{\prime}$ is contained in $q^{+}$,
(2) $q^{\prime}$ is proper cyclic, and
(3) for any two domains $R_{i} . A_{i}$ and $R_{j} . A_{j}$ that appear in clauses in $q^{\prime}$, if ( $R_{i}, A_{i}=R_{j} . A_{j}$ ) is in $q^{+}$, then $\left(R_{i} . A_{i}=R_{j} . A_{j}\right)$ is in $q$.
Notice that $q$ is a proper cyclic subqualification of $q$ iff $q$ is proper cyclic, since (3) is irrelevant when $q$ is proper cyclic.

Part (3) of the above definition guarantees that the cycle corresponding to $q^{\prime}$ cannot be broken up into two shorter cycles, for if ( $R_{i} . A_{i}=R_{j} . A_{j}$ ) were in $q^{\prime}$ but not in $q$ ', then a "path" of clauses from $R_{i} . A_{i}$ to $R_{j} . A_{j}$ in $q$ ' could be replaced by ( $R_{i}, A_{i}=R_{j}, A_{j}$ ), thereby shortening or destroying the cycle.

Lemma 8. Query $Q$ with qualification $q$ is in $C Q$ iff $q$ has a proper cyclic subqualification.

The proof of Lemma 8 is most conveniently presented after we have developed more machinery for manipulating equivalent queries. Section 5 develops this machinery for the purpose of testing membership in $T Q$, so we defer the proof to the end of that section.

Since every cyclic query has a proper cyclic subquery, Lemma 7 holds at least for the proper cyclic subquery (when it is treated as a proper cyclic query). We will show that Lemma 7 actually holds for the entire query. Since by Lemma 8 every cyclic query has a proper cyclic subquery, we can use the database of Figure 4 to produce the desired effect. However, there is a technical problem here; we must assign data values to all domains outside the proper cyclic subquery. We will show that this can be handled by making the rest of the database a Cartesian product.

Let $A_{i}$ be the domains of $R_{i}$. If $B \subseteq A_{i}$, then $S_{i}[B]$ denotes the projection of $S_{i}$ on $B$, where $S_{i}[B]=\left\{\left(s . B_{1}, \ldots, s . B_{m}\right\rangle \mid s \in S_{i}\right.$ and $B_{j} \in B$ for $\left.1 \leq j \leq m\right\}$. Let $\left\langle R_{i} \times R_{j}\right\rangle$ denote the Cartesian product operation. Now if $S_{i}=S_{i}[B] \times S_{i}\left[A_{i}-B\right]$, then for any $A_{i, k} \in B$ we have

$$
\begin{aligned}
& \left(R_{i A_{i, j}} \propto R_{j}\right\rangle\left(\left\{S_{i}, S_{j}\right\}\right) \\
& \quad=\left\langle R_{i}[B] \times R_{i}\left[A_{i}-B\right]\right\rangle\left(\left\langle R_{i}[B] \propto R_{j}\right\rangle\left(\left\{S_{i}[B], S_{i}\left[A_{i}-B\right], S_{j}\right\}\right)\right) .
\end{aligned}
$$

In words, if $S_{i}$ is a Cartesian product of two subrelations, then any semi-join on $S_{i}$ can be considered as a semi-join on the database consisting of the subrelations of $S_{i}$. A Cartesian product operation can always reconstruct the database to its original form-either before or after the semi-join.

