

Uniqueness of Update Strategies for Database Views

Stephen J. Hegner

Umeå University
Department of Computing Science
SE-901 87 Umeå, Sweden
hegner@cs.umu.se
<http://www.cs.umu.se/~hegner>

Abstract. The problem of supporting updates to views of a database schema has been the focus of a substantial amount of research over the years. Since the mapping from base schema to view schema is seldom injective, there is usually a choice of possibilities for the reflection of view updates to base-schema updates. This work presents a solution to this problem which augments the constant-complement strategy of Bancilhon and Spyratos with order-theoretic properties to guarantee unique reflection of view updates. Specifically, most database formalisms endow the database states with a natural order structure, under which update by insertion is an increasing operation, and update by deletion is decreasing. Upon augmenting the original constant-complement strategy with compatible order-based notions, the reflection to the base schema of any update to the view schema which is an insertion, a deletion, or a modification which is realizable as a sequence of insertions and deletions is shown to be unique and independent of the choice of complement.

1. Introduction

Database management systems are typically large and complex, and it is seldom the case that an individual user is granted full access to the entire system. Rather, user access is via windows to the entire system, called *views*. Consequently, a great deal of research has been devoted to the various issues surrounding such views. Among these issues, perhaps none is thornier than the update problem. In general, the mapping from states in the main (or *base*) schema to states in the view schema is not injective; rather, a single view state can be the image of many different base states. At the very least, this leads to the question of which of the many alternatives should be selected as the appropriate reflection of an update to the view. However, more comprehensively, it leads to serious questions about providing a view in which the family of supported updates is systematic, appropriate, and effectively manageable. This, in turn, suggests that not all possible updates to a view should be allowed.

1.1 Open vs. closed update strategies Not surprisingly, a wide selection of approaches to the view update problem has evolved. In [Heg90], the extremes

have been termed open and closed update strategies. Roughly speaking, an *open* strategy is very liberal; as many updates as possible are allowed, with the user expected to be aware of the consequences of such a strategy. A *closed* strategy, on the other hand, is very conservative and systematic. The view appears as a schema unto itself, and the family of updates which is allowed looks exactly as would a family of updates to a base schema.

As the focus of this paper is support for updates under closed strategies, it is fair to ask why one should consider such strategies at all, given that open strategies invariably allow a wider range of updates to the view. The short answer is that open strategies inevitably give rise to certain anomalies which render them less than suitable in certain situations. A few examples will illustrate this point. Let \mathbf{C}_0 denote the schema with the single relation symbol $P[\text{Name}, \text{Dept}, \text{Proj}]$, and the initial instance $M_0 = \{(\text{Smith}, 1, \text{A}), (\text{Jones}, 2, \text{A}), (\text{Jones}, 2, \text{B})\}$. The informal semantics of a tuple such as $(\text{Smith}, 1, \text{A})$ is that employee Smith works in department 1 and on project A. It is furthermore assumed that the functional dependency $\text{Name} \rightarrow \text{Dept}$ holds; i.e., that an individual works in only one department. On the other hand, an individual may work on several projects.

First, let $\Pi_{\text{NP}} = (P[\text{Name}, \text{Proj}], \pi_{(\text{Name}, \text{Proj})})$ denote the view which retains just the **Name** and **Proj** attributes; the image of M_0 under this view is $\{(\text{Smith}, \text{A}), (\text{Jones}, \text{A}), (\text{Jones}, \text{B})\}$. Consider deletion of the tuple (Smith, A) . Under an open strategy, it might be argued that such an operation should be allowed, as the only reasonable reflection is to delete $(\text{Smith}, 1, \text{A})$ from M_0 . While this reflection is unique and unambiguous, there are nonetheless at least two points of caution. First of all, this update involves a *hidden trigger*, in that information about the department in which Smith works is also deleted. The user must have access to information beyond the view Π_{NP} itself in order to be aware of this. Furthermore, this update is *irreversible*, in the sense that it is not possible to undo its effect by re-inserting (Smith, A) into the view, since the information regarding the department of Smith has been lost. For these reasons, this update would not be allowed in a closed view. On the other hand, these difficulties do not arise in the deletion of (Jones, A) ; in this case the view update may be realized via deletion of $(\text{Jones}, 2, \text{A})$ in M_0 . Since information about the department of Jones is retained in the tuple $(\text{Jones}, 2, \text{B})$, and since $\text{Name} \rightarrow \text{Dept}$ holds, re-insertion of (Jones, A) is unambiguous and restores the initial state M_0 . Thus, deletion of (Jones, A) would be allowed, even in a closed strategy, from the view state resulting from M_0 .

Next, consider the schema \mathbf{C}_1 , a slight modification of \mathbf{C}_0 in which nulls are allowed for the attribute **Proj**. A state such as $M_1 = \{(\text{Smith}, 1, \text{A}), (\text{Jones}, 2, \text{A}), (\text{Jones}, 2, \text{B}), (\text{Wilson}, 1, \text{NULL})\}$ is now allowed. The interpretation is that Wilson does not work on any projects. Let $\Pi_{\text{N}\tilde{\text{P}}} = (P[\text{Name}, \text{Proj}], \pi_{(\text{Name}, \tilde{\text{Proj}})})$ denote the view which projects only those tuples on **(Name, Proj)** which do not contain any null values. Thus, the image of M_1 under this view is $\{(\text{Smith}, \text{A}), (\text{Jones}, \text{A}), (\text{Jones}, \text{B})\}$; the tuple $(\text{Wilson}, \text{NULL})$ is not included. In comparison to the previous example, many of the complications surrounding deletion of the tuple (Smith, A) disappear. Indeed, this deletion may be realized in M_1 by modify-

ing the tuple (Smith, 1, A) to (Smith, 1, NULL). This update is reversible, and involves no hidden triggers. However, there remains a difficulty; namely, the admissibility of an insertion into the state of the view depends upon information not visible within that view. More specifically, it is clear that the tuple (Wilson, C) may be inserted into the view state by replacing (Wilson, 1, NULL) with (Wilson, 1, C) in M_1 . On the other hand, the tuple (Young, C) may not be so inserted, since no department information is available for Young. This situation involves a *hidden dynamic constraint*, in the sense that whether or not the insertion is permitted depends upon information not contained in the view itself. For this reason, under a closed interpretation, insertion of (Wilson, C) would be disallowed, although it might be allowed under some open strategies.

Finally, consider the view $\Pi_{\widehat{\mathbf{NP}}} = (P[\text{Name}, \text{Proj}], \pi_{(\text{Name}, \widehat{\text{Proj}})})$ of \mathbf{C}_1 which does allow for the projection of null values; the image of M_1 under this view is {(Smith, A), (Jones, A), (Jones, B), (Wilson, NULL)}. Now, even under a closed interpretation, insertion of (Wilson, C) is allowed; the tuple (Wilson, NULL) is simply replaced with this new value. Insertion of (Young, C) is still not allowed; however, this information necessary to make this distinction is now embodied in the view.

It is not the purpose of this paper to argue that closed update strategies are better than open ones, or that one should be used to the exclusion of the other. Indeed, no such argument seems possible, as they each have their place. Open strategies are useful as tools for experienced users who have wide access to the database, but wish, for convenience, to represent certain operations within a view. Such users can and will understand the more global implications of view updates, and so issues such as hidden triggers, irreversibility, and hidden dynamic constraints pose no real difficulties. On the other hand, there is also a clear need for views which are totally encapsulated. Many users, for reasons of security, lack of expertise, or simply effective management of complex schemata, need to work with views which, for operational purposes, are self contained. In particular, in such situations, the admissibility and effects of updates must be understandable within, and restricted entirely to, the view itself. In this context, anomalies such as hidden triggers, hidden dynamic constraints, and irreversibility must be avoided. It is to such closed update strategies which this paper is directed.

1.2 The constant-complement strategy The preceding discussion provides only an anecdotal characterization of the notion of a closed view; it remains to formalize this concept. Fortunately, to a large degree, this has already been done. The seminal work on closed strategies is the *constant-complement* approach, developed some twenty years ago by Bancilhon and Spyrtos [BS81]. The idea is quite simple. To support updates to the view I_1 of the main schema \mathbf{D} , a view I_2 which is complementary to I_1 is identified; i.e., $\{I_1, I_2\}$ forms a lossless decomposition of \mathbf{D} . Then, the only updates to I_1 which are allowed are those which hold the state of I_2 fixed. Intuitively, the changes are isolated in I_1 ; the “rest” of the schema, which is I_2 , cannot change. This freezing of the state of the complement under view update eliminates the possibility of hidden triggers

and hidden dynamic constraints. Remarkably, it also eliminates the possibility of anomalies such as irreversibility.

An example will help illustrate. Let the base schema \mathbf{E}_1 have the single relation $R[ABC]$, governed by the functional dependency $B \rightarrow C$. Let the view to be updated be $\Pi_{AB} = (R[AB], \pi_{AB})$, the projection of the main relation onto the attributes AB . (In the context of the example of 1.1, just take $A = \text{Proj}$, $B = \text{Name}$, and $C = \text{Dept.}$) A natural complement to Π_{AB} is the view $\Pi_{BC} = (R[BC], \pi_{BC})$. It is easy to see that updates to Π_{AB} which keep Π_{BC} constant are precisely those which hold the projection $\Pi_B = (R[B], \pi_B)$ of the relation $R[AB]$ constant. The common view Π_B is called the *meet* of Π_{AB} and Π_{BC} , and the latter two views are called Π_B -complements. In [Heg90, 2.10], it is established that every constant-strategy in the sense of [BS81] is in fact based upon a meet complement, and the updates which are allowed are precisely those which hold that meet constant.

