

13

Notions of Dependency Satisfaction

MARC H. GRAHAM

Georgia Institute of Technology, Atlanta, Georgia

ALBERTO O. MENDELZON

University of Toronto, Toronto, Ontario, Canada

AND

MOSHE Y. VARDI

Stanford University, Stanford, California

Abstract. Two notions of dependency satisfaction, *consistency* and *completeness*, are introduced. Consistency is the natural generalization of weak-instance satisfaction and seems appropriate when only equality-generating dependencies are given, but disagrees with the standard notion in the presence of tuple-generating dependencies. Completeness is based on the intuitive semantics of tuple-generating dependencies but differs from the standard notion for equality-generating dependencies. It is argued that neither approach is the *correct* one, but rather that they correspond to different policies on constraint enforcement, and each one is appropriate in different circumstances. Consistency and completeness of a state are characterized in terms of the tableau associated with the state and in terms of logical properties of a set of first-order sentences associated with the state. A close relation between the problems of testing for consistency and completeness and of testing implication of dependencies is established, leading to lower and upper bounds for the complexity of consistency and completeness. The possibility of formalizing dependency satisfaction without using a universal relation scheme is examined.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*complexity of proof procedures*; H.2.1 [Database Management]:

The research of A. O. Mendelzon was supported by the National Science and Engineering Research Council of Canada. The research of M. Y. Vardi was supported by a Weizmann Post-Doctoral Fellowship and AFSOR grant 80-0212.

A preliminary version of this paper, by M. H. Graham and A. O. Mendelzon, appeared in the *Proceedings of the 1st ACM Symposium on Principles of Database Systems* (Los Angeles, Calif., Mar. 29–31). ACM, New York, 1982, 177–188.

Author's addresses: M. H. Graham, School of Information and Computer Science, Georgia Institute of Technology, Atlanta, GA 30332; A. O. Mendelzon, Computer Systems Research Institute, University of Toronto, Toronto, Ontario, Canada M5S 1A4; M. Y. Vardi, IBM, Almaden Research Center K55/802, 650 Harry Road, San Jose, CA 95120-6099.

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.

© 1986 ACM 0004-5411/86/01000-0105 \$00.75

Our first objection to this proposal is that weak satisfaction now becomes different from standard satisfaction for single-relation databases. For example, if all the dependencies are total tuple-generating dependencies (such as multivalued and join dependencies), then any database state satisfies any set of dependencies. In particular, every single relation can be made into a weak instance for itself by adding a finite number of tuples to it [6, 19].

A second, related, objection is that this proposal does not seem to capture the intuitive semantics of tuple-generating dependencies in multirelation databases either. Consider, for example, the following database:

Example 1

R_1		R_2		
<i>Student</i>	<i>Course</i>	<i>Course</i>	<i>Room</i>	<i>Hour</i>
<i>Jack</i>	<i>CS378</i>	<i>CS378</i>	<i>B215</i>	<i>M10</i>
		<i>CS378</i>	<i>B213</i>	<i>W10</i>

R_3		
<i>Student</i>	<i>Room</i>	<i>Hour</i>
<i>Jack</i>	<i>B215</i>	<i>M10</i>

and dependencies $\{SH \rightarrow R, RH \rightarrow C, C \twoheadrightarrow S \mid RH\}$.

The multivalued dependency $C \twoheadrightarrow S \mid RH$ is intended to express the fact that a student is associated with every (r, h) pair such that some course that the student takes meets at room r at time h . This constraint is intuitively violated in the example, since Jack is not associated with room B213 on Wednesdays at 10, although he is taking CS378.

An alternative approach to defining satisfaction is through the notion of *complete states*, introduced for a different purpose in [22]. We say that a state is complete if it contains each tuple that appears in the projections of every weak instance for the state. The state in Example 1 is not complete, because every weak instance for it contains the subtuple $\langle \text{Jack}, B213, W10 \rangle$, which does not appear in R_3 . When only tuple-generating dependencies are given, completeness coincides with the standard notion of satisfaction on single-relation databases.

An objection to the completeness criterion is that it seems unnatural for equality-generating dependencies such as functional dependencies. Consider the following example:

Example 2

R_1		R_2		
<i>Student</i>	<i>Course</i>	<i>Course</i>	<i>Room</i>	<i>Hour</i>
<i>Jack</i>	<i>CS378</i>	<i>CS378</i>	<i>B215</i>	<i>M10</i>

R_3		
<i>Student</i>	<i>Room</i>	<i>Hour</i>
<i>John</i>	<i>B320</i>	<i>F12</i>

with the only dependency $C \rightarrow RH$.

This is not a complete state, since the subtuple $\langle \textit{Jack}, \textit{B215}, \textit{M10} \rangle$ will be forced by $C \rightarrow RH$ to appear in every weak instance, but it does not appear in R_3 . However, it is hard to argue that this state *violates* the $C \rightarrow RH$ dependency, which simply requires that each course be associated with a unique room and time.

In sum, we have described two notions of dependency satisfaction, *consistency* and *completeness*. Consistency is the natural generalization of weak satisfaction and seems appropriate when only equality-generating dependencies are given, but disagrees with the standard notion in the presence of tuple-generating dependencies. Completeness is based on the intuitive semantics of tuple-generating dependencies but appears unnatural for equality-generating dependencies. It is our thesis that neither approach is the *correct* one, but rather than they correspond to different policies on constraint enforcement, and each one is appropriate in different circumstances.

After introducing definitions and notation in Section 2, Section 3 presents the notions of consistency and completeness. We show how to construct for a database state ρ two sets of first-order sentences, C_ρ and K_ρ , such that ρ is consistent with the given dependencies if and only if C_ρ is finitely satisfiable and ρ is complete with respect to the given dependencies if and only if K_ρ is finitely satisfiable. In Section 4, we characterize consistency and completeness in terms of the chase of the associated tableaux [22]. When all dependencies are total, our results provide a decision procedure for testing consistency and completeness of a state. However, we show that testing whether a state is inconsistent with a typed equality-generating dependency and testing whether a state is incomplete with respect to a join dependency are both NP-complete problems. Furthermore, the general problems of testing consistency and completeness under full dependencies are shown to be complete in exponential time. In Section 5, we study the decision problem for consistency and completeness when embedded dependencies are present. By relating consistency and completeness to the well-studied question of dependency

implication, we show that both are undecidable in the general case. Finally, in Section 6 we examine the construction of sets of sentences similar to C_ρ and K_ρ , but without using a predicate that stands for the universal relation scheme. We show that this can be done when the database scheme is weakly cover-embedding, a necessary condition for independence of the database scheme [15].

Our results deal with *untyped* relations and dependencies, that is, a value may appear in different columns of a relation. However, all of the results, except for Theorems 8, 9, and 15, can be specialized to the typed case.

2. Definitions and Notation

2.1 RELATIONS, DATABASE STATES, AND TABLEAUX. We fix a finite set of *attributes* called the *universe*, $U = \{A_1, \dots, A_n\}$. Each attribute A_i has an associated infinite set called its *domain* and denoted $\text{dom}(A_i)$. Since we deal with *untyped* databases, we let all the domains be the same, say the integers. A *relation scheme* R is a subset of U . A *database scheme* $\mathbb{R} = \{R_1, \dots, R_k\}$ is a collection of relation schemes such that the union of the R_i s is U . A *tuple* defined on relation scheme R is a function that maps each attribute in R to a value. The value can be either an integer or a *variable* taken from an infinite set of uninterpreted symbols. A tuple can be visualized as a row of a table where the columns are labeled by the attributes. A *tableau* on R is a finite set of tuples defined on R . If t is a tuple on R and X is a subset of R , $t[X]$ denotes the restriction of t to X . If $t(A)$ is an integer for every A in X , we say that t is *total on* X . A *relation* on R is a tableau on R such that every tuple is total on R . A relation on U is called a *universal relation*.

For r a tableau on R and X a subset of R , the *projection of r on X* is

$$\pi_X(r) = \{t[X] \mid t \in r \text{ and } t \text{ is total on } X\}.$$

Note that our definition of projection corresponds to what is sometimes called *total projection* in the literature, so the projection of any tableau on any set of attributes is always a relation. When $\mathbb{R} = \{R_1, \dots, R_k\}$, we write $\pi_{\mathbb{R}}(r)$ for $\langle \pi_{R_1}(r), \dots, \pi_{R_k}(r) \rangle$.

A state of a database scheme \mathbb{R} is a function ρ that maps every relation scheme R in \mathbb{R} to a relation on R . We write

$$\rho = \langle r_1, \dots, r_k \rangle = \langle \rho(R_1), \dots, \rho(R_k) \rangle.$$

We associate with each state ρ of \mathbb{R} a tableau T_ρ defined on the universe. T_ρ contains exactly one tuple for each tuple in each relation of ρ . The tuple r corresponding to tuple t in $\rho(R)$ is constructed by letting $r[R] = t[R]$ and filling the rest of r with distinct variables that appear nowhere else in T_ρ .

Example 3. Let $\mathbb{R} = \{AB, BCD, AD\}$, and ρ is the state below.

$\rho(AB)$		$\rho(BCD)$			$\rho(AD)$	
A	B	B	C	D	A	D
1	2	2	5	8	1	9
1	3	4	6	7		

Then T_ρ is

A	B	C	D
1	2	b_1	b_2
1	3	b_3	b_4
b_5	2	5	8
b_6	4	6	7
1	b_7	b_8	9

A *valuation* ν for a tableau R is a mapping from the symbols in the tableau into variables and constants such that $\nu(c) = c$ for every constant c that appears in T .