Let $q^{\prime}=\Lambda_{i=1}^{p}\left(R_{i} . A_{i, 2}=R_{i+1} . A_{i+1,1}\right)$ be a proper cyclic subqualification of some query $Q$. We say that a database state $\mathbf{S}$ is decomposable with respect to $q^{\prime}$ if
(dl) for $1 \leq i \leq p, S_{i}=S_{i}\left[A_{i, 1}, A_{i, 2}\right] \times\left(\times_{k \neq 1,2} S_{i}\left[A_{i, k}\right]\right)$;
(d2). for $p<i \leq n, S_{i}=x_{k} S_{i}\left[A_{i, k}\right]$;
(d3) for all $i$ and $k, S_{i}\left[A_{i, k}\right] \neq \varnothing$; and
(d4) for $1 \leq i \leq p$ and $k=1,2, S_{i}\left[A_{i, k}\right] \subseteq \cap_{R_{i}, A_{j, \ell} \in J\left(A_{i, k}\right)} S_{j}\left[A_{j, 1}\right]$;
where $J\left(A_{i, k}\right)$ is the set of domains that join with $A_{i, k}$ in $q^{+}$.
Let DOMAINS $(q)$ be the domains referenced in qualification $q$, and let RELATIONS $(q)$ be the relations referenced in qualification $q$. Intuitively, by making $S$ decomposable with respect to $q^{\prime}$ we have made every relation not in RELATIONS $\left(q^{\prime}\right)$ a Cartesian product (by (d2)); and for each relation in RELATIONS $\left(q^{\prime}\right)$ we have made those domains not in DOMAINS $\left(q^{\prime}\right)$ a Cartesian product (by (dl)). Furthermore, (d3) and (d4) imply that no domain not in DOMAINS $\left(q^{\prime}\right)$ can reduce any
domain in DOMAINS $\left(q^{\prime}\right)$ by means of semi-joins, since the former are supersets of the latter. Finally, by the above observation about semi-joins applied to Cartesian products, we can decompose $\mathbf{S}$ so that domains not in DOMAINS $\left(q^{\prime}\right)$ become singledomain relations; semi-joins applied to this decomposed database have the same effect as if they were applied to $S$.

Let $q^{\prime}$ be a proper cyclic subqualification of $q$, and let $S$ be a database state decomposable with respect to $q^{\prime}$. We can decompose $\mathbf{S}$ by applying a decomposition mapping which maps $\mathbf{S}$ into $\mathbf{S}_{d}$ so that each domain not in DOMAINS $\left(q^{\prime}\right)$ becomes a single-domain relation. Having decomposed $\mathbf{S}$, we can apply a renaming mapping to $q$ (and $q^{\prime}$ ), which maps $q$ (and $q^{\prime}$ ) into $q_{d}$ (and $q_{d}^{\prime}$ ) by mapping domain references in $q$ and $q^{\prime}$ on $\mathbf{S}$ into corresponding domain references in $q_{d}$ and $q_{d}^{\prime}$ on $\mathbf{S}_{d}$. We make two important observations regarding the decomposed database and renamed query:
(dm1) Each relation $R$ in RELATIONS $\left(q_{d}^{\prime}\right)$ is a binary relation containing only the two domains of $R$ in DOMAINS $\left(q^{\prime}\right)$. All other relations in $S_{d}$ are unary.
(dm2) For each domain (= unary relation) $R_{k} . A_{k}$ not in DOMAINS $\left(q_{d}^{\prime}\right)$, either there is no domain $R_{i} \cdot A_{i}$ in DOMAINS $\left(q_{d}^{\prime}\right)$ such that $\left(R_{k} . A_{k}=\right.$ $\left.R_{i} . A_{i}\right)$ is in $q_{d}^{+}$, or there is exactly one clause in $q_{d}^{\prime},\left(R_{i}, A_{i}=R_{j} . A_{j}\right)$, such that ( $R_{i} . A_{i}=R_{k} \cdot A_{k}$ ) and ( $R_{j} . A_{j}=R_{k}, A_{k}$ ) are in $q_{d}^{+}$.
(dm2) follows from the fact that $q_{d}^{\prime}$ is a proper cyclic subqualification; if there are two (or more) distinct clauses of $q_{d}^{\prime},\left(R_{i 1} . A_{i 1}=R_{j 1} . A_{j 1}\right)$ and $\left(R_{i 2} . A_{i 2}=R_{j 2} . A_{j 2}\right)$, such that ( $R_{k} \cdot A_{k}=R_{i 1} . A_{i 1}$ ) and ( $R_{k} . A_{k}=R_{i 2} . A_{i 2}$ ) are in $q_{d}^{+}$, then part (3) of the definition of proper cyclic subqualification is violated. Thus, the clauses of $q^{\prime}{ }_{d}$ partition the relations that are not in RELATIONS $\left(q_{d}^{\prime}\right)$ but are connected to $\operatorname{RELATIONS}\left(q_{d}^{\prime}\right)$ in the query graph of $Q_{d}$.

The following lemma extends the first part of Lemma 7 to arbitrary cyclic queries. It implies that if $q^{\prime}$ is a proper cyclic subqualification of $q$ and $\mathbf{S}^{*} \neq \mathbf{S}^{\star}$ with respect to $q^{\prime}$, then $\mathbf{S}^{*} \neq \mathbf{S}^{\star}$ with respect to $q$.