This situation is illustrated diagrammatically in Fig. 1 to the right. The area with up-to-right crosshatch is to be held constant; this includes in particular the meet $\Pi_B = \Pi_{AB} \wedge \Pi_{BC}$. The rest of Π_{AB} , with down-to-right crosshatching, may be updated without restriction.




Fig. 1: Visualization of the constant complement strategy

1.3 Overview and scope of this work Part of the appeal of the constant-complement strategy is its generality; it is formulated in a framework in which database schemata are sets and database mappings are functions. Thus, it is applicable, in principle, to almost any data model. With this generality comes a certain lack of uniqueness. In general, there are many complements to choose from, and distinct complements yield distinct update strategies. As illustrated in [BS81, Sec. 6], this ability to choose can sometimes be useful. If Γ_2 and Γ_3 are both meet complements of the view Γ_1 to be updated, and the meet of $\{\Gamma_1, \Gamma_2\}$ is different than that of $\{\Gamma_1, \Gamma_3\}$, then selecting Γ_2 as the constant complement for an update strategy will result in a different set of allowable updates than will choosing Γ_3 as the complement. On the other hand, if the pairs $\{\Gamma_1, \Gamma_2\}$ and $\{\Gamma_1, \Gamma_3\}$ have the same meet, then the allowable updates to Γ_1 will be the same for each complement; the difference will lie solely in how they are reflected into the base schema \mathbf{D} . This is almost never a desirable situation; a continuation of the above relational example will illustrate the problems which can ensue. Let $\text{Dom}(C)$ denote the set of allowable domain values for attribute C , let $\alpha : \text{Dom}(C) \rightarrow \text{Dom}(C)$ be any permutation of $\text{Dom}(C)$, and define the view $\Pi'_{BC} = (R'[BC], \pi'_{BC})$ as follows. States of Π'_{BC} are binary tuples on the domains of B and C , just as are states of the projection $R[BC]$. However, the tuples which are in the instance of $R'[BC]$ have their C -

position values translated by α . Specifically, let M be a state of E_1 . For any $b \in \text{Dom}(B)$, let $\#_A(b)$ denote the number of distinct values of for attribute A associated with b in the relation $\pi_{AB}(M)$. Define $\pi'_{BC}(M) = \{(b, c) \mid (b, c) \in \pi_{BC}(M) \text{ and } \#_A(b) \text{ is odd}\} \cup \{(b, \alpha(c)) \mid (b, c) \in \pi_{BC}(M) \text{ and } \#_A(b) \text{ is even}\}$. It is easy to see that Π'_{BC} is a Π_B -complement of Π_{AB} , and the updates permitted to Π_{AB} under constant complement Π'_{BC} are exactly the same as those permitted under constant complement Π_{BC} . However, the translation is quite different. As a specific example, suppose that $\text{Dom}(C) = \{c_0, c_1, c_2\}$, and let $\alpha(c_i) = c_{(i+1) \bmod 3}$. Let $M = \{(a_0, b_0, c_0), (a_1, b_1, c_1), (a_2, b_1, c_1)\}$. Then $\pi'_{BC}(M) = \{(b_0, c_0), (b_1, \alpha(c_1))\} = \{(b_0, c_0), (b_1, c_2)\}$. Now suppose that the tuple (a_1, b_0) is to be added to the state of Π_{AB} . Under constant complement Π_{BC} , the translation is simply to insert (a_1, b_0, c_0) to M . However, under constant complement Π'_{BC} , the translation is to insert the tuple $(a_1, b_0, \alpha^{-1}(c_0)) = (a_1, b_0, c_2)$, and to change the tuple (a_0, b_0, c_0) to $(a_0, b_0, \alpha^{-1}(c_0)) = (a_0, b_0, c_2)$, so that the state of E_1 after the update is $\{(a_0, b_0, c_2), (a_1, b_0, c_2), (a_1, b_1, c_1), (a_2, b_1, c_1)\}$.

It is difficult to imagine a circumstance under which update to Π_{AB} with constant complement Π'_{BC} would be desirable. It is clear that Π_{BC} is *the* complement to be kept constant. Yet, the constant-complement approach itself gives no preference to Π_{BC} over Π'_{BC} . The goal of the work reported here is to augment that approach so that only updates with respect to the “natural” complement (Π_{BC} in the example) are permitted, while preserving, to the greatest extent possible, the generality and data-model independence of the original work of Bancilhon and Spyratos. Principally, such a theory is useful not because one might be led to select Π'_{BC} over Π_{BC} as a complement to Π_{AB} (clearly, based upon aesthetics alone, one would not), but rather because it is useful to have a theory which provides a sound foundation as to why Π_{BC} is preferable. Furthermore, it eliminates the possibility that some other view I_1 of some schema \mathbf{S} might have two apparently equally aesthetically reasonable complementary views I_2 and I'_2 from which to choose.

The solution utilizes natural order structure. Most data models admit an order on the states under which insertions increase the position of the state in the order, while deletions decrease it. Furthermore, basic data morphisms are typically monotonic under this ordering. For example, in the relational model, the natural ordering is defined by relation-by-relation inclusion, while the basic morphism operations of projection, selection, and join are monotonic with respect to this ordering. The context of this work is that of ordered schemata and monotonic morphisms. The central result is that, within this restricted context, the translation of insertions and deletions is unique, and independent of the particular choice of complement. This is a very versatile result, in that it does not depend upon establishing the uniqueness of a complement; indeed, it holds even across complements with different meets.

Relative to the relational example above, with relation-by-relation inclusion defining the ordering, it is easy to see that the view mappings π_{AB} and π_{BC} are monotonic, while π'_{BC} is not. Thus, Π'_{BC} is not a suitable complement to Π_{AB} for constant-complement update within the framework presented here. Further-

more, since every update to Π_{AB} which holds Π_B constant may be realized as a sequence of insertions and deletions (the family of updates is *order realizable* in the terminology of 4.1), translation of these updates via constant complement Π_{BC} is the only strategy allowed under order-based constant complement.

The results presented here are limited in that they apply only to updates which are insertions, deletions, or realizable as sequences of such. They do not apply, in general, to modifications. For example, let \mathbf{E}_2 be the relational schema which is identical to \mathbf{E}_1 , save that the additional functional dependency $B \rightarrow A$ holds. With constant complement Π_{BC} , the only updates which are allowed to Π_{AB} are those which replace the A component of a tuple (a, b) with a new value. It is not possible to realize such an update as an insertion followed by a deletion, since the insertion would result in a violation of the functional dependency $B \rightarrow A$, which embeds in $R[AB]$. Thus, with the natural relation-by-relation inclusion, the theory developed here has nothing to say about updates to this view. It is an order-based theory, and updates which sidestep the order are not covered. Fortunately, there is still a way to establish uniqueness of the update strategy. In many cases, including this one, it is possible to find another order which will render these updates to be order-based, while retaining the monotonicity of the underlying view mappings. A full solution is presented in 4.5.

1.4 Applicability and relationship to previous work Over the years, there has been a substantial amount of research on the topic of view updates, a majority of which has focused upon the relational model; some quite recent [BL98]. The problem has also been studied, to a limited extent, in the context of the ER model [Tha00, Sec. 8.3] and the object-oriented model [Bel00]. However, except for the seminal work of Bancilhon and Spyratos [BS81] and that of the author [Heg90], most of this research has the flavor of open strategies.

The work reported here depends only upon the data model possessing a natural order structure, and so is potentially applicable to any of the contexts identified above. Although the examples used in this paper are all relational, the theory is not tied to that model in any way. Nonetheless, it does presume a situation in which the state of the database is represented by a single structure, rather than as a set of constraints, so applicability to deductive models is not clear. In any case, research on updates to deductive databases has taken a direction quite its own; see [MT99] for a recent survey.

The choice of order as a tool to establish uniqueness results is motivated by earlier work [Heg94], in which such tools were used to show that decompositions into independent components are unique when databases are suitably ordered. However, the techniques employed in this paper are vastly different than those of [Heg94], since independence is not a property which most user views have. This paper is in some way part of a long overdue sequel to [Heg90], which laid out principles for the support of closed views, but never unified the results to fruition.

1.5 The results and the relational model Although the results are not specific to the relational model, it remains not only the most widely used data model,

but the one with, by far, the most extensive theoretical foundation. Therefore, it is important to show applicability to that model. To this end, several examples set within the classical relational context which illustrate the use of the results developed here are presented. Unfortunately, it is not practical to summarize even the full notation and terminology from that model; rather, it must be assumed that the reader is familiar with the standard terminology, notation, and results within that context. The monograph [AHV95] provides a relatively recent survey and the necessary background. Only a few key points which are necessary to show how this model fits into the order-based framework are presented here. Furthermore, due to space limitations, proofs of results which are specific to the relational context are only sketched, with many details omitted.

2. Database Concepts in the Order-Based Context

In this section, the fundamental ideas of database schemata, morphisms, views, and complements are formulated within the context of order. That is, database schemata are presumed to have an underlying order to their states, and database mappings are assumed to preserve this order. A few of the ideas are related to those presented in [Heg94]; however, the details are quite different, since that work dealt with a much more specialized context.