2.2 DEPENDENCIES. Following [6], we use tableaux to represent implicational dependencies. A *template dependency* (TD) is a pair $d = \langle T, w \rangle$, where T is a tableau containing no constants and w is a tuple containing no constants. We say d is *full* or *total* if $w[A]$ appears in T for every attribute A . Otherwise, d is said to be *embedded* or *partial*. A relation I satisfies d if for every valuation ν such that $\nu(T) \subseteq I$, there exists an extension ν' of ν to all the symbols of w such that $\nu'(w) \in I$. Informally, a template dependency says that, if certain tuples appear in I , then some other tuple must also appear. Template dependencies are a special case of the *tuple-generating dependencies* (TGDs), where a set of tuples is allowed instead of the single tuple w . For total dependencies, one can assume without loss of generality that a single tuple appears on the right-hand side [6]; thus total template dependencies are no less general than total tuple-generating dependencies. *Join dependencies* (JDs) [1, 24] are a special case of total TDs. An *equality-generating dependency* (EGD) is a pair $d = \langle T, (a_1, a_2) \rangle$, where T is a tableau containing no constants, and a_1, a_2 are variables that appear in T for some A . A tableau S satisfies an EGD $d = \langle T, (a_1, a_2) \rangle$ if for every valuation ν such that $\nu(T) \subseteq S$, $\nu(a_1) = \nu(a_2)$. *Functional dependencies* (FDs) are a special case of EGDs. We use the term *dependency* in the generic sense, meaning either a TGD or an EGD.

EGDs also act like TGDs, since by generating new equalities they generate new tuples. This action can be simulated by total TDs. Beeri and Vardi [5, 6] show how to construct, given a set D of dependencies, a set \bar{D} of TGDs that has the following properties:

- (1) \bar{D} is obtained from D by replacing each EGD by some TDs.
- (2) $D \models \bar{D}$.
- (3) Let d be a TGD. If $D \models d$, then $\bar{D} \models d$.

We call \bar{D} the *EGD-free version* of D .

3. Consistency and Completeness

In this section, we define two properties of a database state with respect to a set of dependencies, consistency, and completeness, which we consider to be two different aspects of dependency satisfaction. We characterize these properties in terms of the satisfiability of two first-order theories associated with the state.

Let us fix a database scheme $\mathbb{R} = \{R_1, \dots, R_n\}$. Let $\text{WEAK}(D, \rho)$ be the set of all weak instances for a database state ρ under a set of dependencies D . That is, $\text{WEAK}(D, \rho)$ is the set of all universal relations satisfying D such that their projections contain each relation in ρ . We say that a state ρ is *consistent* with set of dependencies D if $\text{WEAK}(D, \rho) \neq \emptyset$. The *completion* of a state ρ , ρ^+ , is defined by

$$\rho^+ = \bigcap_{I \in \text{WEAK}(\bar{D}, \rho)} \{\pi_R(I)\},$$

where the intersection is taken relationwise. Note that $\rho \subseteq \rho^+$ for any ρ . We say that a state ρ is *complete* with respect to set of dependencies D if ρ equals its completion, that is, $\rho = \rho^+$.

Intuitively, a state ρ is consistent if there is some way of adding tuples to relations of ρ that will transform it into the set of projections of some satisfying universal instance. If ρ is consistent, then there are many different sets of tuples that can be added to it to demonstrate its consistency. However, there are certain tuples that must appear in every such extension of ρ . If all these necessary tuples are already in ρ , then we say ρ is complete. Note that the definition of completeness is based on the EGD-free version of D , \bar{D} . This is done to allow consistency and completeness to be independent notions. Although $\text{WEAK}(D, \rho)$ could be empty (for inconsistent ρ), $\text{WEAK}(\bar{D}, \rho)$ is never empty. We show later that, for consistent states, it does not matter whether D or \bar{D} is used.

Several workers [11, 23] have advocated the use of first-order logic to express dependencies and other constraints. As we explained in the introduction, their approach does not easily generalize to dependencies in multirelation databases. Given a dependency statement such as $X \rightarrow Y$ in a database scheme where X and Y do not appear together in one relation scheme, or appear in more than one, it is not clear how satisfaction of the dependency can be formalized as satisfaction of some first-order sentence.

The notions of satisfaction proposed above do provide a means of using first-order logic to formalize dependencies; a rather drastic shift in point of view is, however, required. It is no longer possible to write down a sentence for each dependency and ask whether the database provides a model for each of these sentences. Consider for example the notion of consistency as satisfaction, and let $d_1 = A \rightarrow C$, $d_2 = B \rightarrow C$, with the database scheme $\{AB, BC\}$. Let $\rho(AB) = \langle 00, 01 \rangle$ and $\rho(BC) = \langle 01, 12 \rangle$. It is easy to see that ρ is consistent with d_1 and with d_2 , but it is not consistent with $\{d_1, d_2\}$.

Our approach is to construct two sets of sentences, C_ρ and K_ρ , for each state ρ . We show that ρ is consistent exactly when C_ρ is finitely satisfiable and that ρ is complete exactly when K_ρ is finitely satisfiable. Thus we reduce both notions of satisfaction to the standard logical notion of finite satisfiability. Note that, in this approach, consistency and completeness of a state are not first-order notions; they are statements about first-order theories rather than statements in these theories.

Before constructing the sets of sentences C_ρ and K_ρ , we fix a linear ordering on the elements of the universe U . We now write U as the sequence $\langle A_1, \dots, A_n \rangle$. Each relation scheme R in \mathbb{R} may be written as the sequence $\langle A_{i_1}, \dots, A_{i_m} \rangle$, where $i_j < i_k$ for $j < k$.

C_ρ and K_ρ each contain two subsets of sentences, the *database scheme axioms* and the *state axioms*. The scheme axioms depend only on the database scheme \mathbb{R} and the set of dependencies D ; the state axioms depend on the state ρ .

The scheme axioms of both C_ρ and K_ρ include the *containing instance* axioms. For each relation scheme R , there is one containing instance axiom that says that

every tuple in the R -relation of the state must be the projection on R of some tuple of the universal relation. In other words, the containing instance axioms assert the existence of a containing instance for the state. Formally, for each relation scheme $R = \langle A_1, \dots, A_m \rangle$ in \mathbb{R} , there is a sentence of the form

$$\forall y \exists x (R(y_1, \dots, y_m) \rightarrow U(u_1, \dots, u_n)),$$

where y is the sequence y_1, y_2, \dots, y_m , x is the sequence x_1, \dots, x_{n-m} , and u_1, \dots, u_n is a sequence without repetitions of variables such that u_{ij} is y_j for all $1 \leq i \leq m$, and the rest of the u s are x s.

The scheme axioms of C_ρ also include the *dependency axioms*, which are just the dependencies in D encoded as implicational sentences as described by Fagin [10]. K_ρ also contains dependency axioms, but in this case we use the EGD-free version of D , \bar{D} , rather than D itself.

The state axioms of both C_ρ and K_ρ contain the state ρ encoded as a set of quantifier-free sentences. For each tuple $\langle a_1, \dots, a_m \rangle$ in $\rho(R)$, we include the sentence $R(a_1, \dots, a_m)$, where the a_i 's are constants.

Finally, the state axioms of C_ρ include the *distinctness axioms*, and those of K_ρ include the *completeness axioms*. The distinctness axioms are the set of inequalities $c \neq d$, where c and d are distinct constants appearing in ρ . The completeness axioms contain, for each tuple $\langle a_1, \dots, a_m \rangle$ such that each a_i appears in ρ but the tuple itself does not appear in $\rho(R)$, the sentence

$$\forall y (\neg U(y_0, a_1, \dots, a_m, y_m)),$$

where y is constructed as in the containing instance axioms. Intuitively, the completeness axioms say that only tuples appearing in $\rho(R)$ can be in the projection of the universal relation on R .

Example 4. We construct C_ρ and K_ρ for the state shown in Example 1. U is the sequence $\langle S, C, R, H \rangle$; \mathbb{R} contains the schemes $R_1 = \langle S, C \rangle$, $R_2 = \langle C, R, H \rangle$, and $R_3 = \langle S, R, H \rangle$. The dependencies are the functional dependencies $SH \rightarrow R$, $RH \rightarrow C$, and the multivalued dependency (MVD) $C \twoheadrightarrow S \mid RH$.

—*Containing instance axioms*

$$\begin{aligned} \forall s, c \quad \exists r, h (R_1(s, c) \rightarrow U(s, c, r, h)) \\ \forall c, r, h \quad \exists s (R_2(c, r, h) \rightarrow U(s, c, r, h)) \\ \forall s, r, h \quad \exists c (R_3(s, r, h) \rightarrow U(s, c, r, h)) \end{aligned}$$

—*Dependencies*

$$\begin{aligned} (\forall s_1 c_1 c_2 h_1 r_1 r_2) (U(s_1, c_1, r_1, h_1) \wedge U(s_1, c_2, r_2, h_1) \rightarrow r_1 = r_2) \\ (\forall s_1 s_2 c_1 c_2 h_1 r_1) (U(s_1, c_1, r_1, h_1) \wedge U(s_2, c_2, r_1, h_1) \rightarrow c_1 = c_2) \\ (\forall s_1 s_2 c_1 r_1 r_2 h_1 h_2) (U(s_1, c_1, r_1, h_1) \wedge U(s_2, c_1, r_2, h_2) \rightarrow U(s_2, c_1, r_1, h_1)) \end{aligned}$$

—*EGD-free dependency axioms*

$$\begin{aligned} (\forall s_1 s_2 c_1 c_2 c_3 h_1 h_2 r_1 r_2) (U(s_1, c_1, r_1, h_1) \wedge U(s_1, c_2, r_2, h_1) \wedge U(s_2, c_3, r_1, h_2) \\ \rightarrow U(s_2, c_3, r_2, h_2)) \\ (\forall s_1 s_2 s_3 c_1 c_2 h_1 h_2 r_1 r_2) (U(s_1, c_1, r_1, h_1) \wedge U(s_2, c_2, r_1, h_1) \wedge U(s_3, c_1, r_2, h_2) \\ \rightarrow U(s_3, c_2, r_2, h_2)) \\ (\forall s_1 s_2 c_1 r_1 r_2 h_1 h_2) (U(s_1, c_1, r_1, h_1) \wedge U(s_2, c_1, r_2, h_2) \rightarrow U(s_2, c_1, r_1, h_1)) \end{aligned}$$

