

Philip A. Bernstein
Nathan Goodman

Aiken Computation Laboratory
Harvard University
Cambridge, MA 02138

Abstract

Normalization research has concentrated on defining normal forms for database schemas and developing efficient algorithms for attaining these normal forms. It has never been *proved* that normal forms are *good*, i.e. that normal forms are beneficial to database users. This paper considers one of the earliest normal forms (Boyce Codd normal form [Cod2]) whose benefits are intuitively understood. We formalize these benefits and attempt to prove that the normal form attains them. Instead we prove the opposite: *Boyce-Codd normal form fails to meet its goals except in trivial cases*. This counterintuitive result is a consequence of the "universal relation assumption" upon which normalization theory rests. Normalization theory will remain an isolated theoretical area, divorced from database practice, until this assumption is circumvented.

1. Introduction

Normalization is a schema design process. The input is a schema (i.e., database description) that describes an application; the output is another schema that describes the same application, but in a "better" way. Normalization theory provides mathematical tools for defining schemas, for characterizing the application that a schema represents, and for specifying the form a "good" schema must satisfy. Most normalization research has focused on defining good forms for schemas (*normal forms*) and developing efficient algorithms for attaining these normal forms. However, normalization theory does not provide tools for explaining *why* normal forms are good. Presumably, normal forms help database users in some way, but these benefits have never been characterized within the formal framework of normalization theory. This paper attempts such a characterization for one particular normal form, namely Boyce Codd normal form.

The benefits of Boyce Codd normal form (BCNF) are well understood at an intuitive level. BCNF and its precursor, third normal form (3NF), were introduced by Codd to eliminate "anomalous" update* behavior of relations. BCNF accomplishes this by "placing independent relationships into independent relations." (A more precise characterization appears later.) BCNF is widely recommended in practice for this purpose [Date, Mar], and there is little doubt that BCNF is a good normal form in many practical situations.

This paper formalizes the intuitive benefits of BCNF and attempts to prove that these benefits are attained. The attempt fails. Instead we prove that *BCNF does not meet its goals, except in trivial cases*. This result is quite disturbing because it contradicts BCNF's intuitive appeal and empirically observed benefits. Either the theory is wrong, or database practitioners are making a serious error by using BCNF. Our opinion is that the theory is at fault--that normalization theory, as currently constituted, is inadequate to characterize the practical use and benefits of normal forms. The source of inadequacy can be traced to the "universal relation assumption" upon which normalization theory is based.

On the surface, this paper is an analysis of Boyce Codd normal form. Our deeper purpose is to draw attention to a missing link in normalization theory. For normalization theory to be well grounded in the database problems it intends to solve, it must be possible to use the theory to prove the benefits of normal forms. Unfortunately, it appears that normalization theory, as currently formulated, cannot do this. Until this missing link is filled in, normalization theory will be incapable of explaining what Boyce Codd normal form does.

[†]This work was supported by the National Science Foundation under Grant Number MCS-77-05314 and by the Advanced Research Projects Agency of the Department of Defense, Contract Number N00039-78-G-0020.

*The word "update" is used generically to denote insertions, deletions, or replacements (i.e., modification) of tuples in a relation.

2. Terminology

2.1 Relations and Databases

In formalizing relational databases we distinguish the time-invariant description of a relation vs. its contents at a particular moment. The description of a relation is called a *relation schema* and consists of a set of *attributes*, and a set of *functional dependencies*. $\underline{R} = \langle \text{ATTR}, F \rangle$ denotes a relation schema \underline{R} with attributes ATTR, and functional dependencies F. Figure 2.1 illustrates a relation schema. The contents of a relation is called a *relation state*, and may be visualized as a table of data. Figure 2.2 illustrates a relation state.

Relation Schema:

$$\underline{R} = \langle \{E\#, JC, D\#, M\#, CT\}, \{E\# \rightarrow JC, D\#, M\#, CT; D\# \rightarrow M\#, CT; M\# \rightarrow D\#, CT\} \rangle$$

Description of Attributes:

E# employee number
 JC job code
 D# department number
 M# employee number of manager
 CT contract type

Figure 2.1. An Example Relation Schema.

Relation Schema: as defined in Fig. 2.1

Relation: R

| E# | JC | D# | M# | CT |
|----|----|----|----|----|
| 1 | a | x | 11 | g |
| 2 | c | x | 11 | g |
| 3 | a | y | 12 | n |
| 4 | b | x | 11 | g |
| 5 | b | y | 12 | n |
| 6 | c | y | 12 | n |
| 7 | a | z | 13 | n |
| 8 | c | z | 13 | n |

Figure 2.2. An Example Relation.

Notationally, states of relation schema \underline{R} are represented by \underline{R} with possible superscript. Elements of \underline{R} (called *tuples*) are denoted r with possible superscript. If $X \subseteq \text{ATTR}$, we use $r[X]$ to denote the projection of r on the attributes in X , and $R[X] = \{r[X] \mid r \in R\}$.

A *database schema* is a set of relation schemas and is denoted $\underline{D} = \{\underline{R}_i = \langle \text{ATTR}_i, F_i \rangle \mid i = 1, \dots, n\}$.

A database state D for \underline{D} is an assignment of relation states to the relation schemas of \underline{D} . $D(\underline{R}_i)$ denotes the relation state assigned by D to \underline{R}_i , for $i = 1, \dots, n$. If D and D' are both database states for database schema \underline{D} , we define

$D \leq D'$ iff $D(\underline{R}_i) \subseteq D'(\underline{R}_i)$ for $i = 1, \dots, n$; clearly, \leq is a partial order.

2.2 Functional Dependencies

A *functional dependency* (abbr. FD) is a statement of the form $f: X \rightarrow Y$, where f is the name of the FD, and X, Y are sets of attributes. f is defined on relation schema $\underline{R} = \langle \text{ATTR}, F \rangle$ if $X, Y \subseteq \text{ATTR}$. We shall assume throughout this paper that every $f \in F$ is defined on \underline{R} . If f is defined on \underline{R} we interpret it as a predicate on states of \underline{R} ; $f(R)$ is true if for all $r, r' \in R$, whenever $r[X] = r'[X]$, $r[Y] = r'[Y]$ as well.

A relation state R of schema $\underline{R} = \langle \text{ATTR}, F \rangle$ is *consistent* if every $f \in F$ is true in R . For example, the state depicted in Fig. 2.2 is consistent. The set of all consistent states of \underline{R} is denoted $\text{DOM}(\underline{R})$.

Given that a set of FDs, F , holds in state R , one can infer other FDs that must hold in R as well. The set of all such FDs can be constructed syntactically using a system of *inference rules* [DC], see Figure 2.3. The set of all FDs derivable from F using the inference rules of Fig. 2.3 is called the *closure* of F and is denoted F^+ .

If $Y \subseteq X$ then $X \rightarrow Y$
 If $Z \subseteq W$ and $X \rightarrow Y$ then $XW \rightarrow YZ$
 If $X \rightarrow Y$ and $Y \rightarrow Z$ then $X \rightarrow Z$

Figure 2.3. Inference Rules for Functional Dependencies.

Within the confines of a single relation schema, F^+ is precisely the set of FDs *implied* by F : i.e., $f \in F^+$ iff f holds in every state of \underline{R} in which F holds [Arm]. This fact is called the *completeness property* of FD-closure [BFH].

FDs also enjoy the *uniqueness property* within a single relation schema, which states that syntactically identical FDs are semantically equivalent as well [Bern].

2.3 The Universal Relation Assumption: Schema Equivalence

Intuitively, two schemas are equivalent if they represent the same external application. For relation schemas this notion may be formalized in terms of consistent states: relation schema $\underline{R} = \langle \text{ATTR}, F \rangle$ and $\underline{R}' = \langle \text{ATTR}', F' \rangle$ are *equivalent* if $\text{DOM}(\underline{R}) = \text{DOM}(\underline{R}')$. As a consequence of completeness and uniqueness, $\text{DOM}(\underline{R}) = \text{DOM}(\underline{R}')$ iff $\text{ATTR} = \text{ATTR}'$ and $F^+ = F'^+$. Hence, when normalizing \underline{R} one can substitute F' for F whenever $F'^+ = F^+$, since the resulting schema $\langle \text{ATTR}, F' \rangle$ remains equivalent to \underline{R} .

All normalization algorithms depend on this fact. Bernstein [Bern], for example, exploits it

to substitute a "reduced" set of FDs for the ones supplied by the user. Without this substitution, many schemas cannot be normalized due to redundancies in the way FDs are expressed in the schema [Bern].