2.1 Posets Familiarity with the fundamental ideas of posets, such as presented in [DP90], is presumed; only a few notational and terminological points are reviewed here. A *partially ordered set (poset)* is pair $\mathbf{P} = (P, \leq)$ in which P is a set and \leq is a reflexive, transitive, and antisymmetric relation on P . Given posets $\mathbf{P} = (P, \leq)$ and $\mathbf{Q} = (Q, \leq)$, a *morphism* $f : \mathbf{P} \rightarrow \mathbf{Q}$ is a *monotone* function $f : P \rightarrow Q$; i.e., $p_1, p_2 \in P$ with $p_1 \leq p_2$ implies that $f(p_1) \leq f(p_2)$. The morphism f is *open* if, for any $q_1, q_2 \in Q$ with $q_1 \leq q_2$, there are $p_1 \in f^{-1}(q_1)$, $p_2 \in f^{-1}(q_2)$ with $p_1 \leq p_2$. In other words, f is open if Q carries the least order which renders f a morphism. Following the standard terminology of category theory [Mac98], the morphism f is an *isomorphism* iff it has both left and right inverses. It is easily verified that this is equivalent to being an open bijection.


2.2 Schemata and morphisms Mathematically, a database schema with order is just a partially ordered set, and a morphism of such schemata is a poset morphism. However, to emphasize the database aspects, a special notation is employed in this work. Specifically, a *database schema with order* is a poset $\mathbf{D} = (\text{LDB}(\mathbf{D}), \leq_{\mathbf{D}})$ in which $\text{LDB}(\mathbf{D})$ is a set, called the set of *legal databases* of \mathbf{D} .

In the case that the order $\leq_{\mathbf{D}}$ is the identity order in which $M \leq_{\mathbf{D}} N$ iff $M = N$, \mathbf{D} is called a *flat* or *unordered* schema. In this case, the order relation $\leq_{\mathbf{D}}$ plays no rôle whatever, and the schemata and morphisms are just sets and functions, respectively.

To simplify terminology, throughout the rest of this paper, the term *schema* shall mean *schema with order*, unless specifically stated to the contrary.

Given database schemata $\mathbf{D}_1 = (\text{LDB}(\mathbf{D}_1), \leq_{\mathbf{D}_1})$ and $\mathbf{D}_2 = (\text{LDB}(\mathbf{D}_2), \leq_{\mathbf{D}_2})$, a *morphism* $h : \mathbf{D}_1 \rightarrow \mathbf{D}_2$ is a function $h : \text{LDB}(\mathbf{D}_1) \rightarrow \text{LDB}(\mathbf{D}_2)$ which is a poset morphism with respect to the orders $\leq_{\mathbf{D}_1}$ and $\leq_{\mathbf{D}_2}$. It is called an *open surjection* precisely in the case that it has that property as a poset morphism.

2.3 Views and congruences An *order view* of the schema \mathbf{D} is a pair $\Gamma = (\mathbf{V}, \gamma)$ in which \mathbf{V} is a schema and $\gamma : \mathbf{D} \rightarrow \mathbf{V}$ is an open surjection. Given order views $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ of \mathbf{D} , a *morphism* $h : \Gamma_1 \rightarrow \Gamma_2$ is a schema morphism $h : \mathbf{V}_1 \rightarrow \mathbf{V}_2$ with the property that $h \circ \gamma_1 = \gamma_2$, that is, such that the diagram to the right commutes. Since the morphisms γ_1 and γ_2 are open surjections, so too is h . The order views Γ_1 and Γ_2 are *isomorphic* just in case they satisfy the standard categorical notion of isomorphism [Mac98]; that is, there are morphisms $h_1 : \Gamma_1 \rightarrow \Gamma_2$ and $h_2 : \Gamma_2 \rightarrow \Gamma_1$ such that $h_1 \circ h_2$ and $h_2 \circ h_1$ are both identities. The *congruence* of order view $\Gamma = (\mathbf{V}, \gamma)$ is the equivalence relation $\text{Congr}(\Gamma)$ on $\text{LDB}(\mathbf{V})$ defined by $(M, N) \in \text{Congr}(\Gamma)$ iff $\gamma(M) = \gamma(N)$. Given $M \in \text{LDB}(\mathbf{D})$, the notation $[M]_\Gamma$ is shorthand for the more cumbersome $[M]_{\text{Congr}(\Gamma)}$. Both denote the equivalence class of M under the equivalence relation $\text{Congr}(\Gamma)$.



Throughout the remainder of this work, unless specifically stated to the contrary, the term *view* will mean *order view*, as defined above.

Before proceeding further, it is important to establish that the general notions introduced here are applicable to the classical relational setting. The result 2.5 below shows that the common *Select-Project-Join* mappings of the relational theory define order based views in the sense of 2.2 above. First, some clarification of terminology and notation is in order.

2.4 Order-based schemata and morphisms in the classical relational context The *named perspective* [AHV95, Sec. 3.2] is used; this means that the columns of relations are identified by attribute names from a universe \mathbf{U} . For a given attribute A , $\text{Dom}(A)$ denotes the set of all allowable values for that attribute. A *relational schema* \mathbf{D} consists of a finite set of relational symbols $\text{Rel}(\mathbf{D})$, each with an arity $\text{Arity}(R) \subseteq \mathbf{U}$. An *unconstrained database* M over a relational schema \mathbf{D} consists of a set of relations $\{M^R \mid R \in \text{Rel}(\mathbf{D})\}$ of the appropriate arities. The set of all unconstrained databases on \mathbf{D} is denoted $\text{DB}(\mathbf{D})$. The *natural ordering* $\subseteq_{\mathbf{D}}$ on the databases of \mathbf{D} is defined via relation-by-relation inclusion; i.e., $M_1 \subseteq_{\mathbf{D}} M_2$ iff $R^{M_1} \subseteq R^{M_2}$ for all relation symbols R of \mathbf{D} .

Relational schemata are commonly constrained by Horn sentences. If only universal quantifiers are allowed, such constraints are termed *full dependencies*; if existential quantifiers are allowed on positive atoms, they are termed *embedded dependencies* [AHV95, Chap. 10]. Functional dependencies (*fd*'s) and join dependencies (*jd*'s) are full dependencies. If there is a set Φ of full (resp. embedded) dependencies such that $\text{LDB}(\mathbf{D}) = \{M \in \text{DB}(\mathbf{D}) \mid M \models \text{Dep}(\mathbf{D})\}$, then \mathbf{D} is said to be *constrained by full* (resp. *embedded*) *dependencies*; in both cases the notation $\text{Dep}(\mathbf{D})$ is used to denote Φ .

It is critical to know that the most common types of morphisms within the relational framework define views in the order-based context. The database mappings which will be considered in this work are those which are built up from compositions of the primitives *selection*, *projection*, *join*, and *renaming*. Such compositions are termed *SPJR-morphisms*. [AHV95, Sec. 4.4]. It is clear that *SPJR*-mappings are monotonic with respect to the natural orderings, and so define morphisms in the order-based sense. It is equally important to note that database mappings which involve negation, such as difference and division, are not monotonic.

2.5 Proposition — *SPJR*-morphisms define views *Let \mathbf{D} and \mathbf{V} be relational schemata, and suppose furthermore that \mathbf{D} is constrained by full dependencies. Let $\gamma : \mathbf{D} \rightarrow \mathbf{V}$ be a surjective *SPJR*-morphism. Then \mathbf{V} is also constrained by full dependencies, and γ is an open poset morphism with respect to the natural orderings $\subseteq_{\mathbf{D}}$ and $\subseteq_{\mathbf{V}}$, so that $\Gamma = (\mathbf{V}, \gamma)$ is a view of \mathbf{D} in the order-based sense.*

PROOF OUTLINE: First of all, that the class of full dependencies is closed under projection and join is established in [Hul84, Thm. 5.4]. Selection may be added easily using the submodel characterization of universal Horn sentences [Mon76, Thm. 25.13]. Thus, \mathbf{V} is also constrained by full dependencies.

To establish that γ is open, break the problem into three cases, one each for projection, selection, and join, each with possible renaming. Let $N_1, N_2 \in \text{LDB}(\mathbf{V})$ with $N_1 \subseteq_{\mathbf{V}} N_2$. Now, reflect N_1 and N_2 back into \mathbf{D} as minimal structures P_1 and P_2 , respectively, possibly with variables, and apply the chase inference procedure [AHV95, Sec. 10.2] to generate legal databases of \mathbf{D} . When variables occur (this will happen with projection), they must be the same in tuples which are in both N_1 and N_2 , and their eventual bindings to domain values must be the same in each case. Let $\hat{P}_1, \hat{P}_2 \in \text{LDB}(\mathbf{D})$ be the models generated from the chase on P_1 and P_2 , respectively. Then $\hat{P}_1 \subseteq_{\mathbf{D}} \hat{P}_2$, with $\gamma(\hat{P}_1) = N_1$ and $\gamma(\hat{P}_2) = N_2$. \square

2.6 Relational projections The key relational examples used in this work are projections. Although this framework has already been used in the introduction, it is important to crystallize the notation for the formal part of the paper. Let \mathbf{D} be a relational schema which is constrained by full dependencies, let $R \in \text{Rel}(\mathbf{D})$, let $\mathbf{A} = \text{Arity}(R)$, and let $\mathbf{B} \subseteq \mathbf{A}$. The \mathbf{B} -*projection* of R is the view $\Pi_{\mathbf{B}} = (R[\mathbf{B}], \pi_{\mathbf{B}})$ with $R[\mathbf{B}]$ the schema with a single relational symbol R on attribute set \mathbf{B} , and $\pi_{\mathbf{B}} : \mathbf{D} \rightarrow R[\mathbf{B}]$ the morphism which sends $M \in \text{LDB}(\mathbf{D})$ to the projection onto attributes in \mathbf{B} of the relation R in \mathbf{D} . $\text{Dep}(R[\mathbf{B}])$ is taken to be the set of full dependencies which render $\pi_{\mathbf{A}}$ an open surjection. In view of 2.5 above, $\text{Dep}(R[\mathbf{B}])$ is guaranteed to exist, and $\Pi_{\mathbf{B}}$ is guaranteed to be a view in the order-based sense.