—State axioms

$$\begin{aligned}
 &R_1(\text{Jack}, \text{CS378}) \\
 &R_2(\text{CS378}, \text{B215}, \text{M10}) \\
 &R_2(\text{CS378}, \text{B213}, \text{W10}) \\
 &R_3(\text{Jack}, \text{B215}, \text{M10})
 \end{aligned}$$

—Distinctness axioms

$$\begin{aligned}
 &\text{B215} \neq \text{B213} \\
 &\text{M10} \neq \text{W10} \\
 &\text{Jack} \neq \text{CS378} \\
 &\text{Jack} \neq \text{B215} \quad \text{etc.}
 \end{aligned}$$

—Completeness axioms

For R_1 :
$$\begin{aligned}
 &\forall r, h \neg U(\text{CS378}, \text{CS378}, r, h) \\
 &\forall r, h \neg U(\text{B215}, \text{CS378}, r, h) \quad \text{etc.}
 \end{aligned}$$
For R_2 :
$$\begin{aligned}
 &\forall s \neg U(s, \text{CS378}, \text{B213}, \text{M10}) \\
 &\forall s \neg U(s, \text{CS378}, \text{B215}, \text{W10}) \quad \text{etc.}
 \end{aligned}$$
For R_3 :
$$\forall c \neg U(\text{Jack}, c, \text{B213}, \text{W10}) \quad \text{etc.}$$

C_ρ consists of the containing instance axioms, the dependency axioms, the state axioms, and the distinctness axioms. K_ρ consists of the containing instance axioms, the EGD-free dependency axioms, the state axioms, and the completeness axioms.

Before proceeding to our results, we introduce some basic definitions and notation of model theory. A *structure* for a language L consists of a *domain* of elements and an *interpretation* of each predicate and constant of L . A structure with a finite domain is *finite*. If \mathbb{M} is a structure with domain A and P is a k -ary predicate symbol, then $\mathbb{M}(P) \subseteq A^k$ is the interpretation of P in \mathbb{M} , and if c is a constant, $\mathbb{M}(c) \in A$ is the interpretation of c in \mathbb{M} .

A structure \mathbb{M} for a language L is a *model* of a set of sentences Σ if for every $\sigma \in \Sigma$, σ is true in \mathbb{M} , written $\mathbb{M} \models \sigma$ (we assume the reader to be familiar with the notion of truth in a model). A set of sentences is *finitely satisfiable* if it has a finite model.

The next two theorems establish that finite satisfiability of C_ρ and K_ρ are identical respectively to consistency and completeness of ρ with respect to D .

THEOREM 1. C_ρ is finitely satisfiable if and only if ρ is consistent with D .

PROOF. If ρ is consistent, it is clear that for every $I \in \text{WEAK}(D, \rho)$ the structure \mathbb{M} with $\mathbb{M}(R) = \rho(R)$ and $\mathbb{M}(U) = I$ is a finite model of C_ρ . For the converse, let \mathbb{M} be a finite model of C_ρ . For each pair of constants c, d in the language of C_ρ , we have $\mathbb{M} \models c \neq d$, so $\mathbb{M}(c) \neq \mathbb{M}(d)$. Thus we can assume, without loss of generality, that constants are interpreted as themselves, that is, $\mathbb{M}(c) = c$. It is easy to verify that $\mathbb{M}(U)$ is a containing instance for ρ that satisfies D . That is, $\mathbb{M}(U) \in \text{WEAK}(D, \rho)$, so ρ is consistent with D . \square

THEOREM 2. K_ρ is finitely satisfiable if and only if ρ is complete with respect to D .

PROOF. Suppose K_ρ is satisfiable with a finite model M . We claim that K_ρ has a finite model M' where no two constants have the same interpretation. To construct M' , we replace each element of the domain of M with many distinct copies of that element. For example, if $\langle a, b, c \rangle \in M(R)$, then $M'(R)$ would contain tuples $\langle a_1, b_1, c_1 \rangle, \langle a_1, b_1, c_2 \rangle, \langle a_1, b_2, c_3 \rangle$. If two constants are interpreted as the same element in M , they would be interpreted as two distinct copies of that element in M' . Since K_ρ does not have equality in it, M' is still a model for K_ρ . Thus, without loss of generality, we can assume that in M' constants are interpreted as themselves, so $M'(U) \in \text{WEAK}(\bar{D}, \rho)$. By the completeness axioms, $\pi_R(I)$ does not contain any tuple constructed from values appearing in ρ but not itself appearing in ρ . Hence, ρ^+ cannot contain any such tuple, so ρ is complete.

For the converse, suppose that ρ is complete. Consider the set S of all tuples t on some relation scheme R such that t is constructed from values appearing in ρ but t is not in ρ . Since ρ is complete, if t is an R -tuple in S , then there is a universal relation $I_t \in \text{WEAK}(\bar{D}, \rho)$ such that $t \notin \pi_R(I_t)$. We use now the *direct product* construction [10] to produce a universal relation $I \in \text{WEAK}(\bar{D}, \rho)$ such that if t is an R -tuple from S , then $t \notin \pi_R(I)$. I is the direct product $\times_{t \in S} I_t$, constructed as follows. Let $S = \{t_1, \dots, t_m\}$. The values in I are m -sequences $c = \langle c_1, \dots, c_m \rangle$ of constants from ρ , where we identify the m -sequence $\langle c, c, \dots, c \rangle$ with the constant c . Now a tuple s is in I if and only if the tuple s_i , which is obtained from s by projecting each m -sequence on its i th component, is in I_{t_i} . It is straightforward to verify that I is a containing instance for ρ . Furthermore, since dependencies are preserved under direct product [10], I must satisfy \bar{D} . The finite structure M with $M(R) = \rho(R)$ for each R and $M(U) = I$ is a model of K_ρ . \square

4. Testing Satisfaction under Full Dependencies

In this section we show that both consistency and completeness of a state can be tested by chasing the associated tableau, when embedded dependencies are not present. Thus the upper bounds on complexity known for chasing tableaux under various special kinds of full dependencies apply also to testing satisfaction. Throughout this section, D is a set of full dependencies. We start by defining the chase of a tableau under full dependencies and establishing preliminary results about chasing with the EGD-free version of D .

The *chase* of a tableau T with respect to a set of dependencies D , denoted by $\text{CHASE}_D(T)$, is the result of applying the following two transformation rules exhaustively to T .

TD-rule. If $\langle S, w \rangle$ is a TD in D , and there exists a valuation ν such that $\nu(S) \subseteq T$, add $\nu(w)$ to the rows of T .

EGD-Rule. Suppose $\langle S, (a_1, a_2) \rangle$ is an EGD in D , and there exists a valuation ν such that $\nu(S) \subseteq T$, and $\nu(a_1), \nu(a_2)$ are not both constants. If only one of $\nu(a_1), \nu(a_2)$ is a constant, rename all occurrences of the other one in T to that constant. If both are variables, rename all occurrences of the higher numbered variable to the lower numbered one.

Given a tableau T_ρ and set of dependencies D , we write

$$T_\rho^* = \text{CHASE}_D(T_\rho) \quad \text{and} \quad T_\rho^+ = \text{CHASE}_{\bar{D}}(T_\rho)$$

where \bar{D} is the EGD-free version of D .

We need some preliminary results before characterizing consistent states.

LEMMA 1. *For each $I \in \text{WEAK}(D, \rho)$, there is a valuation ν such that $\nu(T_\rho^*) \subseteq I$.*

PROOF. By induction on the computation of T_ρ^* . \square

LEMMA 2. *For any consistent state ρ , $\bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\} = \pi_R(T_\rho^*)$.*

PROOF. Let $I \in \text{WEAK}(D, \rho)$. Let ν be the valuation of the previous lemma, $\nu(T_\rho^*) \subseteq I$. Then

$$\pi_R(T_\rho^*) \subseteq \pi_R(\nu(T_\rho^*)) \subseteq \pi_R(I).$$

It follows that $\pi_R(T_\rho^*) \subseteq \bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$. For the other inclusion, let t be a tuple in the R_i -component of $\bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$. Let ν be an injective valuation for T_ρ^* that maps each variable to a constant not appearing in t . Since $\nu(T_\rho^*) \in \text{WEAK}(D, \rho)$, $t \in \pi_{R_i}(\nu(T_\rho^*))$. By construction of ν , this implies that t must come from some R_i -total tuple of T_ρ^* , so $t \in \pi_{R_i}(T_\rho^*)$. \square

Consistency and completeness of a state can be characterized in terms of the associated tableau T_ρ as follows.

THEOREM 3. *The following are equivalent.*

- (a) ρ is consistent with D .
- (b) T_ρ^* satisfies D .

PROOF

(a) *implies (b).* Let $I \in \text{WEAK}(D, \rho)$, and let ν be the valuation of Lemma 1. Suppose that T_ρ^* does not satisfy some $d \in D$. d cannot be a TD, since in that case a TD-rule is applicable to T_ρ^* . Thus d must be an EGD $\langle S, (a_1, a_2) \rangle$, and there is a valuation ν' such that $\nu'(S) \subseteq T_\rho^*$ and $\nu'(a_1) \neq \nu'(a_2)$. Both $\nu'(a_1)$ and $\nu'(a_2)$ must be constants, otherwise an EGD-rule is applicable to T_ρ^* . But now $\nu(\nu'(S)) \subseteq I$ and $\nu(\nu'(a_1)) \neq \nu(\nu'(a_2))$, so I does not satisfy D —a contradiction.

(b) *implies (a).* Let ν be an injective valuation for T_ρ^* that maps each variable to a constant not appearing in ρ . Then $\nu(T_\rho^*)$ satisfies D and it is a containing instance for ρ . Thus $\text{WEAK}(D, \rho) \neq \emptyset$. \square

We need two more preliminary results before characterizing complete states.

