

Characterization of Optimal Complements of Database Views Defined by Projection

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. A complement to a database view Γ is a second view Γ' which provides the information necessary to reconstruct the entire state of the main schema. View complementation is central in situations in which a view is to be updated, since the complement Γ' embodies the information not contained in Γ . In general, Γ may have many complements. In this work, an approach to identifying and constructing optimal relational complements for relational views defined by projections, including not only single projections but also sets of projections, is initiated. The approach is based upon the idea of identifying when the main schema has a governing join dependency; that is, a join dependency which implies all others. Four distinct classes of such governing dependencies are identified, corresponding to ordinary complements as well as three distinct types of dependency-preserving complements.

1 Introduction

A view of a database schema embodies partial, but in general not total, information about the state of the schema. To recover the remaining information, a second view, called a complement, may be used. To illustrate the surrounding ideas, a set of four closely related schemata and some of their views will be considered. Table 1 provides a summary of the associated notation and of their properties. Each of the four schemata has the same, single relation symbol

Schema	Relation	Constraints	Views for $\mathbf{W} \subseteq ABCD$
\mathbf{E}_0	$R[ABCD]$	$\mathcal{F}_0 = \{C \rightarrow D\}$	$\Pi_{\mathbf{W}}^{\mathbf{E}_0} = (\mathbf{E}_0^{\mathbf{W}}, \pi_{\mathbf{W}}^{\mathbf{E}_0})$
\mathbf{E}_1	$R[ABCD]$	$\mathcal{F}_1 = \{B \rightarrow D, C \rightarrow D\}$	$\Pi_{\mathbf{W}}^{\mathbf{E}_1} = (\mathbf{E}_1^{\mathbf{W}}, \pi_{\mathbf{W}}^{\mathbf{E}_1})$
\mathbf{E}_2	$R[ABCD]$	$\mathcal{F}_2 = \{B \rightarrow C, C \rightarrow D\}$	$\Pi_{\mathbf{W}}^{\mathbf{E}_2} = (\mathbf{E}_2^{\mathbf{W}}, \pi_{\mathbf{W}}^{\mathbf{E}_2})$
\mathbf{E}_3	$R[ABCD]$	$\mathcal{F}_3 = \{B \rightarrow C, C \rightarrow D, D \rightarrow B\}$	$\Pi_{\mathbf{W}}^{\mathbf{E}_3} = (\mathbf{E}_3^{\mathbf{W}}, \pi_{\mathbf{W}}^{\mathbf{E}_3})$

Table 1. Example schemata and their views

$R[ABCD]$; they differ only in the set of constraints which govern them, which in each case consists of a set of functional dependencies (FDs). Sets of attributes are

represented in the conventional relational format; $ABCD$ denotes $\{A, B, C, D\}$. The view $\Pi_{\mathbf{W}}^{\mathbf{E}_i}$ recaptures the projection of the relation $R[ABCD]$ of \mathbf{E}_i onto the subset \mathbf{W} of $ABCD$. The relation symbol of the view schema $\mathbf{E}_i^{\mathbf{W}}$ is denoted $R_{\mathbf{W}}[\mathbf{W}]$. The symbol $\pi_{ABC}^{\mathbf{E}_i}$ denotes the associated view morphism (i.e., the projection mapping itself). For example, $\Pi_{ABC}^{\mathbf{E}_0} = (\mathbf{E}_0^{ABC}, \pi_{ABC}^{\mathbf{E}_0})$ is the projection of $R[ABCD]$ of \mathbf{E}_0 onto ABC ; the relation symbol of this view is $R_{ABC}[ABC]$, and the view morphism is $\pi_{ABC}^{\mathbf{E}_0}$ projects $R[ABCD]$ onto $R_{ABC}[ABC]$.

In this simple context, for $\Pi_{\mathbf{W}_2}^{\mathbf{E}_i}$ to be a complement of $\Pi_{\mathbf{W}_1}^{\mathbf{E}_i}$, it is necessary and sufficient that the join dependency (JD) $\bowtie[\mathbf{W}_1, \mathbf{W}_2]$ hold. This follows immediately from the classical result [24, Thm. 1], which establishes further that, in the context of FDs, this JD holds iff at least one of the FDs $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_1$, $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_2$ holds; i.e., iff $\mathbf{W}_1 \cap \mathbf{W}_2$ is a key for at least one of the projections.

In the setting of \mathbf{E}_0 , it is thus the case that $\Pi_{CD}^{\mathbf{E}_0}$ is a complement of $\Pi_{ABC}^{\mathbf{E}_0}$. Furthermore, amongst projections, it is a minimal complement, in the sense that no projection whose attribute set is a proper subset of CD is also a complement. Of course, there are other complements; indeed $\Pi_{BCD}^{\mathbf{E}_0}$, $\Pi_{ACD}^{\mathbf{E}_0}$, and even $\Pi_{ABCD}^{\mathbf{E}_0}$ are each complements of $\Pi_{ABC}^{\mathbf{E}_0}$ as well. However, these are all larger than $\Pi_{CD}^{\mathbf{E}_0}$ in that they recapture more attributes and hence more information about the state of \mathbf{E}_0 .