However, completeness and uniqueness do not hold in multi-relation databases generally. Consider Figure 2.4. The schemas D^1 , D^2 , and D^3 "contain" the same FDs--i.e., $F1^+ = F2^+ = F3^+$ --yet these three schemas differ qualitatively in representational capability. D^1 can only represent universities in which professors and teaching assistants must teach courses in their own department; D^2 permits teaching assistants to be drawn from other departments, but not professors; and D^3 permits complete flexibility in the assignment of teaching assistants and professors. These schemas do not represent the same external application, and hence are not equivalent, even though they contain the same attributes and FDs.

$$\begin{aligned}
 D^1 &= \{ \text{TEACH}^1 = \langle \{ \text{COURSE, PROF, TA, DEPT} \}, \\
 &\quad \{ \text{COURSE} \rightarrow \text{PROF}; \text{COURSE} \rightarrow \text{TA}; \\
 &\quad \text{COURSE} \rightarrow \text{DEPT}; \text{PROF} \rightarrow \text{DEPT}; \\
 &\quad \text{TA} \rightarrow \text{DEPT} \} \rangle \\
 D^2 &= \{ \text{TEACH}^2 = \langle \{ \text{COURSE, PROF, TA} \}, \\
 &\quad \{ \text{COURSE} \rightarrow \text{PROF}; \text{COURSE} \rightarrow \text{TA} \} \rangle, \\
 &\quad \text{FACULTY}^2 = \langle \{ \text{COURSE, PROF, DEPT} \}, \\
 &\quad \{ \text{COURSE} \rightarrow \text{DEPT}; \text{PROF} \rightarrow \text{DEPT} \} \rangle, \\
 &\quad \text{STUDENTS}^2 = \langle \{ \text{TA, DEPT} \}, \{ \text{TA} \rightarrow \text{DEPT} \} \rangle \\
 D^3 &= \{ \text{TEACH}^3 = \langle \{ \text{COURSE, PROF, TA} \}, \\
 &\quad \{ \text{COURSE} \rightarrow \text{PROF}; \text{COURSE} \rightarrow \text{TA} \} \rangle, \\
 &\quad \text{FACULTY}^3 = \langle \{ \text{PROF, DEPT} \}, \{ \text{PROF} \rightarrow \text{DEPT} \} \rangle, \\
 &\quad \text{STUDENTS}^3 = \langle \{ \text{TA, DEPT} \}, \{ \text{TA} \rightarrow \text{DEPT} \} \rangle, \\
 &\quad \text{CLASSES}^3 = \langle \{ \text{COURSE, DEPT} \}, \{ \text{COURSE} \rightarrow \text{DEPT} \} \rangle \}
 \end{aligned}$$

Figure 2.4. Completeness and Uniqueness Do Not Hold in Multi-Relation Databases.

This problem is quite serious. It means that the definition of schema equivalence used to construct a normalized schema \underline{D} is not semantically valid. To circumvent this problem the *universal relation assumption* is invoked.

Let $\underline{D} = \{ R_i = \langle \text{ATTR}_i, F_i \rangle \mid i = 1, \dots, n \}$ be a database schema and let D be a database state for \underline{D} . D satisfies the *universal relation assumption* (or is *UR-consistent*) if there exists a relation state U for schema $\underline{U} = \langle \text{ATTR}, \emptyset \rangle$ such that $U[\text{ATTR}_i] = D(R_i)$ for $i = 1, \dots, n$. D is *consistent* if it is UR-consistent, and $D(R_i) \in \text{DOM}(R_i)$ for $i = 1, \dots, n$. $\text{DOM}(\underline{D})$ denotes the set of all consistent states of \underline{D} .

By requiring that database states satisfy the UR-assumption, we can attain the desired notion of schema equivalence. Any UR-consistent state D is precisely representable by $U = \ast_{i=1}^n D(R_i)$, since $U[\text{ATTR}_i] = D(R_i)$ for $i = 1, \dots, n$. Therefore $\text{DOM}(\underline{D})$

is precisely representable by the set $\text{REP}(\underline{D}) = \{ U \mid \exists D \in \text{DOM}(\underline{D}), U = \ast_{i=1}^n D(R_i) \}$. Two schemas \underline{D} and \underline{D}' represent the same external application, i.e., are *equivalent under the UR-assumption* if $\text{REP}(\underline{D}) = \text{REP}(\underline{D}')$. But $\text{REP}(\underline{D}) = \text{DOM}(U = \langle \text{ATTR}, F \rangle)$ where $\text{ATTR} = (U_{i=1}^n \text{ATTR}_i)$ and $F = (U_{i=1}^n F_i)$, and $\text{REP}(\underline{D}') = \text{DOM}(U' = \langle \text{ATTR}', F' \rangle)$ where $\text{ATTR}' = (U_{i=1}^{n'} \text{ATTR}'_i)$ and $F' = (U_{i=1}^{n'} F'_i)$. Consequently $\text{REP}(\underline{D}) = \text{REP}(\underline{D}')$ iff $\text{DOM}(U) = \text{DOM}(U')$ iff $\text{ATTR} = \text{ATTR}'$ and $F^+ = F'^+$, which is precisely the notion of schema equivalence required by normalization theory.

Finally we observe that the FDs that hold in $D(R_i)$ for every $D \in \text{DOM}(\underline{D})$ are $F^+[\text{ATTR}_i]$, not simply F_i . We say that $\underline{D}(R_i)$ is *FD-consistent* relative to \underline{D} if every FD in $F^+[\text{ATTR}_i]$ is true in $D(R_i)$.

2.4 Representation

A special case of schema equivalence arises frequently in this paper. Let $U = \langle \text{ATTR}, F \rangle$ be a (universal) relation schema, and let $\underline{D} = \{ R_i = \langle \text{ATTR}_i, F_i \rangle \mid i = 1, \dots, n \}$. Using the terminology of [BBG], we say \underline{D} *Rep2-represents* U if $(U_{i=1}^n \text{ATTR}_i) = \text{ATTR}$ and $(U_{i=1}^n F_i)^+ = F^+$. Note that \underline{D} Rep2-represents U iff $\text{REP}(\underline{D}) = \text{REP}(U)$. Rep2 is essentially the notion of schema equivalence used in [Bern].

Given that \underline{D} Rep2-represents U we define mappings between their consistent states:

$$\begin{aligned}
 J: \text{DOM}(\underline{D}) &\rightarrow \text{DOM}(U), \text{ where } J: D \mapsto U = \ast_{i=1}^n D(R_i); \text{ and} \\
 P: \text{DOM}(U) &\rightarrow \text{DOM}(\underline{D}), \text{ where } P: U \mapsto D, \text{ st. } D(R_i) = \\
 &\quad U[\text{ATTR}_i] \text{ for } i = 1, \dots, n.
 \end{aligned}$$

The following facts are well known:

- (i) J and P are total functions.
- (ii) $J(P(U)) \supseteq U$.
- (iii) $P(J(D)) = D$.
- (iv) $U \subseteq U'$ implies $P(U) \subseteq P(U')$.
- (v) $D \leq D'$ iff $J(D) \subseteq J(D')$.

A stronger form of representation is used by some authors [Cod1], [Ris]. \underline{D} *Rep4-represents* U if \underline{D} Rep2-represents U and for all $U \in \text{DOM}(U)$, $J(P(U)) = U$. Rep4-representation implies that $\text{DOM}(\underline{D})$ and $\text{DOM}(U)$ are isomorphic, with J and P the respective isomorphisms. Hence every consistent state of U can be exactly reconstructed from the corresponding state of \underline{D} .

2.5 Boyce Codd Normal Form

Let $R = \langle \text{ATTR}, F \rangle$, and let $X \subseteq \text{ATTR}$. X is called a *superkey* of R if $X \rightarrow \text{ATTR} \in F^+$. X is a *key* of R if X is a superkey but does not properly contain a superkey. An FD $X \rightarrow X'$ is called *trivial*

if $X' \subseteq X$, because it holds in every state of a relation independent of F .

R is in *Boyce Codd Normal Form* (abbr. BCNF) if for all non-trivial FDs $X \rightarrow Y$ in F^+ , X is a superkey. In extending this notion to database schemas, we must be conscious of the UR-assumption. We say that $R_i = \langle \text{ATTR}_i, F_i \rangle$ is in BCNF if the schema $\langle \text{ATTR}_i, F^+[\text{ATTR}_i] \rangle$ is in BCNF, and D is in BCNF if each R_i is.

3. Goals of Boyce Codd Normal Form

Third normal form and later Boyce Codd normal form were introduced to eliminate anomalous update behavior. Let us quote from [Cod1] the motivation for these normal forms.

Looking at a sample instantaneous tabulation of R (Figure 2.2) the undesirable properties of the R schema become immediately apparent. We observe, for example, that, if the manager of department y should change, more than one tuple has to be updated. The actual number of tuples to be updated can, and usually will, change with time. A similar remark applies if department x is switched from government work (contract type g) to non-government work (contract type n).

Deletion of the tuple for an employee has two possible consequences: deletion of the corresponding department information if his tuple is the sole one remaining just prior to deletion, and non-deletion of the department information otherwise.

Inserting an employee tuple has the complementary problem: if the employee is the first member of a department, we must simultaneously insert new $M\#$ and CT information for that department; inserting the second (third, etc.) employee into that department has no such requirement.

Abstracting from the above discussion, we find R to be an undesirable schema because the precise effects of a given update operation cannot be predicted simply by examining the schema. One must examine the state of R to see the effects that a given update produces. When all of the effects of an update can be determined by examining the schema alone, we say that the update is *syntactically predictable* (more precise definitions will appear later).

In state R for example, replacing department y 's manager by another value implies updating a (syntactically) unpredictable number of tuples. Inserting an employee into (a new) department v requires creating new department-related data, while inserting an employee into (an existing) department x does not have this effect.

In this section we formalize this motivation, and prove that BCNF *relation schemas* attain it. In Sections 4 and 5 we will consider multi-relation database schemas.

3.1 Update Operations

Since the goal of normalization is to improve update behavior, our first step is to define update operations.

Let $R = \langle \text{ATTR}, F \rangle$ be a schema. The update operations we define on R are $+$ (insertion), $-$ (deletion), and $\&$ (replacement).

If $R \in \text{DOM}(R)$, then

$$+(R, r) = \begin{cases} R \cup \{r\}, & \text{if consistent} \\ R & , \text{ otherwise} \end{cases}$$

$$-(R, r) = R - \{r\}$$

$$\&(R, r, r') = \begin{cases} (R - \{r\}) \cup \{r'\}, & \text{if consistent} \\ R & , \text{ otherwise.} \end{cases}$$

3.2 Insertion and Deletion Anomalies

Intuitively, R exhibits an *insertion anomaly* if for some states R and tuples r , the operation $+(R, r)$ affects one collection of FDs, while for other states and tuples the insertion affects a different collection of FDs. The goal of normalization (with respect to insertions) is to eliminate such anomalies.