Lemma 9. Let $Q$ be a cyclic query with qualification $q$, let $q^{\prime}$ be a proper cyclic subqualification of $q$, and let $\mathbf{S}$ be a database state decomposable with respect to $q^{\prime}$. Let $\mathbf{S}_{d}, q_{d}$, and $q_{d}^{\prime}$ be the results of decomposing $\mathbf{S}$ and renaming $q$ and $q^{\prime}$. Then the full semi-join reduction of $\mathbf{S}_{d}$ with respect to $q_{d}$ (denoted $\mathbf{S}_{d, q_{d}}^{\star}$ ) equals the full semi-join reduction of $\mathbf{S}_{d}$ with respect to $q_{d}^{\prime}\left(\right.$ denoted $\left.\mathbf{S}_{d, q_{d}^{\prime}}^{\times}\right)$on each relation in RELATIONS $\left(q_{d}^{\prime}\right)$.

Proof. It suffices to show that no relation $\mathbf{S}_{d \cdot q_{d}^{\prime}}^{\times}$[RELATIONS $\left.\left(q_{d}^{\prime}\right)\right]$ can be reduced by a semi-join corresponding to a clause in $q_{d}^{+}$. However, this follows directly from (d4) and (dm2).

To extend the remainder of Lemma 7, we introduce some additional notation. Given a query $Q$ and a database state $S$, we define $M(Q, S)$ to be the length of the shortest semi-join program to produce a full semi-join reduction of $S$ with respect to $Q$; that is,

$$
M(Q, \mathbf{S})=\min _{\substack{\sigma \in \operatorname{PrOGRAMSIQ}) \\ \sigma(\mathbf{S})=\mathbf{S}^{\times}}}|\sigma| .
$$

Lemma 10. Let $Q$ be a cyclic query with qualification $q$, let $q^{\prime}$ be a proper cyclic subqualification of $q$, and let $\mathbf{S}$ be a database state decomposable with respect to $q^{\prime}$. Let $\mathbf{S}_{d}, q_{d}$, and $q_{d}^{\prime}$ be the results of decomposing $\mathbf{S}$ and renaming $q$ and $q^{\prime}$. Then $M\left(Q_{d}^{\prime}, \mathbf{S}_{d}\right) \leq M\left(Q_{d}, \mathbf{S}_{d}\right)$.

Proof. Let $\sigma$ be a semi-join program of length $M\left(Q_{d}, \mathbf{S}_{d}\right)$ such that $\sigma\left(\mathbf{S}_{d}\right)=$ $\mathbf{S}_{d, q_{d}}^{\kappa}$. To prove the lemma, we will transform $\sigma$ into $\sigma^{\prime}$ such that $M\left(Q_{d}^{\prime}, \mathbf{S}_{d}\right) \leq$ $\left|\sigma^{\prime}\right| \leq|\sigma|$ and $\sigma^{\prime}\left(Q_{d}^{\prime}, S_{d}\right)=S_{d, q_{d}^{\prime}}^{\star}$. Begin by finding the first semi-join in $\sigma$, say ( $\left.R_{i A_{i}} \times_{A_{j}} R_{j}\right\rangle$, that reduces a relation in RELATIONS $\left(q_{d}^{\prime}\right)$. If $\left(R_{i}, A_{i}=R_{j}, A_{j}\right)$ is in $q_{d,}^{\prime}$ then leave the semi-join as is. Otherwise, remove $\left\langle R_{i A_{i}} \times_{A_{j}} R_{j}\right\rangle$ from $\sigma$. By (d4) and (dm2), $\left\langle R_{i A_{i}} \propto_{A_{j}} R_{j}\right\rangle$ could only reduce $S_{d, i}$ if there were a chain of semi-joins preceding it of the form $\left.\left\langle R_{j 1 A_{j 1}} \propto_{A_{k}} R_{k}\right\rangle,\left(R_{j 2 A_{j 2}} \propto_{A_{j 1}} R_{j 1}\right\rangle, \ldots, R_{j A_{j}} \propto_{A_{j m}} R_{j m}\right\rangle$, where ( $R_{k} A_{k}=R_{i} A_{i}$ ) is the unique clause in $q_{d}^{\prime}$ referencing $R_{i}, A_{i}$. If the chain exists, remove it from $\sigma$ and insert ( $R_{i A_{i}} \propto_{A_{k}} R_{k}$ ) in $\sigma$ where $\left\langle R_{i A_{i}} \propto_{A_{j}} R_{j}\right.$ ) used to be; this has exactly the same effect on $R_{i}$ as the chain did. Now find the second such semijoin, perform the replacement if appropriate, etc., until all semi-joins reducing relations in RELATIONS $\left(q_{d}^{\prime}\right)$ have been examined and/or replaced. Now all semijoins in $\sigma$ chat reduce relations in RELATIONS $\left(q_{d}^{\prime}\right)$ correspond to clauses of $q_{d}^{\prime}$. So remove from $\sigma$ all semi-joins that do not reference relations in RELATIONS $\left(q_{d}^{\prime}\right)$, since none of these semi-joins can affect relations in RELATIONS $\left(q^{\prime}{ }_{d}\right)$. The resulting semi-join program $\sigma^{\prime}$ is no longer than $\sigma$, and it produces the same effect on relations in RELATIONS $\left(q_{d}^{\prime}\right)$ as $\sigma$. Since $\sigma\left(Q_{d}, \mathbf{S}_{d}\right)$ [RELATIONS $\left.\left(q_{d}^{\prime}\right)\right]=\mathbf{S}_{d, q_{d}^{\prime}}^{\aleph}$ by Lemma $9, \sigma^{\prime}\left(Q_{d}^{\prime}, \mathbf{S}_{d}\right)=\mathbf{S}_{d, q_{d}^{\prime}}^{\kappa}$ as desired.