To see why this minimality is important, suppose that it is desired to update the state of $\Pi_{ABC}^{\mathbf{E}_0}$. In general, there are many ways to reflect such an update back to \mathbf{E}_0 . For example, let the state of \mathbf{E}_0 be $M_{00} = \{R(a_1, b_1, c_1, d_1), R(a_2, b_2, c_2, d_2)\}$, so that the state of \mathbf{E}_0^{ABC} is $N_{00} = \{R(a_1, b_1, c_1), R(a_2, b_2, c_2)\}$. Suppose further that the desired new state of \mathbf{E}_0^{ABC} is $N_{01} = \{R(a_3, b_1, c_1), R(a_2, b_2, c_2)\}$. Two possibilities for translation of this view update to the main schema \mathbf{E}_0 include $M_{01} = \{R(a_3, b_1, c_1, d_1), R(a_2, b_2, c_2, d_2)\}$ and $M_{01'} = \{R(a_3, b_1, c_1, d_3), R(a_2, b_2, c_2, d_2)\}$. Of these two, M_{01} seems the more natural, in that it does not alter the CD projections of the tuples. $M_{01'}$ makes the additional and rather arbitrary change of d_1 to d_3 . This preference is formalized by the classical constant-complement strategy [3]. The update (M_{00}, M_{01}) on the main schema is the only one which is a reflection of the update (N_{00}, N_{01}) on $\Pi_{ABC}^{\mathbf{E}_0}$ and which holds the state of the complement $\Pi_{CD}^{\mathbf{E}_0}$ constant. Indeed, since the states of $\Pi_{ABC}^{\mathbf{E}_0}$ and the complement $\Pi_{CD}^{\mathbf{E}_0}$ together determine the state of \mathbf{E}_0 , there can be only one such reflection. The appeal of the constant-complement strategy is that the reflection is limited to that part of the state of the main schema which is determined by the view to be updated. The rest of the state of the main schema, as captured by the complement, is left unchanged. Although a classical approach, the constant-complement strategy has seen renewed interest in recent years [20] [21] [14], and is also related to other modern approaches to view updates, such as those based upon lenses [7] [12].

It was noted already in the initial seminal paper on the subject that a view almost always has many complements [3, Thm. 4.4]. A question which arises immediately is whether by choosing a complement other than $\Pi_{CD}^{\mathbf{E}_0}$, a reflection other than (M_{00}, M_{01}) , such as (M_{00}, M'_{01}) , is possible for the view update

(N_{00}, N_{01}) on $\Pi_{ABC}^{\mathbf{E}_0}$. The answer is a qualified yes, but a rather “unnatural” complement is required. An example of such an unnatural complement for a similar schema and view may be found in [14, Sec. 1.3]. However, it has been shown that for a very wide range of schemata and views which occur in practice, in which the schemata have a natural order structure and the morphisms are order preserving, the reflection of the view update is in fact independent of the choice of complement [14, Thm. 4.4], [17, Thm. 53]. Under these conditions, (M_{00}, M_{01}) is the only reasonable translation of (N_{00}, N_{01}) from initial state M_{00} . Nevertheless, the set of view updates which are allowed under the constant-complement strategy does in fact depend upon the choice of complement. As noted above, $\Pi_{BCD}^{\mathbf{E}_0}$, the projection onto $R[BCD]$, is also a complement of $\Pi_{ABC}^{\mathbf{E}_0}$, as is $\Pi_{ACD}^{\mathbf{E}_0}$. The view update (N_{00}, N_{01}) is also supported via constant complement $\Pi_{BCD}^{\mathbf{E}_0}$, with the same reflection, but not with constant complement $\Pi_{ACD}^{\mathbf{E}_0}$. On the other hand, with $N_{01'} = \{R(a_1, b_3, c_1), R(a_2, b_2, c_2)\}$, the view update $(N_{00}, N_{01'})$ is supported with constant complement $\Pi_{CD}^{\mathbf{E}_0}$, as well as with constant complement $\Pi_{ACD}^{\mathbf{E}_0}$, via the reflection (M_{00}, M_{01}'') with $M_{01}'' = \{R(a_3, b_3, c_1, d_1), R(a_2, b_2, c_2, d_2)\}$, but not with constant complement $\Pi_{BCD}^{\mathbf{E}_0}$. The reason is clear. $\Pi_{CD}^{\mathbf{E}_0}$ is a smaller complement than either $\Pi_{BCD}^{\mathbf{E}_0}$ or $\Pi_{ACD}^{\mathbf{E}_0}$, and the smaller the complement, the less that must be held constant, and so the larger the set of possible view updates. Thus, in order to maximize the set of view updates which are supported, the complement must be minimized.

It is easy to see that $\Pi_{CD}^{\mathbf{E}_0}$ is the unique minimal complement to $\Pi_{ABC}^{\mathbf{E}_0}$ amongst projections; it is thus called an *optimal* projective complement. Unfortunately, there are simple examples for which there are incomparable minimal complements, so no such optimal complement exists. The schema \mathbf{E}_1 provides one such setting, in which both $\Pi_{CD}^{\mathbf{E}_1}$ and $\Pi_{BD}^{\mathbf{E}_1}$ are minimal complements to $\Pi_{ABC}^{\mathbf{E}_1}$ within the family of projections. That both $\Pi_{CD}^{\mathbf{E}_1}$ and $\Pi_{BD}^{\mathbf{E}_1}$ are complements of $\Pi_{ABC}^{\mathbf{E}_1}$ follows from the classical result [24, Thm. 1] noted above. Furthermore, since the setting is completely symmetric with respect to attributes B and C , there is no way to prefer one over the other. Thus, it is not possible, in general, to expect an optimal projective complement, even in the basic setting of FDs and views which are projections. Consequently, there is no complement which supports all updates which any other complement does. For example, if the state of \mathbf{E}_1 is M_{00} as given above, then the update on $\Pi_{ABC}^{\mathbf{E}_1}$ from $N_{10} = \{R(a_1, b_1, c_1), R(a_2, b_2, c_2)\}$, to $N_{11} = \{R(a_1, b_3, c_1), R(a_2, b_2, c_2)\}$ is supported with constant complement $\Pi_{CD}^{\mathbf{E}_1}$ but not with constant complement $\Pi_{BD}^{\mathbf{E}_1}$, while the update from $N_{10'}$ = $\{R(a_1, b_1, c_1), R(a_2, b_2, c_2)\}$, to $N_{11'}$ = $\{R(a_1, b_1, c_3), R(a_2, b_2, c_2)\}$ is supported with constant complement $\Pi_{BD}^{\mathbf{E}_1}$ but not with constant complement $\Pi_{CD}^{\mathbf{E}_1}$.