The focus now returns to the more general context. The following proposition characterizes views and their isomorphism classes in terms of their congruences. The proof is completely straightforward, and so omitted.

2.7 Proposition Let \mathbf{D} be a schema, and let $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ be views of \mathbf{D} .

- (a) Every view morphism $\Gamma_1 \rightarrow \Gamma_2$ is an open surjection of posets.
- (b) There is at most one morphism $\Gamma_1 \rightarrow \Gamma_2$.
- (c) There is a morphism $\Gamma_1 \rightarrow \Gamma_2$ iff $\text{Congr}(\Gamma_1) \subseteq \text{Congr}(\Gamma_2)$.
- (d) The views Γ_1 and Γ_2 are isomorphic iff $\text{Congr}(\Gamma_1) = \text{Congr}(\Gamma_2)$. \square

2.8 View equivalence and order The (equivalence) class of all views which are isomorphic to Γ is denoted $[\Gamma]$. In view of 2.7(d) above, $[\Gamma_1] = [\Gamma_2]$ iff $\text{Congr}(\Gamma_1) = \text{Congr}(\Gamma_2)$. The notation $[\Gamma_2] \leq [\Gamma_1]$ denotes that there is a morphism $\Gamma_1 \rightarrow \Gamma_2$, or, equivalently, that $\text{Congr}(\Gamma_1) \subseteq \text{Congr}(\Gamma_2)$. As a shorthand, $\Gamma_2 \leq \Gamma_1$ shall also be written to denote this fact, with the understanding that \leq is not a partial order on the views themselves, but only on the underlying equivalence classes.

If $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ are views of the schema \mathbf{D} with $\Gamma_2 \leq \Gamma_1$, then the unique morphism $\Gamma_1 \rightarrow \Gamma_2$ is denoted $\lambda\langle\Gamma_1, \Gamma_2\rangle$. In this case, \mathbf{V}_2 may be regarded as a view of \mathbf{V}_1 under the view mapping $\lambda\langle\Gamma_1, \Gamma_2\rangle$. Specifically, the *relativization* of Γ_2 to Γ_1 is the view $\Lambda(\Gamma_1, \Gamma_2) = (\mathbf{V}_2, \lambda\langle\Gamma_1, \Gamma_2\rangle)$ of \mathbf{V}_1 .

2.9 Order-compatible congruences In the context of a flat schema \mathbf{D} with no order relation, every equivalence relation R on $\text{LDB}(\mathbf{D})$ gives rise to a view whose states are the equivalence classes of R , with the view mapping the natural projection of an element to its equivalence class. In the context of order-based schemata, an additional constraint mandating that the equivalence respect the order on the database states must be imposed.

Specifically, let \mathbf{D} be an order-based schema, and let R be an equivalence relation on $\text{LDB}(\mathbf{D})$. Call R *order compatible* for \mathbf{D} if for every pair $(M_1, M_2) \in R$ with $M_1 \leq_{\mathbf{D}} M_2$ and every $M_3 \in \text{LDB}(\mathbf{D})$ with $M_1 \leq_{\mathbf{D}} M_3 \leq_{\mathbf{D}} M_2$, $(M_1, M_3) \in R$ as well. Now, define $\Theta_R = (\text{LDB}(\mathbf{D})/R, \theta_R)$ to be the view of \mathbf{D} with $\text{LDB}(\mathbf{D})/R$ the schema whose underlying set is $\text{LDB}(\mathbf{D})/R$, the set of equivalence classes of R , with $\theta_R : \mathbf{D} \rightarrow \text{LDB}(\mathbf{D})/R$ the morphism which sends each $M \in \text{LDB}(\mathbf{D})$ to its equivalence class $[M]_R$ under R . The order \leq_{Θ_R} is given by $[M]_R \leq_{\Theta_R} [N]_R$ iff $(\exists M_1 \in [M]_R)(\exists N_1 \in [N]_R)(M_1 \leq_{\mathbf{D}} N_1)$. The order compatibility of R ensures that \leq_{Θ_R} is a partial order, and the construction itself ensures that θ_R is an open surjection. Thus, Θ_R is indeed an order-based view of \mathbf{D} .

Note, conversely, that for any view $\Gamma = (\mathbf{V}, \gamma)$, $\text{Congr}(\Gamma)$ is an order-based congruence, since otherwise the order $\leq_{\mathbf{V}}$ would not be well defined.

2.10 Products and complements of views Let $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ be views of the schema \mathbf{D} . The product view $\Gamma_1 \times \Gamma_2 = (\mathbf{V}_1 \times_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2, \gamma_1 \otimes \gamma_2)$ has $\text{LDB}(\mathbf{V}_1 \times_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2) = \{(\gamma_1(M), \gamma_2(M)) \mid M \in \text{LDB}(\mathbf{D})\}$. The morphism $\gamma_1 \otimes \gamma_2 : \mathbf{D} \rightarrow \mathbf{V}_1 \times_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2$ is given on elements by $M \mapsto (\gamma_1(M), \gamma_2(M))$. The order on $\mathbf{V}_1 \times_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2$ is that which renders $\gamma_1 \otimes \gamma_2$ an open poset morphism.

The pair $\{\Gamma_1, \Gamma_2\}$ of views is said to form a *subdirect complementary pair* just in case $\gamma_1 \otimes \gamma_2 : \mathbf{D} \rightarrow \mathbf{V}_1 \times_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2$ is a poset isomorphism.

It is entirely possible for $\gamma_1 \otimes \gamma_2$ to be a bijection and a poset morphism without being a poset isomorphism, since the order on $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$ induced by the product order on $\mathbf{V}_1 \times \mathbf{V}_2$ is, in general, strictly stronger than that induced by the morphism $\gamma_1 \otimes \gamma_2$. Thus, two order-based views may be complementary in a “flat” environment in which order is ignored, without being complements in the order-based sense.

In the examples of 1.2 and 1.3, the common view Π_B was referred to as the *meet* of Π_{AB} and Π_{BC} . This notion of the common component of two views is extremely important in a number of contexts involving views and database decomposition, including the generalization of acyclic decompositions [Heg93] and decomposition into independent components [Heg94]. Not surprisingly, it is also central to the constant complement strategy. For it to be well defined, the congruences of the two views must commute; the definition follows.

2.11 Fully commuting views and meet complements The pair $\{\Gamma_1, \Gamma_2\}$ of views of \mathbf{D} is called a *fully commuting pair* if $\text{Congr}(\Gamma_1) \circ \text{Congr}(\Gamma_2) = \text{Congr}(\Gamma_2) \circ \text{Congr}(\Gamma_1)$, with “ \circ ” denoting ordinary relational composition. A subdirect complementary pair $\{\Gamma_1, \Gamma_2\}$ which is fully commuting is called a *meet-complementary pair*, and Γ_1 and Γ_2 are called *meet complements* of one another.

2.12 Generalized dependencies Let $\{\Gamma_1, \Gamma_2\}$ be a subdirect complementary pair.

- (a) The $\{\Gamma_1, \Gamma_2\}$ -*reconstruction dependency* on $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$, denoted $\otimes[\Gamma_1, \Gamma_2]$, is satisfied iff for any $M_1, N_1 \in \text{LDB}(\mathbf{V}_1)$ and $M_2, N_2 \in \text{LDB}(\mathbf{V}_2)$, if any three of the elements of the set $\{(M_1, M_2), (M_1, N_2), (N_1, M_2), (N_1, N_2)\}$ is in $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$, then so too is the fourth.
- (b) Let $\Gamma_3 = (\mathbf{V}_3, \gamma_3)$ be a view of \mathbf{D} , with $[\Gamma_3] \leq [\Gamma_1]$ and $[\Gamma_3] \leq [\Gamma_2]$. The Γ_3 -*independence dependency* on $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$, denoted \otimes_{Γ_3} , is satisfied iff for any $M_1 \in \text{LDB}(\mathbf{V}_1)$ and $M_2 \in \text{LDB}(\mathbf{V}_2)$, $((M_1, M_2) \in \text{LDB}(\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2)) \Leftrightarrow (\lambda\langle\Gamma_1, \Gamma_3\rangle(M_1) = \lambda\langle\Gamma_2, \Gamma_3\rangle(M_2))$

The following characterization first appeared in [Heg90, 1.13]. Unfortunately, due to space constraints, no proof was presented. Because of its importance, the proof is sketched here. Note that this result is essentially independent of additional constraints imposed by the order-based context.

2.13 Theorem – characterization of meet-complementary pairs *Let $\{\Gamma_1, \Gamma_2\}$ be a subdirect complementary pair. Then the following conditions are equivalent.*

- (a) $\{\Gamma_1, \Gamma_2\}$ is a meet-complementary pair.
- (b) $\text{Congr}(\Gamma_1) \circ \text{Congr}(\Gamma_2)$ is an equivalence relation.
- (c) $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$ satisfies $\otimes[\Gamma_1, \Gamma_2]$.
- (d) $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$ satisfies \otimes_{Γ_3} , with Γ_3 the view (unique up to equivalence) whose congruence is the smallest equivalence relation on $\text{LDB}(\mathbf{D})$ containing both $\text{Congr}(\Gamma_1)$ and $\text{Congr}(\Gamma_2)$.

