

A Semideterministic Approach to Object Creation and Nondeterminism in Database Queries*

Jan Van den Bussche[†]

Department of Mathematics & Computer Science, University of Antwerp (UIA), Universiteitsplein 1, B-2610 Antwerp, Belgium

and

Dirk Van Gucht[‡]

Computer Science Department, Indiana University, Bloomington, Indiana 47405-4101

Received May 6, 1996

We introduce and study the concept of semideterminism. A non-deterministic, generic query is called semideterministic if any two possible results of the query to a database are isomorphic. Semideterminism is a generalization of determinacy, proposed by Abiteboul and Kanellakis in the context of object-creating query languages. The framework of semi-deterministic queries is less restrictive than that of the determinate queries and avoids the problem of copy elimination connected with determinacy. We argue that semideterminism is also interesting in its own right and show that it is natural and desirable, although hard to achieve in general. Nevertheless, we exhibit two major applications where semideterministic computations are possible. First, we show that there is a universal procedure to compute any semideterministic query in a semideterministic manner. Second, we show that the polynomial-time counting queries can be efficiently expressed semideterministically. © 1997 Academic Press

1. INTRODUCTION

The theory of queries in the context of the conventional relational database model is well understood. For a survey, see [3]. The theory started with the work of Chandra and Harel [10], who formally defined a *query* as a function from databases to databases which is partial recursive and preserves isomorphisms. They also presented a query language called QL, an extension of the relational algebra with unbounded looping and the possibility to store intermediate results in relations with unbounded arity. It was furthermore shown that QL is complete in the sense that it can express exactly all queries.

* An extended abstract of a preliminary version of this paper was presented at the *11th ACM Symposium on Principles of Database Systems*.

[†] Current address: Dept. WNI, University of Limburg (LUC), B-3590 Diepenbeek, Belgium. E-mail: vdbuss@luc.ac.be.

[‡] E-mail: vgucht@cs.indiana.edu.

Abiteboul and Vianu [5] extended the notion of query by allowing nondeterminism. They redefined a query as a binary relationship between databases which is recursively enumerable and preserves isomorphisms.¹ If the binary relationship happens to be a function, we call the query *deterministic* and the deterministic queries are exactly those of the previous paragraph. A language called TL [5, 6] was then proved to be complete for all (possibly nondeterministic) queries. TL is equivalent to the extension of the relational algebra with (i) unbounded looping, (ii) an operation to choose arbitrary tuples from a relation, and (iii) an operation to create new objects by tagging each tuple of a relation with a different new object (which can be thought of as an identifier.) It was also shown in [5] that detTL, a deterministic version of TL obtained by disallowing the choice operation and the appearance of new objects in the final result, is complete for the deterministic queries. In particular, this showed that object creation in intermediate relations is an alternative to the intermediate relations with unbounded arity of QL.

Motivated by applications in object-oriented database systems [19], the need arose for queries where new objects *do* appear in the final result. This led Abiteboul and Kanellakis to the study of IQL [4], a query language roughly comparable in expressive power to detTL, but without the prohibition of new objects in the result. In an attempt to capture the queries expressible in IQL, they defined the intuitively appealing class of *determinate* queries.² A nondeterministic query is called determinate if any two different results of the query applied to a database

¹ The work of Abiteboul and Vianu was not restricted to queries only, but also included updates. In this paper we will focus on the query aspect.

² Determinacy was also implicit in Kuper's thesis [21].

are isomorphic through an isomorphism that is the identity on the domain of the input database. Hence, the isomorphism can only permute the new objects. Determinacy thus intends to isolate the very weak form of nondeterminism which is needed to accommodate new objects.

It turned out that IQL is not complete for the determinate queries: an extra operation called *copy elimination* [4] has to be added in order to obtain determinate-completeness. Nevertheless, most likely all determinate queries arising in practice can already be expressed in IQL (without copy elimination.) So, one could conclude that the class of determinate queries is perhaps not so natural after all. This philosophically unsatisfactory situation left open two natural directions for further research: restrict or extend the class of determinate queries to find more natural classes.

The first direction was explored in [27], where the precise expressive power of IQL was characterized. It was shown that this characterization is obtained when an additional requirement, besides determinacy, is made which simply expresses that the creation of new objects can be interpreted as the deterministic construction of hereditarily finite sets. The feeling expressed above that the class of IQL-expressible determinate queries, coined the *constructive* queries in [27], indeed arises very naturally was thus confirmed.

In the present paper, we explore the other direction. We propose the notion of *semideterminism*. A nondeterministic query is called semideterministic if any two different results of the query to a database are isomorphic, by an isomorphism which maps the input database onto itself (i.e., an automorphism.) Note that determinate queries are very restricted semideterministic queries, for which this automorphism is the identity. Thus, semideterminism is a natural generalization of determinacy and, at the same time, a natural restriction on arbitrary nondeterminism; the only sources of nondeterminism are the symmetries (automorphisms) present in the input database.

The contents of this paper can be summarized as follows. In Section 2, we review the necessary preliminary notions.

In Section 3, we introduce and motivate semideterminism. Given that semideterminism is an extension of determinacy, we provide a necessary and sufficient algebraic condition for when a semideterministic query is actually determinate. We also present an alternative characterization for semideterminism which demonstrates its naturalness. Another motivation for semideterminism which we prove is that a nondeterministic loop program which makes its choices in a semideterministic manner, has the good property that if one of its possible computations halts, then all possible computations will halt, and this after the same number of iterations of the loop. We also observe that semideterministic queries which yield a yes–no answer must in fact be deterministic, which is quite desirable.

While semideterminism has good properties, it is hard to achieve in general. Indeed, we prove that the problem of

syntactically verifying whether a nondeterministic program expresses a semideterministic query is undecidable. Furthermore, checking a computation for semideterminism at run-time is shown to be polynomial-time equivalent to checking graph isomorphism.

Nevertheless, in Sections 4 and 5, we exhibit two major applications where semideterministic computations are possible: *completeness* and *counting*. In both of these applications, object creation plays a crucial role.

In Section 4, we prove that every semideterministic query can be expressed by a program in TL that contains only one single application of the choice operation. This application is semideterministic and is a generalization of the copy elimination operation to the semideterministic context. This result indicates that the obvious candidate language for semideterministic completeness, namely the semideterministic TL programs, is indeed semideterministic complete. It also suggests that the notion of copy elimination naturally falls in the framework of semideterministic queries.

In Section 5, we show that all polynomial-time counting queries can be expressed efficiently in a semideterministic manner. Queries involving counting are computationally simple but cannot be expressed in the extension of the relational algebra with looping. Using nondeterminism, this situation can be remedied: it is well known [3] that all polynomial-time computable queries can be expressed efficiently in the extension of the relational algebra with looping and nondeterministic choice. Our result shows that, as far as counting is concerned, the advantages of nondeterminism can also be obtained semideterministically. To this end, we demonstrate a new, semideterministic technique for constructing restricted types of orderings that are sufficient for counting purposes.

Concluding remarks are in Section 6.

2. PRELIMINARIES

In this section, we review the necessary preliminary notions. We will work in the following version of the well-known relational database model [3]. Formally, assume the existence of sufficiently many *relation names*. Every relation name R has an associated *arity* $a(R)$, a natural number. For each natural number n there are sufficiently many relation names R with $a(R) = n$. A *database scheme* is a finite set of relation names. Furthermore, let \mathbf{U} be an infinite, recursively enumerable universe of data elements called *atomic objects* (or *objects* for short). An *instance* I over a scheme \mathcal{S} is a mapping, assigning to each relation name R of \mathcal{S} a finite relation $I(R) \subset \mathbf{U}^{a(R)}$. The *active domain* of a relation r (or an instance I) is the set of all objects occurring in it, and is denoted by $\text{adom}(r)$ ($\text{adom}(I)$). The set of all instances over a scheme \mathcal{S} is denoted by $\text{inst}(\mathcal{S})$. If $\mathcal{S} = \{R_1, \dots, R_n\}$, then an instance $I \in \text{inst}(\mathcal{S})$ will be denoted as $I = (R_1 : I(R_1), \dots, R_n : I(R_n))$.

In its most general form, a query can be thought of as a possibly nondeterministic process, augmenting databases with derived information. Formally, let \mathcal{S} be a scheme, and let A_1, \dots, A_n be relation names not in \mathcal{S} . The following definition is adapted from [5].

DEFINITION 2.1. A query of type $\mathcal{S} \rightarrow A_1, \dots, A_n$ is a recursively enumerable, binary relationship $Q \subseteq \text{inst}(\mathcal{S}) \times \text{inst}(\mathcal{S} \cup \{A_1, \dots, A_n\})$ such that:

1. if $Q(I, J)$ then J is equal to I on \mathcal{S} ; and
2. if $Q(I, J)$ and f is a permutation of \mathbf{U} then $Q(f(I), f(J))$.

The latter requirement is now commonly called *genericity* and traces back to Refs. [8, 10].

We will sometimes use the following notation: if ρ is a binary relationship, then $\rho(x)$ stands for the set $\{y \mid \rho(x, y)\}$. For a query Q and instance I , $Q(I)$ is called the set of *possible results* of Q for I . If $Q(I) = \emptyset$ then the result of Q is undefined on input I .