We can now extend Lemma 7 to all cyclic queries.
Theorem 2. For any query $Q \in C Q$, there exists a database state $\mathbf{S}$ such that $\mathbf{S}^{\kappa} \neq \mathbf{S}^{*}$ and the fewest number of semi-joins required to obtain $\mathbf{S}^{\star}$ is $O(m)$, where $m$ is the number of tuples of some relation in $\mathbf{S}$.

Proof. By Lemma 8, $q$ has a proper cyclic subqualification $q^{\prime}$. Let $\mathbf{S}$ be a database state decomposable with respect to $q^{\prime}$ such that if $\mathbf{S}_{d}$ is the decomposed version of $\mathbf{S}$ and $q_{d}, q_{d}^{\prime}$ are the renamed versions of $q, q^{\prime}$, then $\mathbf{S}_{d}$ [RELATIONS $\left(q_{d}^{\prime}\right)$ ] is a state satisfying Lemma 7. That is, $\mathbf{S}_{d, q_{d}^{\prime}}^{\kappa} \neq \mathbf{S}_{d, q_{d}^{\prime}}^{*}$ and $M\left(Q_{d}^{\prime}, \mathbf{S}_{d}\right)$ is of size $O(m)$. By Lemma 9, $\mathbf{S}_{d, q_{d}^{\prime}}^{\kappa}\left[\operatorname{RELATIONS}\left(q_{d}^{\prime}\right)\right]=\mathbf{S}_{d, q_{d}}^{\kappa}\left[\operatorname{RELATIONS}\left(q_{d}^{\prime}\right)\right]$, so $\mathbf{S}_{d, q_{d}}^{\kappa} \neq \mathbf{S}_{d, q_{d}}^{*}$. By Lemma $10, M\left(Q_{d}^{\prime}, \mathbf{S}_{d}\right) \leq M\left(Q_{d}, \mathbf{S}_{d}\right)$, so $M\left(Q_{d}, \mathbf{S}_{d}\right)$ is at least of size $O(m)$. Since all of these results hold on $S$ by taking the inverse of the decomposition and renaming maps, the theorem is proved.

## 5. A Fast Tree Query Membership Test

To make use of our result that shows tree queries to be well behaved with respect to semi-join, we need a procedure that tests if a given query is in $T Q$. In this section we present such a test that runs in linear time. The test is constructive. If the given query is in $T Q$, an equivalent query is produced whose query graph is a tree.

Given a query with a cyclic query graph, we cannot immediately tell if the query is cyclic. For example, if the qualification consists of the cycle $\left(\left(R_{1} . A=R_{2} \cdot B\right)\right.$ and $\left(R_{2} \cdot B=R_{3} \cdot C\right)$ and ( $\left.R_{3} . C=R_{1} . A\right)$ ), we can drop any one of the clauses; the result is a tree query that is obviously equivalent to the given query. The property of the cycle that permits us to drop one clause is that each relation participates in the cycle with only one joining domain.