PROOF: The implications (a) \Leftrightarrow (b) and (a) \Rightarrow (c) are straightforward and left to the reader, while (a) \Leftrightarrow (d) is a special case of the Chinese-remainder characterization for schema decomposability [Heg93, 2.1.5]. To show the implication, (c) \Rightarrow (a) let $(M, N) \in \text{Congr}(I_1) \circ \text{Congr}(I_2)$. Then, there is a $P \in \text{LDB}(\mathbf{D})$ with $(M, P) \in \text{Congr}(I_1)$ and $(P, N) \in \text{Congr}(I_2)$. Let (M_1, M_2) , (N_1, N_2) , and (P_1, P_2) denote the images, under $\gamma_1 \otimes \gamma_2$, of M , N , and P , respectively. Since $(M, P) \in \text{Congr}(I_1)$, $M_1 = P_1$. Similarly, since $(P, N) \in \text{Congr}(I_2)$, $N_2 = P_2$. Thus, $(P_1, P_2) = (M_1, N_2)$. In particular, $(M_1, N_2) \in \text{LDB}(\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2)$. So, applying $\otimes[I_1, I_2]$ to the set $\{(M_1, M_2), (M_1, N_2), (N_1, N_2)\}$, it follows that $(N_1, M_2) \in \text{LDB}(\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2)$. Thus, $((\gamma_1 \otimes \gamma_2)^{-1}(M_1, M_2), (\gamma_1 \otimes \gamma_2)^{-1}(N_1, M_2)) \in \text{Congr}(I_2)$ and $((\gamma_1 \otimes \gamma_2)^{-1}(N_1, M_2), (\gamma_1 \otimes \gamma_2)^{-1}(N_1, N_2)) \in \text{Congr}(I_1)$, so that $(M, N) \in \text{Congr}(I_2) \circ \text{Congr}(I_1)$, whence $\text{Congr}(I_1) \circ \text{Congr}(I_2) \subseteq \text{Congr}(I_2) \circ \text{Congr}(I_1)$. The opposite inclusion is proved in an analogous fashion, so that $\text{Congr}(I_1) \circ \text{Congr}(I_2) = \text{Congr}(I_2) \circ \text{Congr}(I_1)$. \square

2.14 Meets and I -complements Let $I_1 = (\mathbf{V}_1, \gamma_1)$ and $I_2 = (\mathbf{V}_2, \gamma_2)$ be views of \mathbf{D} which form a subdirect complementary pair. In the case that the equivalent conditions of 2.13 above are satisfied, the view I_3 guaranteed by (d) is called the *meet* of $\{I_1, I_2\}$, and is denoted $I_1 \wedge I_2 = (\mathbf{V}_1 \gamma_1 \wedge \gamma_2 \mathbf{V}_2, \gamma_1 \wedge \gamma_2)$. The set $\{I_1, I_2\}$ is then said to form a *meet-complementary pair*, or *I_3 -complementary pair*, and I_1 and I_2 are called *I_3 -complements* of one another.

The meet view I_3 is defined only up to isomorphism of views, but this is no limitation for the work presented here. Also, since $\mathbf{V}_1 \gamma_1 \wedge \gamma_2 \mathbf{V}_2$ factors through both I_1 and I_2 , it is immediate that the order on $\text{LDB}(\mathbf{V}_1 \gamma_1 \wedge \gamma_2 \mathbf{V}_2)$ induced by either $\leq_{\mathbf{V}_1}$ or $\leq_{\mathbf{V}_2}$ is the same as that induced by $\leq_{\mathbf{D}}$, so there is no question that this order is well defined.

2.15 Products and meets in the relational context Let \mathbf{D} be a relational schema, and let $I_1 = (\mathbf{V}_1, \gamma_1)$ and $I_2 = (\mathbf{V}_2, \gamma_2)$ be views of \mathbf{D} defined by SPJR-morphisms. The relation names of $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$ will be the disjoint union of those in \mathbf{V}_1 and \mathbf{V}_2 . Furthermore, if the pair $\{I_1, I_2\}$ forms a subdirect decomposition, then, in general, there will interrelational constraints. For example, let \mathbf{E}_1 be the schema introduced in 1.2, with the single relation symbol $R[ABC]$, governed by the fd $B \rightarrow C$. Let $I_1 = \Pi_{AB}$ and $I_2 = \Pi_{BC}$. On $\mathbf{V}_1 \gamma_1 \otimes \gamma_2 \mathbf{V}_2$, which contains relation symbols $R_1[AB]$ and $R_2[BC]$ (the subscripts on the R 's are used here to avoid name collision), the fd $B \rightarrow C$ holds on $R_2[BC]$, and the jd $R_1[AB] \bowtie R_2[BC]$ binds the two relations together. In this case, it is easy to see that $\{I_1, I_2\}$ are Π_B -complements, as condition (d) of 2.13 applies. On the other hand, if the fd $A \rightarrow C \in \text{Dep}(\mathbf{D})$ as well, as in the schema \mathbf{E}_2 of 1.3, then while $\{I_1, I_2\}$ remains a subdirect complementary pair, it is no longer a meet-complementary pair, as condition (d) of 2.13 fails [Heg90, 1.11].

The generalization of this idea makes use of the idea of dependency preservation. A decomposition $\{I_1 = (\mathbf{V}_1, \gamma_1), I_2 = (\mathbf{V}_2, \gamma_2)\}$ of a relational schema \mathbf{D} is *dependency preserving* if $\text{Dep}(\mathbf{D})$ is recoverable from the reconstruction dependency (usually a join dependency), together with $\text{Dep}(\mathbf{V}_1) \cup \text{Dep}(\mathbf{V}_2)$. If $\text{Dep}(\mathbf{D})$ consists of fd's, this means that a cover of $\text{Dep}(\mathbf{D})$ embeds in the relations of the

view schemata. For details, see [AHV95, Sec. 11.2]. The formalization is then the following.

2.16 Proposition — meets of projections in the relational context *Let \mathbf{D} be a relational schema consisting of a single relation $R[\mathbf{U}]$, and let $\mathbf{A}, \mathbf{B} \subseteq \mathbf{U}$. Assume further that $\text{Dep}(\mathbf{D})$ is generated by full dependencies, and that $\bowtie[\mathbf{A}, \mathbf{B}] \in \text{Dep}(\mathbf{D})$. Then the pair $\{\Pi_{\mathbf{A}}, \Pi_{\mathbf{B}}\}$ forms a meet-complementary pair iff the decomposition is dependency preserving. In this case, $\Pi_{\mathbf{A}} \wedge \Pi_{\mathbf{B}} = \Pi_{\mathbf{A} \cap \mathbf{B}}$.*

PROOF OUTLINE: It is straightforward to verify that condition (d) of 2.13 is satisfied for these projections. \square

3. Updates in the Order-Based Context

Using the context established in the previous section, the ideas of constant-complement update within an order-based framework are now established. These results extend those of Bancilhon and Spyrtos [BS81] to the order-based case. Because this extension is far from trivial, and because the formalism employed here is different than that of [BS81], all results are proven directly. The results presented here strengthen those of [BS81], not only in extension to the order-based case, but also in that meet complementation is used to provide a complete bijective correspondence between update strategies and complements.

3.1 Update strategies Let \mathbf{D} be a database schema. A *closed update family* for \mathbf{D} is an order-compatible equivalence relation U on $\text{LDB}(\mathbf{D})$. Think of a pair $(M_1, M_2) \in U$ as describing an update of the database from state M_1 to state M_2 . The equivalence-relation requirement implies that the identity update is always allowable, that all updates are reversible, and that updates may be composed.

Now let $\Gamma = (\mathbf{V}, \gamma)$ be a view of \mathbf{D} , and let T be a closed update family for \mathbf{V} . An *update strategy* for T with respect to U is a partial function $\rho : \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{V}) \rightarrow \text{LDB}(\mathbf{D})$ which has the eight properties listed below. The first five, (upt:1)-(upt:5), constitute an alternative formulation of those in the original work [BS81], while the last three, (upt:6)-(upt:8), are specific to the order-based context. The notation $\rho(M, N) \downarrow$ means that ρ is defined on the argument (M, N) . If a formula involving ρ appears in a formula, then it is implicitly assumed that it is defined.

- (upt:1) $\rho(M, N) \downarrow$ iff $(\gamma(M), N) \in T$.
- (upt:2) If $\rho(M, N) \downarrow$, then $(M, \rho(M, N)) \in U$ and $\gamma(\rho(M, N)) = N$.
- (upt:3) For every $M \in \text{LDB}(\mathbf{D})$, $\rho(M, \gamma(M)) = M$. [Identity updates are reflected as identities.]
- (upt:4) If $\rho(M, N) \downarrow$, then $\rho(\rho(M, N), \gamma(M)) = M$. [Every view update is globally reversible.]
- (upt:5) If $\rho(M, N_1) \downarrow$ and $\rho(\rho(M, N_1), N_2) \downarrow$, then $\rho(M, N_2) = \rho(\rho(M, N_1), N_2)$. [View update reflection is transitive.]