There is a further, central issue. A cover (i.e., a logical representation) of the constraints $\mathcal{F}_0 = \{C \rightarrow D\}$ of \mathbf{E}_0 embeds into the pair $\{\Pi_{ABC}^{\mathbf{E}_0}, \Pi_{CD}^{\mathbf{E}_0}\}$; indeed, $C \rightarrow D$ itself embeds into $\Pi_{CD}^{\mathbf{E}_0}$. This implies that whether an update to $\Pi_{ABC}^{\mathbf{E}_0}$ is supported by constant complement $\Pi_{CD}^{\mathbf{E}_0}$ is determined by the state of \mathbf{E}_0^{ABC} alone; it is not necessary to know the state of \mathbf{E}_0^{CD} . In the classical setting, it is

said that $\Pi_{ABC}^{\mathbf{E}_0}$ and $\Pi_{CD}^{\mathbf{E}_0}$ are *independent* [24, Thm. 2]. The updates to $\Pi_{ABC}^{\mathbf{E}_0}$ which are allowed with constant complement $\Pi_{CD}^{\mathbf{E}_0}$ are exactly those which keep the view $\Pi_C^{\mathbf{E}_0}$ defined by the common column of $\Pi_{ABC}^{\mathbf{E}_0}$ and $\Pi_{CD}^{\mathbf{E}_0}$ constant. In this case, $\{\Pi_{ABC}^{\mathbf{E}_0}, \Pi_{CD}^{\mathbf{E}_0}\}$ is said to form a *meet-complementary pair* [14, Def. 2.12], with $\Pi_C^{\mathbf{E}_0}$ their *meet*.

On the other hand, in the context of \mathbf{E}_1 , neither $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{CD}^{\mathbf{E}_1}\}$ nor $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{BD}^{\mathbf{E}_1}\}$ forms a meet-complementary pair. Indeed, $B \rightarrow D$ cannot be embedded into the first pair while $C \rightarrow D$ cannot be embedded into the second. Consequently, it cannot be determined in general whether a given update to $\Pi_{ABC}^{\mathbf{E}_1}$ is supported via constant complement without knowing the state of the complementary view. For example, the update on $\Pi_{ABC}^{\mathbf{E}_1}$ from N_{10} to $N_{12} = \{R(a_1, b_2, c_1), R(a_2, b_2, c_2)\}$ is supported with constant complement $\Pi_{CD}^{\mathbf{E}_1}$ if the state of \mathbf{E}_1 is $M_{11} = \{R(a_1, b_1, c_1, d_1), R(a_2, b_2, c_2, d_1)\}$, but not if it is M_{00} . Thus, requiring a meet complement for the constant-complement strategy is critical. Interestingly, $\Pi_{ABC}^{\mathbf{E}_1}$ does have an optimal projective meet complement, namely, $\Pi_{BCD}^{\mathbf{E}_1}$, which is obtained by combining $\Pi_{BD}^{\mathbf{E}_1}$ and $\Pi_{CD}^{\mathbf{E}_1}$.

The schema \mathbf{E}_2 illustrates yet another possibility, in which both $\Pi_{BD}^{\mathbf{E}_2}$ and $\Pi_{CD}^{\mathbf{E}_2}$ are minimal projective complements of $\Pi_{ABC}^{\mathbf{E}_2}$, yet only $\{\Pi_{ABC}^{\mathbf{E}_2}, \Pi_{CD}^{\mathbf{E}_2}\}$ forms a meet-complementary pair. Thus, it is possible for a schema to have several minimal projective complements, of which only one is a meet complement.

In view of these three examples, it might be conjectured that an optimal meet complement always exists, at least in the context of projections of a universal schema which is constrained by FDs. That this is not the case is illustrated by the situation surrounding the schema \mathbf{E}_3 . In that context, both $\Pi_{CD}^{\mathbf{E}_3}$ and $\Pi_{BD}^{\mathbf{E}_3}$ are minimal meet complements of $\Pi_{ABC}^{\mathbf{E}_3}$. Indeed, \mathcal{F}_3 is equivalent to $\mathcal{F}'_3 = \{B \rightarrow C, C \rightarrow B, C \rightarrow D, D \rightarrow C\}$, which embeds into $\{\Pi_{ABC}^{\mathbf{E}_3}, \Pi_{CD}^{\mathbf{E}_3}\}$, as well as equivalent to $\mathcal{F}''_3 = \{B \rightarrow C, C \rightarrow B, B \rightarrow D, D \rightarrow B\}$, which embeds into $\{\Pi_{ABC}^{\mathbf{E}_3}, \Pi_{BD}^{\mathbf{E}_3}\}$. Consequently, in this case, there is no optimal meet complement.