Queries that are functions are called *deterministic*. A basic language for expressing deterministic queries is RA. This is a version of the relational algebra [3], an algebraization of first-order predicate logic (a.k.a. relational calculus), consisting of the relational operators union (\cup), difference ($-$), Cartesian product (\times), selection ($\sigma_{i=j}$, selects from a relation those tuples for which the i th and the j th component are equal), and projection (π_{i_1, \dots, i_k}). Let \mathcal{S} be a scheme and let R_1, \dots, R_m be a sequence of relation names not in \mathcal{S} . They will be used as relation-valued variables, holding intermediate results of the computation. The sequence P of *relation assignments*:

$$R_1 := E_1; \dots; R_m := E_m,$$

where each E_i is a relational algebra expression using only relation names in $\mathcal{S} \cup \{R_1, \dots, R_{i-1}\}$, is a RA *program*. If we choose a subset of *answer relation* names $\{A_1, \dots, A_n\}$ among the R_i 's, then this program P computes a query of type $\mathcal{S} \rightarrow A_1, \dots, A_n$ in the obvious and well-known way. We will not always explicitly indicate the answer relations in a program; they will often be clear from the context.

A query Q is called *object-creating* if there exist I, J such that $Q(I, J)$ and $\text{adom}(I) \subsetneq \text{adom}(J)$ (note that for any query Q , $Q(I, J)$ implies $\text{adom}(I) \subseteq \text{adom}(J)$). Object-creating queries are necessarily nondeterministic, because of genericity. In proof, assume $Q(I, J)$ and $o \in \text{adom}(J) - \text{adom}(I)$. Take an arbitrary $o' \in \mathbf{U}$ such that $o' \notin \text{adom}(J)$, and consider the permutation f of \mathbf{U} which transposes o and o' : $f = (o \ o')$. Then $f(I) = I$ and $f(J) \neq J$ and, by genericity, $Q(I, f(J))$.

A natural object-creating operation which can be added to RA is *new*. Let \mathcal{S} be a scheme, $I \in \text{inst}(\mathcal{S})$, and $R \in \mathcal{S}$ with $\text{a}(R) = a$. A relation r of arity $a + 1$ is a possible result

of $\text{new}(R)$ applied to I if r is obtained from $I(R)$ by extending each tuple t of $I(R)$ with a different new object $t(a + 1)$ that is not yet in $\text{adom}(I)$.

EXAMPLE 2.2. Let $\mathcal{S} = \{R\}$ with $\text{a}(R) = 2$ and

$$I(R) = \{\langle a, b \rangle, \langle a, c \rangle, \langle b, d \rangle\}.$$

A possible result of $\text{new}(R)$ applied to I is

$$\{\langle a, b, \alpha_1 \rangle, \langle a, c, \alpha_2 \rangle, \langle b, d, \alpha_3 \rangle\},$$

where the α_i are three arbitrary new objects from \mathbf{U} .

So, in a sense, *new* is the converse of projection. We can add *new* to the operations of the relation algebra, thus obtaining the query language RA + new.

The language RA + new provides a basic and general object creation mechanism in function of tuples. Essentially this same mechanism is provided by many other query languages with object-creating capabilities that have been proposed in the literature [4–6, 16–18, 22]. Some languages [4, 16, 22, 24] provide also object creation in function of sets. We will not explicitly need this capability to prove our results, but we will return to it in Section 6.

While we know that an object-creating query must be nondeterministic, it is clear that the queries expressed by programs of RA + new are “nearly” deterministic: different results of these queries to a database differ only in the particular choice of the new objects that have been created. To capture this intuition, the notion of *determinacy* was defined in [4].

DEFINITION 2.3. Query Q is called *determinate* if whenever $Q(I, J_1)$ and $Q(I, J_2)$, then J_1 and J_2 must be *I-isomorphic*; i.e., there must be a permutation f of \mathbf{U} which is the identity on $\text{adom}(I)$, such that $f(J_1) = J_2$.

We now define an operation, W (Witness, [3]), which allows for arbitrary nondeterminism. Let \mathcal{S} be a scheme, $I \in \text{inst}(\mathcal{S})$, and $R \in \mathcal{S}$ with $X \subseteq \{1, \dots, \text{a}(R)\}$. A relation r of arity $\text{a}(R)$ is a possible result of $W_X(R)$ applied to I if r is a subset of $I(R)$ obtained by choosing for each class of X -equivalent tuples of $I(R)$ exactly one representative. Here, two tuples are called *X-equivalent* if they are equal outside X . In particular, if $X = \{1, \dots, \text{a}(R)\}$, then any two tuples are X -equivalent.

EXAMPLE 2.4. Let \mathcal{S} and I be as in Example 2.2. We interpret $I(R)$ as a parent–child relation, then $W_2(R)$ amounts to choosing a child for each parent. The two possible results of $W_2(R)$ applied to I are

$$r_1 = \{\langle a, b \rangle, \langle b, d \rangle\}$$

and

$$r_2 = \{\langle a, c \rangle, \langle b, d \rangle\}.$$

On the other hand, $W_{1,2}(R)$ chooses an arbitrary tuple from $I(R)$, so the three possible results of $W_{1,2}(R)$ applied to I are $\{\langle a, b \rangle\}$, $\{\langle a, c \rangle\}$, and $\{\langle b, d \rangle\}$.

We can extend RA and RA + new with nondeterminism by allowing assignments of the form $A := W_X(E)$, where E is a relational algebra expression. The resulting language is denoted by RA + W or RA + new + W. Note that we allow applications of the Witness operation only at the end of a relational algebra expression in an assignment, not within expressions. This is no real restriction and will make it easier to define semideterministic programs in Section 3.

Finally, all languages introduced so far must often be enriched with a looping construct to increase expressiveness. We will use here a loop with partial fixpoint semantics [3]. This is a “repeat-while-change” looping semantics, but special care must be taken to define it in the presence of nondeterminism. Let P be a program in RA + new + W. Let Q be the query expressed by P and let Q^n be the query expressed by P^n , the n -fold concatenation of P . Then the query Q' expressed by the program $\text{loop}[P]$ is defined as follows: On input instance I , $J \in Q'(I)$ if there exists an n such that $J \in Q^n(I) \cap Q(J)$. We can extend any query language \mathcal{L} to $\mathcal{L} + \text{loop}$ as follows: (i) if P is an \mathcal{L} -program, then $\text{loop}[P]$ is also a program; (ii) programs can be composed using “;”.

If P is determinate, then the above amounts to the typical fixpoint semantics. So either $\text{loop}[P]$ is undefined on I since it does not halt, or the result equals $Q^n(I)$ for the least n such that $Q^n(I) = Q^{n+1}(I)$. However, if P is arbitrarily nondeterministic, $\text{loop}[P]$ can have a much more complicated behavior. Some possible computations may halt, while others may not halt. Furthermore, not all possible results of the loop will be computed after the same number of iterations.

EXAMPLE 2.5. Let R be a binary relation, viewed as (the set of edges of) a directed graph. Let Q be the query of type $\{R\} \rightarrow A$, where A is binary, defined by $Q(I, J)$ if J^A lists each node x in I together with an arbitrary node reachable from x in R^I . This query is expressible by the following program:

$$D := \pi_1(R) \cup \pi_2(R);$$

$$E := \sigma_{1=2}(D \times D);$$

$$A := W_2(R);$$

$$\text{loop}[A := W_2(\pi_{1,4}\sigma_{2=3}(A \times (R \cup E)))].$$

Possible computations of the loop in this program, when applied to a fixed instance, may halt after any number of iterations, or may loop indefinitely.

We will show later (Theorem 3.13) that nondeterministic loop programs do have a nice behavior if the nondeterminism is in fact semideterministic.

The language which includes all features introduced in this section, RA + new + W + loop, will often be referred to as TL, since it is equivalent to the language TL of [5] (see also [6]). TL is complete; it can express all queries.

3. SEMIDETERMINISM

In this section, we introduce a motivate semideterminism and also show that it is hard to achieve in general.

3.1. Definition and General Properties

Semideterminism is a natural restriction on the amount of nondeterminism of a query Q :

DEFINITION 3.1. Query Q is called *semideterministic* if whenever $Q(I, J_1)$ and $Q(I, J_2)$, then J_1 and J_2 are isomorphic, i.e., there is a permutation f of \mathbf{U} such that $f(J_1) = J_2$.

By definition of the query, J_1 and J_2 are extensions of I , and it follows that the stated isomorphism f from J_1 to J_2 is an automorphism of I . (In this paper, with an automorphism of I we mean a permutation f of \mathbf{U} such that $f(I) = I$.) So intuitively, the only sources of nondeterminism in a semideterministic query are the symmetries (automorphisms) in the input database.

Clearly, determinate queries (Definition 2.3) are very restricted semideterministic queries, for which the above-mentioned automorphism is actually the identity on I . The difference between determination and semideterminism can also be characterized by the next theorem. We say that an automorphism f of an instance I can be *extended* to an automorphism of an instance J , with $\text{adom}(J) \supseteq \text{adom}(I)$, if there is a permutation g of \mathbf{U} such that $g|_{\text{adom}(I)} = f|_{\text{adom}(I)}$ and $g = (J) = J$.