In formalizing this notion a technical problem surfaces. If $R = \emptyset$, the insertion of any tuple must affect *every* FD, while if $R \neq \emptyset$, the insertion of any tuple already in R does not affect *any* FD. Consequently, every schema exhibits insertion anomalies in this trivial sense (unless $\text{ATTR} = \emptyset$). A similar problem arises with deletions: if $R = \{r\}$, the deletion of r must affect every FD, while for any R , the deletion of a tuple not in R cannot affect any FD.

A second problem arises with respect to *trivial* FDs. Since every attribute functionally determines itself, the "collection of FDs" affected by an update will in general vary by these trivial FDs at least. We therefore choose to discard trivial FDs in this context.

Formally, we say that $+(R, r)$ *affects* $f: X \rightarrow Y$ if $R[XUY] \neq (+(R, r))[XUY]$, and we define

$$\text{Affect}(+R) = \{f \in F^+ \mid f \text{ is non-trivial, and for some } R \neq \emptyset, \text{ and some } r, +(R, r) \text{ affects } f\}$$

$$\text{NoAffect}(+R) = \{f \in F^+ \mid f \text{ is non-trivial, and for some } R \neq \emptyset, \text{ and some } r, +(R, r) \neq R, \text{ yet } +(R, r) \text{ does not affect } f\}$$

R is free of *insertion anomalies* iff $\text{Affect}(+R) \cap \text{NoAffect}(+R) = \emptyset$. The definition of *deletion anomalies* is similar and appears in Figure 3.1.

$Affect(-R) = \{f \in F^+ \mid f \text{ is non-trivial, and for some } R \text{ and } r, R \neq \{r\}, \text{ yet } -(R,r) \text{ affects } f\}.$

$NoAffect(-R) = \{f \in F^+ \mid f \text{ is non-trivial, and for some } R \text{ and } r, R \neq \{r\} \text{ and } -(R,r) \neq R, \text{ yet } -(R,r) \text{ does not affect } f\}.$

R is free of *deletion anomalies* iff $Affect(-R) \cap NoAffect(-R) = \emptyset.$

Figure 3.1. Definition of Deletion Anomaly.

It is conjectured in [Cod2] that BCNF schemas are the only ones that can avoid insertion and deletion anomalies. Now that we have formalized these anomalies, we can prove this conjecture true.

Theorem 1. (i) R is free of insertion anomalies iff R is in BCNF.

(ii) R is free of deletion anomalies iff R is in BCNF.

Proof. (i) $Affect(+R) = \{\text{non-trivial FDs in } F^+\}$ whether or not R is in BCNF. So it suffices to prove that $NoAffect(+R) = \emptyset$ iff R is in BCNF.

if: Let $f: X \rightarrow Y$ be any non-trivial FD in $F^+.$ We will construct an insertion that affects $f.$ By definition of BCNF, X is a superkey. Hence, all tuples in any consistent state of R have different X -projections. Let R' be any consistent state containing two or more tuples; such a state must exist for any $R,$ e.g., the state $\{<1, \dots, 1>, <2, \dots, 2>\}$ is such a state. Let $r \in R',$ and let $R = R' - \{r\}.$ $R \neq \emptyset, + (R,r) = R' \neq R,$ yet $+ (R,r)$ affects $f.$ Consequently $NoAffect(+R) = \emptyset$ as desired.

only if: We prove that if R is not in BCNF, then $NoAffect(+R) \neq \emptyset.$ If R is not in BCNF, then there exists a non-trivial FD $f: X \rightarrow Y$ in F^+ for which X is not a superkey. We will construct an insertion that does not affect $f.$

Let R' be any consistent state containing two (or more) tuples r and r' such that $r[X] = r'[X].$ Such a state must exist by the completeness of $F^+:$ since $X \rightarrow ATTR \notin F^+,$ there must exist a state R' in which F^+ holds, but $X \rightarrow ATTR$ does not hold [Theorem 3, BFH]; for $X \rightarrow ATTR$ not to hold in $R',$ distinct tuples r and r' must be present in R' with $r[X] = r'[X].$

Let $R = R' - \{r\}.$ $R \neq \emptyset, + (R,r) = R' \neq R,$ yet $+ (R,r)$ does not affect $f.$ Hence, $NoAffect(+R) \neq \emptyset,$ as desired.

(ii) The result for deletions follows by similar proof. \square

3.3 Replacement Anomalies

Replacement anomalies are concerned with non-predictability of a different nature. The replacement depicted in the Introduction to Section 3 sought to change the manager of department $y,$

leaving all other values unchanged. This replacement was judged to be anomalous because an unpredictable number of tuples might have to be replaced to achieve this effect. If we scrutinize this motivation, however, difficulties emerge.

The operation used in the example is *not* a single-tuple replacement of the type defined in Section 3.1. Instead, it attempts to modify a set of tuples that satisfy a boolean qualification, " $D\# = y$ " in this case. Arbitrarily many tuples may satisfy an arbitrary qualification. So, nonpredictability of this sort can hardly be called "anomalous". Apparently, Codd did not have such general operations in mind.

The operation that seems to be intended is the following. Let $R = \langle ATTR, F \rangle,$ let $R \in DOM(R),$ and let $f: X \rightarrow Y$ be a non-trivial FD in $F^+.$ For any value $x,$ and for each r in $\{r \in R \mid r[X] = x\}$ the user may replace $r[Y]$ with a new value. The replacement has *predictable size* if $|\{r \in R \mid r[X] = x\}| \leq 1.$ R is free of *replacement anomalies* if for all non-trivial $f: X \rightarrow Y$ in $F^+,$ all $R \in DOM(R),$ and all $x,$ the replacement has predictable size.

Theorem 2. R is free of replacement anomalies iff R is in BCNF.

Proof. All replacements have predictable size iff for all non-trivial $f: X \rightarrow Y$ in F^+, X is a superkey. \square

It is not apparent to us that multi-tuple replacements of this type warrant special consideration in the schema design process. Replacement anomalies will not be studied further in this paper.

4. Update Operations Under UR-Assumption

In the following two sections we attempt to extend Theorem 1 to multi-relation database schemas. This attempt will *fail,* because of the UR-assumption. Instead, we shall prove that BCNF database schemas are not free of insertion and deletion anomalies except in trivial cases.

4.1 Problems in Preserving UR-Consistency

Recall the example of Section 3. R is not in BCNF and therefore exhibits update anomalies; the remedy recommended in [Cod1] is to decompose R into two relation schemas, R_1 and $R_2,$ which are in BCNF (see Fig. 4.1). Notice that the database schema $D = \{R_1, R_2\}$ Rep4-represents $R;$ therefore D and R have equivalent representational power, provided the database system only permits UR-consistent states of D to occur.

This leaves us the problem of designing update operations that preserve UR-consistency. For example, to insert $\langle v, 14, n \rangle$ into $R_2,$ a tuple r_1 with $r_1[D\#] = v$ must be inserted into $R_1.$ Similarly, to delete $\langle x, 11, g \rangle$ from $R_2,$ all tuples r_1 in R_1 with $r_1[D\#] = x$ must be

deleted too. In these two cases, the requirement of UR-consistency dictates a "natural" semantics for the update. However, in other cases, the choice of a reasonable semantics is less constrained.

For example, suppose we want to insert $\langle 9, d, v \rangle$ into R_1 . One way to preserve UR-consistency is to insert $\langle 9, d, v \rangle$ into R_1 and $\langle v, \text{null}, \text{null} \rangle$ into R_2 , where "null" denotes a blank or uncommitted value that is distinguishable from "real" values, such as 11. To adopt this interpretation, we must develop a complete semantics for update operations when null values are present. We attempted such a development (see Appendix I) but discovered that for each semantics that we tried, certain bizarre behavior was forced by the UR-assumption. For now, we are defeated in taking this route.

The choices that remain are to use D# values already in R_2 or to invent new ones. Consider the first choice. To insert $\langle 9, d, v \rangle$ into R_1 , we can preserve UR-consistency by also replacing $\langle x, 11, g \rangle$ by $\langle v, 11, g \rangle$ in R_2 , and replacing each instance of 'x' in R_1 by 'v'. There are several problems with this semantics. First, the choice of $\langle x, 11, g \rangle$ as the victim to be replaced was arbitrary; $\langle y, 12, n \rangle$ or $\langle z, 13, n \rangle$ would have worked as well. Second, the interpretation fails when R_2 is empty, because there are no tuples in R_2 to replace. Third, although we have not inserted new D# values, we have produced new relationships both in R_1 and R_2 that were not part of the insertion request. Finally, this approach effectively asserts that the tuple being inserted is "more reliable" than data already in the database, since we have chosen to modify existing relationships instead of rejecting the insertion as an integrity violation. For these reasons we have rejected this "replacement" approach to insertion semantics.

The only remaining choice is to insert a "real" valued tuple into R_2 , such as $\langle v, 14, g \rangle$. This approach, like the preceding one, arbitrarily selects values to be placed into the database, and produces a relationship that was not part of the insertion request. However, it has desirable properties not shared by the replacement interpretation. First, it succeeds even when R_2 is empty. Second, the effect on R_1 of "insert $\langle 9, d, v \rangle$ into R_1 " is to insert that tuple into that relation, and to make no other changes to R_1 . Third, an insertion never has the effect of deleting a value or relationship; this is consonant with the intuitive understanding that an insertion creates data and does not destroy it. For these reasons, we have selected this interpretation of insertion in the UR-environment.

In Section 4.2 we formalize this semantics and generalize it to deletions and replacements. Basic properties of these operations are investigated in Appendix II.