In summary, it the case that, even in the context of simple universal relational schemata constrained only by FDs, there are many possibilities regarding the existence of optimal complements and optimal meet complements. The goal of this paper is to address the issues surrounding such optimal complements in a systematic fashion.

The organization of the paper is as follows. Section 2 provides an overview of the basic terminology and notation required for the following sections, Section 3 comprises the main part of the paper. Four fundamental notions of optimal complement are developed, first for the case of views on a universal relational schema defined by single projections, and then for the more general context of views defined by sets of projections. Finally, Sec. 4 provides conclusions and sketches some further directions.

2 Database Schemata, Views, and Complements

Although the main focus in this paper is upon views defined by projection on the classical relational model, the underlying theory is not based upon any particular data model. Rather, the set-based model described in [14] and [15] is more than adequate. It is in fact simpler to present the underlying theory in a general context, and only afterwards to specialize it to the relational context. In this section, the key ideas of the underlying theoretical model are outlined. Much of the material is based upon earlier papers, such as [14], although the classification scheme of Definition 2.5 is new.

Summary 2.1 (Database schemata and views). In the underlying theory of this work, a database schema \mathbf{D} is just a set. This set, denoted $\text{LDB}(\mathbf{D})$, consists of the *legal databases* of \mathbf{D} . Thus, a database schema is modelled by its instances alone; constraints, schema structure, and the like are not represented explicitly.

A *morphism* $f : \mathbf{D}_1 \rightarrow \mathbf{D}_2$ of database schemata is a function $f : \text{LDB}(\mathbf{D}_1) \rightarrow \text{LDB}(\mathbf{D}_2)$. A *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 a surjective database morphism. In the relational context, a morphism is usually defined using the relational algebra or calculus, but the abstraction to a function on states is adequate for the framework developed here.

The *congruence* of Γ is the equivalence relation on $\text{LDB}(\mathbf{D})$ defined by $(M_1, M_2) \in \text{Congr}(\Gamma)$ iff $\gamma(M_1) = \gamma(M_2)$. Views with identical congruences are considered “abstractly the same” for the purposes of this work. The views $\Gamma_1 = (\mathbf{V}_1, \gamma_1)$ and $\Gamma_2 = (\mathbf{V}_2, \gamma_2)$ of \mathbf{D} are *equivalent* or (*congruence*) *isomorphic* iff $\text{Congr}(\Gamma_1) = \text{Congr}(\Gamma_2)$. In general, views with vastly different representations may have the same congruence. However, in the context of the examples of this paper, which are all relational and defined by projections, unless two views are “obviously” equivalent, they will generally have distinct congruences. Such equivalences occur in particular in the context of views defined by sets of projections, where a view with a single relation is equivalent to one with two relations connected by a join dependency. See Definition 3.22 and Example 3.24 for examples. The bottom line is that in this paper $\text{Congr}(\Gamma)$ may be regarded as identifying the view Γ up to an obvious relational equivalence.

A congruence on $\text{LDB}(\mathbf{D})$ may be represented by the partition which it induces [22, Sec. 1]. The partition of $\text{LDB}(\mathbf{D})$ induced by $\text{Congr}(\Gamma)$ is denoted $\text{Partition}(\text{Congr}(\Gamma))$.

There is a natural partial order on the equivalence classes of the views of \mathbf{D} , given by $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1$ iff $\text{Congr}(\Gamma_1) \subseteq \text{Congr}(\Gamma_2)$. If $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1$, then every pair $\{M_1, M_2\} \subseteq \text{LDB}(\mathbf{D})$ which is distinguished by Γ_2 (in the precise sense that $\gamma_2(M_1) \neq \gamma_2(M_2)$) is also distinguished by Γ_1 . Less formally, Γ_1 preserves at least as much information about the state of \mathbf{D} as does Γ_2 . The notation $\Gamma_2 \prec_{\mathbf{D}} \Gamma_1$ means that $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1$ holds but $\Gamma_1 \preceq_{\mathbf{D}} \Gamma_2$ does not.

It is easy to see that $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1$ iff every block of $\text{Partition}(\text{Congr}(\Gamma_2))$ is the union of blocks of $\text{Partition}(\text{Congr}(\Gamma_1))$. If $\Gamma_2 \preceq_{\mathbf{D}} \Gamma_1$, there is thus a unique view morphism $\lambda(\Gamma_1, \Gamma_2) : \Gamma_1 \rightarrow \Gamma_2$, called the *relativization morphism* from Γ_1 to

Γ_2 , which sends a block of $\text{Partition}(\text{Congr}(\Gamma_1))$ to the block which contains it in $\text{Partition}(\text{Congr}(\Gamma_2))$.

Summary 2.2 (Complements and Optimality). The pair $\{\Gamma_1 = (\mathbf{V}_1, \gamma_1), \Gamma_2 = (\mathbf{V}_2, \gamma_2)\}$ of views of the database schema \mathbf{D} is said to be *complementary* if the *decomposition mapping* $\gamma_1 \times \gamma_2 : \text{LDB}(\mathbf{D}) \rightarrow \text{LDB}(\mathbf{V}_1) \times \text{LDB}(\mathbf{V}_2)$ given on elements by $M \mapsto (\gamma_1(M), \gamma_2(M))$ is injective. This is nothing more than a repackaging of the classical idea of a lossless decomposition. Γ_2 is said to be a *complement* of Γ_1 , and vice versa. The inverse $(\gamma_1 \times \gamma_2)^{-1} : (\gamma_1 \times \gamma_2)(\text{LDB}(\mathbf{D})) \rightarrow \text{LDB}(\mathbf{D})$ of the decomposition mapping is called the *reconstruction mapping*.