- (upt:6) If $\rho(M, N) \downarrow$ and $\gamma(M) \leq_{\mathbf{V}} N$, then $M \leq_{\mathbf{D}} \rho(M, N)$. [View update reflects order.]
- (upt:7) If $\rho(M_1, N_1) \downarrow$ with $M_1 \leq_{\mathbf{D}} \rho(M_1, N_1)$, then for all $M_2 \in \text{LDB}(\mathbf{D})$ with $M_1 \leq_{\mathbf{D}} M_2 \leq_{\mathbf{D}} \rho(M_1, N_1)$, there is an $N_2 \in \text{LDB}(\mathbf{V})$ with $\rho(M_1, N_2) = M_2$. [This condition is called *chain reflection*.]
- (upt:8) If $M_1, M_2 \in \text{LDB}(\mathbf{D})$ with $\gamma(M_1) \leq_{\mathbf{V}} \gamma(M_2)$ and $(\exists N_1, N_2 \in \text{LDB}(\mathbf{V}))(\rho(M_1, N_1) \leq_{\mathbf{D}} \rho(M_2, N_2))$, then $M_1 \leq_{\mathbf{D}} M_2$. [This condition is called *order inheritance*.]

The *induced update family* on \mathbf{D} is the smallest subset of U which will support the updates in T . It is denoted \equiv_{ρ} and is given by $\{(M_1, M_2) \in \text{LDB}(\mathbf{D}) \mid (\exists N \in \text{LDB}(\mathbf{V}))(\rho(M_1, N) = M_2)\}$.

While conditions (upt:7) and (upt:8) are somewhat technical, it will be seen in the following that they are exactly what is required to ensure that the update strategy defines an order-based complement, and conversely.

3.2 Notational convention For 3.3 through 3.6 below, \mathbf{D} will be a database schema, $\Gamma = (\mathbf{V}, \gamma)$ will be a view of \mathbf{D} , U and T will be closed update families for \mathbf{D} and \mathbf{V} , respectively, and ρ will be an update strategy for T with respect to U .

3.3 Proposition \equiv_{ρ} is an order-compatible equivalence relation for \mathbf{D} .

PROOF: The reflexivity, symmetry, and transitivity of \equiv_{ρ} follow from conditions (upt:3), (upt:4), and (upt:5), respectively. Order compatibility follows from (upt:7). \square

3.4 The complementary view for an update strategy The ρ -complement of Γ , denoted $\tilde{\Gamma}^{\rho} = (\tilde{\mathbf{V}}^{\rho}, \tilde{\gamma}^{\rho})$, is defined to have $\text{LDB}(\tilde{\mathbf{V}}^{\rho}) = \text{LDB}(\mathbf{D}) / \equiv_{\rho}$, with the morphism $\tilde{\gamma}^{\rho} : \mathbf{D} \rightarrow \tilde{\mathbf{V}}^{\rho}$ given by $M \mapsto [M]_{\equiv_{\rho}}$. The order $\leq_{\tilde{\mathbf{V}}^{\rho}}$ is just that which makes $\tilde{\gamma}^{\rho}$ an open surjection. More specifically, $[M_1]_{\equiv_{\rho}} \leq_{\tilde{\mathbf{V}}^{\rho}} [M_2]_{\equiv_{\rho}}$ iff there are $M_3 \in [M_1]_{\equiv_{\rho}}$, $M_4 \in [M_2]_{\equiv_{\rho}}$, with the property that $M_3 \leq_{\mathbf{D}} M_4$.

The *reflection of T along γ* is defined to be the relation $\text{Refl}_{\gamma}(T) = \{(M_1, M_2) \in \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{D}) \mid (\gamma(M_1), \gamma(M_2)) \in T\}$.

3.5 Proposition $\{\Gamma, \tilde{\Gamma}^{\rho}\}$ forms a meet-complementary pair, with meet $\theta_{\text{Refl}_{\gamma}(T)} = (\text{LDB}(\mathbf{D}) / \text{Refl}_{\gamma}(T), \theta_{\text{Refl}_{\gamma}(T)})$.

PROOF: First, it will be shown that $\gamma \otimes \tilde{\gamma}^{\rho} : \mathbf{D} \rightarrow \mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho}} \tilde{\mathbf{V}}^{\rho}$ is a bijection. It suffices to establish that it is an injection, i.e., that $\text{Congr}(\Gamma) \cap \text{Congr}(\tilde{\Gamma}^{\rho}) = \{(M, M) \mid M \in \text{LDB}(\mathbf{D})\}$. Let $(M_1, M_2) \in \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{D})$. If $(M_1, M_2) \in \text{Congr}(\tilde{\Gamma}^{\rho})$, then $(\exists N \in \text{LDB}(\mathbf{V}))(\rho(M_1, N) = M_2)$; in particular, $N = \gamma(M_2)$. Since $(M_1, M_2) \in \text{Congr}(\Gamma)$ iff $\gamma(M_1) = \gamma(M_2)$, condition (upt:3) ensures that $(M_1, M_2) \in \text{Congr}(\Gamma)$ iff $M_1 = M_2$.

Next, it will be shown that $\gamma \otimes \tilde{\gamma}^{\rho}$ is open. Let $M_1, M_2 \in \text{LDB}(\mathbf{D})$ be such that $([M_1]_{\Gamma}, [M_1]_{\tilde{\Gamma}^{\rho}}) \leq_{\Gamma \otimes \tilde{\Gamma}^{\rho}} ([M_2]_{\Gamma}, [M_2]_{\tilde{\Gamma}^{\rho}})$; i.e., $\gamma(M_1) \leq_{\mathbf{V}} \gamma(M_2)$ and $\tilde{\gamma}^{\rho}(M_1) \leq_{\tilde{\mathbf{V}}^{\rho}} \tilde{\gamma}^{\rho}(M_2)$. The latter inequality implies that there are $M_3 \in [M_1]_{\tilde{\Gamma}^{\rho}}$ and $M_4 \in [M_2]_{\tilde{\Gamma}^{\rho}}$ with the property that $M_3 \leq_{\mathbf{D}} M_4$. By the definition of $\text{Congr}(\tilde{\Gamma}^{\rho})$,

there are $N_1, N_2 \in \text{LDB}(\mathbf{V})$ with the property that $M_3 = \rho(M_1, N_1)$ and $M_4 = \rho(M_2, N_2)$, whence condition (upt:8) mandates that $M_1 \leq_{\mathbf{D}} M_2$, as required.

Finally, it will be established that $\text{Congr}(I) \circ \text{Congr}(\tilde{I}^\rho) = \text{Congr}(\tilde{I}^\rho) \circ \text{Congr}(I) = \text{Refl}_\gamma(T)$. It is immediate that $\text{Congr}(I) \subseteq \text{Refl}_\gamma(T)$ and $\text{Congr}(\tilde{I}^\rho) \subseteq \text{Refl}_\gamma(T)$. To establish the converse, let $(M_1, M_2) \in \text{Refl}_\gamma(T)$. Then there is a $N_1 \in \text{LDB}(\mathbf{D})$ with $\gamma(M_1) = \gamma(N_1)$ and $\rho(M_2, \gamma(M_1)) = N_1$. In other words, $(M_1, N_1) \in \text{Congr}(I)$ and $(N_1, M_2) \in \text{Congr}(\tilde{I}^\rho)$. Thus, $\text{Refl}_\gamma(T) \subseteq \text{Congr}(I) \circ \text{Congr}(\tilde{I}^\rho)$, and so $\text{Congr}(I) \circ \text{Congr}(\tilde{I}^\rho) \subseteq \text{Refl}_\gamma(T) \circ \text{Refl}_\gamma(T) = \text{Refl}_\gamma(T)$ i.e., $\text{Refl}_\gamma(T) = \text{Congr}(I) \circ \text{Congr}(\tilde{I}^\rho)$. Since $\text{Refl}_\gamma(T)$ is an equivalence relation, condition (b) of 2.13 is satisfied, and so $\{I, \tilde{I}^\rho\}$ is a meet-complementary pair with meet $\Theta_{\text{Refl}_\gamma(T)}$. \square

The following theorem is the order-based analog of the characterization first reported by Bancilhon and Spyratos in their seminal paper [BS81, Thm. 7.3].

3.6 Theorem – constant complement representation of update *For every $(N_1, N_2) \in T$ and $M \in \text{LDB}(\mathbf{D})$ with $\gamma(M) = N_1$, $\rho(M, N_2) = (\gamma \otimes \tilde{\gamma}^\rho)^{-1}(N_2, [M]_{\tilde{I}^\rho})$.*

PROOF: Follows from 3.5 and the definition of \tilde{I}^ρ . \square

Bancilhon and Spyratos also present a result which associates an update strategy for a closed update family T of $I_1 = (\mathbf{V}_1, \gamma_1)$ to each complement view I_2 which admits a translation of T (I_2 -*translatable* in their terminology) [BS81, Thm. 7.1]. However, their characterization does not provide conditions under which I_2 admits such a translation. The result 3.8 below makes this characterization precise in terms of the existence and character of the meet $I_1 \wedge I_2$.

3.7 The update strategy associated with a meet complement Let $\{I_1 = (\mathbf{V}_1, \gamma_1), I_2 = (\mathbf{V}_2, \gamma_2)\}$ be a meet-complementary pair of the schema \mathbf{D} .

- (a) Define $\text{UpdStr}\langle I_1, I_2 \rangle : \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{V}_1) \rightarrow \text{LDB}(\mathbf{D})$ by $(M, N) \mapsto (\gamma_1 \otimes \gamma_2)^{-1}(N, \gamma_2(M))$. $\text{UpdStr}\langle I_1, I_2 \rangle$ is called the *update strategy* for I_1 with respect to I_2 .
- (b) Define $\text{UpdFam}\langle I_1, I_2 \rangle = \{(N_1, N_2) \in \text{LDB}(\mathbf{V}_1) \times \text{LDB}(\mathbf{V}_1) \mid \lambda\langle I_1, I_1 \wedge I_2 \rangle(N_1) = \lambda\langle I_1, I_1 \wedge I_2 \rangle(N_2)\}$. $\text{UpdFam}\langle I_1, I_2 \rangle$ is called the *update family* induced by I_2 on I_1 .