THEOREM 3.2. *A semideterministic query Q is determinate if and only if whenever $Q(I, J)$ and f is an automorphism of I , then f can be extended to an automorphism of J_2 .*

Proof. (If) Assume $Q(I, J_1)$ and $Q(I, J_2)$. Since Q is semideterministic, there is a permutation f of \mathbf{U} such that $f(J_1) = J_2$. f is an automorphism of I , so it can be extended to an automorphism of J_2 . This means there is a permutation g such that $g|_{\text{adom}(I)} = f|_{\text{adom}(I)}$ and $g(J_2) = J_2$. Let p be the order of $g|_{\text{adom}(I)}$ in the permutation group of $\text{adom}(I)$. I.e., p is the least such that $p > 0$ and $(g|_{\text{adom}(I)})^p$ is the identity. Define $h := g^{(p-1)}f$. Then $h(J_1) = J_2$, and h is the identity on $\text{adom}(I)$. Therefore, Q is determinate.

(Only if) Let f be an automorphism of I . By genericity, $Q(I, f(J))$. By determinacy, there is permutation g such that

$g(f(J)) = J$ and g is the identity on $\text{adom}(I)$. So, $h := gf$ satisfies $h|_{\text{adom}(I)} = f|_{\text{adom}(I)}$ and $h(J) = J$, as desired. ■

The above theorem also yields the following.

COROLLARY 3.3. *A semideterministic query Q is deterministic if and only if whenever $Q(I, J)$ and f is an automorphism of I , then f is also an automorphism of J .*

Proof. (Only if) Suppose for the sake of contradiction that f is not an automorphism of J . This means that $f(J) \neq J$. But by genericity, $Q(I, f(J))$, whence Q is not deterministic, a contradiction.

(If) If $Q(I, J)$ and f is an automorphism of I , then surely f can be trivially extended to an automorphism of J . Hence, by Theorem 3.2, Q is determinate. So to prove that Q is deterministic, it suffices to prove that Q is not object-creating. Suppose for the sake of contradiction that Q is object-creating. Then there exist I, J such that $Q(I, J)$ and $\text{adom}(J) \not\subseteq \text{adom}(I)$. Let $o \in \text{adom}(J) - \text{adom}(I)$. Let f be a permutation of \mathbf{U} such that $f(I) = I$ and $f(o) \notin \text{adom}(J)$. But the $f(J) \neq J$, so f is an automorphism of I but not of J , a contradiction. ■

The qualification *semideterministic* cannot be omitted from Theorem 3.2 or Corollary 3.3. To show this for Theorem 3.2, consider the query Q of type $\{V\} \rightarrow W$, where $a(V) = a(W) = 1$, defined as follows: $Q(I, J)$ if $J(W) = I(V)$ or if $J(W) = I(V) \cup \{o\}$ for some arbitrary $o \notin I(V)$. Although every automorphism of I can be extended to an automorphism of J , Q is not determinate. For Corollary 3.3, consider the query Q of type $\{V\} \rightarrow R$, where $a(V) = 1$ and $a(R) = 2$, defined as follows: $Q(I, J)$ if $I(V)$ is of the form $\{a, b\}$ and $J(W)$ is either of the form:

$$\{\langle a, a \rangle, \langle b, b \rangle\}$$

or of the form:

$$\{\langle a, b \rangle, \langle b, a \rangle\}.$$

Although every automorphism of I is also an automorphism of J , Q is not deterministic.

We also point out that, since we defined queries as augmentations, the following straightforward but fundamental closure property holds:

PROPOSITION 3.4. *The composition of two semideterministic queries is again semideterministic.*

We now present an interesting characterization of the semideterministic queries [14]. To do this, we need the following auxiliary notion:

DEFINITION 3.5. Let $\mathcal{S}, A_1, \dots, A_n$ be as in Definition 2.1. A *prequery* of type $\mathcal{S} \rightarrow A_1, \dots, A_n$ is a recursively enumerable, binary relationship $Q_0 \subset \text{inst}(\mathcal{S}) \times \text{inst}(\mathcal{S} \cup \{A_1, \dots, A_n\})$ such that:

1. if $Q_0(I, J)$ then J is equal to I on \mathcal{S} ; and
2. if $Q_0(I, J)$ and $Q_0(I', J')$ then either $I = I'$ and $J = J'$, or I and I' are not isomorphic.

So, a prequery is an arbitrary partial recursive function from instances to instances with the property that it is defined on at most one representative of each isomorphism type. In particular, prequeries are never generic. Therefore, it makes sense to define the closure of a prequery under the genericity requirement:

DEFINITION 3.6. The *closure* of prequery Q_0 is the query

$$Q_0^* := \{(f(I), f(J)) \mid Q_0(I, J) \text{ and } f \text{ permutation of } \mathbf{U}\}.$$

We observe:

PROPOSITION 3.7. *For any prequery Q_0 , its closure Q_0^* is recursively enumerable.*

Proof. The standard approach to enumerate Q_0^* would be to enumerate the Cartesian product of Q_0 with the set of all permutations f of \mathbf{U} . The only problem with this approach is that the latter set is uncountable. However, since we are working with finite instances, it suffices to consider only the permutations having finite support, i.e., those permutations that are the identity on all but a finite number of elements. Since every such permutation is the composition of a finite number of transpositions, it thus suffices to enumerate the Cartesian product of Q_0 with the set of all finite sequences of transpositions. ■

Hence, Q_0^* is the minimal query containing Q_0 . We now establish:

THEOREM 3.8. *Query Q is semideterministic if and only if $Q = Q_0^*$ for some prequery Q_0 .*

Proof. (If) Assume $Q(I, J_1)$ and $Q(I, J_2)$. By the definition of Q_0^* , there are instances I_0, J_0 and permutations f_1, f_2 of \mathbf{U} such that $Q_0(I_0, J_0)$, $I = f_1(I_0)$ and $J_1 = f_1(J_0)$, and $I = f_2(I_0)$ and $J_2 = f_2(J_0)$. Hence, $f_2 f_1^{-1}$ is the desired isomorphism from J_1 to J_2 .

(Only if) Let Q_0 be any maximal prequery contained in Q . Since Q is recursively enumerable, such a Q_0 exists. We show that $Q = Q_0^*$. Since Q is generic, it suffices to show that $Q \subseteq Q_0^*$. Assume $Q(I, J)$. Since Q_0 is maximal, there exists an I_0 isomorphic to I such that $Q_0(I_0, J_0)$ for some J_0 . Let f be a permutation of \mathbf{U} such that $f(I_0) = I$. Since $Q(I_0, J_0)$, we have $Q(I, f(J_0))$, by genericity of Q . Since Q is semideterministic there exists a permutation g of \mathbf{U} such that $g(f(J_0)) = J$. Since $g(I) = I$ and, hence, $g(f(I_0)) = I$, it follows that $Q_0^* = (I, J)$ through application of gf to $Q_0(I_0, J_0)$, as desired. ■

Before we move to semideterministic *programs*, we make one final observation concerning semideterministic queries in general. A *boolean* query is a query of type $\mathcal{S} \rightarrow A$, with

$a(A)=0$. So a boolean query can have only two possible answers: the empty zero-ary relation \emptyset , which is interpreted as false, and the nonempty zero-ary relation $\{\langle \rangle\}$, which is interpreted as True. Since True and False are non-isomorphic, we immediately obtain:

PROPOSITION 3.9. *Every boolean semideterministic query is deterministic.*

3.2. Semideterministic Programs

Given a TL program P , when would we call it semideterministic? We could call it semideterministic simply if the query expressed by P is semideterministic. But this is not a very interesting definition. It would be more interesting if the whole of P 's computation is semideterministic, not just merely the end result. Formally, the semideterministic programs are defined as follows:

DEFINITION 3.10. 1. We first define the *unfoldings* of a TL program P :

(a) If P is a program not containing any loops, then the only unfolding of P is P itself.

(b) If P is of the form $\text{loop}[P']$, then the unfoldings of P are all n -fold concatenations $(P')^n$ of P' with itself, for all natural numbers n .

(c) If P is of the form $P_1; P_2$ then the unfoldings of P are all programs of the form $P'_1; P'_2$, with $P'_1(P'_2)$ an unfolding of $P_1(P_2)$.

2. An unfolding of a program is a sequence of assignments $\sigma_1; \dots; \sigma_n$. We call such a sequence *semideterministic* if for each $i \leq n$, the program $\sigma_1; \dots; \sigma_i$ expresses a semideterministic query.

3. Finally, a program is called *semideterministic* if all its unfoldings are.

In item (2) of the above definition, it is required that every intermediate stage is semideterministic. In principle it is sufficient to verify the requirement only for those i where σ_i contains an application of the Witness operation, since these are the only places where things can go wrong.

Clearly, the query expressed by a semideterministic program is semideterministic. Of course, the converse does not hold. In the next example, we will give examples of non-semideterministic queries and of semideterministic queries expressed by programs that are semideterministic and by programs that are not.