Given a set \mathcal{V} of views of \mathbf{D} , call a complement $\Gamma_2 \in \mathcal{V}$ of Γ_1 *minimal* relative to \mathcal{V} if there is no complement $\Gamma'_2 \in \mathcal{V}$ with the property that $\Gamma'_2 \prec_{\mathbf{D}} \Gamma_2$. Call Γ_2 *optimal* relative to \mathcal{V} if for every other complement $\Gamma'_2 \in \mathcal{V}$, it must be the case that $\Gamma_2 \preceq_{\mathbf{D}} \Gamma'_2$. Less formally, a complement is minimal relative to \mathcal{V} if no other complement in \mathcal{V} is smaller, and it is optimal relative to \mathcal{V} if every other complement in \mathcal{V} is at least as large. Thus, there may be many incomparable minimal complements, but there is at most one optimal complement, or, better put, all optimal complements must have the same congruence.

If each member of $\{\Gamma_1, \Gamma_2\}$ is a minimal (resp. optimal) complement of the other relative to \mathcal{V} , then it is called a *minimal complementary pair* (resp. *optimal complementary pair*) relative to \mathcal{V} .

Summary 2.3 (Fully commuting pairs and meet complements). The pair $\{\Gamma_1 = (\mathbf{V}_1, \gamma_1), \Gamma_2 = (\mathbf{V}_2, \gamma_2)\}$ of views of \mathbf{D} is said to be *fully commuting* [14, Def. 2.12] if $\text{Congr}(\Gamma_1) \circ \text{Congr}(\Gamma_2) = \text{Congr}(\Gamma_2) \circ \text{Congr}(\Gamma_1)$, with “ \circ ” denoting the ordinary composition of relations. If $\{\Gamma_1, \Gamma_2\}$ furthermore forms a complementary pair, then it is called *meet-complementary pair*, and Γ_1 and Γ_2 are called *meet complements* of one another. In the case that $\{\Gamma_1, \Gamma_2\}$ forms a fully commuting pair, the view whose congruence is $\text{Congr}(\Gamma_1) \circ \text{Congr}(\Gamma_2)$ is called the *meet* of Γ_1 and Γ_2 , and is denoted $\Gamma_1 \wedge \Gamma_2$.

The importance of full commutativity is that it recaptures dependency preservation in an abstract fashion. Specifically, if $\{\Gamma_1, \Gamma_2\}$ forms a fully commuting pair, with $N_i \in \text{LDB}(\mathbf{V}_i)$ for $i \in \{1, 2\}$, then there is an $M \in \text{LDB}(\mathbf{D})$ with $\gamma_i(M) = N_i$ for $i \in \{1, 2\}$ iff $\lambda\langle \Gamma_1, \Gamma_1 \wedge \Gamma_2 \rangle(N_1) = \lambda\langle \Gamma_2, \Gamma_1 \wedge \Gamma_2 \rangle(N_2)$. If $\{\Gamma_1, \Gamma_2\}$ is furthermore a complementary pair, then this M is necessarily unique [14, Thm. 2.14].

In the context of projections on universal schemata constrained by full dependencies (including FDs in particular), the meet, when it exists, is always the projection on the columns common to the two views [14, Prop. 2.17]. For example, in the context of \mathbf{E}_0 of Sec. 1, $\Pi_{ABC}^{\mathbf{E}_0} \wedge \Pi_{CD}^{\mathbf{E}_0} = \Pi_C^{\mathbf{E}_0}$.

Minimal and optimal meet complements are defined in the natural way. Let Γ be a view of the \mathbf{D} , and let \mathcal{V} be any set of views of \mathbf{D} . Define $\text{MeetSet}(\mathcal{V}; \Gamma)$ to be the subset of \mathcal{V} which identifies just those views which are meet complements of Γ . A *minimal* (resp. *optimal*) meet complement of the view Γ relative to the set \mathcal{V} of views is a minimal (resp. optimal) complement relative to $\text{MeetSet}(\mathcal{V}; \Gamma)$.

Summary 2.4 (The constant-complement approach to view update).

While the results of this paper are presented and may be understood without any reference to the constant-complement view-update strategy, the latter forms the motivation for exploring the notion of optimal complement in the first place. Therefore, a brief sketch of this approach is presented. For details, consult [3] and [14].

An *update* on the database schema \mathbf{D} is a pair $(M_1, M_2) \in \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{D})$. M_1 is the current state, and M_2 the new state. To describe the situation surrounding an update request on the view $\Gamma = (\mathbf{V}, \gamma)$, it is sufficient to specify the current state M_1 of the main schema and the desired new state N_2 of the view schema \mathbf{V} . The current state of the view can be computed as $\gamma(M_1)$; it is only the new state M_2 of the main schema (subject to $N_2 = \gamma(M_2)$) which must be obtained from an update strategy. Formally, an *update request* from Γ to \mathbf{D} is a pair (M_1, N_2) in which $M_1 \in \text{LDB}(\mathbf{D})$ (the old state of the main schema) and $N_2 \in \text{LDB}(\mathbf{V})$ (the new state of the view schema). A *realization* of (M_1, N_2) along Γ is an update (M_1, M_2) on \mathbf{D} with the property that $\gamma(M_2) = N_2$. The update (M_1, M_2) is called a *reflection* (or *translation*) of the view update $(\gamma(M_1), N_2)$. Thus, the realization (M_1, M_2) tells how to reflect the view update $(\gamma(M_1), N_2)$ to the main schema \mathbf{D} when the state of \mathbf{D} is M_1 .