3.8 Theorem *Let $\{I_1 = (\mathbf{V}_1, \gamma_1), I_2 = (\mathbf{V}_2, \gamma_2)\}$ be a meet-complementary pair of the schema \mathbf{D} . Then $\text{UpdFam}\langle I_1, I_2 \rangle$ is a closed update family for I_1 , and $\text{UpdStr}\langle I_1, I_2 \rangle$ is an update strategy for $\text{UpdFam}\langle I_1, I_2 \rangle$ with respect to $\text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{D})$.*

PROOF: First of all, $\text{UpdFam}\langle I_1, I_2 \rangle$ must be an order-compatible equivalence relation, since it is the congruence of a view. Next, it is completely straightforward to verify that $\text{UpdStr}\langle I_1, I_2 \rangle$ satisfies conditions (up1:1)-(upt:6) of 3.1. To show (upt:7), let $M_1, M_2 \in \text{LDB}(\mathbf{D})$, $N_1 \in \text{LDB}(\mathbf{V}_1)$, be such that

$\text{UpdStr}\langle I_1, I_2 \rangle \downarrow$ with $M_1 \leq_{\mathbf{D}} M_2 \leq_{\mathbf{D}} \text{UpdStr}\langle I_1, I_2 \rangle(M_1, N_1)$. Then $(\gamma_1 \otimes \gamma_2)(M_1)$ is of the form $(\gamma(M_1), P)$ for some $P \in \text{LDB}(\mathbf{V}_2)$, and $(\gamma_1 \otimes \gamma_2)(\text{UpdStr}\langle I_1, I_2 \rangle(M_1, N_1)) = (N, P)$. Thus, $(\gamma_1 \otimes \gamma_2)(M_2) = (\gamma(M_2), P)$, whence $(\gamma(M_1), \gamma(M_2)) \in T$.

Finally, to show that (upt:8) holds, let $M_1, M_2 \in \text{LDB}(\mathbf{D})$ with $\gamma_1(M_1) \leq_{\mathbf{V}} \gamma_2(M_2)$. Then $(\gamma_1 \otimes \gamma_2)^{-1}(M_1)$ is of the form $(\gamma_1(M_1), P_1)$ and $(\gamma_1 \otimes \gamma_2)^{-1}(M_2)$ is of the form $(\gamma_1(M_2), P_2)$ for some $P_1, P_2 \in \text{LDB}(\mathbf{V}_1)$. Now if $(\exists N_1, N_2 \in \text{LDB}(\mathbf{V}_1))(\text{UpdStr}\langle I_1, I_2 \rangle(M_1, N_1) \leq_{\mathbf{D}} \text{UpdStr}\langle I_1, I_2 \rangle(M_2, N_2))$, then $(\gamma_1 \otimes \gamma_2)^{-1}(\text{UpdStr}\langle I_1, I_2 \rangle(M_1, N_1)) = (Q_1, P_1)$ and $(\gamma_1 \otimes \gamma_2)^{-1}(\text{UpdStr}\langle I_1, I_2 \rangle(M_2, N_2)) = (Q_2, P_2)$, with $Q_1, Q_2 \in \text{LDB}(\mathbf{V}_1)$ and $(Q_1, P_1) \leq_{\mathbf{V}_1 \otimes_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2} (Q_2, P_2)$. In particular, $P_1 \leq_{\mathbf{V}_2} P_2$, whence $(\gamma_1(M_1), P_1) \leq_{\mathbf{V}_1 \otimes_{\gamma_1 \otimes \gamma_2} \mathbf{V}_2} (\gamma_1(M_2), P_2)$, and so $M_1 \leq_{\mathbf{D}} M_2$. \square

Note carefully the definition of $\text{UpdFam}\langle I_1, I_2 \rangle$ in the above. The allowable updates to I_1 under constant complement I_2 are precisely those which hold the meet $I_1 \wedge I_2$ constant. Thus, only $I_1 \wedge I_2$, and no further properties of I_2 , influence which updates are allowed. However, it may very well influence how those updates are reflected, as illustrated by the Π_B -complement views Π_{BC} and Π'_{BC} to Π_{AB} in the example schema \mathbf{E}_1 in 1.2 and 1.3. This uniqueness issue will be addressed in the next section.

To close, the following corollary, which identifies explicitly the natural association between meet complements and update strategies, is presented. It follows from the combination of 3.6 and 3.8.

3.9 Corollary *Let \mathbf{D} be a database schema, and let Γ be a view of \mathbf{D} . There is natural bijective correspondence between update strategies for the view Γ and (equivalence classes of) meet complements that view. Specifically:*

- (a) *For any update strategy ρ , $\text{UpdStr}\langle \Gamma, \tilde{\Gamma}^\rho \rangle = \rho$.*
- (b) *For any meet complement Γ_1 of Γ , $\tilde{\Gamma}^{\text{UpdStr}\langle \Gamma, \Gamma_1 \rangle} = \Gamma_1$.* \square

4. Uniqueness Results for Update Strategies

The necessary background having been established, the main results on the uniqueness of view update strategies and meet complements in the order-based context are presented. First, it is necessary to make precise the relationship between the order properties and types of updates.

4.1 Types of updates Let \mathbf{D} be a database schema, and let U be a closed update family for \mathbf{D} .

- (a) A pair $(M_1, M_2) \in U$ is called:
 - (i) a *formal insertion* with respect to U if $M_1 \leq_{\mathbf{D}} M_2$;
 - (ii) a *formal deletion* with respect to U if $M_2 \leq_{\mathbf{D}} M_1$;
 - (iii) an *order-based update* with respect to U if there exists a nonempty sequence $(N_1, N_2), (N_2, N_3), \dots, (N_{k-2}, N_{k-1}), (N_{k-1}, N_k)$ of elements of U with the properties that $N_1 = M_1$, $N_k = M_2$, and each pair (N_i, N_{i+1}) ,

$1 \leq i \leq k - 1$, is either a formal insertion or else a formal deletion with respect to U .


- (b) The update family U is called *order realizable* if every pair in U is an order-based update.

The main theorem of this paper states that within the constant-complement order-based framework, the reflection of an order-based update of a view to the base schema is unique, period. It does not depend upon the choice of complement, and it does not matter whether other updates are or are not order based. In short, the reflection of order-based updates is unique in a strong and global sense.

4.2 Theorem – uniqueness of reflection of order-based view updates

Let \mathbf{D} be a database schema, let $\Gamma = (\mathbf{V}, \gamma)$ be a view of \mathbf{D} , and let U and T be a closed update families for \mathbf{D} and \mathbf{V} , respectively. Let ρ_1 and ρ_2 be update strategies for T with respect to U . Then, for any $M \in \text{LDB}(\mathbf{D})$ and $N \in \text{LDB}(\mathbf{V})$ with $(\gamma(M), N) \in T$ an order-based update, it must be the case that $\rho_1(M, N) = \rho_2(M, N)$. In particular, if T is order realizable, then $\rho_1 = \rho_2$.


PROOF: The diagram to the right provides a general view of how updates are processed under the two different strategies. From top to bottom, the first leg of each path corresponds to translation from \mathbf{D} to the corresponding sub-direct complement representation, the second leg (μ_1 and μ_2) corresponds to the actual update, and the final leg corresponds to translation back to \mathbf{D} . The horizontal



arrows marked with “ \cong ” indicate the natural isomorphism between the two decompositions of the schema \mathbf{D} . Suppose that the state of the schema \mathbf{D} is M , and that it is desired to update the state of the view Γ from $\gamma(M)$ to N , with $(\gamma(M), N) \in T$. Relative to the diagram above and to the right, the two diagrams on the next page show how this update proceeds on elements. The diagram to the left below corresponds to ρ_1 or constant complement \tilde{I}^{ρ_1} , and the diagram to the right below to ρ_2 or constant complement \tilde{I}^{ρ_2} . Thus, in the diagram to the left, μ_1 corresponds to update under constant complement \tilde{I}^{ρ_1} , while μ_2 gives the corresponding update in the schema $\mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho_2}} \tilde{\mathbf{V}}^{\rho_2}$. In the diagram to the right, the situation is reversed; μ_2 gives the update under constant complement \tilde{I}^{ρ_2} , while μ_1 gives the corresponding update in the schema $\mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho_1}} \tilde{\mathbf{V}}^{\rho_1}$.