EXAMPLE 3.11. Let \mathcal{S} and I be as in Example 2.2. Consider the (RA + W)-program P_1 :

$$A_1 := \pi_1(R) \cup \pi_2(R); \quad A_2 := W_1(A_1).$$

In a possible result of applying this program to I , the A_1 -relation will hold $\text{adom}(I)$, and the A_2 -relation will hold

one of the four singleton subsets of $\text{adom}(I)$. Among these four possible outcomes, the two possible results ($R: I(R), A_1: \text{adom}(I), A_2: \{a\}$) and ($R: I(R), A_1: \text{adom}(I), A_2: \{b\}$) are not isomorphic: there is no automorphism of I mapping a to b . So, the program does not express a semideterministic query, and so it is certainly not semideterministic.

To illustrate the difference between semideterministic programs and semideterministic queries, consider the very simple scheme $\{V\}$, with $a(V) = 1$, and the programs P_2 :

$$A_1 := W_1(V); \quad A_2 := W_1(V)$$

and P_3 :

$$A_1 := W_1(V); \quad A_2 := W_1(V - A_1).$$

Let us restrict attention to instances I for which $I(V)$ has at least two elements. If in both programs, A_2 is the answer relation, the query Q expressed by P_2 and P_3 is the same. Applied to an instance I , a possible result of Q will hold (in the A_2 relation) an arbitrary element if $I(V)$. Since the elements of $I(V)$ are indistinguishable, Q is semideterministic. However, to check whether the programs P_2 and P_3 are semideterministic, we also have to take the temporary A_1 relation into account. So we see that P_2 is not semideterministic; two possible intermediate results could be the instance J_1 with $J_1(A_1) = \{\langle a \rangle\}$ and $J_1(A_2) = \{\langle a \rangle\}$, and the instance J_2 with $J_2(A_1) = \{\langle b \rangle\}$ and $J_2(A_2) = \{\langle c \rangle\}$. Clearly, J_1 and J_2 are not isomorphic. On the other hand, P_3 is semideterministic, since any possible intermediate result of P_3 will contain an element of V in A_1 and *another* element of V in A_2 . All such configurations are clearly isomorphic.

EXAMPLE 3.12. Finally, to give a less abstract example, consider the scheme $\{\text{Prof}, \text{Stud}\}$ with $a(\text{Prof}) = a(\text{Stud}) = 1$. The intended meaning of this scheme is that the Prof-relation holds a set of professor names, and the Stud-relation holds a set of student names. We assume that the sets of student names and professor names are disjoint. The reader is invited to verify that the following program is semideterministic:

```

Prof_chosen := ∅;
Stud_chosen := ∅;
Advisor := ∅;
loop[
  C_stud := W1(Stud - Stud_chosen);
  C_prof := W1(Prof - Prof_chosen);
  Advisor := Advisor ∪ (C_stud × C_prof);
  Stud_chosen := Stud_chosen ∪ C_stud;
  Prof_chosen := Prof_chosen ∪ C_prof;
]

```

The query expressed by this program produces in the answer relation *Advisor* an arbitrary one-to-one student-advisor assignment. All possible such assignments are isomorphic by an isomorphism mapping profs to profs and students to students; hence, the query is indeed semideterministic.

The above examples do not involve object creation, but it will become clear later that object creation is actually crucial to performing semideterministic computations in less contrived situations. For now, we demonstrate a good property of semideterministic loop programs which we can prove in the most general setting of all TL programs.

THEOREM 3.13. *Let P be a TL program such that $\text{loop}[P]$ is semideterministic, and let I be an input instance. Then either the result of $\text{loop}[P]$ on I is undefined since no possible computation halts, or every possible computation halts after the same number of iterations of the loop.*

Proof. Assume J is a possible result of $\text{loop}[P]$ on I . By definition, there exists n such that $J \in P^n(I) \cap P(J)$. So $J \in P^{n+1}(I)$ and J is output after $n+1$ iterations. Now let J' be another element of $P^{n+1}(I)$, arbitrarily chosen. So, J' is the preliminary result, after $n+1$ iterations, of an arbitrary possible computation of the program. We must show that $J' \in P^n(I) \cap P(J')$. Indeed, if this is the case then J' is actually a final result of $\text{loop}[P]$, or in other words, its corresponding computation halts after $n+1$ iterations, as desired.

Since $\text{loop}[P]$ is semideterministic, P^{n+1} is semideterministic. Since both $P^{n+1}(I, J)$ and $P^{n+1}(I, J')$, J' must be isomorphic to J . But then since $P^n(I, J)$, by genericity also $P^n(I, J')$. Moreover, since $P(J, J)$, by genericity, also $P(J, J')$. ■

It is known [3] that it is undecidable whether an $(\text{RA} + \text{W})$ -program, and more specifically, an $(\text{RA} + \text{W})$ -program of the form: $R := E; A := W_X(R)$ (where E is an arbitrary relational algebra expression) expresses a deterministic query.³ Although semideterminism is less restrictive than plain determinism, the analogue of this result still holds.

THEOREM 3.14. *It is undecidable whether a program in $\text{RA} + \text{W}$ expresses a semideterministic query.*

Proof. Let P be a program of the form $R := E; A_1 := W_X(R)$. Let P' be the program $P; A_2 := W_X(R)$;

³ In short, the reason is the following: The Witness operation $W_X(R)$ is deterministic iff the complement X^c of X is a key for its argument relation R . Since R is the result of algebra expression E , it thus follows that the program is deterministic iff E implies the key dependency $\text{key}(X^c)$. This is undecidable because it is undecidable whether a functional dependency is implied by a first-order sentence; a much stronger version of this statement was proven in [11].

$A := A_1 - A_2$, with A the answer relation. Since the assignments to A_1 and A_2 might be the same, we have that for any instance I , there is a possible result J of P' on I for which $J(A) = \emptyset$. Since the only relation isomorphic to \emptyset is \emptyset itself, it follows that the query Q' expressed by P' is semideterministic iff Q' is deterministic. But Q' is deterministic iff the query Q expressed by P is deterministic, and this is undecidable. ■

COROLLARY 3.15. *It is undecidable whether a program in $\text{RA} + \text{W}$ is semideterministic.*

Proof. Follows from the observation that programs P' of the special form exhibited in the proof of Theorem 3.14 are semideterministic iff the query they express is semideterministic, which was just shown to be undecidable. ■

Theorem 3.14 shows that “compile-time” checking for semideterminism is infeasible. As an alternative, we can check for semideterminism “at run-time.” Given a TL program P and an input instance I , execute P on I and add, a posteriori, an extra checking phase to see that all possible results are pairwise isomorphic. If this check fails, the result of P on I is overruled to some default value, e.g., all empty answer relations. We have thus defined an alternative semantics for TL programs, which we naturally call the *semideterministic* semantics. If we want not only the query expressed by the program, but also the program itself to be semideterministic, we further add a similar check after each application of the Witness operation. This stronger version of the semideterministic semantics is called the *uniformly semideterministic* semantics.

We now show that run-time checking can be performed in TL. That is, the just mentioned extra checking phases can be simulated in TL itself. This should come as no surprise, in view of the computational completeness of TL. Therefore, we only give a sketch of the proof.

PROPOSITION 3.16. *For every TL program P there is another TL program P^{sd} (resp. P^{usd}) such that the ordinary semantics of P^{sd} corresponds to the (resp. uniformly) semideterministic semantics of P .*

Proof. P^{sd} consists of two parts. The first part is a deterministic (or rather determinate) simulation of P . Applications of Witness are simulated by generating all possible results and keeping track of them in subsequent computations. After the end of the simulation, there is a second part that checks whether the accumulated possible results are pairwise isomorphic. (Between each pair of possible results, all possible bijections are generated by a powerset-like construction using object creation. Then it is verified whether at least one of these bijections is an isomorphism.) If the test fails, the answer relations are assigned empty by default. Otherwise, using Witness, an arbitrary possible result is chosen. The proof for P^{usd} is analogous. ■

We have just shown that run-time checking for semideterminism, in contrast to compile-time checking, is decidable. However, the tests for isomorphism that are involved are similar to testing graph isomorphism, for which no polynomial-time algorithm is known. Indeed, we next show that the two problems are polynomial-time equivalent.

PROPOSITION 3.17. *Run-time checking for semideterminism is polynomial-time equivalent to checking graph isomorphism.*

Proof. Run-time checking for semideterminism can be reduced in polynomial time to checking graph isomorphism, because checking isomorphism of relational structures (i.e., database instances) reduces to checking graph isomorphism [23].

For the converse direction, consider the scheme $\mathcal{S} = \{R, V\}$ with $a(R) = 2$ and $a(V) = 1$, and consider the class \mathcal{I} of instances I over \mathcal{S} for which $I(V)$ contains exactly two elements of $\text{adom}(I(R))$. Consider the following one-line program $P \equiv A := W_1(V)$. Checking P for semideterminism on an instance I of the class \mathcal{I} amounts to checking whether two given nodes o_1, o_2 of a graph (binary relation) are *autoequivalent*, meaning that there is an automorphism of the graph mapping o_1 to o_2 . We conclude the proof by showing that graph isomorphism can be reduced in polynomial time to autoequivalence. Assume we are given two graphs, G_1, G_2 , which we may assume are connected and disjoint. In order to test whether there is an isomorphism between G_1 and G_2 , it suffices to test whether there is a pair (o_1, o_2) , where o_i is a node of G_i , such that o_1 and o_2 are autoequivalent within $G_1 \cup G_2$. ■

This section can be concluded by saying that semideterminism is a desirable and natural notion, but at the same time it is hard to achieve. Nevertheless, in the next two sections we will exhibit two major applications where semideterministic computations are possible.