There is a second situation in which a cycle can be broken without changing the meaning of the query. Consider the qualification $\left(\left(R_{1} . A=R_{2} \cdot B\right)\right.$ and $\left(R_{2} . B=R_{3} . C\right)$ and $\left(R_{3} . C=R_{1} . D\right)$, which produces a cyclic query graph with edges $\left\{\left(R_{1}, R_{2}\right)\right.$, $\left(R_{2}, R_{3}\right),\left(R_{3}, R_{1}\right)$ ). This qualification does not fit the form of our previous example, because $R_{1}$ has two joining domains. Yet we can transform the qualification into ( $R_{1} . A=R_{2} \cdot B$ ) and ( $R_{2} \cdot B=R_{3} . C$ ) and ( $\left.R_{1} \cdot A=R_{1} \cdot D\right)$ ); since $R_{1} . A$ must equal $R_{3} . C$ in the first qualification, we can substitute $\left(R_{1}, A=R_{1} \cdot D\right)$ for $\left(R_{3}, C=R_{1} \cdot D\right)$. The


Fig. 5. A join graph and its corresponding query graph for query $Q$ with qualification $q=\left(\left(R_{1} A=R_{2} . B\right) \wedge\left(R_{2} . C=R_{3} . D\right) \wedge\left(R_{3} . E=R_{1} A\right) \wedge\left(R_{2} . F=R_{4} . G\right)\right)$. (a) Join graph $J_{Q .}$ (b) Query graph $G_{Q}$.
resulting query graph contains edges $\left\{\left(R_{1}, R_{2}\right),\left(R_{2}, R_{3}\right)\right\}$ and is now acyclic; we have replaced an interrelation clause by an intrarelation clause that does not produce a query graph edge. Intuitively, we have replaced a join between two relations by a restriction within a single relation, thereby breaking the cycle. As we will now show formally, the transformations described by the above two examples are the only ones needed to map a tree query whose query graph is cyclic into an equivalent query whose query graph is a tree.
To perform the above transformations on clauses, we will use another graph model of a query. For a query $Q$ with qualification $q$, we define the join graph for $Q$, $J_{Q}(V, E)$, to be a node-labeled undirected graph where

$$
V=\left\{R_{i} \cdot A \mid A \text { is a domain of } R_{i}\right\}
$$

and
*

$$
E=\left\{\left(R_{i} \cdot A, R_{j} \cdot B\right) \mid\left(R_{i} \cdot A=R_{j} \cdot B\right) \text { is a join clause in } Q\right\} .
$$

The join graph simply represents the joins in a qualification by edges in the graph (e.g., see Figure 5). We say that a join graph $J$ corresponds to a query graph $G$ if the query represented by $J$ has (the unique) query graph $G$. Unlike query graphs, join graphs are not multigraphs.
The transitive closure of $J_{Q}$, denoted $J_{Q}^{\dagger}$, represents all join clauses that are logically implied by $Q^{\prime}$ 's qualification. A spanning forest of $J_{Q}^{+}$is defined as a minimal subgraph of $J_{Q}^{+}$whose closure is $J_{Q}^{+}$. A spanning forest of $J_{Q}^{+}$is in some sense a minimal representation of a query. The following lemma says that to test if $Q \in T Q$ we need only look at spanning forests of $J_{Q}^{+}$.

Lemma 11. $Q \in T Q$ iff there exists a spanning forest of $J_{Q}^{+}$that corresponds to an acyclic query graph.

Proof. Since a spanning forest of $J_{Q}^{+}$has the same transitive closure as $J_{Q}$, the query represented by the spanning forest is equivalent to $Q$. If the spanning forest corresponds to an acyclic query graph, then $Q$ is obviously equivalent to a query whose query graph is a tree, and hence $Q \in T Q$. If $Q \in T Q$, then it is equivalent to some $Q^{\prime}$ whose query graph is a tree. Any spanning forest of its join graph $J_{Q^{\prime}}$ is a spanning forest of $J_{Q}^{+}$and therefore satisfies the conditions of the lemma.