LEMMA 3. *For each $I \in \text{WEAK}(\bar{D}, \rho)$, there is a valuation ν such that $\nu(T_\rho^+) \subseteq I$.*

PROOF. By induction on the computation of T_ρ^+ . \square

The next lemma shows that the completion of a state can be obtained from T_ρ^+ .

LEMMA 4. *For any state ρ , $\rho^+ = \pi_R(T_\rho^+)$.*

PROOF. Let $I \in \text{WEAK}(\bar{D}, \rho)$. Let ν be the valuation of the previous lemma, $\nu(T_\rho^+) \subseteq I$. Then

$$\pi_R(T_\rho^+) \subseteq \pi_R(\nu(T_\rho^+)) \subseteq \pi_R(I).$$

It follows that $\pi_R(T_\rho^+) \subseteq \rho^+$. For the other inclusion, let t be a tuple in the R_i -component of ρ^+ . Let ν be an injective valuation for T_ρ^+ that maps each variable to a constant not appearing in t . Since $\nu(T_\rho^+) \in \text{WEAK}(\bar{D}, \rho)$, $t \in \pi_{R_i}(\nu(T_\rho^+))$. By construction of ν , this implies that t must come from some R_i -total tuple of T_ρ^+ , so $t \in \pi_{R_i}(T_\rho^+)$. \square

THEOREM 4. *The following are equivalent:*

- (a) ρ is complete with respect to D .
- (b) ρ is complete with respect to \bar{D} .
- (c) $\rho = \pi_R(T_\rho^+)$.

PROOF. The equivalence of (a) and (b) is immediate from the fact that $\bar{\bar{D}} = \bar{D}$. The equivalence of (a) and (c) follows directly from Lemma 4. \square

We have defined completeness and consistency to be independent notions. However, it is interesting to note that for consistent states the notion of completeness can be simplified.

THEOREM 5. *For state ρ consistent with dependencies D , the following are equivalent:*

- (a) ρ is complete with respect to D .
- (b) $\rho = \pi_R(T_\rho^*)$.
- (c) $\rho = \bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$.

PROOF

(b) equivalent to (c). Follows from Lemma 2.

(a) equivalent to (b). By Theorem 4, ρ is complete with respect to D iff $\rho = \pi_R(T_\rho^+)$. We claim that $\pi_R(T_\rho^*) = \pi_R(T_\rho^+)$ for consistent states. Since ρ is consistent, by Theorem 1, T_ρ^* satisfies D . By property (2) of \bar{D} , we also have that T_ρ^* satisfies \bar{D} . Hence, by Lemma 3, there is a valuation ν such that $\nu(T_\rho^+) \subseteq T_\rho^*$. Consequently, $\pi_R(T_\rho^+) \subseteq \pi_R(T_\rho^*)$. For the other inclusion, let t be a tuple in the R_i -component of $\pi_R(T_\rho^*)$. Let ν be a map for T_ρ that sends distinct constants to distinct variables. Let $T = \nu(T_\rho)$, and let s be a tuple such that $s[R_i] = \nu(t)$ and the rest of s consists of distinct new variables. We claim that $D \models \langle T, s \rangle$. Indeed, there is a tuple t_1 in T_ρ^* such that $t_1[R_i] = t$. Therefore, there is a tuple s_1 in $\text{CHASE}_D(T)$ such that $s_1[R_i] = s[R_i]$. By the results in [6], it follows that $D \models \langle T, s \rangle$. But then also $\bar{D} \models \langle T, s \rangle$, by property (3) of \bar{D} . Thus, by the results in [6] there is a tuple s_2 in $\text{CHASE}_{\bar{D}}(T)$ such that $s_2[R_i] = s[R_i]$. It follows that $t \in \pi_R(T_\rho^+)$. \square

COROLLARY 1. *The following are equivalent.*

- (a) ρ is consistent and complete with respect to D .
- (b) T_ρ^* satisfies D and $\rho = \pi_R(T_\rho^*)$.
- (c) $\rho = \bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$.

The next theorem relates consistency and completeness to standard satisfaction for single relations.

THEOREM 6. *For $R = \{U\}$, $\rho(U)$ satisfies D if and only if ρ is consistent and complete with respect to D .*

PROOF

(Only if) Let $J = \rho(U) \in \text{SAT}(D)$. Since $J \in \text{WEAK}(D, \rho)$, ρ is consistent with D . We claim also that $\rho = \bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$. Since $\rho = \pi_R(J)$, clearly $\bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\} \subseteq \rho$. Furthermore, for every $I \in \text{WEAK}(D, \rho)$, $J \subseteq I$, that is, $\rho \subseteq \pi_R(I)$, so $\rho \subseteq \bigcap_{I \in \text{WEAK}(D, \rho)} \{\pi_R(I)\}$. Therefore ρ is complete with respect to D .

(If) Let ρ be consistent and complete with respect to D . Then T_ρ^* satisfies D by Theorem 3, and $\rho = \pi_R(T_\rho^*)$ by Theorem 5. Since $R = \{U\}$ and all dependencies are total, clearly $\pi_R(T_\rho^*) = T_\rho^*$, hence $\rho = T_\rho^*$ satisfies D . \square

As a consequence of Theorems 3 and 4, the chase is a decision procedure for consistency and completeness under full dependencies. In the rest of this section we shall give upper and lower complexity bounds for these problems.

We first give NP-completeness results that follow from Theorem 6.

THEOREM 7

- (1) For $\mathbb{R} = \{U\}$, testing whether a state ρ is inconsistent with a typed EGD or whether it is incomplete with respect to a JD is NP-complete.
- (2) For $\mathbb{R} = \{U\}$, testing whether a state ρ is not complete with respect to a set D of full dependencies is NP-complete.
- (3) Testing whether a state ρ is inconsistent with a set D of EGDs is NP-complete.

PROOF

- (1) In [20] it is shown that testing whether a relation violates a JD is NP-complete, and in [4] it is shown that testing whether a relation violates a typed EGD is NP-complete. The claim now follows by Theorem 6.
- (2) The claim follows by the above-mentioned NP-completeness results in [4] and [20] and Theorem 6.
- (3) NP-hardness follows from the first claim; we have to show that the problem is in NP. To test for inconsistency, one constructs T_ρ and chases it by D . If at any stage the chase requires identifying two constants, then ρ is inconsistent with D . By [4] chasing by EGDs can be done in nondeterministic polynomial time. \square

We now refer to the general case and prove lower and upper exponential time bounds.

THEOREM 8. *Testing whether a state ρ is consistent with a set D of full dependencies is EXPTIME-complete.*

PROOF. As observed before, to test for inconsistency, one constructs T_ρ and chases it by D . If at any stage the chase requires identifying two constants, then ρ is inconsistent with D . Otherwise, it is consistent. An analysis of the chase in [4] shows that it can be done in exponential time. It remains to show that the problem is EXPTIME-hard. We show it by reduction from the implication problem for full TDs, which was shown in [8] to be EXPTIME-complete. That is, given a set D of full TDs and a full TD d , we construct in polynomial time a set D' of full dependencies and a state ρ such that $D \models d$ iff ρ is inconsistent with D' .

Let U be the relation scheme for the dependencies $D \cup d$. Let d be $\langle T, w \rangle$, with $T = \{w_1, \dots, w_m\}$. Without loss of generality, assume that there are at least two variables in T . To test whether $D \models d$, we chase T by D and see whether w is generated. The idea of the reduction is to have a state ρ that “looks like” T and a set D' of dependencies that simulate D and in addition force identification of two constants if w is generated. In order to do that, we need to mark the tuples in the original state and the tuples that are generated by the chase. The marking is done by equalities by the tuples. For that we add new attributes: the database scheme is $\mathbb{R} = \{U'\}$, where

$$U' = U \cup \{A, A_1, \dots, A_m, B, B_1, \dots, B_m\}.$$

Let α be a one-to-one valuation that maps the variables in T to constants. $\rho(U')$ has tuples u_1, \dots, u_m that correspond to the tuples w_1, \dots, w_m in T in the

following way:

- (1) $u_i[U] = \alpha(w_i)$,
- (2) $u_i[A]$ and $u_i[A_i]$ are the same new constant, and
- (3) u_i has distinct new constants elsewhere.

Note that, since \mathbb{R} has a single relation scheme, T_p is just $\rho(U')$.

Let now $\langle S, v \rangle$ be a full TD in D . We construct a full TD $\langle S', v' \rangle$ on U' and put it in D' . For each tuple v_i in S we have a tuple v'_i in S' defined as follows:

- (1) $v'_i[U] = v_i$,
- (2) v'_i has distinct new variables elsewhere.

v' is defined as follows:

- (1) $v'[U] = v$, and
- (2) $v'[A, A_1, \dots, A_m] = v'[B, B_1, \dots, B_m] = v_1[B, B_1, \dots, B_m]$.

For example, if $\langle S, v \rangle$ is

	<i>F</i>	<i>G</i>	<i>H</i>
<i>v</i> :	<i>f</i>	<i>g</i>	<i>h</i>
<i>v</i> ₁ :	<i>f</i>	<i>g</i>	<i>h1</i>
<i>v</i> ₂ :	<i>f</i>	<i>g1</i>	<i>h</i>
<i>v</i> ₃ :	<i>f1</i>	<i>g</i>	<i>g</i>

then $\langle S', v' \rangle$ is

	<i>A</i>	<i>A</i> ₁	<i>A</i> ₂	<i>A</i> ₃	<i>B</i>	<i>B</i> ₁	<i>B</i> ₂	<i>B</i> ₃	<i>F</i>	<i>G</i>	<i>H</i>
<i>v'</i> :	<i>b</i>	<i>b1</i>	<i>b2</i>	<i>b3</i>	<i>b</i>	<i>b1</i>	<i>b2</i>	<i>b3</i>	<i>f</i>	<i>g</i>	<i>h</i>
<i>v'</i> ₁ :	<i>b</i>	<i>b1</i>	<i>b2</i>	<i>b3</i>	<i>f</i>	<i>g</i>	<i>h1</i>
<i>v'</i> ₂ :	<i>f</i>	<i>g1</i>	<i>h</i>
<i>v'</i> ₃ :	<i>f1</i>	<i>g</i>	<i>h</i>

(Dots represent variables with unique occurrences.)