Given a second view $\Gamma' = (\mathbf{V}', \gamma')$ of \mathbf{D} which is a complement of Γ , the realization (M_1, M_2) of (M_1, N_2) is *defined by constant complement Γ'* if $\gamma'(M_1) = \gamma'(M_2)$. It is easy to see that there can be only one realization of (M_1, N_2) which is defined by constant complement Γ' , given explicitly by $(M_1, (\gamma \times \gamma')^{-1}(N_2, \gamma'(M_1)))$ if and only if the latter exists, so the translation truly is defined by Γ' . The family $S \subseteq \text{LDB}(\mathbf{D}) \times \text{LDB}(\mathbf{V})$ of update requests on Γ is *defined by constant complement Γ'* if for every $(M_1, N_2) \in S$, $(M_1, (\gamma \times \gamma')^{-1}(N_2, \gamma'(M_1)))$ exists and is a realization of (M_1, N_2) . Update reflections which are defined by constant complement have many desirable properties; see in particular [14, Sec. 1] for an overview with examples.

If the pair $\{\Gamma, \Gamma'\}$ is meet complementary, an important additional property is guaranteed; namely, that whether an update to the state of Γ is allowed with constant complement Γ' is independent of the state of Γ' . Specifically, the view update (N, N') to Γ is allowed iff $\lambda\langle \Gamma, \Gamma \wedge \Gamma' \rangle(N) = \lambda\langle \Gamma, \Gamma \wedge \Gamma' \rangle(N')$; that is, iff the update keeps the meet view constant. Examples have already been given in introduction. This property is crucial because it provides a form of closure to Γ ; all allowable update operations may be identified locally.

Definition 2.5 (Types of complements and complementary pairs). As observed in the discussion surrounding the examples of Sec. 1, there are several distinct levels of optimality of a complementary view which arise when combining notions of ordinary and meet complements. Because it is precisely such notions of optimality which will be addressed in Sec. 3, it is important to identify and label these systematically. Relative to a base set \mathcal{V} of views of the schema \mathbf{D} , there are four principal types of complements of a given view Γ , defined as follows.

Type 0: An optimal complement of Γ of *type 0* relative to \mathcal{V} is an ordinary optimal complement of Γ relative to \mathcal{V} , as described in Summary 2.2 above.

Type 1: An optimal complement of Γ of *type 1* relative to \mathcal{V} is an optimal meet complement of Γ relative to \mathcal{V} , as described in Summary 2.3 above.

Type 2: An optimal complement of Γ of *type 2* relative to \mathcal{V} is simultaneously an optimal meet complement of Γ (and hence of type 1) and a minimal complement relative to \mathcal{V} .

Type 3: An optimal complement of Γ of *type 3* relative to \mathcal{V} is simultaneously an optimal meet complement of Γ (and hence of type 1) and an optimal complement of Γ relative to \mathcal{V} (and hence of type 0).

For $i \in \{0, 1, 2, 3\}$, if each member of $\{\Gamma_1, \Gamma_2\}$ is an optimal complement of the other of type i , then $\{\Gamma_1, \Gamma_2\}$ is called an *optimal complementary pair of type i* relative to \mathcal{V} .

These four types are summarized in Table 2.

Type of complement	Set of possible complements	
	\mathcal{V}	$\text{MeetSet}(\mathcal{V}; \Gamma)$
0	optimal	—
1	—	optimal
2	minimal	optimal
3	optimal	optimal

Table 2. Types of optimal complements for the view Γ

Note that this is not a linear hierarchy. While type 3 \Rightarrow type 0 and type 3 \Rightarrow type 2 \Rightarrow type 1, no other subsumptions hold.

To illustrate these classifications in the context of the examples of Sec. 1, let $\Pi\text{-Views}(\mathbf{E}_i)$ denote the set of all views of the schema \mathbf{E}_i (for $i \in \{0, 2, 2, 3\}$) which are defined by projections. This notion is formalized more carefully in Definition 3.2, but the main idea should already be clear. In this context, the following hold.

- (\mathbf{E}_0): The set $\{\Pi_{ABC}^{\mathbf{E}_0}, \Pi_{CD}^{\mathbf{E}_0}\}$ is an optimal complementary pair for all four types relative to $\Pi\text{-Views}(\mathbf{E}_0)$.
- (\mathbf{E}_1): The set $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{BCD}^{\mathbf{E}_1}\}$ is an optimal complementary pair of type 1 relative to $\Pi\text{-Views}(\mathbf{E}_1)$. It is not an optimal complementary pair for any of the other three types. Neither $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{CD}^{\mathbf{E}_1}\}$ nor $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{BD}^{\mathbf{E}_1}\}$ is an optimal complementary pair for any of the four types, although in both cases each element of the pair is a minimal complement of the other relative to $\Pi\text{-Views}(\mathbf{E}_1)$.
- (\mathbf{E}_2): The set $\{\Pi_{ABC}^{\mathbf{E}_2}, \Pi_{CD}^{\mathbf{E}_2}\}$ is an optimal complementary pair of types 1 and 2 relative to $\Pi\text{-Views}(\mathbf{E}_2)$. It is not an optimal complementary pair for types 0 or 3; indeed, there is no view which is an optimal complement of $\Pi_{ABC}^{\mathbf{E}_2}$ of type 0 relative to $\Pi\text{-Views}(\mathbf{E}_2)$.
- (\mathbf{E}_3): The pair $\{\Pi_{ABC}^{\mathbf{E}_3}, \Pi_{CD}^{\mathbf{E}_3}\}$ is not an optimal complementary pair relative to $\Pi\text{-Views}(\mathbf{E}_3)$ for any of the four types.