4.2 Update Operations that Preserve UR-Consistency

Let $\underline{D} = \{R_1, \dots, R_n\}$ be a database schema. The update operations defined on \underline{D} are named $+R_i$, $-R_i$, and $\&R_i$ for $i=1, \dots, n$. (We include the schema name, R_i , as part of the operation name for notational clarity.) If $D \in \text{DOM}(\underline{D})$, then

$+R_i(D, r_i)$ denotes insertion of r_i into $D(R_i)$,

$-R_i(D, r_i)$ denotes deletion of r_i from $D(R_i)$,

$\&R_i(D, r_i, r'_i)$ denotes replacement of r_i by r'_i in $D(R_i)$.

As indicated in Section 4.1, each operation may have to update other relations to preserve UR-consistency; to formalize the semantics of these operations we must also specify these other effects. As we have illustrated there is no unique, "natural" semantics. Instead we choose to base the semantics of these operations on properties that we judge to be desirable for update operations in general. These properties are closely related to the correctness criteria for "view updates" postulated in [DB]. We state the properties for insertions; the definitions for deletions and replacements are analogous.

Property 1--Consistency

The result of an update operation must be a consistent state.

Property 2--Exact Performance

Given $+R_i(D, r_i)$, let $R'_i = +(D(R_i), r_i)$. State D' exactly performs $+R_i(D, r_i)$ if $D'(R_i) = R'_i$.

Property 3--Intent Performance

State D' performs the intent of $+R_i(D, r_i)$ if D' can be obtained from D by applying insertion operations to individual relations. (This property is stated more precisely in Appendix II.)

If D' satisfies Properties 1-3 for an update operation, we say that D' performs that update. In general, many states may perform an update. Our final property provides a metric choosing certain of these states over others.

Property 4--Nonextraneous Performance

Suppose D' performs $+R_i(D, r_i)$ and let Ω_+ be a sequence of insertions that maps D into D' . Since every operation in Ω_+ (except for the one that inserts r_i into R_i) causes a change to the database not specified by the user, all unnecessary operations should be eliminated. D' nonextraneously performs $+R_i(D, r_i)$ if (a) D' performs the insertion, and (b) for all subsets* Ω'_+ of Ω_+ , Ω'_+ either maps D into D' , or it maps D

* Technically, we mean "projection" here. In Appendix II, however, we demonstrate that Ω_+ can be treated as a set with no loss of generality; with this understanding, the word "subset" is correct.

into a state that does not perform the insertion. (A more precise statement is given in Appendix II.)

The semantics of $+R_i$, $-R_i$, and IR_i can now be defined. Let $D \in \text{DOM}(\underline{D})$. $+R_i(D, r_i) = \text{any}$ state that nonextraneously performs $+R_i(D, r_i)$. $-R_i$ and $\&R_i$ have analogous definitions. In general, $+R_i(D, r_i)$ and $\&R_i(D, r_i, r_i')$ have many possible meanings and we treat these operations as non-deterministic functions. By contrast, $-R_i(D, r_i)$ is uniquely specified by this definition, a fact we prove in Appendix II.

Database Schema:

$$\underline{D} = \{R_1 = \langle \{E\#, JC, D\# \}, \{E\# \rightarrow JC, D\#\} \rangle, \\ R_2 = \langle \{D\#, M\#, CT\}, \{D\# \rightarrow M\#, CT, M\# \rightarrow D\#, CT\} \rangle\}$$

Database:

| E# | JC | D# |
|----|----|----|
| 1 | a | x |
| 2 | c | x |
| 3 | a | y |
| 4 | b | x |
| 5 | b | y |
| 6 | c | y |
| 7 | a | z |
| 8 | c | z |

$$R_1 =$$

| D# | M# | CT |
|----|----|----|
| x | 11 | g |
| y | 12 | n |
| z | 13 | n |

$$R_2 =$$

R_1 and R_2 are in BCNF and \underline{D} Rep4-represents \underline{R} of Fig. 2.1.

Figure 4.1. A Normalized Database Schema.

5. BCNF Does Very Little

Having developed an update semantics that preserves UR-consistency, we are ready to study normalization in this context. There are two issues: (1) characterizing database schemas that avoid insertion and deletion anomalies; and (2) the schema design question--characterizing relation schemas \underline{U} for which there exists a database schema \underline{D} that represents \underline{U} , yet avoids insertion and deletion anomalies.

5.1 Database Schemas that Avoid Anomalies

Throughout this section let $\underline{D} = \{R_1, \dots, R_n\}$, and $\underline{U} = \langle \text{ATTR}, F \rangle$, such that \underline{D} Rep2-represents \underline{U} . If \underline{D} is to avoid update anomalies it is necessary for \underline{D} to be in BCNF: otherwise the "internal effects" of updates would exhibit anomalies as proved in Section 3. However, we shall prove that BCNF is not sufficient to prevent anomalies caused by the "external effects" of these operations.

The definitions of insertion and deletion anomalies require technical changes in the UR-environment. First, the definitions of Affect and NoAffect given in Section 3 quantify over all

consistent states of a schema, but it is technically better to quantify over *reachable* states only; in light of Theorem II.1 (see Appendix II), though, this distinction has no impact. Second, since we may assume that \underline{D} is in BCNF, we may interpret Affect and NoAffect as sets of *relation schemas* instead of FDs: e.g., for any $R_i, R_j \in \underline{D}$, $R_j \in \text{Affect}(+R_i)$ iff every non-trivial FD in $F^+[\text{ATTR}_j] \in \text{Affect}(+R_i)$, because all external effects of $+R_i$ are themselves insertions. Figure 5.1 presents the definitions of Affect, NoAffect, and anomalies as used in this section.

$\text{Affect}(+R_i) = \{R_j \in \underline{D} \mid \text{for some } D \text{ other than the empty state, for some tuple } r_i, \text{ and for some meaning, } +R_i(D, r_i) \text{ affects } R_j\}$

$\text{NoAffect}(+R_i) = \{R_j \in \underline{D} \mid \text{for some } D \text{ other than the empty state, for some tuple } r_i, \text{ and for some meaning, } +R_i(D, r_i) \neq D, \text{ yet } +R_i(D, r_i) \text{ does not affect } R_j\}$

$\text{Affect}(-R_i) = \{R_j \in \underline{D} \mid \text{for some } D \text{ and } r_i, D(R_i) \neq \{r_i\}, \text{ yet } -R_i(D, r_i) \text{ affects } R_j\}$

$\text{NoAffect}(-R_i) = \{R_j \in \underline{D} \mid \text{for some } D \text{ and } r_i, D(R_i) \neq \{r_i\}, \text{ and } -R_i(D, r_i) \neq D, \text{ yet } -R_i(D, r_i) \text{ does not affect } R_j\}$.

\underline{D} is free of *insertion anomalies* iff for all $R_i \in \underline{D}$, $\text{Affect}(+R_i) \cap \text{NoAffect}(+R_i) = \emptyset$

\underline{D} is free of *deletion anomalies* iff for all $R_i \in \underline{D}$, $\text{Affect}(-R_i) \cap \text{NoAffect}(-R_i) = \emptyset$

Figure 5.1. Definitions of Anomalies for UR Environment.

A graphic representation of \underline{D} helps characterize the effects of insertions. Let $G(\underline{D})$ be an undirected graph whose *vertex set* equals \underline{D} , and whose *edge set* contains (R_i, R_j) iff $\text{ATTR}_i \cap \text{ATTR}_j \neq \emptyset$.

Lemma 1. (i) $\text{Affect}(+R_i) = \{R_j \in \underline{D} \mid R_i \text{ and } R_j \text{ are connected by a path in } G(\underline{D})\}$.

(ii) $\text{NoAffect}(+R_i) = \{R_j \in \underline{D} \mid \text{ATTR}_j \rightarrow \text{ATTR}_i \in F^+\}$.

Proof. See Appendix II.4. □

Lemma 2. (i) $\text{Affect}(-R_i) = \text{Affect}(+R_i)$

(ii) $\text{NoAffect}(-R_i) = \text{NoAffect}(+R_i)$.

Proof. See Appendix II.4. □

Combining these results we obtain a precise characterization of database schemas that are free of insertion and deletion anomalies under the UR-assumption.

Theorem 3. *D is free of insertion anomalies (or equivalently deletion anomalies) iff D is in BCNF, and for all R_i, R_j , if $ATTR_i \cap ATTR_j \neq \emptyset$ then $ATTR_i \rightarrow ATTR_j \in F^+$ and $ATTR_j \rightarrow ATTR_i \in F^+$.*

The conditions placed on \underline{D} by this result are quite restrictive; for example, the normalized schema suggested in Figure 4.1 does not satisfy them, nor for that matter do most of the normalized schemas commonly illustrated in the literature. This bodes poorly for the schema design aspect of normalization.

5.2 Attainability of Normalized Schemas

The remaining issue is one of schema design. Let $\underline{U} = \langle ATTR, F \rangle$ be any relation schema; we prove that \underline{U} can be normalized into a schema \underline{D} that is free of insertion and deletion anomalies iff \underline{U} is "almost" in BCNF itself.

A database schema $\underline{D}_p = \{D_{pi} = \langle ATTR_{pi}, F_{pi} \rangle \mid i = 1, \dots, n\}$ is said to partition \underline{U} if D_{pi} Rep2-represents \underline{U} and $ATTR_{pi} \cap ATTR_{pj} = \emptyset$ for $i \neq j$.