In addition we put in D' an EGD $\langle T', (a_1, a_2) \rangle$, where a_1 and a_2 are two distinct variables from T . T' has tuples w'_1, \dots, w'_m, w' that correspond to w_1, \dots, w_m, w .

w'_i is defined as follows:

- (1) $w'_i[U] = w_i$,
- (2) $w'_i[A]$ and $w'_i[A_i]$ are the same new variable, and
- (3) w'_i has new distinct variables elsewhere.

w' is defined as follows:

- (1) $w'[U] = w$, and
- (2) w' has new distinct variables elsewhere.

Notions of Dependency Satisfaction

119

For example, if $\langle T, w \rangle$ is

	<i>F</i>	<i>G</i>	<i>H</i>
<i>w</i> :	<i>f</i>	<i>g</i>	<i>h</i>
<i>w</i> ₁ :	<i>f</i>	<i>g</i>	<i>h</i> 1
<i>w</i> ₂ :	<i>f</i>	<i>g</i> 1	<i>h</i>
<i>w</i> ₃ :	<i>f</i> 1	<i>g</i>	<i>g</i>

then the constructed EGD is $\langle T', (f, f1) \rangle$, where T' is

	<i>A</i>	<i>B</i>	<i>A</i> ₁	<i>A</i> ₂	<i>A</i> ₃	<i>B</i> ₁	<i>B</i> ₂	<i>B</i> ₃	<i>F</i>	<i>G</i>	<i>H</i>
<i>w</i> ' ₁ :	<i>a</i> 1	.	<i>a</i> 1	<i>f</i>	<i>g</i>	<i>h</i> 1
<i>w</i> ' ₂ :	<i>a</i> 2	.	.	<i>a</i> 2	<i>f</i>	<i>g</i> 1	<i>h</i>
<i>w</i> ' ₃ :	<i>a</i> 3	.	.	.	<i>a</i> 3	.	.	.	<i>f</i> 1	<i>g</i>	<i>h</i>
<i>w</i> ' ₄ :	<i>f</i>	<i>g</i>	<i>h</i>

(Dots represent variables with unique occurrences.)

To prove that $D \models d$ iff ρ is inconsistent with D' we show that a chase of T by D can be simulated by a chase of T_p by D' , and vice versa.

Consider first a chase of T by D . We claim that for any tuple t generated by a TD $\langle S, v \rangle$ in a chase of T by D , one can generate a tuple t' by the TD $\langle S', v' \rangle$ in a chase of T_p by D' such that

- (1) $t'[U] = \alpha(t)$, and
- (2) $t'[A] \neq t'[A_i]$ for $1 \leq i \leq m$.

We leave the verification of this claim to the reader. If $D \models d$, then w is generated by the chase of T by D . Therefore, a chase of T_p by D' generates a tuple u such that $u[U] = \alpha(w)$. Let us now apply the EGD $\langle T', (a_1, a_2) \rangle$ with a valuation β that maps w'_i to u_i and w' to u . β agrees with α on the variables of T , so we are forced to identify $\alpha(a_1)$ and $\alpha(a_2)$. That means that ρ is inconsistent with D' .

Consider now a chase of T_p by D' . We claim that for any tuple t' generated by a TD $\langle S', v' \rangle$ in a chase of T_p by D' , one can generate a tuple t by the TD $\langle S, v \rangle$ in a chase of T by D such that

- (1) $t'[U] = \alpha(t)$, and
- (2) $t'[A] \neq t'[A_i]$ for $1 \leq i \leq m$.

We leave the verification of this claim to the reader. If ρ is inconsistent with D' then the EGD $\langle T', (a_1, a_2) \rangle$ must be applied with some valuation β . But since $w'_i[A] = w'_i[A_i]$, w'_i cannot be mapped by β to any other tuple but u_i . Thus β agrees with α on the variables of T . In particular, $\beta(w')[U] = \alpha(w)$. That is, w' must be mapped to a tuple u generated by the chase of T_p such that $u[U] = \alpha(w)$. But then w is generated by a chase of T by D , so $D \models d$.

To complete the proof we note that the reduction from D and d to ρ and D' can be done in polynomial time. \square

COROLLARY 2. For $\mathbb{R} = \{U\}$, testing whether a state ρ is consistent with a set D of full dependencies is EXPTIME-complete.

The corollary should be contrasted with clause (2) of Theorem 7. Although consistency is EXPTIME-complete even for database schemes with a single relation

scheme, for completeness there is a complexity gap between the case of database schemes with a single relation scheme and the case of database schemes with two relation schemes.

THEOREM 9. *Testing whether a state ρ is complete with respect to a set D of full TDs is EXPTIME-complete.*

PROOF. To test for incompleteness, one constructs T_ρ and chases it by D . If at any stage a tuple t is generated such that $t[R_i]$ has no variables and $t[R_i]$ is not in $\rho(R_i)$ for some of the relation schemes R_i in the database scheme \mathbf{R} , then ρ is incomplete with respect D . Otherwise, it is complete. An analysis of the chase in [4] shows that it can be done in exponential time. It remains to show that the problem is EXPTIME-hard. We show it by reduction from the implication problem for full TDs which was shown in [8] to be EXPTIME-complete. That is, given a set D of full TDs and a full TD d , we construct in polynomial time a set D' of full TDs and a state ρ such that $D \models d$ iff ρ is incomplete with D' .

Let U be the relation scheme for the dependencies $D \cup d$. Let d be $\langle T, w \rangle$, with $T = \{w_1, \dots, w_m\}$. Without loss of generality, assume that w is not in T . To test whether $D \models d$, we chase T by D and see whether w is generated. The idea of the reduction is to have a state ρ that “looks like” T and a set D' of dependencies that simulate D and in addition generate a “forbidden” tuple if w is generated. Unlike the reduction in the proof of Theorem 8, we have to be careful not to generate “forbidden” tuples too early. For that we add new attributes: the database scheme is $\mathbf{R} = \{R_1, R_2\}$, where $R_1 = U \cup \{A, B, A_1, \dots, A_m\}$ and $R_2 = \{C, D\}$. The new universe is $U' = R_1 \cup R_2$.

Let α be a one-to-one valuation that map the variables in T to constants. $\rho(R_1)$ has tuples u_1, \dots, u_m that correspond to the tuples w_1, \dots, w_m in T in the following way:

- (1) $u_i[U] = \alpha(w_i)$,
- (2) $u_i[A]$, $u_i[B]$, and $u_i[A_i]$ are the same new constant, and
- (3) u_i has distinct constants elsewhere.

$\rho(R_2)$ has a single tuple u_0 such that $u_0[C]$ and $u_0[D]$ are the same new constant.

In T_ρ , the tuples u_1, \dots, u_m are extended with distinct new variables for the attributes C and D , and u_0 is extended with unique variables for all attributes other than C and D . The idea is that for every tuple t generated by a chase of T_ρ , $t[A_i]$ and $t[D]$ are variables, so no “forbidden” tuple is generated until the very last step.

Let now $\langle S, v \rangle$ be a full TD in D . We construct a full TD $\langle S', v' \rangle$ on U' and put it in D' . For each tuple v_i in S we have a tuple v'_i in S' defined as follows:

- (1) $v'_i[U] = v_i$,
- (2) $v'_i[A]$ and $v'_i[B]$ are the same new variable, and
- (3) v'_i has distinct new variables elsewhere.

In addition S' has a tuple v'_0 defined as follows:

- (1) $v'_0[C]$ and $v'_0[D]$ are the same new variable, and
- (2) v'_0 has distinct new variables elsewhere.

Finally, v' is defined as follows:

- (1) $v'[U] = v$,
- (2) $v'[A_1, \dots, A_m] = v'_0[A_1, \dots, A_m]$,

- (3) $v'[A]$ and $v'[B]$ are the same old variable (any variable from v will do), and
 (4) $v'[CD] = v'[CD]$.

For example, if $\langle S, v \rangle$ is

	<i>F</i>	<i>G</i>	<i>H</i>
v :	<i>f</i>	<i>g</i>	<i>h</i>
v_1 :	<i>f</i>	<i>g</i>	<i>h1</i>
v_2 :	<i>f</i>	<i>g1</i>	<i>h</i>
v_3 :	<i>f1</i>	<i>g</i>	<i>h</i>

then $\langle S', v' \rangle$ is

	<i>A</i>	<i>B</i>	<i>A</i> ₁	<i>A</i> ₂	<i>A</i> ₃	<i>C</i>	<i>D</i>	<i>F</i>	<i>G</i>	<i>H</i>
v' :	<i>f</i>	<i>f</i>	<i>a1</i>	<i>a2</i>	<i>a3</i>	<i>c1</i>	<i>d1</i>	<i>f</i>	<i>g</i>	<i>h</i>
v'_0 :	.	.	<i>a1</i>	<i>a2</i>	<i>a3</i>	<i>c</i>	<i>c</i>	<i>f2</i>	<i>g2</i>	<i>h2</i>
v'_1 :	<i>a4</i>	<i>a4</i>	.	.	.	<i>c1</i>	<i>d1</i>	<i>f</i>	<i>g</i>	<i>h1</i>
v'_2 :	<i>a5</i>	<i>a5</i>	<i>f</i>	<i>g1</i>	<i>h</i>
v'_3 :	<i>a6</i>	<i>a6</i>	<i>f1</i>	<i>g</i>	<i>h</i>

(Dots represent variables with unique occurrences.)