There is a particular kind of spanning forest of $J_{Q}^{+}$which is easy to compute and which has the property that $Q \in T Q$ iff such a spanning forest corresponds to an acyclic query graph. This kind of canonical spanning forest is constructed as follows


Fig. 6. A canonical spanning forest for query $Q$ with qualification $q=\left(\left(R_{1}, A=R_{2} . C\right) \wedge\left(R_{2 .} . C=R_{2} . D\right)\right.$ $\wedge\left(R_{2} . C=R_{1} \cdot B\right) \wedge\left(R_{2} . D=R_{3 .} . F\right) \wedge\left(R_{1} \cdot B=R_{3} . F\right) \wedge\left(R_{3} \cdot E=R_{4} \cdot G\right)$. (a) Join graph $J_{Q}($ b) A canonical spanning forest of $J_{Q}{ }^{+}$
(see Figure 6 for an example):
do for each connected component of $J_{Q}$;
partition the nodes according to the relations that name them; for each partition, connect all the nodes in the partition by a chain; choose one node from each partition and connect them by a chain; end;

Each tree in the forest spans a connected component of $J_{Q}^{+}$. Since there is some choice in constructing each tree, canonical spanning forests are not unique for each $J_{Q}$. However, the following lemma shows that any one of them can tell if $Q \in T Q$.

Lemma 12. For every query $Q$ and for every canonical spanning forest $J^{C}$ of $J_{Q}^{+}$, $Q \in T Q$ iff $J^{C}$ corresponds to an acyclic query graph.

Proof. The if part follows from Lemma 11. To prove the only-if part, suppose $J^{\mathrm{C}}$ corresponds to a cyclic query graph $G^{\mathrm{C}}$. We will show that every spanning forest of $J_{Q}^{+}$corresponds to a cyclic query graph, so that by Lemma 11, $Q \notin T Q$. Therefore $Q \in T Q$ must imply that every $J^{C}$ corresponds to an acyclic query graph.

Consider a simple cycle in $G^{\text {C }}$, i.e., a cycle in which each node is incident with at most two edges in the cycle. The edges of this cycle map directly into edges of $J^{\mathrm{c}}$; suppose this corresponding sequence of edges of $J^{\mathrm{C}}$ is $S=\left[\left(v_{0}, v_{1}\right),\left(v_{2}, v_{3}\right), \ldots\right.$, ( $v_{2 n-2}, v_{2 n-1}$ )]. Clearly, $v_{2 i-1}$ and $v_{2 i}$ are labeled by domains of the same relation for $i=1, \ldots, n($ arithmetic is $\bmod 2 n)$. For each $i, 1 \leq i \leq n$, if $v_{2 i-1}$ and $v_{2 i}$ are in the same connected component, then by the construction of canonical spanning forests, $v_{2 i-1}=v_{2 i}$. Thus $S$ can be represented as a sequence of $p+1$ paths, $p \geq 0$, $\left[\left[\left(v_{0,0}, v_{0,1}\right),\left(v_{0,1}, v_{0,2}\right), \ldots,\left(v_{0, m_{0}-1}, v_{0, m_{0}}\right)\right], \ldots,\left[\left(v_{p, 0}, v_{p, 1}\right), \ldots,\left(v_{p, m_{p}-1}, v_{p, m_{p}}\right)\right]\right]$, where each path is entirely contained in a connected component of $J_{Q}^{+}$.

We claim $p \geq 1$. Suppose $p=0$. Since $\nu_{0,0}$ and $\nu_{0, m_{0}}$ are in the same connected component of the join graph and are in the same relation, the construction of canonical spanning forests requires that $v_{0,0}=v_{m, 0}$. But then the path in the canonical spanning forest is a cycle, a contradiction. So $p>0$.

Consider one of the paths, $\left[\left(v_{i, 0}, v_{i, 1}\right), \ldots,\left(v_{i, m_{i}-1}, v_{i, m_{i}}\right)\right]$. Since the transitive closure of any spanning forest of $J_{Q}^{+}$equals the transitive closure of $J^{\mathrm{C}}$, there must be a path from $v_{i, 0}$ to $\nu_{i, m_{i}}$ in all spanning forests of $J_{Q}^{+}$. This holds for $0 \leq i \leq p$. Hence for each spanning forest $F$ there exists in $F$ a sequence of $p+1$ paths $S_{F}$ with the same endpoints as the $p+1$ paths of $S$.