Theorem 4. *There exists a database schema \underline{D} that Rep2-represents \underline{U} yet is free of insertion and deletion anomalies iff there exists a BCNF database schema \underline{D}_p that partitions \underline{U} .*

Proof. if: If \underline{D}_p exists, it Rep2-represents \underline{U} and satisfies the conditions of Theorem 3.

only if: Let P_1, \dots, P_n be the connected components of $G(\underline{D})$, and let $\{R_{k1}, \dots, R_{km_k}\}$ be the schemas in component P_k . The schemas in each component are "equivalent" in the sense that $\forall R_{ki}, R_{kj}$ in P_k , $ATTR_{ki} \rightarrow ATTR_{kj} \in F^+$ and $ATTR_{kj} \rightarrow ATTR_{ki} \in F^+$. Also each schema is in BCNF. Therefore, if we merge all schemas in P_k into one schema $R_{-pk} = \langle \bigcup_{i=1}^{m_k} ATTR_{ki}, \bigcup_{i=1}^{m_k} F_{ki} \rangle$, R_{-pk} is in BCNF also. The database schema $\underline{D}_p = \{R_{-pk} \mid k=1, \dots, m\}$ is thus a BCNF schema which partitions \underline{U} . \square

Figure 5.2 illustrates two typical relation schemas that are BCNF-partitionable, and database schemas that Rep2-represent them while satisfying the conditions of Theorem 3. In both cases the database schemas do not Rep4-represent the original schema.

If Rep4-representation is required, \underline{U} can be "normalized" if and only if it already is in BCNF!

Corollary to Theorem 4. *There exists a database schema \underline{D} that Rep4-represents \underline{U} yet is free of insertion and deletion anomalies iff \underline{U} is in BCNF.*

- (a) $\underline{U} = \langle \{E\#, D\#, M\#, PROJ\}, \{E\# \rightarrow D\#, M\# \rightarrow PROJ\} \rangle$
 $\underline{D} = \{R_1 = \langle \{E\#, D\#\}, \{E\# \rightarrow D\#\} \rangle,$
 $R_2 = \langle \{M\#, PROJ\}, \{M\# \rightarrow PROJ\} \rangle\}$
- (b) $\underline{U} = \langle \{SUPPLIER, PART, CITY\}, \{SUPPLIER \rightarrow CITY\} \rangle$
 $\underline{D} = \{R_1 = \langle \{SUPPLIER, CITY\}, \{SUPPLIER \rightarrow CITY\} \rangle,$
 $R_2 = \langle \{PART\}, \{ \} \rangle\}$

Figure 5.2. Schemas That Can Be Normalized.

Proof. If \underline{U} is in BCNF there is nothing to prove. To prove the converse, observe that if \underline{D} Rep4-represents \underline{U} , the J operator must denote a lossless join $[ABU]$, hence $G(\underline{D})$ must be a connected graph. If we merge all relation schemas as in the proof of Theorem 4, the result is a BCNF relation schema $\underline{R} = \langle ATTR, F' \rangle$ where $F' = F^+$. Hence $\underline{U} = \langle ATTR, F \rangle$ is in BCNF as well. \square

6. Conclusion

BCNF was invented to prevent anomalous side effects of relational updates. The question we have asked is, "does BCNF attain this goal?" We have given two conflicting answers: (1) In the context of a single relation schema, BCNF is successful. But (2) in a multirelation database schema, BCNF fails.

In spite of (2), we believe BCNF to be an important schema design goal.

The failure of BCNF in the multirelation context is caused by the universal relation assumption. The UR-assumption is adopted in normalization theory to formalize notions of FD completeness and uniqueness, and schema equivalence. Recent work on normal forms, e.g., fourth and fifth normal forms [Fagl,2], leans even more heavily upon the UR-assumption; these normal forms cannot even be defined without the UR-assumption. Recent work on schema equivalence [BMSU] depends strongly on this assumption as well.

Yet update operations are clumsy to define when the UR-assumption is present, and behave quite badly. In addition to the problems noted in this paper, the UR-assumption introduces complexity problems as well, since testing whether an insertion preserves the UR-assumption is NP-complete [HLY]. Apparently, the UR-assumption is incompatible with databases that are updated.

At present, normalization theory is an isolated theoretical area divorced from database practice. This separation will persist until normalization theory is made adequate to prove the benefits of normal forms.

APPENDIX I

Pitfalls in Null-Value Update Semantics

Let us return to the example database of Section 4.1. The schema is

$$D = \{R_1 = \langle \{E\#, JC, D\#\}, \{E\# \rightarrow JC, D\#\} \rangle, \\ R_2 = \langle \{D\#, M\#, CT\}, \{D\# \rightarrow M\#, CT; M\# \rightarrow D\#, CT\} \rangle\}$$

and the database state is

$$D:R_1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ 2 & c & x \\ 3 & a & y \\ 4 & b & x \\ 5 & b & y \\ 6 & c & y \\ 7 & a & z \\ 8 & c & z \\ \hline \end{array}, \quad R_2 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ y & 12 & n \\ z & 13 & n \\ \hline \end{array}$$

If the user says to insert $\langle 9, d, v \rangle$ into R_1 , the most plausible way to preserve UR-consistency is to simultaneously insert $\langle v, \text{null}, \text{null} \rangle$ into R_2 . All other interpretations, we saw in 4.1, share the dubious property of inventing data values and/or relationships among data values in the database.

In this appendix we explore the consequences of using null values to preserve UR-consistency.

I.1 The Second-Insertion Problem

The database state that results from the preceding insertion is

$$D^1:R_1^1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ 8 & c & z \\ 9 & d & v \\ \hline \end{array}, \quad R_2^1 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ y & 12 & n \\ z & 13 & n \\ v & \text{null} & \text{null} \\ \hline \end{array}$$

Suppose the user now says to insert $\langle v, 14, g \rangle$ into R_2 . We cannot simply add $\langle v, 14, g \rangle$ to R_2^1 , because the resulting relation state would be inconsistent: the FD $D\# \rightarrow M\#, CT$ would be violated. To overcome this inconsistency, the natural interpretation is to replace $\langle v, \text{null}, \text{null} \rangle$ by $\langle v, 14, g \rangle$. The rationale for this interpretation is that $\langle v, 14, g \rangle$ signifies "more complete" information than $\langle v, \text{null}, \text{null} \rangle$, hence this new tuple makes the old one obsolete.

Applying this rationale to another case, consider the insertion of the same tuple in the state

$$D^2:R_1^2 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ 9 & d & \text{null} \\ \hline \end{array}, \quad R_2^2 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \text{null} & \text{null} & g \\ \hline \end{array}$$

since $\langle v, 14, g \rangle$ is "more complete" than $\langle \text{null}, \text{null}, g \rangle$ the correct action apparently is to replace $\langle \text{null}, \text{null}, g \rangle$ by $\langle v, 14, g \rangle$. But then UR-consistency is violated, since $\langle 9, d, \text{null} \rangle$ in R_1^2 no longer "matches up" with any tuple in R_2^2 , and further correction is needed. The correction would seem to be the replacement of $\langle 9, d, \text{null} \rangle$ by $\langle 9, d, v \rangle$. The resulting state is

$$D^3:R_1^3 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ 9 & d & v \\ \hline \end{array}, \quad R_2^3 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ v & 14 & g \\ \hline \end{array}$$

D^3 asserts that employee 9 works in department 'v', but this relationship is in no sense implied by the user level insertion operation. The database system has *invented* a relationship between employee 9 and department 'v', which is precisely the problem we thought null values would prevent!

This problem is not caused by the use of null values per se; it is caused by the interpretation of insertion operations as replacements when null values are present. The alternative, of course, is to treat insertions as insertions whether or not nulls are present. However, this interpretation begs the issue, because it assigns the same semantics to null values as real ones.

I.2 What Do Deletions Mean?

Let us return to state D and consider possible meanings of the delete operation. Suppose the user says to delete $\langle z, 13, n \rangle$ from R_2 . Two plausible meanings for this are

1. replace $\langle z, 13, n \rangle$ by $\langle z, \text{null}, \text{null} \rangle$;
2. delete $\langle z, 13, n \rangle$ from R_2 and delete employee in department 'z' from R_1 .

Meaning (1) is the inverse of the insertion semantics suggested in I.1. The rationale for this interpretation is that deletions are too drastic for the system to undertake automatically--the database system should never delete more data than the user specifies. Meaning (1) supports this reasoning while meaning (2) deletes far more data than the user specified.

But consider what happens if we add relation schemas

$$R_3 = \langle \{CT, VP\}, \{CT \rightarrow VP\} \rangle$$

and

$$R_4 = \langle \{PROJ, M\#\}, \{PROJ \rightarrow M\#\} \rangle$$

to the database with states

$$R_3 = \begin{array}{|c|c|} \hline CT & VP \\ \hline g & MONDALE \\ n & FORD \\ \hline \end{array}, \quad \text{and} \quad R_4 = \begin{array}{|c|c|} \hline PROJ & M\# \\ \hline p1 & 11 \\ p2 & 12 \\ p3 & 13 \\ \hline \end{array}$$

If we adopt meaning (1) now, the deletion of $\langle z,13,n \rangle$ has no effect whatever! Given this semantics, the following procedure is required for a user to actually remove $\langle z,13,n \rangle$ from the database:

1. delete $\langle p3,13 \rangle$ from R_4 , and delete every employee from R_1 with $D\# = 'z'$. The resulting state is

$$D^1:R_1^1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 6 & c & y \\ \hline null & null & z \\ \hline \end{array}, R_2^1 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline y & 12 & n \\ \hline z & 13 & n \\ \hline \end{array}$$

$$R_3^1 = \begin{array}{|c|c|} \hline CT & VP \\ \hline g & MONDALE \\ \hline n & FORD \\ \hline \end{array}, R_4^1 = \begin{array}{|c|c|} \hline PROJ & M\# \\ \hline p1 & 11 \\ \hline p2 & 12 \\ \hline null & 13 \\ \hline \end{array}$$

2. delete $\langle z,13,n \rangle$ from R_2 ; to implement this deletion the database system must delete $\langle null,null,z \rangle$ from R_1^1 , $\langle z,13,n \rangle$ from R_2^1 , and $\langle null,13 \rangle$ from R_4^1 . While in a strict sense this implementation deletes more data than the user specified, we judge these extra deletions to be acceptable since all unspecified data values are nulls.