4. SEMIDETERMINISM AND COMPLETENESS

In this section, we prove that every semideterministic query can be expressed by a semideterministic TL program.

4.1. Determinate-Completeness up to Copies

The language RA + new + loop is a very powerful, determinate language, which is roughly comparable in expressive power to the language IQL [4]. The only difference is that IQL supports set values. This difference is irrelevant to the issues discussed in the main body of this paper, but we will return to set values in Section 6.

If we restrict the use of object creation syntactically, so that new objects appear only in intermediate relations and not in the final answer relations, we obtain a sublanguage of RA + new + loop, which we call detTL since it is equivalent

to the language detTL of [5, 6]. The queries expressed by detTL programs are not just determinate, but actually deterministic, exactly because no new objects appear in the result. In fact, it is known [3] that detTL is deterministic-complete: *all* deterministic queries can be expressed in the language.

Therefore, it came much as a surprise that RA + new + loop is *not* determinate-complete. This is illustrated by the following example.

EXAMPLE 4.1. Consider a relation name A of arity 1. Let B, C be two relation names of arity 2. Consider any determinate query Q of type $\{A\} \rightarrow B, C$ containing a pair $Q(I^{\text{diff}}, J^{\text{diff}})$. Here, I^{diff} is an instance of the very simple form $(A: \{a_1, a_2\})$, and the corresponding result J^{diff} has the form:

$$A: \{a_1, a_2\},$$

$$B: \{\langle b_1, a_1 \rangle, \langle b_3, a_1 \rangle, \langle b_2, a_2 \rangle, \langle b_4, a_2 \rangle\},$$

$$C: \{\langle b_1, b_2 \rangle, \langle b_2, b_3 \rangle, \langle b_3, b_4 \rangle, \langle b_4, b_1 \rangle\}.$$

So the b_i 's are new objects. This can be easily visualized using graphs: starting from a discrete graph containing two isolated A -nodes a_1, a_2 , a C -cycle of four new nodes b_1, \dots, b_4 is created such that two opposite b -nodes are associated to a common a -node through the B -relation. It can be shown [2, 9, 27] that such a query Q is not expressible in RA + new + loop.

The noncompleteness of RA + new + loop can be put in a more structured framework using the notion of *instance with copies* [4]. Let \mathcal{S} be a scheme, and let $\mathcal{S}_0 \subseteq \mathcal{S}$. For each $R \in \mathcal{S} - \mathcal{S}_0$, let CR be a relation name not in \mathcal{S} for which $a(CR) = a(R) + 1$. All these CR must be different. Let $\bar{\mathcal{S}} := \mathcal{S}_0 \cup \{CR \mid R \in \mathcal{S} - \mathcal{S}_0\}$. Let $J \in \text{inst}(\mathcal{S})$, and $\bar{J} \in \text{inst}(\bar{\mathcal{S}})$. Then we define:

DEFINITION 4.2. \bar{J} is an instance *with copies* of J w.r.t. \mathcal{S}_0 if there exist:

- (i) An natural number $n > 0$, called the *number of copies*;
- (ii) n instance $J_1, \dots, J_n \in \text{inst}(\mathcal{S})$, called *copies*, such that the sets $\text{adom}(J) - \text{adom}(J|_{\mathcal{S}_0})$, $\text{adom}(J_1) - \text{adom}(J_1|_{\mathcal{S}_0})$, \dots , $\text{adom}(J_n) - \text{adom}(J_n|_{\mathcal{S}_0})$ are pairwise disjoint, and J_k and J are $J|_{\mathcal{S}_0}$ -isomorphic for $k = 1, \dots, n$;⁴
- (iii) n objects $\varepsilon_1, \dots, \varepsilon_n$, called *copy identifiers*, not appearing in J or any J_k , such that

$$\bar{J}(CR) = (J_1(R) \times \{\langle \varepsilon_1 \rangle\}) \cup \dots \cup (J_n(R) \times \{\langle \varepsilon_n \rangle\})$$

for each $R \in \mathcal{S} - \mathcal{S}_0$, and $\bar{J}(R) = J(R)$ for each $R \in \mathcal{S}_0$.

⁴ Recall from Definition 2.3 that an I -isomorphism, for some instance I , is an isomorphism that is the identity on $\text{adom}(I)$.

It follows from (ii) that each J_k agrees with J on \mathcal{S}_0 and, hence, also \bar{J} does.

EXAMPLE 4.3. Recall J^{diff} from Example 4.1 The following instance J^{diff} over scheme $\{A, CB, CC\}$ is an instance with two copies of J^{diff} w.r.t. $\{A\}$:

$$A: \{a_1, a_2\},$$

$$CB: \{\langle b_{11}, a_1, \varepsilon_1 \rangle, \langle b_{31}, a_1, \varepsilon_1 \rangle, \langle b_{21}, a_2, \varepsilon_1 \rangle, \\ \langle b_{41}, a_2, \varepsilon_1 \rangle, \langle b_{12}, a_1, \varepsilon_2 \rangle, \langle b_{32}, a_1, \varepsilon_2 \rangle, \\ \langle b_{22}, a_2, \varepsilon_2 \rangle, \langle b_{42}, a_2, \varepsilon_2 \rangle\},$$

$$CC: \{\langle b_{11}, b_{21}, \varepsilon_1 \rangle, \langle b_{21}, b_{31}, \varepsilon_1 \rangle, \langle b_{31}, b_{41}, \varepsilon_1 \rangle, \\ \langle b_{41}, b_{11}, \varepsilon_1 \rangle, \langle b_{12}, b_{22}, \varepsilon_2 \rangle, \langle b_{22}, b_{32}, \varepsilon_2 \rangle, \\ \langle b_{32}, b_{42}, \varepsilon_2 \rangle, \langle b_{42}, b_{12}, \varepsilon_2 \rangle\}.$$

Although no query expressible in RA + new + loop can contain a pair $(I^{\text{diff}}, J^{\text{diff}})$ as in Example 4.1, it is not difficult to write an RA + new + loop program P such that \bar{J}^{diff} is a possible result of P applied to I^{diff} . More generally, RA + new + loop is complete up to copies:

FACT 4.4 [4]. *For each determinate query Q of type $\mathcal{S}_0 \rightarrow \mathcal{S} - \mathcal{S}_0$ there is an RA + new + loop program expressing a query \bar{Q} of type $\mathcal{S}_0 \rightarrow \bar{\mathcal{S}} - \mathcal{S}_0$ such that $Q(I, J)$ iff $\bar{Q}(I, \bar{J})$, with \bar{J} an instance with copies of J w.r.t. \mathcal{S}_0 .*

Proof. Let us briefly review the proof of this important fact. Since Q is r.e., there is a Turing machine M which enumerates Q . So, M takes a natural number as input and produces a pair (I, J) of (encodings of) instances such that $Q(I, J)$. The range of M is the whole of Q . Then the desired program, on input instance I over \mathcal{S}_0 , visits pairs (k, l) of natural numbers, in some standard order. For each pair (k, l) , k new objects, o_1, \dots, o_k , are created. The collection \mathcal{C} of all instances J over \mathcal{S} that equal I on \mathcal{S}_0 and for which $\text{adom}(J) - \text{adom}(I) = \{o_1, \dots, o_k\}$, is constructed. The subset \mathcal{C}' of \mathcal{C} , consisting of those J for which M on input l produces a pair of instances that is isomorphic to (I, J) , is determined. If \mathcal{C}' is empty, the next pair (k, l) of natural numbers is visited. Otherwise, by the genericity of Q , for each J in \mathcal{C}' , we have $Q(I, J)$. Furthermore, by the determinacy of Q , all J in \mathcal{C}' are pairwise isomorphic by a permutation of $\{o_1, \dots, o_k\}$. Thus, \mathcal{C}' contains all the necessary ingredients from which to construct an instance \bar{J} with copies as desired. ■

It follows that the query CE (for copy elimination) of type $\bar{\mathcal{S}} \rightarrow \mathcal{S} - \mathcal{S}_0$, defined by CE (\bar{J}, J) , if \bar{J} is an instance with copies of $(\mathcal{S}_0: J|_{\mathcal{S}_0}, \mathcal{S} - \mathcal{S}_0: J|_{\mathcal{S} - \mathcal{S}_0})$, is not expressible in RA + new + loop. Copy elimination is a determinate query and, by Fact 4.4, it suffices to add it as a primitive to make RA + new + loop complete for all determinate queries.