In addition we put in D' a full TD $\langle T', w' \rangle$. T' has tuples w'_0, w'_1, \dots, w'_m that correspond to w, w_1, \dots, w_m .

w'_0 is defined as follows:

- (1) $w'_0[U] = w$, and
- (2) w'_0 has new distinct variables elsewhere.

w'_i is defined as follows:

- (3) $w'_i[U] = w_i$,
- (4) $w'_i[A]$ and $w'_i[A_i]$ are the same new variable, and
- (5) w'_i has new distinct variables elsewhere.

Finally, w' is defined as follows:

- (1) $w'[U] = w$, and
- (2) $w'[A, B, A_1, \dots, A_m, C, D] = w'_1[A, B, A_1, \dots, A_m, C, D]$.

For example, if $\langle T, w \rangle$ is

	<i>F</i>	<i>G</i>	<i>H</i>
v :	<i>f</i>	<i>g</i>	<i>h</i>
v_1 :	<i>f</i>	<i>g</i>	<i>h1</i>
v_2 :	<i>f</i>	<i>g1</i>	<i>h</i>
v_3 :	<i>f1</i>	<i>g</i>	<i>h</i>

then $\langle T', w' \rangle$ is

	A	B	A_1	A_2	A_3	D	E	F	G	H
w' :	$a1$	$a2$	$a1$	$a3$	$a4$	$c1$	$d1$	f	g	h
w'_0 :	f	g	h
w'_1 :	$a1$	$a2$	$a1$	$a3$	$a4$	$c1$	$d1$	f	g	$h1$
w'_2 :	$a5$.	.	$a5$.	.	.	f	$g1$	h
w'_3 :	$a6$.	.	.	$a6$.	.	$f1$	g	h

(Dots represent variables with unique occurrences.)

To prove that $D \models d$ iff ρ is incomplete with respect to D' we show that a chase of T by D can be simulated by a chase of T_ρ by D' , and vice versa.

Consider first a chase of T by D . We claim that for any tuple t generated by a TD $\langle S, v \rangle$ in a chase of T by D , one can generate a tuple t' by the TD $\langle S', v' \rangle$ in a chase of T_ρ by D' such that

- (1) $t'[U] = \alpha(t)$,
- (2) $t'[A]$ and $t'[B]$ are the same constant,
- (3) $t'[A_1], \dots, t'[A_m], t'[C]$, and $t'[D]$ are distinct variables.

We leave the verification of this claim to the reader. If $D \models d$, then w is generated by the chase of T by D . Therefore, a chase of T_ρ by D' generates a tuple u such that $u[U] = \alpha(w)$. Let us now apply the TD $\langle T', w' \rangle$ in the chase of T_ρ with the valuation β that maps w'_0 to u and maps w'_i to u_i . This generates the tuple $\beta(w')$ with

- (1) $\beta(w')[U] = \alpha(w)$, and
- (2) $\beta(w')[A, B, A_1, \dots, A_m, C, D] = u_1[A, B, A_1, \dots, A_m, C, D]$.

It follows that $\beta(w')[R_1]$ consists solely of constants and is not in $\rho(R_1)$, since w is not in T . So ρ is incomplete with respect to D' .

Consider now a chase of T_ρ by D' . We claim that for any tuple t' generated by a TD $\langle S', v' \rangle$ in a chase of T_ρ by D' , one can generate a tuple t by the TD $\langle S, v \rangle$ in a chase of T by D such that

- (1) $t'[U] = \alpha(t)$,
- (2) $t'[A]$ and $t'[B]$ are the same constant, and
- (3) $t'[A_1], \dots, t'[A_m], t'[C]$, and $t'[D]$ are distinct variables.

We leave the verification of this claim to the reader. If ρ is incomplete with respect to D' then the TD $\langle T', w' \rangle$ must be applied with some valuation β . But, since $w'_i[A] = w'_i[A_i]$, w'_i cannot be mapped to any other tuple but u_i . Thus β necessarily agrees with α on the variables that are in T . In particular, $\beta(w')[U] = \alpha(w)$, so w'_0 must be mapped to a tuple u generated by the chase of T_ρ such that $u[U] = \alpha(w)$. But then w is generated by a chase of T by D , so $D \models d$.

To complete the proof we note that the reduction from D and d to ρ and D' is polynomial. \square

We note that in the proofs of Theorem 8 and 9 we have untyped dependencies in D' even if the dependencies in D are typed. We believe that the exponential lower bounds for consistency and completeness hold also for typed dependencies.

5. Testing Satisfaction under Embedded Dependencies

In the previous section, we restricted our attention to full dependencies in order to obtain decidability results. In this section we return to arbitrary sets of dependencies. Our main result will be that both consistency and completeness are undecidable in this general setting. To show this, we will prove that consistency and completeness are recursively equivalent to certain dependency implication problems that are known to be undecidable.

The first step is to reduce consistency to the implication problem of EGDs by a set of dependencies. Let ρ be a state and D a set of arbitrary dependencies. Construct a set of EGDs E_ρ as follows. Let $T = \nu(T_\rho)$ be an isomorphic image of T_ρ in which no constants appear. For every pair of distinct constants c and d in T_ρ , $\langle T, (\nu(c), \nu(d)) \rangle$ is an element of E_ρ .

THEOREM 10. ρ is consistent with D if and only if for no EGD $e \in E_\rho$ is it the case that $D \models e$.

PROOF. Suppose ρ is consistent. Let $I \in \text{WEAK}(D, \rho)$. We can construct from ν a homomorphism η with $\eta(T) \subseteq I$ and for each constant c of ρ , $\eta(\nu(c)) = c$. I certainly satisfies D , but it violates each EGD in E_ρ . Therefore no element of E_ρ is implied by D .

For the converse, suppose there is no $e \in E_\rho$ such that $D \models e$. Let $E_\rho = \{e_1, \dots, e_k\}$ and $e_i = \langle T, (\nu(c_i), \nu(d_i)) \rangle$. Construct from T an atomic sentence $\tilde{\tau}$ by letting $\tilde{\tau}$ be the conjunction of all sentences $U(t)$ such that t is a tuple in T . Now consider the sentence d given by

$$\exists x(\tilde{\tau} \wedge \nu(c_1) \neq \nu(d_1) \wedge \dots \wedge \nu(c_k) \neq \nu(d_k)),$$

where x is a sequence of all the variables in T . We claim that $D' = D \cup d$ is finitely satisfiable. Suppose not; then $D \models \neg d$. Now $\neg d$ is a disjunctive EGD of the form

$$\forall x(\tilde{\tau} \rightarrow \nu(c_1) = \nu(d_1) \vee \dots \vee \nu(c_k) = \nu(d_k)).$$

We now rely on a finite version of a theorem of McKinsey [21] due to Graham and Vardi [14] to conclude that for some $1 \leq i \leq k$, $D \models \forall x(\tilde{\tau} \rightarrow \nu(c_i) = \nu(d_i))$, that is, $D \models e_i$, contradicting our assumption that D does not imply e for any $e \in E_\rho$.

Since D' is finitely satisfiable, it has a finite model M . Let $s(c)$ be the domain element assigned to each variable $\nu(c)$ of d to make d true in M . Note that $s(c_i) \neq s(d_i)$ for every c_i, d_i appearing in an inequality in d , so we can assume without loss of generality that $s(c) = c$. Then $M(U)$ is a weak instance for ρ , showing that ρ is consistent. \square

We now reduce the implication problem for EGDs to the consistency problem. Let D be any set of dependencies and let $e = \langle T, (a, b) \rangle$ be an EGD. We form the set R_e of states of the universal scheme $\{U\}$ as follows. For each mapping ν from the symbols of T to constants such that $\nu(a) \neq \nu(b)$, $\nu(T)$ is a member of R_e .

THEOREM 11. $D \models e$ if and only if no state in R_e is consistent with D .

PROOF. Suppose $D \models e$, and let $\nu(T)$ be any state in R_e . Clearly $\nu(T)$ violates e . Any weak instance in $\text{WEAK}(D, \nu(T))$ must satisfy D , and hence e . But $\nu(T)$ would have to be a subset of such a weak instance, which is impossible. Hence no such weak instance exists and $\nu(T)$ is inconsistent with D .

For the converse, suppose D does not imply e . Let I be any relation that satisfies D but not e . Such a relation must contain a homomorphic image $\nu(T)$ of T , such that $\nu(a) \neq \nu(b)$. Hence I is a weak instance for $\nu(T)$, which is an element of R_e . \square

From the last two theorems we obtain the following immediate corollary, relating the decidability of the membership problem for an EGD from a set of dependencies and the decidability of consistency under that set of dependencies.

COROLLARY 3. *Let D be a set of dependencies. Let D_e be the set of EGDs implied by D . The following are equivalent.*

- (a) D_e is recursive.
- (b) For every database scheme \mathbf{R} over the universe on which D is defined, the consistency of every state of \mathbf{R} is decidable.
- (c) The consistency of every state of the universal scheme of D is decidable.

PROOF

- (a) *implies* (b). Follows from Theorem 10.
- (b) *implies* (c). Immediate.
- (c) *implies* (a). Follows from Theorem 11. \square

The development of the last two theorems can be repeated to relate completeness to TD implication. For the analog to Theorem 10, construct a set of exponentially many TDs, G_ρ , from a state ρ . Elements of G_ρ are of the form $\langle T, w \rangle$, where T is the image of T_ρ under an injection ν to variables, and w is constructed as follows. Let R_i be a relation scheme in the given database scheme, and let t be a tuple on R_i such that t consists of constants taken from ρ but $t \notin \rho(R_i)$. Then $w[R_i] = \nu(t)$ and the rest of w consists of distinct new variables. Informally, each element $\langle T, w \rangle$ of G_ρ says that a containing instance for ρ must contain a tuple w such that its projection on some relation scheme is not in ρ . Note that G_ρ is a set of embedded TDs.