As an exercise the reader may devise a similar procedure for the following case:

$$\underline{D} = \{R_1 = \langle \{A,B\}, \{A \rightarrow B\} \rangle, R_2 = \langle \{B,C\}, \{B \rightarrow C\} \rangle, R_3 = \langle \{C,A\}, \{C \rightarrow A\} \rangle\}$$

$$D: R_1 = \begin{array}{|c|c|} \hline A & B \\ \hline a_1 & b_1 \\ \hline a_2 & b_2 \\ \hline \end{array}, R_2 = \begin{array}{|c|c|} \hline B & C \\ \hline b_1 & c_1 \\ \hline b_2 & c_2 \\ \hline \end{array}, R_3 = \begin{array}{|c|c|} \hline C & A \\ \hline c_1 & a_1 \\ \hline c_2 & a_2 \\ \hline \end{array}$$

operation: delete $\langle a_1, b_1 \rangle$ from R_1 .

I.3 Another Interpretation of Delete

Let us return to the original schema \underline{D} and its original state

$$D: R_1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 7 & a & z \\ \hline 8 & c & z \\ \hline \end{array}, R_2 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline y & 12 & n \\ \hline z & 13 & n \\ \hline \end{array}$$

Another plausible meaning of "delete $\langle z,13,n \rangle$ from R_2 " is

Replace $\langle z,13,n \rangle$ by $\langle null,null,null \rangle$ in R_2 , and replace the department of every department 'z' employee by 'null' in R_1 .

The resulting state is

$$D^1:R_1^1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 7 & a & null \\ \hline 8 & c & null \\ \hline \end{array}, R_2^1 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline y & 12 & n \\ \hline null & null & null \\ \hline \end{array}$$

This interpretation eliminates the problems noted in I.2, but its interaction with the insertion semantics of I.1 is strange. Suppose the user says to insert $\langle v,14,g \rangle$ into R_2 . Since $\langle v,14,g \rangle$ is "more complete" than $\langle null,null,null \rangle$, the effect (according to I.1), is to replace $\langle null,null,null \rangle$ by $\langle v,14,g \rangle$. This replacement violates UR-consistency, however, and the corrective action is to replace every 'null' department in R_1^1 by 'v'. This yields

$$D^2:R_1^2 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 7 & a & v \\ \hline 8 & c & v \\ \hline \end{array}, R_2^2 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline y & 12 & n \\ \hline v & 14 & g \\ \hline \end{array}$$

Observe that the deletion of $\langle z,13,n \rangle$ and the insertion of $\langle v,14,g \rangle$ have been *coupled* into a replacement, whether or not this was intended!

I.4 Minimizing the Use of Null Values--A Problem

Additional difficulties surround the question of how many nulls should be inserted when correcting UR violations. Let us agument schema \underline{D} with relation schema $\underline{R}_3 = \langle \{CT,VP\}, \{CT \rightarrow VP\} \rangle$. And consider the state

$$D^1:R_1^1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 9 & d & v \\ \hline \end{array}, R_2^1 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline y & 12 & n \\ \hline z & 13 & n \\ \hline v & null & null \\ \hline \end{array}$$

$$R_3^1 = \begin{array}{|c|c|} \hline CT & VP \\ \hline g & MONDALE \\ \hline n & FORD \\ \hline null & null \\ \hline \end{array}$$

If the user says to insert $\langle 10,d,u \rangle$ into R_1 either of the following interpretations is possible:

1. insert $\langle 10,d,u \rangle$ into R_1^1 and $\langle u,null',null \rangle$ into R_2^1 ; or
2. insert $\langle 10,d,u \rangle$ into R_1^1 , $\langle u,null',null \rangle$ into R_2^1 , and $\langle null',null' \rangle$ into R_3^1 .

The advantage of (1) is that the number of relations affected is minimized. Since null values have no intrinsic value this seems a worthwhile goal. The state corresponding to (1) is

$$D^2:R_1^2 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 9 & d & v \\ \hline 10 & d & u \\ \hline \end{array}, R_2^2 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline v & null & null \\ \hline u & null & null \\ \hline \end{array},$$

$$R_3^2 = \begin{array}{|c|c|} \hline CT & VP \\ \hline g & MONDALE \\ \hline n & FORD \\ \hline null & null \\ \hline \end{array}.$$

Suppose the user subsequently determines that department 'v' has vice-president 'AGNEW'. To place this fact in the database, the usual procedure is

I. Find the R_3 tuple that describes department v's vice-president. I.e.,

- (i) restrict R_2 by $D\# = 'v'$
- (ii) joint the result of (i) with R_3
- (iii) project onto $ATTR_3 = \{CT, VP\}$.

II. Replace the VP-projection of all tuples found in (I) by 'AGNEW'.

Applying this procedure to D^2 , we obtain

$$D^3:R_1^3 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 9 & d & v \\ \hline 10 & d & u \\ \hline \end{array}, R_2^3 = \begin{array}{|c|c|c|} \hline D\# & M\# & CT \\ \hline x & 11 & g \\ \hline v & null & null \\ \hline u & null & null \\ \hline \end{array},$$

$$R_3^3 = \begin{array}{|c|c|} \hline CT & VP \\ \hline g & MONDALE \\ \hline n & FORD \\ \hline null & AGNEW \\ \hline \end{array}.$$

This state asserts, *erroneously*, that AGNEW is vice-president of department 'u'.

To avoid this error, the retrieval step (part I) of the procedure must be cognizant of null join values and must *insert new null values* as it goes. It would certainly be simpler for the insert operator to insert these extra nulls in the first place.

I.5 Maximizing the Use of Null Values--A Problem

On the other hand there are cases where maximizing the use of null values is also wrong. Let us add the following relation schemas to \underline{D} :

$$\underline{R}_4 = \langle \{A\#, D\#\}, \{A\# \rightarrow D\#\} \rangle,$$

and

$$\underline{R}_5 = \langle \{E\#, A\#\}, \{\} \rangle,$$

where $A\#$ denotes an account (in the financial sense), R_4 tells which accounts are used by each department, and R_5 identifies which employees may

post expenses against which accounts. The FD $A\# \rightarrow D\#$ is present in R_4 because the "chart of accounts" normally mirrors the organizational hierarchy; R_5 has no FDs because each employee may generally charge to multiple accounts and vice versa.

Consider the state

$$D^1:R_1^1 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 8 & c & z \\ \hline \end{array}, \dots, R_4^1 = \begin{array}{|c|c|} \hline A\# & D\# \\ \hline a1 & x \\ \hline a2 & x \\ \hline a3 & y \\ \hline a4 & z \\ \hline a5 & z \\ \hline \end{array}, R_5^1 = \begin{array}{|c|c|} \hline E\# & A\# \\ \hline 1 & a1 \\ \hline 2 & a1 \\ \hline 2 & a2 \\ \hline 3 & a3 \\ \hline 8 & a5 \\ \hline \end{array}$$

and suppose two new employees are hired into department 'z' with job-codes 'd'. If we adopt the policy suggested in I.4 of inserting *unique* null values whenever a null value is needed, the resulting state is

$$D^2:R_1^2 = \begin{array}{|c|c|c|} \hline E\# & JC & D\# \\ \hline 1 & a & x \\ \hline 8 & c & z \\ \hline 9 & d & z \\ \hline 10 & d & z \\ \hline \end{array}, \dots, R_4^2 = \begin{array}{|c|c|} \hline A\# & D\# \\ \hline a1 & x \\ \hline a2 & x \\ \hline a3 & y \\ \hline a4 & z \\ \hline null & z \\ \hline null & z \\ \hline \end{array}, R_5^2 = \begin{array}{|c|c|} \hline E\# & A\# \\ \hline 1 & a1 \\ \hline 8 & a5 \\ \hline 9 & null \\ \hline 10 & null \\ \hline \end{array}$$

The values of R_4^2 and R_5^2 are bizarre--it is highly unlikely that the hiring of every new employee would require the establishment of a new account. However, the alternative of letting all new hires share one null account number suffers the problems described in I.4 and is wrong, too.

I.6 Conclusion

Null values are an escape clause from the UR-assumption; they attempt to circumvent UR-consistency in cases where it seems most unrealistic. Since the UR-assumption is central to currently developed database theories, it is not surprising that attempts to circumvent it are so troublesome.

APPENDIX II

Properties of Updates Under UR-Assumption

II.1 Update Maps

Properties 3 and 4 of Section 4.2 are defined in terms of sequences of updates operating on individual relations. This section formalizes this concept, refining it into the notion of "update map".

An *insertion command* is an ordered pair $\omega = \langle +R_j, r_j \rangle$ and is interpreted as a function on FD-consistent database states.

$\omega(D) = D'$ such that

$$D'(R_i) = \begin{cases} D(R_i) & , \text{ for } i \neq j \\ + (D(R_j), r_j) & , \text{ for } i = j. \end{cases}$$

A sequence of insertion commands, $\Omega_+ = \langle \omega_1, \dots, \omega_k \rangle$ is interpreted as the composition of $^+ \omega_1, \dots, \omega_k$. Ω_+ is *minimal* in state D if for all projections Ω'_+ , $\Omega'_+(D) \neq \Omega_+(D)$. It is a fact that if Ω_+ is minimal in state D , every permutation of it is minimal too, and all map D into the same state. Conversely, if two sequences are both minimal and both map D into D' , they are permutations of each other. Moreover, a minimal sequence of insertion commands cannot include duplicate elements. Therefore we can unambiguously consider such sequences to be *sets*. An *insertion map* is such a set. *Deletion maps* are defined analogously.