Although the above examples illustrate situations in which the optimality properties are symmetric between the two views, this need not be the case. For example, $\Pi_{ABC}^{\mathbf{E}_0}$ is an optimal complement of $\Pi_{BCD}^{\mathbf{E}_0}$ of all four types relative to $\Pi\text{-Views}(\mathbf{E}_0)$, while $\Pi_{BCD}^{\mathbf{E}_0}$ is not an optimal complement of $\Pi_{ABC}^{\mathbf{E}_0}$ for any of the four types.

The type classification described in Definition 2.5 has important interpretations within the context of constant-complement update. If Γ' forms a type-3 complement to Γ within the context \mathcal{V} of views under consideration, then every update to Γ which is supported by any complement in \mathcal{V} is supported by Γ' , and whether the update is supported or not may be determined from the state of Γ alone. If Γ' forms a type-1 or type-2 complement, every update to Γ which is supported in an independent fashion, that is, without further knowledge of the state of the main schema, is supported by constant complement Γ' . Other constant complement updates to Γ may be possible, but whether or not they are realizable may depend upon the state of that complement. Type 2 provides the further strength that no other complement supports a superset of those view updates supported by constant-complement Γ' . Type 0 provides no support at all for independence.

From the discussions of Summary 2.4 and Definition 2.5, it is apparent that the smaller the complement Γ' (in the sense of the ordering $\preceq_{\mathbf{D}}$), the larger the set of update requests which are supported. In particular, if there is an optimal complement, then it must support all update requests which are supported by any constant complement. This is made more precise in the following.

Proposition 2.6 (Optimal complements and constant-complement updates). *Let \mathbf{D} be a database schema, \mathcal{V} a set of views of \mathbf{D} , Γ a view of \mathbf{D} , $i \in \{0, 1, 2, 3\}$, and Γ' an optimal complement of Γ of type i relative to \mathcal{V} .*

- (a) *If $i \in \{0, 3\}$, then every view update which is defined by constant complement for some complement $\Gamma'' \in \mathcal{V}$ of Γ is defined by constant complement Γ' as well.*
- (b) *If $i \in \{1, 2, 3\}$, then every view update which is defined by constant complement for some meet complement $\Gamma'' \in \mathcal{V}$ of Γ is defined by constant complement Γ' as well. \square*

In that which follows, in particular for Theorem 3.15 and Theorem 3.23, both of which identify optimal complements in a given setting, there is a natural corollary concerning constant-complement updates which follows from the given theorem and the above proposition. These corollaries will not be stated explicitly, but they should be kept in mind as a fundamental application of these theorems.

3 Optimal Complements in the Context of Projections of a Universal Relation

In the classical relational decomposition theory, the standard setting is that defined by projections on a universal schema which is constrained by a join

dependency [4] [10], often derived from functional dependencies [2]. Although the motivation for studying such decompositions has been largely based upon schema normalization, the ideas apply equally well to the problem of characterizing complements. In this section, a theory of optimal complements is developed within the familiar framework of projections on a universal schema.

In that which follows, it is assumed that the reader is familiar with the standard relational model and its notation, as may be found in [23] and [1].

Notation 3.1 (Universal Relation Schema). Throughout the rest of this section, unless specifically stated to the contrary, take \mathbf{U} to be a finite nonempty set of attributes and $\mathbf{E}_{\mathbf{U}}$ the *universal relation schema* with $R[\mathbf{U}]$ its sole relation symbol. With \mathcal{F} a set of constraints, $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ denotes this same schema constrained by \mathcal{F} ; i.e., $\text{Constr}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}) = \mathcal{F}$. Unless stated to the contrary, \mathcal{F} will be taken to be an arbitrary set of first-order constraints (without non-nullary function symbols).

The symbol \models will be used to denote semantic entailment of constraints.

Definition 3.2 (Projective views). A *projective view* or (*Π -view*) on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ is a projection on some of the attributes of \mathbf{U} . More precisely, for $\mathbf{U}' \subseteq \mathbf{U}$, $\Pi_{\mathbf{U}'}^{\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}}$ is the view which is defined on tuples by $t \mapsto t[\mathbf{U}']$ and which is extended to relations in a tuple-by-tuple fashion. The morphism of this view is denoted $\pi_{\mathbf{U}'}^{\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}}$, and the view schema is $\mathbf{E}_{\mathbf{U}'}^{\pi_{\mathbf{U}'}(\mathcal{F})}$. As this notation quickly becomes quite cumbersome, the superscripts will be dropped when the context is clear. The view which is the projection of $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ onto attributes \mathbf{U}' is then denoted by simply $\Pi_{\mathbf{U}'} = (\mathbf{E}_{\mathbf{U}'}, \pi_{\mathbf{U}'})$. The set of all Π -views on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ is denoted $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$.

Identification of the constraints on $\mathbf{E}_{\mathbf{U}'}$ is not an easy matter in general. Investigations of the behavior of various classes of database dependencies under projection have been reported in [9], [19], and [15]. Since the representation of such constraints is not an issue for the problems considered in this paper, it will not be pursued further.