THEOREM 12. *ρ is complete with respect to D if and only if for no element $g \in G_\rho$ is it the case that $D \models g$.*

PROOF. If there is some $g = \langle T, w \rangle \in G_\rho$ such that $D \models g$, then let R be the relation scheme that led us to include g in G_ρ . There is some tuple t constructed with values from ρ that does not appear in $\rho(R)$. By property (3) of \bar{D} , we know that if $D \models g$ then $\bar{D} \models g$. Let I be an element of $\text{WEAK}(\bar{D}, \rho)$. Since I satisfies \bar{D} , it satisfies g . Let μ be a valuation such that $\mu(T_\rho) \subseteq I$. Since T is the image of T_ρ under an injection ν and I satisfies g , I must contain some tuple whose projection on R is t . It follows that $t \in \bigcap_{I \in \text{WEAK}(\bar{D}, \rho)} \{\pi_R(I)\}$, so ρ is incomplete.

For the converse, suppose that no $g \in G_\rho$ is implied by D . By property (3) of \bar{D} , no $g \in G_\rho$ is implied by \bar{D} . Thus, for every $g \in G_\rho$, there is a universal relation I_g such that I_g satisfies \bar{D} but not g . Let $g = \langle T, w \rangle$ and let t be the tuple on relation scheme R that led to the inclusion of g in G_ρ . There is a valuation ν such that $\nu(T) \subseteq I_g$ and $\nu(w[R]) \notin \pi_R(I_g)$. By the multiple copies construction of Theorem 2, we can assume that ν is injective. Thus we can assume without loss of generality that $I_g \in \text{WEAK}(\bar{D}, \rho)$, and $t \notin \pi_R(I_g)$. It follows that $t \notin \rho^+(R)$. Since this is true for every tuple t constructed from values in ρ but not itself in $\rho(R)$, it follows that ρ is complete. \square

For the analog of Theorem 11, let D be a set of dependencies and $g = \langle T, w \rangle$ a TD. We may assume $w \notin T$, else g is trivial. Let U be the relation scheme of

$D \cup g$, let $R = \{A \mid w[A] \text{ occurs in } T\}$, and let $\mathbb{R} = \{U, R\}$. Let ν be an injection from variables of T to constants. Let S be the set of all relations on U constructed from values in $\nu(T)$ that contain $\nu(T)$. Let K be the set of all states of the form $\pi_{\mathbb{R}}(r)$, where $r \in S$ and $\pi_{\mathbb{R}}(r)$ does not contain $\nu(w)$.

THEOREM 13. $D \models g$ if and only if every state of K is incomplete.

PROOF. Again the proof parallels that of Theorem 11. If $D \models g$, then the completion σ^+ of σ for each $\sigma \in K$ is such that $\nu(w)[R] \in \sigma^+(R)$, but $\nu(w) \notin \sigma(R)$, hence every such σ is incomplete.

Conversely, if D does not imply g , let I be a relation on U that satisfies D but not g . I is a weak instance for the state $\sigma = \pi_{\mathbb{R}}(I)$ in K . Since σ is exactly the projection of one of its weak instances, it must be complete. \square

COROLLARY 4. Let D_i be the set of all TDs implied by a set of dependencies D . D_i is recursive if and only if completeness of any state of any database scheme over the attributes of D is decidable.

We now state the main result of this section, which is a corollary of the four theorems above.

THEOREM 14. There does not exist an algorithm that will determine for every pair $\langle D, \rho \rangle$ whether ρ is consistent nor whether ρ is complete with respect to D .

PROOF. The implication problem of EGDs from arbitrary dependencies was shown undecidable by Vardi [26]. Implication of TDs was shown undecidable by Vardi [28] and Gurevich and Lewis [16]. \square

Since no algorithms exist for deciding either completeness or consistency, we consider solvable subcases. If implication is decidable for D , for example, if D contains only full dependencies, consistency and completeness are decidable, as shown in the previous section. But there may be specific database schemes for which consistency and completeness are decidable, even if implication is not. However, there is no algorithm that will decide, given a database scheme and a set of dependencies, whether consistency and completeness are decidable for that scheme and those dependencies.

THEOREM 15. The set $A = \{\langle D, \mathbb{R} \rangle \mid \text{consistency and completeness of states of } \mathbb{R} \text{ with respect to } D \text{ are decidable}\}$ is not recursive.

PROOF. Vardi [27] showed that it is undecidable whether the implication problem for a set of dependencies D is decidable. Suppose A were recursive. For a fixed set of dependencies D , the predicate of these dependencies has a fixed arity, that is, a fixed set of attributes. There are only finitely many database schemes over this set of attributes. As implication is decidable for D if and only if $\langle D, \mathbb{R} \rangle \in A$ for each of these finitely many \mathbb{R} s, an algorithm for membership in A would yield an algorithm to test decidability of the implication problem. So A cannot be recursive. \square

6. Discarding the Universal Relation Scheme

The sentences in C_p and K_p use a predicate letter corresponding to the universal relation scheme for the database. It is interesting to ask whether we can construct a theory with properties similar to Theorems 4 and 5, but using only predicate symbols R_1, \dots, R_n of the same arities and sorts as the relation schemes, thus avoiding the universal predicate. This amounts to asking whether dependency

satisfaction can be expressed in a “local” way, without having to resort to the existence of a universal relation.

There is a special case, the independent schemes, when the question can clearly be answered in the affirmative. Given a set of dependencies D on a database scheme $\{R_1, \dots, R_n\}$, the *projected dependencies* D_i are all the dependencies that must hold in any relation r_i on R_i such that $r_i = \pi_{R_i}(r)$, where r is a universal relation satisfying D . A state ρ is called *locally satisfying* if every $\rho(R_i)$ satisfies D_i . A database scheme is said to be *independent* if every locally satisfying state is consistent with D . When the database scheme is independent, we can write down the required set of sentences by expressing each dependency in the context of some R_i . For special cases such as functional and multivalued dependencies, projected dependencies can be easily characterized in terms of the original set D , although finding the D_i s is computationally hard [17]. For more general classes of dependencies, we do not even know if the D_i s are finite. In the general case, the results in this section should be viewed as existence proofs for the desired sets of sentences, rather than effective constructions.

Our main result in this section is that in fact it is possible to construct a set of sentences with the desired properties when the database scheme is *weakly cover embedding*. To define this notion, note that the projected dependencies D_i can be viewed as embedded dependencies on U . For D_i defined on R_i , we say a relation on U satisfies D_i if $\pi_{R_i}(I)$ does. We say that a database scheme \mathbb{R} *weakly cover embeds* a set of dependencies D if any state of \mathbb{R} consistent with $\bigcup_{i=1}^n D_i$ is consistent with D . In the framework of Section 4, in a weakly cover embedding scheme it suffices to chase using only dependencies local to some relation scheme of \mathbb{R} .

It is easy to see that the class of weakly cover embedding schemes contains both the cover embedding or dependency preserving schemes [18] and the independent schemes. Since, for cover embedding schemes, we have $\bigcup_{i=1}^n D_i \models D$, such schemes are weakly cover embedding. Since any state consistent with $\bigcup_{i=1}^n D_i$ is locally consistent, the independent schemes are weakly cover embedding. A polynomial time algorithm for testing whether a weakly cover embedding database scheme is independent, in the case where all dependencies are FDs, is given in [15]. Even for this restricted case, no algorithm to test whether a scheme is weakly cover embedding is known.

Given a state ρ of a weakly cover embedding database scheme, we construct a new set of sentences B_ρ as follows. The language of B_ρ is the same as the language of C_ρ , except that we do not use the universal predicate letter U . B_ρ contains four kinds of sentences.

- State axioms.* For each R_i and each tuple $t \in \rho(R_i)$, the sentence $R_i(t)$.
- Join-consistency axioms.* For each R_i in the database scheme, B_ρ contains the sentence

$$\forall \mathbf{x}(R_i(\mathbf{x}) \rightarrow (\exists b_1 \dots b_m)(R_1(\mathbf{v}_1) \wedge \dots \wedge R_n(\mathbf{v}_n)))$$

where $\mathbf{v}_i = \mathbf{x}$ and the \mathbf{v}_s are constructed from values in \mathbf{x} and the b s so that for all $1 \leq p, q \leq n$, if the j th attribute of R_p is the k th attribute of R_q , then $\mathbf{v}_p[j] = \mathbf{v}_q[k]$. Intuitively, the join-consistency axioms, together with the state containment axioms, assert the existence of a join-consistent state that contains ρ .

- Dependencies.* For each i , the set of dependencies D_i can be rewritten as a set of first-order sentences on R_i .
- Distinctness axioms.* As before, these assert that all constants are distinct.

Example 5. We construct B_ρ for the state of Examples 1 and 4. The universe is $U = \{S, C, R, H\}$, the relation schemes $R_1 = SC$, $R_2 = CRH$, $R_3 = SRH$, and the dependencies $SH \rightarrow R$, $RH \rightarrow C$. The projected dependencies are: $D_1 \neq \emptyset$, $D_2 = \{RH \rightarrow C\}$, $D_3 = \{SH \rightarrow R\}$.

—*State axioms*

$R_1(\text{Jack}, \text{CS378})$
 $R_2(\text{CS378}, \text{B215}, \text{M10})$
 $R_2(\text{CS378}, \text{B213}, \text{W10})$
 $R_3(\text{Jack}, \text{B215}, \text{M10})$

—*Join-consistency axioms*

$(\forall x_1 x_2)(R_1(x_1 x_2) \rightarrow (\exists b_1 b_2)(R_2(x_2 b_1 b_2) \wedge R_3(x_1 b_1 b_2)))$
 $(\forall x_1 x_2 x_3)(R_2(x_1 x_2 x_3) \rightarrow (\exists b_1)(R_1(b_1 x_1) \wedge R_3(b_1 x_2 x_3)))$
 $(\forall x_1 x_2 x_3)(R_3(x_1 x_2 x_3) \rightarrow (\exists b_1)(R_1(x_1 b_1) \wedge R_2(b_1 x_2 x_3)))$

—*Dependencies*

$(\forall r_1 h_1 c_1 c_2)(R_2(c_1 r_1 h_1) \wedge R_2(c_2 r_1 h_1) \rightarrow c_1 = c_2)$
 $(\forall s_1 r_1 r_2 h_1)(R_3(s_1 r_1 h_1) \wedge R_3(s_1 r_2 h_1) \rightarrow r_1 = r_2)$

—*Distinctness axioms*

$\text{B215} \neq \text{B213}$
 $\text{M10} \neq \text{W10}$ etc.