Replacement maps require different treatment because the order of replacements can affect the result even in a minimal sequence [BG]. Let Ω_- and Ω_+ be sets of deletion and insertion commands respectively, and let $\Omega_\& = \langle \Omega_-, \Omega_+ \rangle$. We interpret $\Omega_\&$ as the composition of Ω_- followed by Ω_+ ; $\Omega_\&$ is defined in state D iff Ω_- is a deletion map in state D and Ω_+ an insertion map in state $\Omega_-(D)$. We define a partial order over pairs of this form: $\langle \Omega'_-, \Omega'_+ \rangle \leq \langle \Omega_-, \Omega_+ \rangle$ if both $\Omega'_- \subseteq \Omega_-$ and $\Omega'_+ \subseteq \Omega_+$. $\Omega_\&$ is a *replacement map* in state D if (a) $\Omega_\&$ is defined in state D , and (b) for all $\Omega'_\& < \Omega_\&$ either $\Omega'_\&(D)$ is not defined or $\Omega'_\&(D) \neq \Omega_\&(D)$.

Properties 3 and 4 can now be restated.

Property 3--Intent Performance

State D' *performs the intent* of $+R_i(D, r_i)$ if there exists an insertion map from D to D' . The definitions for $-R_i$ and $\&R_i$ are analogous.

Property 4--Nonextraneous Performance

D' *nonextraneously performs* $+R_i(D, r_i)$ if (a) D' performs $+R_i(D, r_i)$, and (b) for all $\Omega'_+ \subsetneq \Omega_+$, where Ω_+ is the insertion map from D to D' , $\Omega'_+(D)$ does not perform $+R_i(D, r_i)$. The definition for $-R_i$ is analogous; for $\&R_i$, change ' \subsetneq ' in part (b) to ' \neq '.

The following facts about update maps are used later. Let D and D' be FD-consistent states.

- (i) there exists an insertion map from D to D' iff $D \leq D'$;

- (ii) there exists a deletion map from D to D' iff $D \geq D'$; and

- (iii) there exists a replacement map from D to D' always.

II.2 Reachability

A database state is *reachable* if it can be attained by applying a sequence of update operations to an initially empty database state. Reachability is complementary to consistency. By assumption, the consistent states of a schema are the states that represent meaningful configurations of information. We require that update operations preserve consistency to ensure that if state D occurs, D represents a meaningful situation. Conversely, it is desirable that if D represents a meaningful situation (i.e., is consistent), then D can occur. This property is called *reachability*. In a single relation using the update operations defined in 3.1, reachability is obviously achieved; in this section we prove that reachability is also achieved in a database using the operations defined in 4.2.

Let $D, D' \in \text{DOM}(D)$. D' is *reachable* from D if there exists a sequence of update operations (as defined in 4.2) that maps D into D' .

LEMMA II.1. Let $D \in \text{DOM}(D)$, $U \in P^{-1}(D)$, and $U' = +(U, u)$. Then $D' = P(U')$ is reachable from D .

Proof. Let Ω_+ be the insertion map from D to D' ; $\Omega_+ = \{ \langle +R_j, r_j \rangle \mid r_j = u[\text{ATTR}_j] \wedge r_j \notin R_j, \text{ for } j = 1, \dots, n \}$. Observe that Ω_+ contains at most one element per relation schema. We shall prove that whenever Ω_+ has this form and $\Omega_+(D)$ is consistent; $\Omega_+(D)$ is reachable from D . The proof is by induction on $|\Omega_+|$.

Basis: If $|\Omega_+| = \emptyset$, the result holds trivially.

Induction: Assume $\Omega_+(D)$ is reachable whenever Ω_+ contains at most one command per relation schema, $|\Omega_+| < N$, and $\Omega_+(D)$ is consistent; prove that $\Omega_+(D)$ remains reachable when $|\Omega_+| = N$.

Let $\omega_j = \langle +R_j, r_j \rangle$ be any element of Ω_+ . We claim that $\Omega_+(D)$ performs $+R_j(D, r_j)$: (1) $\Omega_+(D)$ is consistent; (2) $\Omega_+(D)$ exactly performs the insertion, because ω_j is the only element of Ω_+ that operates on R_j ; and (3) $\Omega_+(D)$ performs the intent by definition. Hence, there exists $\Omega_j \subseteq \Omega_+$ such that $D_j = \Omega_j(D)$ nonextraneously performs $+R_j(D, r_j)$; moreover $\Omega_j \neq \emptyset$ since $\omega_j \in \Omega_j$. Let $\Omega'_+ = \Omega_+ - \Omega_j$; $\Omega'_+(D_j) = \Omega_+(D) = D'$ and $|\Omega'_+| < N$, hence Ω'_+ satisfies the induction hypothesis. Therefore $D' = \Omega'_+(D_j)$ is reachable from D_j by induction hypothesis, while D_j is reachable from D by construction. This proves that D' is reachable from D as desired. \square

THEOREM II.1. Let D be any database schema. Every $D \in \text{DOM}(D)$ is reachable from the empty state.

Proof. Let D' be any consistent state of D and let $U' \in P^{-1}(D')$. We shall prove that D' is

reachable from the empty state, by induction on $|U'|$.

Basis: If $|U'| = \emptyset$, there is nothing to prove.

Induction: Assume D' is reachable from the empty state when $|U'| < N$; prove that D' is reachable when $|U'| = N$.

Let $U = U' - \{u\}$ for any $u \in U'$, and let $D = P(U)$. D is reachable from the empty state by induction hypothesis, while $D' = P(U')$ is reachable from D by Lemma II.1. Thus D' is reachable from the empty state as desired. \square

II.3 Totality and Uniqueness

An update semantics is *total* if it specifies a meaning for every operation applied to every possible configuration of arguments. A related issue is uniqueness: a semantics is *unique* if it specifies a single meaning for each operation whenever that operation is defined.

The example of Section 4.1 suggests that insertions are not uniquely specified in the UR-environment, and one would suspect that replacements are not unique either. In [BG] we demonstrate that this is the case. What is more surprising is that these operations are *not total*.

Consider the database of Figure II.1, and the insertion $+R_1(D, \langle a3, b1, c2 \rangle)$. We shall prove that no state D' exists that performs this insertion, and thus the insertion has no meaning. Suppose such a D' does exist and let $U' \in P^{-1}(D')$. Since $R_1 \cup \{ \langle a3, b1, c2 \rangle \}$ is FD-consistent $D'(R_1)$ must have this value, and U' must include the following three tuples:

- u1: $\langle a1 \ b1 \ c1 \ d1 \rangle$
- u2: $\langle a2 \ b2 \ c2 \ d2 \rangle$
- u3: $\langle a3 \ b1 \ c2 \ dx \rangle$, where dx is a variable.

Since $u1[B] = u3[B]$, the FD $B \rightarrow D$ implies $dx = 'd1'$; but since $u2[C] = u3[C]$, $C \rightarrow D$ implies $dx = 'd2'$. Contradiction! Thus the given insertion is undefined.

$$\begin{aligned} \underline{D} &= \{R_1 = \langle \{A, B, C\}, \{A \rightarrow BC\} \rangle, \\ &R_2 = \langle \{B, D\}, \{B \rightarrow D\} \rangle, \\ &R_3 = \langle \{C, D\}, \{C \rightarrow D\} \rangle\} \end{aligned}$$

$D: R_1 =$

| | | |
|----|----|----|
| A | B | C |
| a1 | b1 | c1 |
| a2 | b2 | c2 |

, $R_2 =$

| | |
|----|----|
| B | D |
| b1 | d1 |
| b2 | d2 |

, $R_3 =$

| | |
|----|----|
| C | D |
| c1 | d1 |
| c2 | d2 |

Figure II.1. Insertion is Not Total

Replacements have greater flexibility in their external effects. Nonetheless, replacements are not defined in all cases.

Consider the database of Figure II.2, and the replacement $\&R_1(D, \langle a1, b1, c2 \rangle, \langle a2, b1, c2 \rangle)$. This operation is defined if and only if a consistent D' exists in which $D'(R_1) = (R_1 - \{ \langle a1, b1, c2 \rangle \}) \cup \{ \langle a2, b1, c2 \rangle \}$. We shall prove that no such state exists. Suppose D' does exist, and let $U' \in P^{-1}(D')$. U' must include the following three tuples:

- u1: $\langle a1 \ b1 \ c1 \ dx1 \ ex1 \rangle$, where $dx1, ex1$ are variables
- u2: $\langle a1 \ b2 \ c2 \ dx2 \ ex2 \rangle$, where $dx2, ex2$ are variables
- u3: $\langle a2 \ b1 \ c2 \ dx3 \ ex3 \rangle$, where $dx3, ex3$ are variables.

Since $u1[A] = u2[A]$, $A \rightarrow D$ implies $dx1 = dx2$; since $u2[C] = u3[C]$, $C \rightarrow D$ implies $dx2 = dx3$; and since $u1[B] = u3[B]$, $B \rightarrow E$ implies $ex1 = ex3$. But this means that $u1[DE] = u3[DE]$, and so $DE \rightarrow A$ implies $'a1' = 'a2'$. Contradiction!

Deletion operations, on the other hand, are totally and uniquely specified.

THEOREM II.2. $-R_i(D, r_i)$ is defined and has a unique value for every $D \in \text{DOM}(D)$ and tuple r_i .

Proof. Let $U = J(D)$, $U' = U - \{u \in U \mid u[\text{ATTR}_i] = r_i\}$, and $D' = P(U')$. We claim that D' performs $-R_i(D, r_i)$: (1) it is consistent; (2) it exactly performs the deletion since $D'(R_i) = U'[\text{ATTR}_i] = U[\text{ATTR}_i] - \{r_i\} = D(R_i) - \{r_i\} = -(D(R_i), r_i)$; and (3) it performs the intent of the deletion since $U' \subseteq U$ implies $P(U') \subseteq P(U) = D$. This establishes that $-R_i(D, r_i)$ is defined.