To complete the proof, we must show that for each spanning forest $F$, the sequence $S_{F}$ produces a cycle in its corresponding query graph. Since we chose a simple cycle in $G^{\mathrm{C}}, v_{m_{i}}$ and $v_{m_{j}}$ are named by domains of different relations if $i \neq j$. Since $p>0$, we have at least two distinct nodes in the query graph path that corresponds to $S_{F}$. So the path is a cycle. Since all spanning forests have a cycle, by Lemma $11 Q \in T Q$, a contradiction.

Theorem 3. Testing if $Q \in T Q$ can be decided in linear time.
Proof. The algorithm for testing if $Q \in T Q$ is

1. Construct a canonical spanning forest $J^{\mathrm{c}}$ of $J_{Q}$.
2. Construct the query graph $G^{\mathrm{C}}$ corresponding to $J^{\mathrm{C}}$.
3. If $G^{\mathrm{C}}$ is acyclic, then answer $Q \in T Q$; else answer $Q \in C Q$.

By Lemma 12, the algorithm is correct. Both spanning forest and cycle detection can be computed in time linear in the number of edges of the graph [1]; so the algorithm has linear worst-case time.

A tree query membership test when multiattribute semi-joins are permitted appears in [3].

We conclude by proving Lemma 8 from Section 4.3 as promised.
Lemma 8. Query $Q$ with qualification $q$ is in $C Q$ iff $q$ has a proper cyclic subqualification.

Proof. If. Let $q^{\prime}$ be a proper cyclic subqualification of $q$. When we construct the canonical spanning forest, in the third step of the do-loop we choose to include a clause of $q^{\prime}$ as an edge in a chain whenever possible. Clearly, every clause of $q^{\prime}$ can be added, since they are all in $q^{+}$. Also, by part (3) of the definition of proper cyclic subqualification, each connected component of the canonical spanning forest can contain at most one clause from $q^{\prime}$; so, incorporating clauses from $q^{\prime}$ can never cause a cycle in a connected component. Since all clauses of $q^{\prime}$ are embedded in the spanning forest, the query graph corresponding to the forest has a cycle. By Lemma 11, $Q \in C Q$.
Only if. Suppose $Q \in C Q$. Construct a canonical spanning forest $F$ for $Q$. By Lemma 11 there is a cycle in the corresponding query graph. Select such a cycle, and let $C$ be the edges of $F$ that correspond to the edges of the cycle. From $C$, construct a qualification $q^{\prime}$ as follows: For each path in $C$ which is not a proper subpath of any other path in $C$ whose endpoints are, say, $R_{i} . A$ and $R_{j} . B$, include ( $R_{i} . A=R_{j} . B$ ) in $q^{\prime}$. Clearly $q^{\prime}$ is in $q^{+}$and $q^{\prime}$ is proper cyclic. Also, if ( $R_{i} . A=R_{j} . B$ ) is in $q^{+}$and $R_{i} . A$ and $R_{j} . B$ are in DOMAINS $\left(q^{\prime}\right)$, then $\left(R_{i} . A=R_{j} . B\right)$ is in $q^{\prime}$; this follows because $R_{i} . A$ and $R_{j} . B$ must be in the same connected component of $F$, and each connected component of $F$ has at most one clause in $q^{\prime}$ by construction, so $q^{\prime}$ is a proper cyclic subqualification of $q$ as desired.

## 6. Conclusions

Semi-joins are commonly used as basic operations in query processing algorithms, especially in a distributed environment. In these algorithms a strategy is simply a semi-join program. Our results show that for tree queries, as soon as a strategy embeds $\sigma_{Q}$, a full reduction and, hence, the full potential of semi-joins are achieved. Consequently the strategies tend to be short; at worst their length is bounded by a linear function of the number of relations referenced. Furthermore, tree-query membership can be tested in linear time. This suggests that searching for optimal strategies is quite likely to be easier for tree queries than for cyclic ones. We therefore recommend that the optimization problem for tree queries be treated as an important (and probably more tractable) special case of the general query processing problem.
For cyclic queries, finding good semi-join programs is likely to be quite difficult. Since cyclic queries are common too, we will either need to find new tactics (other
than semi-joins) for solving them or will probably have to be satisfied with heuristic approaches such as that of [14].