4.2. Semideterministic Completeness

As argued in the Introduction, the determinate-completeness up to copies of RA + new + loop is not very satisfactory. In fact, our original motivation for studying semideterminism was the hope that it could offer a less restrictive setting in which copy elimination could be explained and would appear less ad hoc. Indeed, an alternative way to look at copy elimination is to consider it as a nondeterminate operation which chooses one among several available copies. More formally, we can define a *nondeterminate* version of copy elimination, call it CE^{nd} , as follows: Let C be a relation name not in $\bar{\mathcal{S}}$ with $\text{a}(C) = 1$. Then CE^{nd} is of type $\bar{\mathcal{S}} \rightarrow C$ and defined by $\text{CE}^{\text{nd}}(\bar{J}, K)$, if \bar{J} is an instance with copies as in Definition 4.2 and $K(C) = \{\varepsilon_i\}$ for some arbitrarily chosen i .

Using CE^{nd} as just defined, we can easily simulate CE as originally defined. It is also readily verified that CE^{nd} is semideterministic. And, it can be easily expressed by a semideterministic TL program. This program first checks whether its input is indeed an instance with copies.⁵ Then the Witness operation is applied to the set of all copy identifiers.

This encourages us to generalize the notion of instance with copies to the semideterministic setting, in the hope of being able to prove that every semideterministic query is expressible by a semideterministic TL program. Indeed, the strategy which we could follow is to prove (i) that Fact 4.4 can be generalized to the semideterministic setting and (ii) that nondeterminate copy elimination, adapted to the semideterministic setting, is still expressible by a semideterministic TL program. In the remainder of this section, we will show that this strategy works.⁶

First, we generalize the notion of instance with copies as defined in Definition 4.2 to the semideterministic setting. To do this, it suffices to note that requirement (ii) in that definition, stating that J_k and J must be $J|_{\mathcal{S}_0}$ -isomorphic, is similar to the determinacy condition. Hence, we can generalize the definition of instance with copies in a similar way as we generalized determinacy to semideterminism. Specifically, we now only require in (ii) that J_k and J are isomorphic. Let us refer to this generalized notion of instance with copies as *instance with semideterministic copies*, and to the original notion as *instance with determinate copies*.

We can now observe that the following analogue to Fact 4.4 holds.

⁵ This initial check is computationally expensive, as it requires checking graph isomorphism. However, in the intended application of copy elimination it is guaranteed that the input has the required format (as with the result of the RA + new + loop program of Fact 4.4) so the check can in principle be omitted.

⁶ Contrary to the conjecture expressed in a preliminary version of this paper [25].

PROPOSITION 4.5. *For each semideterministic query Q of type $\mathcal{S}_0 \rightarrow \mathcal{S} - \mathcal{S}_0$ there is an RA + new + loop program expressing a query \bar{Q} of type $\mathcal{S}_0 \rightarrow \bar{\mathcal{S}} - \mathcal{S}_0$ such that $Q(I, J)$ iff $\bar{Q}(I, \bar{J})$, with \bar{J} an instance with semideterministic copies of J w.r.t. \mathcal{S}_0 .*

Proof. The proof of Fact 4.4 goes through verbatim, except for the point where it is stated that all J in \mathcal{C}' are isomorphic by a permutation of $\{o_1, \dots, o_k\}$. The italicized qualification must now be omitted. ■

Suppose Q is a determinate query as in Fact 4.4, $Q(I, J)$, and \bar{J} is an instance with determinate copies of J . Then trivially, any J' such that $Q(I, J')$ is I -isomorphic to a copy (in fact every copy) contained in \bar{J} . The analogue of this property in the semideterministic case is not entirely trivial.

LEMMA 4.6. *Assume Q is a semideterministic query as in Proposition 4.5, $Q(I, J)$, and $\bar{Q}(I, \bar{J})$ as constructed in the proof of said proposition. Then any J' such that $Q(I, J')$ is I -isomorphic to a copy contained in \bar{J} .*

Proof. We use the notation from the proof of Proposition 4.5 (and Fact 4.4). Clearly, J' is I -isomorphic to some J'' in \mathcal{C} . Furthermore, since Q is semideterministic, J' and, hence, also J'' , are isomorphic to J by some automorphism of I . But \mathcal{C}' is computed by a (generic) RA + new + loop program, and therefore if J is in \mathcal{C}' , then so is J'' , as desired. ■

We may thus conclude that also in the semideterministic case, \bar{J} contains “enough” copies. Hence, to express any semideterministic query Q , we can in a first phase express Q up to semideterministic copies, using the RA + new + loop program of Proposition 4.5, and then in a second phase choose one of the copies and assign it to the final answer relation, using an obvious TL program. If we can prove that the two-phase program thus obtained is semideterministic (again something that was trivial in the determinate case), we can conclude:

THEOREM 4.7. *Every semideterministic query can be expressed by a semideterministic TL program.*

Proof. As just observed, we must prove that choosing among semideterministic copies is semideterministic. This follows immediately from the following claim.

CLAIM. *Assume Q is a semideterministic query as in Proposition 4.5, $Q(I, J)$, and $\bar{Q}(I, \bar{J})$ as constructed in the proof of said proposition. Then any two copies, J_1, J_2 , contained in \bar{J} are isomorphic by an automorphism of \bar{J} .*

Proof. We know that J_1 and J_2 are isomorphic by an automorphism of I , say f . We must extend f to an automorphism of \bar{J} . This means we must find a permutation g of \mathbf{U} that equals f on $\text{adom}(J_1)$ such that $g(\bar{J}) = \bar{J}$. Let p

be the order of $f|_{\text{adom}(I)}$ in the permutation group of $\text{adom}(I)$.

If $p = 1$, then J_1 and J_2 are I -isomorphic and the problem is as trivial as in the determinate case (simply define g as f on $\text{adom}(J_1)$, f^{-1} on $\text{adom}(J_2) - \text{adom}(J_1)$, and the identity everywhere else).

Now suppose $p > 1$. By genericity, we have for each $i = 1, \dots, p$ that $Q(I, f^i(J_1))$. By Lemma 4.6, for each i there is a copy J_{i+1} contained in \bar{J} such that J_{i+1} is I -isomorphic to $f^i(J_1)$. (Note that J_2 is compatible with this statement.) Furthermore, since $f^p|_{\text{adom}(I)}$ is the identity, J_{p+1} can be chosen to be J_1 . So we can define the desired permutation g as f on $\text{adom}(J_1)$, and such that $g(J_i) = J_{i+1}$, for $i = 1, \dots, p$, in the obvious way. ■

Note that we have actually obtained a syntactic, semideterministic sublanguage of TL capable of expressing exactly the semideterministic queries: namely, the language consisting of all RA + new + loop programs followed by a semideterministic copy elimination step. Of course, this sublanguage is highly artificial.

5. COUNTING SEMIDETERMINISTICALLY

In this section, we show that every polynomial-time counting query can be computed efficiently in a semideterministic manner. For clarity, we will only consider deterministic queries in this section. So whenever we use the term *query*, we will mean *deterministic query*.

Before coming to the point, we define an intuitive notion of efficiency for programs. There is an obvious implementation of TL programs by nondeterministic Turing machines. We say that a TL program is *efficient* if its thus associated Turing machine is polynomial-time.

5.1. Problem Statement

There are queries that are computationally simple, but not expressible in RA + loop. The most important class of such queries are those involving counting. A simple example in this class is the boolean parity query: Given input instance I , is the cardinality of $\text{adom}(I)$ even? Most queries involving counting cannot be expressed in RA + loop because RA + loop has a 0–1 law [20].

It is well known [3] that using nondeterminism, more deterministic queries can be expressed. Concretely, every polynomial-time computable query is expressible by an efficient RA + W + loop program. This result can be best understood in terms of ordered databases. Ordered databases are databases equipped with a special binary relation containing a linear order of the active domain. It is well known that every polynomial-time computable query is expressible, *on ordered databases*, by an efficient program in RA + loop [3]. Hence, to express an arbitrary polynomial-time query Q in RA + W + loop, it suffices to be able to

construct an ordered list of the active domain, which can be done easily by repeatedly choosing elements. In particular, the counting queries can be expressed efficiently in this manner.

A nondeterministic program which generates an arbitrary ordered list of the input database is not semideterministic in general; only for completely symmetric databases it is true that any two such lists can be mapped into each other by an automorphism. However, in order to express counting queries, constructing an ordering of the entire database domain is not really needed. It is sufficient to generate, for each fragment of the database that has to be counted, a list of *new* objects whose length is equal to the cardinality of the fragment. These new objects are “witness” for the fragment to be counted, and whether we count the original fragment, or the thus-constructed witness list, clearly does not matter. Furthermore, a configuration of witness lists can be computed by an efficient TL program, and this program is semideterministic, intuitively since any two configurations of witness lists are determinate copies of each other.

So, the purpose of this section is to prove formally that the intuition just expressed is correct, i.e., that every counting query can be expressed efficiently by a semideterministic TL program. The qualification *efficiently* is crucial, since otherwise the claim would be trivial. Indeed, we mentioned in Section 4 that RA + new + loop (or, for that matter, detTL) being a sublanguage of TL can already express *all* deterministic queries, in particular the counting queries. And, RA + new + loop programs are trivially semideterministic. However, the counting queries are not *efficiently* expressible in RA + new + loop. For example, it follows from results in generic complexity [7] that there is no efficient RA + new + loop program expressing the parity query.