THEOREM 16. *For a weakly cover embedding database scheme, B_ρ is finitely satisfiable if and only if ρ is consistent with D .*

PROOF

(If) Suppose ρ is consistent, and let $I \in \text{WEAK}(D, \rho)$. Let $r_i = \pi_{R_i}(I)$ for each i . Consider the interpretation for B_ρ where each constant is mapped to itself and R_i is interpreted as r_i for each i . By definition of projected dependencies, each r_i satisfies D_i . Since I is a containing instance for ρ , the r_i s satisfy the state containment axioms. Finally, the r_i s are join-consistent by construction and thus satisfy the join-consistency axioms. It follows that we have a finite model for B_ρ .

(Only if) Suppose B_ρ is finitely satisfiable. Proceeding as in Theorem 1, we can assume that there is a database state ρ' that satisfies B_ρ . That is, ρ' contains ρ , is join consistent, and $\rho'(R_i)$ satisfies D_i for each i . Let I be a universal relation such that $\pi_{\mathbf{R}}(I) = \rho'$. Clearly I satisfies the D_i s, so $I \in \text{WEAK}(\bigcup_i D_i, \rho')$, showing that ρ' is consistent with $\bigcup_i D_i$. Since \mathbf{R} is weakly cover embedding, ρ' is also consistent with D . Since $\rho \subseteq \rho'$, $\text{WEAK}(D, \rho') \subseteq \text{WEAK}(D, \rho)$, so ρ is consistent with D . \square

The following example shows that the construction above does not generalize to noncover embedding schemes; we leave open the question of whether such a set of sentences can be constructed at all for arbitrary schemes.

Example 6. Let $\mathbf{R} = \{AC, BC\}$, $D = \{AB \rightarrow C, C \rightarrow B\}$, $\rho(AC) = \{\langle 01 \rangle, \langle 02 \rangle\}$, $\rho(BC) = \{\langle 31 \rangle, \langle 32 \rangle\}$. Note that $D_1 = \emptyset$, $D_2 = \{C \rightarrow B\}$. It is easy to see that ρ is consistent with $D_1 \cup D_2$, but not consistent with D . However, B_ρ is consistent in this example, since ρ is a join-consistent state that satisfies the local dependencies and hence provides a model for B_ρ .

7. Discussion

We have pointed out that there are two separate sides to the standard notion of dependency satisfaction; consistency and completeness. We view consistency as corresponding to a “lazy evaluation” tactic for constraint maintenance. As long as no violations can be proved, the state is considered legal. The derived tuples not present in the state can be generated on demand, for purposes such as query answering. Note the similarity of this policy to the “deductive databases” approach [11], where any fact deducible from the stored relations is considered part of the database. Requiring both consistency and completeness corresponds to a constraint maintenance policy that guarantees that all derived tuples will be present in the database at all times. There is a storage–computation trade-off in the choice of a policy. This trade-off applies not only to multirelation databases but also to single relations. Consistency of a relation under a set of, say, FDs and MVDs, is strictly weaker than standard satisfaction.

The combination of our notions of satisfaction with the concept of independence leads to interesting questions. For example, what are the database schemes such that every locally consistent state is consistent and complete? Chan and Mendelzon [7] have characterized these schemes when the join dependency for the database scheme and a set of functional dependencies are given.

REFERENCES

[Note: References [2], [3], [9], [12] are not cited in text.]

1. AHO, A. V., BEERI, C., AND ULLMAN, J. D. The theory of joins in relational databases. *ACM Trans. Database Syst.* 4, 3 (Sept. 1979), 297–314.
2. AHO, A. V., SAGIV, Y., AND ULLMAN, J. D. Equivalence among relational expressions. *SIAM J. Comput.* 8, (2) (1979), 218–246.
3. BEERI, C., AND RISSANEN, J. Faithful representations of relational database schemes. Res. Rep. RJ 2722, IBM, 1980.
4. BEERI, C., AND VARDI, M. Y. On the complexity of testing implications of data dependencies. Dept. Comput. Sci., The Hebrew Univ., Jerusalem, Israel, 1980.
5. BEERI, C., AND VARDI, M. Y. The implication problem for data dependencies. In *Proceedings of the 8th International Colloquium on Automata, Languages, and Programming*. Lecture Notes in Computer Science, vol. 115. Springer-Verlag, New York, 1981, pp. 73–85.
6. BEERI, C., AND VARDI, M. Y. A proof procedure for data dependencies. *J. ACM* 31 (1984), 718–741.
7. CHAN, E. P. F., AND MENDELZON, A. O. Independent and separable database schemes. In *Proceedings of the 2nd ACM Symposium on Principles of Database Systems* (Atlanta, Ga., Mar. 21–23). ACM, New York, 1983, pp. 288–296.
8. CHANDRA, A. K., LEWIS, H. R., AND MAKOWSKY, J. A. Embedded implicational dependencies and their inference problem. In *Proceedings of the 13th Annual ACM Symposium on Theory of Computing* (Milwaukee, Wisc., May 11–13). ACM, New York, 1981, pp. 342–354.
9. CODD, E. F. Further normalization of the data base relational model. In *Data Base Systems*, R. Rustin, Ed. Prentice-Hall, N. J., 1972, pp. 33–64.
10. FAGIN, R. Horn clauses and database dependencies. *J. ACM* 29, 4 (Oct. 1982), 952–985.
11. GALLAIRE, H., AND MINKER, J., Eds. *Logic and Data Bases*. Plenum Press, New York, 1978.
12. GAREY, M. R., AND JOHNSON, D. S. *Computers and Intractability*. Freeman, San Francisco, 1979.
13. GRAHAM, M. H. Satisfying database states. Rep. CSRG-137, Univ. of Toronto, Toronto, Ont., Canada, Dec. 1981.
14. GRAHAM, M. H., AND VARDI, M. Y. On the complexity and axiomatizability of consistent database states. In *Proceedings of the 3rd Annual ACM Symposium on Principles of Database Systems* (Waterloo, Ont., Canada, Apr. 2–4). ACM, New York, 1984, pp. 281–289.
15. GRAHAM, M. H., AND YANNAKAKIS, M. Independent database schemas. *J. Comput. Syst. Sci.* 28 (1984), 121–141.
16. GUREVICH, Y., AND LEWIS, H. R. The inference problem for template dependencies. In *Proceedings of the 1st ACM Symposium on Principles of Database Systems* (Los Angeles, Calif., Mar. 29–31). ACM, New York, 1982, pp. 221–229.

17. HONEYMAN, P. Testing satisfaction of functional dependencies. *J. ACM* 29, 3 (July 1982), 668–677.
18. MAIER, D., MENDELZON, A. O., SADRI, F., AND ULLMAN, J. D. Adequacy of decompositions of relational databases. In *Advances in Database Theory*, H. Gallaire, J. Minker, and J. M. Nicolas, Eds. Plenum Press, New York, 1981, pp. 101–114.
19. MAIER, D., MENDELZON, A. O., AND SAGIV, Y. Testing implications of data dependencies. *ACM Trans. Database Syst.* 4, 4 (Dec. 1979), 455–469.
20. MAIER, D., SAGIV, Y., AND YANNAKAKIS, M. On the complexity of testing implications of functional and join dependencies. *J. ACM* 28, 4 (Oct. 1981), 680–695.
21. MCKINSEY, J. C. C. The decision problem for some classes of sentences without quantifiers. *J. Symb. Logic* 8 (1943), 61–76.
22. MENDELZON, A. O. Database states and their tableaux. *ACM Trans. Database Syst.* 9, 2 (June 1984), 264–282.
23. NICOLAS, J. M. First order logic formalization for functional, multivalued and mutual dependencies. In *Proceedings of the ACM Conference on Management of Data*. ACM, New York, 1978, pp. 40–46.
24. RISSANEN, J. Theory of relations for databases—A tutorial survey. In *Proceedings of the 7th Symposium on Mathematical Foundations of Computer Science* (Poland, 1978). Lecture Notes in Computer Science, vol. 64. Springer-Verlag, New York, 1978, pp. 537–551.
25. SAGIV, Y. Can we use the universal instance assumption without using nulls? In *Proceedings of the ACM Conference on Management of Data*. ACM, New York, 1981, 108–120.
26. VARDI, M. The implication problem for data dependencies in relational databases. Ph.D. Dissertation (in Hebrew). The Hebrew University in Jerusalem, Jerusalem, Israel, 1981.
27. VARDI, M. Y. Global decision problems for relational databases. In *Proceedings of the 22nd IEEE Symposium on Foundations of Computer Science*. IEEE, New York, 1981, pp. 198–202.
28. VARDI, M. Y. The implication and finite implication problems for typed template dependencies. *J. Comput. Syst. Sci.* 28 (1984), 3–28.
29. VASSILIOU, Y. A formal treatment of imperfect information in database management. Ph.D. Dissertation. Univ. Toronto, Toronto, Ont., Canada, 1980.
30. YANNAKAKIS, M. Algorithms for acyclic database schemes. In *Proceedings of the 7th Conference on Very Large Databases* (Cannes, France, Sept. 9–11). ACM, New York, 1981, 82–94.

RECEIVED AUGUST 1983; REVISED APRIL 1985; ACCEPTED JUNE 1985