The goal is to show that, in these diagrams, $P_1 = \tilde{\gamma}^{\rho_1}(M)$, and $P_2 = \tilde{\gamma}^{\rho_2}(M)$, which in turn forces $\rho_1(M, N) = \rho_2(M, N)$. First of all, assume that the update $(\gamma(M), N) \in T$ is a formal insertion. Since $(\gamma(M), \tilde{\gamma}^{\rho_1}(M)) \leq_{\mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho_1}} \tilde{\mathbf{V}}^{\rho_1}} (N, \tilde{\gamma}^{\rho_1}(M))$, the order isomorphisms guarantee that $(\gamma(M), \tilde{\gamma}^{\rho_1}(M)) \leq_{\mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho_1}} \tilde{\mathbf{V}}^{\rho_1}} (N, P_2)$ as well. Thus $\rho_2(M, N) = (\gamma \otimes \tilde{\gamma}^{\rho_2})^{-1}((N, \tilde{\gamma}^{\rho_2}(M))) \leq_{\mathbf{V}_{\gamma \otimes \tilde{\gamma}^{\rho_2}} \tilde{\mathbf{V}}^{\rho_2}}$

$(\gamma \otimes \tilde{\gamma}^{\rho_2})^{-1}((N, P_2)) = (\gamma \otimes \tilde{\gamma}^{\rho_1})^{-1}((N, \tilde{\gamma}^{\rho_1}(M))) = \rho_1(M, N)$; i.e., $\rho_2(M, N) \leq_{\mathbf{D}} \rho_1(M, N)$. An analogous argument establishes that $\rho_1(M, N) \leq_{\mathbf{D}} \rho_2(M, N)$, so that $\rho_2(M, N) = \rho_1(M, N)$. The proof is now finished easily. A formal deletion is handled analogously, and a general order-based update is managed by gluing together the insertions and deletions of which it is composed. \square



Although it is primarily uniqueness of view update strategies which is the focus of this work, it is nonetheless worthwhile to note that in the case that the update family in the view is order realizable, the view complement defined by that update family is unique.

4.3 Corollary – uniqueness of meet complements *Let \mathbf{D} be a database schema, let $\Gamma = (\mathbf{V}, \gamma)$ be a view of \mathbf{D} , and T be an order-realizable closed update family for \mathbf{V} . Then there is at most one view of \mathbf{D} which is a Θ_T -complement of Γ .*

PROOF: Follows immediately from 3.9 and the above theorem. \square

4.4 Example Recall the schema \mathbf{E}_1 from 1.2, with relation schema $R[ABC]$, constrained by the fd $B \rightarrow C$. The view to be updated is Π_{AB} . From 2.16, it follows that Π_{BC} is a Π_B -complement of Π_{AB} , and the corresponding update family on Π_{AB} is precisely that which holds Π_B constant. This update family is order realizable, since there are no constraints on Π_{AB} . Any change can be realized by first inserting the new tuples, and then deleting the ones which are no longer wanted. Thus, in view of 4.2, the only update strategy in the sense of 3.1 which supports these updates is that which holds Π_{BC} constant. Furthermore, 4.3 above guarantees that Π_{BC} is the only order-based Π_B -complement of Π_{AB} . In particular, in view of 2.5, Π_{BC} is the only SPJR-view of \mathbf{E}_1 which is a Π_B -complement of Π_{AB} . In short, the theory has very strong things to say about \mathbf{E}_1 .

4.5 Example — using tuple ordering to force order realizability Now, revisit the example schema \mathbf{E}_2 from 1.3. It is identical to \mathbf{E}_1 , save that it is constrained by the additional fd $B \rightarrow A$. Π_{BC} is still a Π_B -complement of Π_{AB} ,

since the dependency $B \rightarrow A$ embeds in Π_{AB} . Thus, the update family in which Π_B is held constant is again that induced on Π_{AB} by the complement Π_{BC} . However, this time the states in Π_{AB} are constrained by the fd $B \rightarrow A$. This has the consequence that none of the allowable updates on Π_{AB} is order based; thus, under the natural order, the theory developed above has nothing to say either about updates to or about Π_B -complements of Π_{AB} .

Fortunately, in this situation, there is a way to establish results similar to those for \mathbf{E}_1 . The trick is to add additional order to the relational states. Let \preceq_A , \preceq_B and \preceq_C be partial orders on $\text{Dom}(A)$, $\text{Dom}(B)$ and $\text{Dom}(C)$, respectively. Define the ordering \preceq_{ABC} on tuples over ABC by $(a_0, b_0, c_0) \preceq_{ABC} (a_1, b_1, c_1)$ iff $a_0 \preceq_A a_1$, $b_0 \preceq_B b_1$, and $c_0 \preceq_C c_1$. Finally, extend the ordering \preceq_{ABC} to relations on ABC by $M_1 \preceq_{ABC} M_2$ iff $(\forall t_0 \in M_1)(\exists t_1 \in M_2)(t_0 \preceq_{ABC} t_1)$. It is easy to see that this induces a partial order on the relations of $R[ABC]$ which is stronger than $\leq_{\mathbf{E}_2}$ in general, and the same as $\leq_{\mathbf{E}_2}$ if each of the three orders \preceq_A , \preceq_B , and \preceq_C is flat (i.e., no ordering at all). Define a similar ordering, using the same base orders \preceq_A , \preceq_B , and \preceq_C , on the states of $R[AB]$ and $R[BC]$.

Now, for the specific problem at hand, let \preceq_A be any *total* order on $\text{Dom}(A)$, and let \preceq_B and \preceq_C be the flat orders on $\text{Dom}(B)$ and $\text{Dom}(C)$, respectively. Using the orders \preceq_{ABC} on $R[ABC]$, \preceq_{AB} on $R[AB]$, and \preceq_{BC} on $R[BC]$, every allowable update to Π_{AB} under constant complement Π_{BC} becomes order realizable. Furthermore, Π_{BC} is completely unchanged. Thus, the above results apply; update with constant complement Π_{BC} is the unique order-based update strategy on Π_{AB} with constant complement Π_B , and Π_{BC} is the unique order-based Π_B -complement of Π_{AB} .

5. Conclusions and Further Directions

It has been shown that the use of order as a central property of database schemata and morphisms leads to strong uniqueness results for view updates under the constant-complement strategy. In many cases, the natural order defined by inclusion is adequate for uniqueness; however, it has also been shown (4.5) that an “artificial” order may be used to achieve uniqueness in certain other situations.

In terms of future theoretical directions for this work, the most immediate one is to examine the circumstances under which update strategies may be combined to yield a single, larger strategy. Ultimately, it would be interesting and useful to identify circumstances under which a given view Γ_1 has a least meet complement Γ_2 , in the sense that for any other meet complement Γ_3 , $\Gamma_1 \wedge \Gamma_2 \leq \Gamma_1 \wedge \Gamma_3$. This would imply that Γ_1 has a largest update family which can be supported via the constant-complement strategy. Other theoretical directions to consider include a more systematic treatment of the kind of artificial orders used in 4.5, as well as study of the complexity of these forms of updates under specific data models.

In terms of more practical follow-up work, certainly a closer examination of how this work might be applied to data models other than the relational is high on the list. However, it seems that there is a more fundamental issue which

must be addressed first. Most of the research on the view update problem has focused upon open strategies, with the goal to support the largest possible family of updates, even at the expense of violating the isolation principles which are implicit in the constant-complement strategy. Indeed, very little work has been done which builds upon the constant complement strategy, and even in the few cases in which it has, the axioms have been weakened so that critical properties, such as reversibility, have been sacrificed [GPZ88]. In the opinion of the author, this situation exists because in traditional database design, the main schema is established first, and only afterwards are views fitted to it. Therefore, *view-centered schema design*, a research direction in which schema design is focused upon the support of views, including view updates based upon the constant-complement strategy, is proposed. The conjecture is that, with the judicious use of null values, most view updates can be accommodated within the constant-complement strategy. The most natural framework in which to pursue these ideas would be within *HERM*, the *Higher-order Entity-Relationship Model* [Tha00]. Since *HERM* is a meta model, conclusions about most modern data models could be extracted from a study based upon it.

References

- [AHV95] Abiteboul, S., Hull, R., and Vianu, V., *Foundations of Databases*, Addison-Wesley, 1995.
- [BS81] Bancilhon, F. and Spyratos, N., "Update semantics of relational views," *ACM Transactions on Database Systems*, **6**(1981), pp. 557–575.
- [Bel00] Bellahsene, Z., "Updates and object-generating views in ODBS," *Data & Knowledge Engr.*, **34**(2000), pp. 125–163.
- [BL98] Bentayeb, F. and Laurent, D., "View updates translations in relational databases," in: *Proc. DEXA '98, Vienna, Sept. 24-28, 1998*, pp. 322–331, 1998.
- [DP90] Davey, B. A. and Priestly, H. A., *Introduction to Lattices and Order*, Cambridge University Press, 1990.
- [GPZ88] Gottlob, G., Paolini, P., and Zicari, R., "Properties and update semantics of consistent views," *ACM Trans. Database Systems*, **13**(1988), pp. 486–524.
- [Heg90] Hegner, S. J., "Foundations of canonical update support for closed database views," in: *Proc. ICDT'90, Paris, Dec. 1990*, Springer-Verlag, 1990.
- [Heg93] Hegner, S. J., "Characterization of desirable properties of general database decompositions," *Ann. Math. Art. Intell.*, **7**(1993), pp. 129–195.
- [Heg94] Hegner, S. J., "Unique complements and decompositions of database schemata," *J. Comput. System Sci.*, **48**(1994), pp. 9–57.
- [Hul84] Hull, R., "Finitely specifiable implicational dependency families," *Journal of the Association for Computing Machinery*, **31**(1984), pp. 210–226.
- [Mac98] Mac Lane, S., *Categories for the Working Mathematician*, Springer-Verlag, second edition, 1998.
- [MT99] Mayol, E. and Teniente, E., "A survey of current methods for integrity constraint maintenance and view updating," in: *Proc. ER '99 Workshops, Paris, Nov. 15-18, 1999*, Springer-Verlag, 1999.
- [Mon76] Monk, J. D., *Mathematical Logic*, Springer-Verlag, 1976.
- [Tha00] Thalheim, B., *Entity-Relationship Modeling*, Springer-Verlag, 2000.