5.2. Counting Semideterministically

We still have to specify precisely what is the class of counting queries. A precisely defined extension of fixpoint logic with counting capabilities has been introduced in [12, 13]. This approach, however, involves a two-sorted extension of the relational signature with a sort for natural numbers to express the results of the counting operations. In our object-creating framework, however, an essentially equivalent counting mechanism can be introduced without at the same time introducing natural numbers as well, since a natural number n can be represented by a linked list of newly created objects of length n . This simply corresponds to writing the natural number in unary notation. Before giving a formal definition, we illustrate this operation by means of an example.

EXAMPLE 5.1. Assume we are given a binary relation r , interpreted as a parent–child relation. Suppose we want to count the number of children of each parent. A uniform,

generic way to do this would be to construct for each parent a list of new objects, the length of which equals the number of his children. This is done by applying the counting operation *count* to r . For example, Fig. 1 shows a relation r , together with $\text{count}(r)$. The Greek letters are newly created objects.

In general, the counting operation is defined as follows.

DEFINITION 5.2. Let \mathcal{S} be a scheme, $I \in \text{inst}(\mathcal{S})$, and $R \in \mathcal{S}$ with $X \subseteq \{1, \dots, a(R)\}$. Let $\pi_X(I(R)) = \{u_1, \dots, u_p\}$. For each $i = 1, \dots, p$, define n_i as the cardinality of the set

$$\{t \in I(R) \mid t|_X = u_i\}.$$

Then each possible result of $\text{count}_X(R)$ applied to I is a relation of arity $|X| + 2$, of the form

$$\begin{aligned} & \{ \langle u_1, \alpha_1^1, \alpha_2^1 \rangle, \dots, \langle u_1, \alpha_{n_1}^1, \alpha_{n_1+1}^1 \rangle \\ & \quad \vdots \\ & \langle u_p, \alpha_1^p, \alpha_2^p \rangle, \dots, \langle u_p, \alpha_{n_p}^p, \alpha_{n_p+1}^p \rangle \}, \end{aligned}$$

where the α_j^i are different new objects not in $\text{adom}(I)$.

The *count* operation is clearly determinate and is readily implementable in polynomial-time.

The counting application in Example 5.1 is formally an application of count_\emptyset . Note that $\text{count}_\emptyset(R)$ counts the total number of tuples in R . Note also that the count operation contains the new operation as a singular case. Indeed,

$$\text{new}(R) \equiv \pi_{1, \dots, a(R)+1} \text{count}_{1, \dots, a(R)}(R).$$

By extending RA + loop with the count operation, we obtain a language which we denote by RA + loop + count. By the above remark and the discussion in the previous subsection, this language has exactly the same expressive

$$r = \begin{array}{|l} a & a_1 \\ b & b_1 \\ b & b_2 \\ c & c_1 \\ c & c_2 \\ c & c_3 \end{array}$$

$$\text{count}(r) = \begin{array}{|l} a & \alpha_1 & \alpha_2 \\ b & \beta_1 & \beta_2 \\ b & \beta_2 & \beta_3 \\ c & \gamma_1 & \gamma_2 \\ c & \gamma_2 & \gamma_3 \\ c & \gamma_3 & \gamma_4 \end{array}$$

FIG. 1. An example of counting.

power as RA + new + loop, but it can express many more queries efficiently. We will refer formally to the class of *counting queries* as consisting of those queries expressed by an efficient RA + loop + count program.⁷

We now show:

THEOREM 5.3. *Every counting query is expressible by an efficient semideterministic TL program.*

Proof. We show how to express the count operation by an efficient semideterministic TL program. It is sufficient to consider applications of the form $\text{count}_1(R)$, where R is a binary relation (as in Example 5.1.) Indeed, if R is not binary and we have to compute $\text{count}_X(R)$, then by two applications of the new operation, $\text{new}(\pi_X(R))$ and $\text{new}(\pi_{\{1, \dots, a(R)\} - X}(R))$, we can easily construct, in the relational algebra, an encoding of R into a binary relation which is equivalent for counting purposes. The following program computes $\text{count}_1(R)$ in its answer relation A . (The relational calculus expressions occurring in the program can easily be translated into the relational algebra [3].)

$$\begin{aligned}
 R' &:= \pi_{1,3} \text{new}(R) \cup \text{new}(\pi_1(R)); \\
 C_1 &:= W_2(R'); \\
 C_2 &:= W_2(R' - C_1); \\
 A &:= \{ \langle u, \alpha, \beta \rangle \mid \langle u, \alpha \rangle \in C_1 \text{ and } \langle u, \beta \rangle \in C_2 \}; \\
 \text{Chosen} &:= C_1 \cup C_2; \\
 \text{loop}[\\
 &C := W_2(R' - \text{Chosen}); \\
 &\text{Chosen} := \text{Chosen} \cup C; \\
 &A := A \cup \{ \langle u, \beta, \gamma \rangle \mid \langle u, \gamma \rangle \in C \text{ and} \\
 &\quad \exists \alpha: \langle u, \alpha, \beta \rangle \in A \text{ and } \neg \exists \gamma': \langle u, \beta, \gamma' \rangle \in A \}; \\
 &]
 \end{aligned}$$

Using the familiar parent–child terminology, the R' relation defined in the first assignment can be viewed as being obtained from the R relation by replacing each child by a different, newly created object (and one more.) The remainder of the program then simply orders these “witnesses” (recall the discussion in the first subsection) into lists, as required by the definition of $\text{count}_1(R)$. This is accomplished by initializing the lists to length 1, and then completing them by repeatedly choosing the remaining witnesses and appending them to the lists. Hence, the program is correct, and it is clearly also efficient.

It remains to show that the program is semideterministic. Observe that in R' , any two witnesses of a common parent are logically interchangeable. More precisely, for any two tuples $\langle u, \alpha_1 \rangle$ and $\langle u, \alpha_2 \rangle$ in R' , the transposition $(\alpha_1 \alpha_2)$ is an automorphism of R' . From this it follows immediately

⁷ This is not a syntactic definition. The usual approach to syntactically approximate the efficient RA + loop programs is to give an inflationary semantics to the loops [3]. This approach does not work in the presence of object creation.

that the first two applications of the Witness operation, in the second and third statements of the program, are semideterministic. To show that the loop is semideterministic, we similarly observe that before each application of the assignment $C := W_2(R' - \text{Chosen})$, any two witnesses of a common parent remaining in $R' - \text{Chosen}$ are logically interchangeable with respect to the whole intermediate result of the computation at that moment. ■

The class of counting queries includes a wide variety of useful queries. In particular, it can be verified that every query expressible in the extension of fixpoint logic with counting defined in [12, 13] is a counting query in our sense. Rather than proving this claim formally (which is tedious but straightforward), we illustrate it in the remainder of this section.

For example, we can express all “global” boolean counting queries, such as the parity query. This class of queries could be defined as follows: Let E be a relational algebra expression. Let $F(n)$ be a property of natural numbers which can be decided by a Turing machine in polynomial time, given n in unary notation. Then the boolean counting query associated with E and F yields True for an input instance I iff $F(|E(I)|)$ is true. To express this query efficiently and semideterministically, first compute $\text{count}_\emptyset(E)$, which produces an ordered list whose length equals the cardinality of $E(I)$. Then invoke the well-known fact [3] that every polynomial-time property of ordered instances can be expressed efficiently in RA + loop.

But we can express much more counting queries than the global boolean ones. The result of a general application of the count operation can be interpreted as a relation containing tuples having one entry whose value is a natural number, encoded as a (unary) string. For example, the result relation shown in Fig. 1 can be interpreted as the relation

$$\{ \langle a, 1 \rangle, \langle b, 2 \rangle, \langle c, 3 \rangle \}.$$

By combining the results of different counting operations, we can also construct relations with tuples having multiple entries with natural numbers as values. For example, we can count the number of children of each parent in one relation and we can count the number of grandchildren of each parent in another relation. These two relations can then be joined, yielding a relation containing, for each parent, the number of children and the number of grandchildren.

One could now imagine a selection operation for relations containing natural numbers, based on an arbitrary k -ary property $F(n_1, \dots, n_k)$ of natural numbers computable by a polynomial-time Turing machine, given n_1, \dots, n_k in unary notation on k input tapes. An example with $k = 1$ would be: *Give all parents with an even number of children.* In this case, F would be the property of being even. An

example with $k=2$ would be: *Give all parents having an equal number of children and grandchildren.* In this case, F would be the property of being equal. It is possible to show that these selection operations are efficiently expressible in $\text{RA} + \text{loop}$. Indeed, as already mentioned above, one can simulate in $\text{RA} + \text{loop}$ any polynomial-time Turing machine operating on a unary string (actually, any string). The generalization to a tuple of strings is straightforward. It finally suffices to run the simulations in parallel for each tuple in the relation. The techniques required to do this are straightforward adaptations of the basic simulation technique described in [3].