We emphasize, however, that our results only describe worst-case behavior of semijoin strategies that try to achieve full reductions. The pathological case that shows cyclic queries to behave badly may rarely occur. More important, full reductions are not always profitable, and a query optimizer should only produce the most profitable semi-join strategies. So while we have learned much about the strategy space of semijoin programs, the query optimization problem on "average" databases using the semi-join tactic remains open.
acknowledgments. We thank Wing S. Wong for simplifying our first version of the $T Q$ membership algorithm, which resulted in the one presented in Section 5. Many notational improvements and clarifications over earlier drafts were contributed by Marco Casanova, Nat Goodman, and the referees. Finally, we thank M. Clarke, R. D'Arcangelo, and C. Louis for their expert preparation of this manuscript.

## REFERENCES

1. Aho, A., Hopcroft, J., And Ullman, J. The Design and Analysis of Computer Algorithms. AddisonWesley, Reading, Mass., 1974.
2. Aho, A., Sagiv, Y., and Ullman, J.D. Equivalence among relational expressions. SIAM J. Comput. 2 (May 1979), 218-246.
3. Bernstein, P.A., and Goodman, N. Full reducers for relational queries using multi-attribute semijoins. Proc. Computer Networking Symposium, Gaithersburg, Md., 1979, pp. 206-215.
4. Chamberlin, D.D., Astrahan, M.M., Eswaran, K.P., Griffiths, P.P., Lorie, R.A., Mehl, J.W., Reisner, P., and Wade, B.W. SEQUEL 2: A unified approach to data definition, manipulation, and control. IBM J. Res. Dev. 20, 6 (Nov. 1976), 560-575.
5. Chandra, A.K., and Merlin, P.M. Optimal implementation of conjunctive queries in relational data bases. Proc. 9th Ann. ACM Symp. on Theory of Computing, Boulder, Colo., May 1976, pp. 7790.
6. CoDd, E.F. A relational model of data for large shared data banks. Commun. ACM 13, 6 (June 1970), 377-387.
7. Codd, E.F. Relational completeness of data base sublanguages. In Data Base Systems, R. Rustin, Ed., Courant Computer Science Symposia Series, Vol. 6, Prentice-Hall, Englewood Cliffs, N.J., 1972, pp. 65-90.
8. Held, G.D., Stonebraker, M., and Wong, E. INGRES-A relational data base management system. Proc. AFIPS 1975 National Computer Conf., AFIPS Press, Arlington, Va., 1975, pp. 409-416.
9. Ozkarahan, E.A., Schuster, S.A., and Sevcik, K.C. Performance evaluation of a relational associative processor. ACM Trans. Database Syst. 2, 2 (June 1977), 175-196.
10. Rothnie, J.B. Evaluating inter-entry retrieval expressions in a relational database management system. Proc. AFIPS 1975 National Computer Conf., AFIPS Press, Arlington, Va., 1975, pp. 417-423.
11. Rothnie, J.B., and Goodman, N. A survey of research and development in distributed database management. Proc. 3rd Inter. Conf. on Very Large Data Bases, Tokyo, Japan, 1977, pp. 48-61.
12. Smith, J.M., and Chang, P.Y.-T. Optimizing the performance of a relational algebra database interface. Commun. ACM 18, 10 (Oct. 1975), 568-579.
13. Wong, E., and Youssefi, K. Decomposition-A strategy for query processing. ACM Trans. Database Syst. 1, 3 (Sept. 1976), 223-241.
14. Wong, E. Retrieving dispersed data from SDD-1: A system for distributed databases. Proc. 1977 Berkeley Workshop on Distributed Data Management and Computer Networks, Berkeley, Calif., May 1977.

RECEIVED DECEMBER 1978; REVISED SEPTEMBER 1979; ACCEPTED SEPTEMBER 1979


[^0]:    ${ }^{1}$ The join of relations $R(A, B)$ and $S(B, C)$ on domain $B$ is the set of tuples $\{\langle a, b, c\rangle \in A \times B \times$ $C \mid(a, b) \in R$ and $\langle b, c\rangle \in S\}$. The projection of relation $T(A, B, C)$ on domains $A$ and $B$ is the set $\{(a, b) \in A \times B \mid \exists c \in C((a, b, c\rangle \in T)\}$, that is, the join of $R$ and $S$ projected back on the domains of $R$.

[^1]:    2 The lack of negation, an existential quantifier, and tuple variables imply that equi-join queries are not
    "relationally complete" in the sense of $[7]$ "relationally complete" in the sense of [7].

[^2]:    ${ }^{3}$ For notational convenience, we shall use $i$ in place of $R_{i}$.