To prove uniqueness, observe that by construction, U' is the *unique subset* of U for which $U'[\text{ATTR}_i] = U[\text{ATTR}_i] - \{r_i\}$. Let D'' be any state that performs $-R_i(D, r_i)$, and let $U'' = J(D'')$. To satisfy the intent of the deletion $D'' \leq D$, hence $U'' \subseteq U$, while to exactly perform the deletion, $D''(R_i) = -(R_i, r_i)$, hence $U''[\text{ATTR}_i] = U[\text{ATTR}_i] - \{r_i\}$. But U' is the unique subset of U whose ATTR_i projection has that value. Therefore $J(D'') = U'' = U' = J(D')$, and since J is one-one, $D'' = D'$. Thus D' is the unique state that performs $-R_i(D, r_i)$. \square

$$\begin{aligned} \underline{D} &= \{R_1 = \langle \{A, B, C\}, \{BC \rightarrow A\} \rangle, \\ &R_2 = \langle \{A, D, E\}, \{A \rightarrow D, DE \rightarrow A\} \rangle, \\ &R_3 = \langle \{B, E\}, \{B \rightarrow E\} \rangle, \\ &R_4 = \langle \{C, D\}, \{C \rightarrow D\} \rangle\} \end{aligned}$$

$D: R_1 =$

| | | |
|----|----|----|
| A | B | C |
| a1 | b1 | c1 |
| a1 | b2 | c2 |
| a1 | b1 | c2 |

, $R_2 =$

| | | |
|----|----|----|
| A | D | E |
| a1 | d1 | e1 |
| a1 | d1 | e2 |

, $R_3 =$

| | |
|----|----|
| B | E |
| b1 | e1 |
| b2 | e2 |

$R_4 =$

| | |
|----|----|
| C | D |
| c1 | d1 |
| c2 | d1 |

Figure II.2. Replacement is Not Total.

II.4 Proofs of Lemmas 1 and 2

LEMMA 1. (i) $\text{Affect}(+R_i) = \{R_j \in D \mid R_i \text{ and } R_j \text{ are connected by a path in } \bar{G}(D)\}$.

(ii) $\text{NoAffect}(+R_i) = \{R_j \in D \mid \text{ATTR}_j \rightarrow \text{ATTR}_i \notin F^+\}$.

Proof. (i) We first prove that if R_i and R_j are not connected, $R_j \notin \text{Affect}(+R_i)$. Let D be any non-empty consistent state, and consider an arbitrary (defined) insertion, $+R_i(D, r)$. Let D' be any meaning of that insertion, let Ω_r be the insertion map from D to D' , and let $\Omega_i = \{\langle +R_k, r_k \rangle \in \Omega \mid R_i \text{ and } R_k \text{ are connected}\}$. Observe that $\Omega_i(D)$ performs the given insertion, and $\Omega_i \subseteq \Omega_r$, consequently $\Omega_i = \Omega_r$. $R_j \notin \text{Affect}(+R_i)$ since Ω_i contains no command of the form $\langle +R_j, r_j \rangle$ by construction.

We now prove that if R_i and R_j are connected, $R_j \in \text{Affect}(+R_i)$. Let D be any non-empty consistent state, and let u be a "thoroughly distinct" tuple relative to D , meaning that for all $A \in \text{ATTR}$, $u[A] \notin J(D)[A]$. Let $\Omega_+ = \{\langle +R_j, u[\text{ATTR}_j] \rangle \mid R_i, R_j \text{ are connected}\}$. $\Omega_+(D)$ performs the insertion $+R_i(D, u[\text{ATTR}_i])$ and affects every R_j connected to R_i , while for all $\Omega' \subseteq \Omega$, $\Omega'(D)$ does not perform the insertion.

(ii) If $\text{ATTR}_j \rightarrow \text{ATTR}_i \in F^+$, every $+R_i(D, r_i)$ that has any effect must certainly affect R_j ; otherwise $J(+R_i(D, r_i))$ would be inconsistent. This establishes that $\text{NoAffect}(+R_i) \subseteq \{R_j \in D \mid \text{ATTR}_j \rightarrow \text{ATTR}_i \notin F^+\}$.

To prove inclusion in the opposite direction, suppose $\text{ATTR}_j \rightarrow \text{ATTR}_i \notin F^+$. We will construct an insertion that does not affect R_j . By the completeness property of FDs, there exists a consistent state U' in which $\text{ATTR}_j \rightarrow \text{ATTR}_i$ does not hold. For this FD not to hold, U' must contain distinct tuples u and u' such that $u[\text{ATTR}_j] = u'[\text{ATTR}_j]$, yet $u[\text{ATTR}_i] \neq u'[\text{ATTR}_i]$. Let $U = U' - \{u\}$, let $D = P(U)$, and consider the insertion $+R_i(D, u[\text{ATTR}_i])$. This insertion has an effect and $P(U')$ performs it. Thus there exists a state that nonextraneously performs the insertion such that $D < D' \leq P(U')$. Since $U[\text{ATTR}_j] = U'[\text{ATTR}_j]$ by construction, it follows that $D(R_j) = D'(R_j)$, and so the insertion does not affect R_j as claimed. \square

LEMMA 2. (i) $\text{Affect}(-R_i) = \text{Affect}(+R_i)$

(ii) $\text{NoAffect}(-R_i) = \text{NoAffect}(+R_i)$.

Proof. Observe that $-R_i(+R_i(D, r_i), r_i) = D$ whenever $+R_i(D, r_i) \neq D$. This establishes $\text{Affect}(-R_i) \supseteq \text{Affect}(+R_i)$ and $\text{NoAffect}(-R_i) \supseteq \text{Affect}(-R_i)$.

To prove inclusion in the opposite direction, let $-R_i(D, r_i)$ be any deletion that has an effect, and let $D' = -R_i(D, r_i)$. Observe that for all $u \in J(D) - J(D')$, $P(J(D') \cup \{u\})$ is a meaning of $+R_i(D, r_i)$. Hence any R_j affected by the deletion is also affected by one of these meanings of the insertion, establishing $\text{Affect}(-R_i) \subseteq \text{Affect}(+R_i)$. Conversely, any R_j not affected by the deletion

is also not affected by these meanings for the insertion, establishing $\text{NoAffect}(-R_i) \subseteq \text{NoAffect}(+R_i)$. \square

References

- [ABU] A.V. Aho, C. Beeri, and J.D. Ullman, "The Theory of Joins in Relational Databases," *ACM Trans. on Database Syst.*, 4,3 (Sept., 1979), 297-314.
- [AM] M.A. Arbib, and E.G. Manes, *Arrows, Structures, and Functors--The Categorical Imperative*, Academic Press, New York, 1975.
- [Arm] W.W. Armstrong, "Dependency Structures of Database Relationships," *Proc. IFIP 74* (1974), 580-583.
- [BBG] C. Beeri, P.A. Bernstein, and N. Goodman, "A Sophisticate's Introduction to Database Normalization Theory," *Proc. 4th Int'l Conf. on Very Large Databases* (1978), 113-124.
- [BFH] C. Beeri, R. Fagin, and J.H. Howard, "A Complete Axiomatization for Functional and Multivalued Dependencies," *Proc. ACM SIGMOD Conf.* (Aug. 1977), 47-61.
- [Bern] P.A. Bernstein, "Synthesizing Third Normal Form Relations from Functional Dependencies," *ACM Trans. on Database Sys.*, 4 (Dec. 1976), 277-298.
- [BG] P.A. Bernstein, and N. Goodman, "What Does Boyce-Codd Normal Form Do?" Tech. Rept. 07-79, Center for Research in Computing Technology, Harvard University (May 1979).
- [BMSU] C. Beeri, A. Mendelzon, Y. Sagiv, and J.D. Ullman, "Equivalence of Relational Database Schemes," *Proc. ACM Symp. on Theory of Computing* (May 1979).
- [Cod1] E.F. Codd, "Further Normalization of the Data Base Relational Models," in *Data Base Systems* (R. Rustin, ed.), Prentice-Hall, Englewood Cliffs, N.J. (1972), 33-64.
- [Cod2] E.F. Codd, "Recent Investigations in Relational Data Base Systems," *Proc. IFIP 74* (1974), 1017-1021.
- [Date] C.J. Date, *Introduction to Database Systems*, Addison-Wesley, Reading, MA (1977).
- [DB] U. Dayal, and P.A. Bernstein, "The Updatability of Relational Views," *Proc. 4th Int'l Conf. on Very Large Databases* (1978).
- [DC] C. Delobel, and R.C. Casey, "Decomposition of a Data Base and the Theory of Boolean Switching Functions," *IBM J. of Res. and Dev.*, 17, 5 (Sept. 1972), 370-386.

- [Fag1] R. Fagin, "Multivalued Dependencies and a New Normal Form for Relational Databases," *ACM Trans. on Database Systems*, 2, 3 (Sept. 1977), 262-278.
- [Fag2] R. Fagin, "Normal Forms and Relational Database Operators," *Proc. ACM SIGMOD Conf.* (May 1979).
- [HLY] P. Honeyman, R.E. Ladner, and M. Yannakakis, "Testing the Universal Instance Assumption," *Inf. Proc. Letters*, 10, 1 (Feb. 12, 1980), 14-19.
- [Mar] J. Martin, *Computer Database Organization*, Prentice Hall, Englewood Cliffs, NJ (1975).
- [Ris] J. Rissanen, "Independent Components of Relations," *ACM Trans. on Database Sys.*, 2, 4 (Dec. 1977), 317-325.