6. CONCLUDING REMARKS

1. We would like to have a better understanding of why semideterministic computing is desirable, i.e., why it is important to express a query with a semideterministic program instead of with just an arbitrarily nondeterministic one. One of the most important characteristics of semideterministic programs is that all possible computations behave similarly, as illustrated by Theorem 3.13. This property could be crucial when having to combine different parts of a nondeterministic program which have been distributed on different computers.

2. In Section 5 we have shown that every polynomial-time counting query can be expressed by an efficient, semideterministic TL program. Can we do better? More precisely, exactly which polynomial-time computable queries can be thus expressed? In this respect, we point out that on the class of *rigid* instances (i.e., having no nontrivial automorphisms), fixpoint logic is strictly weaker than polynomial time.⁸ On rigid structures, semideterministic choices among elements are impossible. This strongly indicates that not *all* the polynomial-time computable queries will be expressible by an efficient, semideterministic TL program.

3. In the present paper, we have focused on object creation in the function of tuples (the *new* operation.) Object creation in function of sets has also been considered in the literature; a simple and uniform operation providing this functionality is the *abstraction* operation [16, 24]. One among the many equivalent ways to define this operation is as follows: Let $I \in \text{inst}(\mathcal{S})$ and $R \in \mathcal{S}$ with $\text{a}(R) = 2$. If the binary relation $I(R)$ is not an equivalence relation, then the application of $\text{abstr}(R)$ to I is undefined. Otherwise, assume Z_1, \dots, Z_k is an enumeration of the equivalence classes, and let $\alpha_1, \dots, \alpha_k$ be distinct new objects not in $\text{adom}(I)$. Then the binary relation

$$\{ \langle o, \alpha_i \rangle \mid o \in \text{adom}(I(R)) \text{ and } [o] = Z_i \}$$

⁸This is because fixpoint logic has a 0–1 law [20] and almost all instances are rigid.

is the result of $\text{abstr}(R)$ applied to I ($[o]$ denotes o 's equivalence class). So, while new tags tuples, abstr tags sets.

Using abstraction, it seems that more queries can be expressed efficiently and semideterministically. For example, consider the following semideterministic query: Given an equivalence relation R , choose for each equivalence class one representative. We conjecture that this query is not expressible by an efficient semideterministic TL program. It can, however, be expressed efficiently and semideterministically by the following simple program in $\text{RA} + \text{W} + \text{abstr}$:

$$\begin{aligned} C &:= W_1 \text{abstr}(R); \\ A &:= \pi_2(C). \end{aligned}$$

By Theorem 4.7, abstraction is of course semideterministically expressible in TL, but, we conjecture, not efficiently so. We point out that abstraction is not expressible in $\text{RA} + \text{new} + \text{loop}$ [24].

4. Not every TL program is semideterministic or expresses a semideterministic query. TL programs are in general not semideterministic, and, even more, in general do not express a semideterministic query. An important issue that has not been considered in the present paper is the design of (preferably efficient) query languages that are guaranteed to express semideterministic queries. In investigating this problem, one might take inspiration from the semideterministic counting program of Theorem 5.3. A semideterministic operation which is implicit in this program is what could be called *swap-choice*, the choice of one representative for each equivalence class of objects that are logically interchangeable. Swap-choice has been studied in [15], where it was shown that swap-choice and the count operation are polynomial-time equivalent within $\text{RA} + \text{new} + \text{loop}$.

5. In the present paper, we have focused on queries. In [26], we have conducted an initial study on semideterminism in the context of arbitrary database transformations, including updates. The preliminary conclusion of this study was that there seems to be a fundamental difference between queries and updates in this respect. Several soundness problems arise and it remains to be seen whether semideterminism is compatible at all with updates.

ACKNOWLEDGMENTS

We thank Serge Abiteboul, Tony Bonner, Anuj Dawar, Marc Gyssens, and Jan Paredaens for interesting discussions and Moshe Vardi for bringing up the question whether every polynomial-time deterministic query is expressible by an efficient semideterministic TL program. We also thank the referees for their thorough and insightful comments on an earlier draft of this article, in particular Tova Milo for suggesting a more natural semantics for the looping construct.

REFERENCES

1. "Proceedings, Ninth ACM Symposium on Principles of Database Systems," Assoc. Comput. Mach., New York, 1990.
2. S. Abiteboul, personal communication, 1990.
3. S. Abiteboul, R. Hull, and V. Vianu, "Foundations of Databases," Addison-Wesley, Reading, MA, 1994.
4. S. Abiteboul and P. Kanellakis, "Object-Identity as a Query Language Primitive," Technical Report 1022, INRIA-Rocquencourt, 1989; *SIGMOD Record* **18**, No. 2 (1989), 159–173.
5. S. Abiteboul and V. Vianu, Procedural languages for database queries and updates, *J. Comput. System Sci.* **41**, No. 2 (1990), 181–229.
6. S. Abiteboul and V. Vianu, Datalog extensions for database queries and updates, *J. Comput. System Sci.* **43**, No. 1 (1991).
7. S. Abiteboul and V. Vianu, Generic computation and its complexity, in "Proceedings, 23rd ACM Symposium on Theory of Computing," pp. 209–218, Assoc. Comput. Mach., New York, 1991.
8. A. V. Aho and J. D. Ullman, Universality of data retrieval languages, in "Proceedings, Sixth ACM Symposium on Principles of Programming Languages, 1979," pp. 110–120.
9. M. Andries and J. Paredaens, On instance-completeness of database query languages involving object creation, *J. Comput. System Sci.* **52**, No. 2 (1996), 357–373.
10. A. Chandra and D. Harel, Computable queries for relational database systems, *J. Comput. System Sci.* **21**, No. 2 (1980), 156–178.
11. A. K. Chandra and M. Y. Vardi, The implication problem for functional and inclusion dependencies is undecidable, *SIAM J. Comput.* **14**, No. 3 (1985), 671–677.
12. E. Grädel and M. Otto, Inductive definability with counting on finite structures, in "Computer Science Logic" (E. Börger *et al.*, Ed.), Lect. Notes in Comput. Sci., Vol. 702, pp. 231–247, Springer-Verlag, New York, Berlin, 1993.
13. S. Grumbach and C. Tollu, Query languages with counters, in "Database Theory—ICDT'92" (J. Biskup and R. Hull, Eds.), Lect. Notes in Comput. Sci., Vol. 646, pp. 124–139, Springer-Verlag, New York/Berlin, 1992.
14. M. Gyssens, personal communications, 1992.
15. M. Gyssens, J. Van den Bussche, and D. Van Gucht, Expressiveness of efficient semideterministic choice constructs, in "Automata, Languages, and Programming" (S. Abiteboul and E. Shamir, Eds.), Lect. Notes in Comput. Sci., Vol. 80, pp. 106–117, Springer-Verlag, New York/Berlin, 1994.
16. M. Gyssens, J. Paredaens, J. Van den Bussche, and D. Van Gucht, A graph-oriented object database model, *IEEE Trans. Knowledge Data Eng.* **6**, No. 4 (1994), 572–586.
17. R. Hull and M. Yoshikawa, ILOG: Declarative creation and manipulation of object identifies, in "Proceedings, 16th International Conference on Very Large Data Bases" (D. McLeod, R. Sacks-Davis, and H. Schek, Eds.), pp. 455–468, Morgan Kaufmann, San Mateo, CA, 1990.
18. M. Kifer and J. Wu, A logic for programming with complex objects, *J. Comput. System Sci.* **47**, No. 1 (1993), 77–120.
19. W. Kim, A model of queries for object-oriented databases, in "Proceedings, Fifteenth International Conference on very Large Data Bases" (P. Apers and G. Wiederhold, Eds.), pp. 423–432, Morgan Kaufmann, San Mateo, CA, 1989.
20. P. G. Kolaitis and M. Y. Vardi, Infinitary logics and 0–1 laws, *Inform. and Comput.* **98**, No. 2 (1992), 258–294.
21. G. Kuper, "The Logical Data Model: A New Approach to Database Logic," Ph.D. thesis, Stanford University, 1985.
22. G. Kuper and M. Vardi, The logical data model, *ACM Trans. Database Systems* **18**, No. 3 (1993), 379–413.
23. G. L. Miller, Graph isomorphism, general remarks, *J. Comput. System Sci.* **13**, No. 3 (1976), 300–317.
24. J. Van den Bussche and J. Paredaens, The expressive power of complex values in object-based data models, *Inform. and Comput.* **120** (1995), 220–236.
25. J. Van den Bussche and D. Van Gucht, Semi-determinism, in "Proceedings, 11th ACM Symposium on Principles of Database Systems," pp. 191–201, Assoc. Comput. Mach., New York, 1992.
26. J. Van den Bussche and D. Van Gucht, Non-deterministic aspects of database transformations involving object creation, in "Modeling Database Dynamics" (U. Lipeck and B. Thalheim, Eds.), pp. 3–16, Workshops in Computing, Springer-Verlag, New York/Berlin, 1993.
27. J. Van den Bussche, D. Van Gucht, M. Andries, and M. Gyssens, On the completeness of object-creating query languages, in "Proceedings, 33rd Symposium on Foundations of Computer Science," pp. 372–379, IEEE Comput. Soc. Press, New York, 1992.