In the notation of Summary 2.2, it is easy to see that for $\mathbf{U}_1, \mathbf{U}_2 \subseteq \mathbf{U}$, $\mathbf{U}_1 \subseteq \mathbf{U}_2$ implies that $\Pi_{\mathbf{U}_1} \preceq_{\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}} \Pi_{\mathbf{U}_2}$. The converse, that $\Pi_{\mathbf{U}_1} \preceq_{\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}} \Pi_{\mathbf{U}_2}$ implies that $\mathbf{U}_1 \subseteq \mathbf{U}_2$, holds in the context of usual data dependencies, but can fail in certain situations. $\Pi_{AB} \preceq_{\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}} \Pi_A$ if there is only one possible value for attribute A . A less trivial example arises when a conditional functional dependency in the sense of [11] holds for $A \rightarrow B$. This means that not only does the usual FD $A \rightarrow B$ hold, but also that the value of attribute B determined for each value of attribute A is defined by the dependency itself.

Definition 3.3 (Attribute-set collections and join dependencies). An *attribute-set collection* over \mathbf{U} is any set J of nonempty subsets of \mathbf{U} . Let $J = \{\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_k\}$ be an attribute-set collection over \mathbf{U} ; i.e., $\mathbf{U}_i \subseteq \mathbf{U}$ for $1 \leq i \leq k$. A database $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}})$ satisfies the *join dependency* (or *JD*) $\bowtie[J]$ if for every sequence $\langle t_1, t_2, \dots, t_k \rangle$ of k (not necessarily distinct) tuples for $R[\mathbf{U}]$, there is a tuple t for $R[\mathbf{U}]$ with the property that $t[\mathbf{U}_i] = t_i[\mathbf{U}_i]$ for $1 \leq i \leq k$.

If $\mathbf{U} = \mathbf{U}_1 \cup \mathbf{U}_2 \cup \dots \cup \mathbf{U}_k$, then both J and the associated join dependency $\bowtie[J]$ are called *full*. If $\bowtie[J]$ is not full, it is said to be *embedded*. Say that $\bowtie[J]$ is *entailed by* \mathcal{F} if $\bowtie[J]$ holds on every $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ (i.e., if $\mathcal{F} \models \bowtie[J]$).

Thus, if $\bowtie[J]$ is entailed by \mathcal{F} , the schema $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ decomposes into the views in $\{\Pi_{\mathbf{U}_1}, \Pi_{\mathbf{U}_2}, \dots, \Pi_{\mathbf{U}_k}\}$ in the precise sense that any $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}})$ is recoverable from the k -tuple $(\pi_{\mathbf{U}_1}(M), \pi_{\mathbf{U}_2}(M), \dots, \pi_{\mathbf{U}_k}(M))$ of view states via the reconstruction mapping which sends $\langle t_1, t_2, \dots, t_k \rangle$ to the \mathbf{U} -tuple t which agrees with t_i on all attributes of the latter. This reconstruction mapping is also called the *join*.

Properties of JDs are often described in terms of associated graphs and/or hypergraphs. A full development of these ideas is not needed here, but a few definitions will prove useful. Say that \mathbf{U}_i and \mathbf{U}_j are *directly connected* if $\mathbf{U}_i \cap \mathbf{U}_j \neq \emptyset$. A subset $J' \subseteq J$ is *connected* if there is an ordering $\langle \mathbf{U}_{m_1}, \mathbf{U}_{m_2}, \dots, \mathbf{U}_{m_\ell} \rangle$ of the elements of J' such that each element except the first is directly connected to at least one which precedes it. A *connected block* is a maximal connected subset of J . This definition accommodates JDs with disjoint blocks. For example, if $J = \{AB, BC, CD, EF, FG, GH\}$, then $\{AB, BC, CD\}$ and $\{DE, EF, FG\}$ are maximal connected subsets. The blocks are related via the *cartesian-product* dependency $\bowtie[ABC, DEF]$.

A central focus of classical research on JDs surrounds the property of *acyclicity* of the underlying hypergraph of a JD. This property of the hypergraph of a JD has been shown to be equivalent to a long list of “desirable” properties [4]. Although such properties are not the subject of this work, two of the equivalent characterizations of such “desirability” do arise indirectly. The first asserts that pairwise consistency implies global consistency. For this definition, assume that $\mathcal{F} = \bowtie[J]$; that is, that $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ is constrained by $\bowtie[J]$ alone. Call $\bowtie[J]$ *pairwise definable* if for any $M_1, M_2, \dots, M_k \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$, if $\pi_{\mathbf{U}_j \cap \mathbf{U}_{j'}}(M_j) = \pi_{\mathbf{U}_j \cap \mathbf{U}_{j'}}(M_{j'})$ for all $j, j' \in \{1, \dots, k\}$, then there is an $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ with $\pi_{\mathbf{U}_i}(M_i) = \pi_{\mathbf{U}_i}(M)$ for $i \in \{1, \dots, k\}$. (In the cases that $\mathbf{U}_j \cap \mathbf{U}_{j'} = \emptyset$, this condition is trivially satisfied.)

A second “desirable” condition involves representation of $\bowtie[J]$ by simple JDs. Specifically, the JD $\bowtie[J]$ is *binary* (or a *BJD*) if J consists of exactly two elements. Binary JDs are an alternate representation of *multivalued dependencies* or *MVDs*. Call $\bowtie[J]$ *BJD-representable* if it is equivalent to a set of binary join dependencies, called a *binary basis* for $\bowtie[J]$. Then, $\bowtie[J]$ is has an acyclic hypergraph iff it is pairwise definable and iff it has a binary basis [4, Thm. 3.4].

In accordance with customary terminology, a JD which satisfies any of these “desirable” properties will be called *acyclic*.

Definition 3.4 (Constructive join dependencies). The classification scheme of Definition 2.5 involves not only minimality and optimality of ordinary complements, but optimality of meet complements as well. As illustrated in the examples of the introduction, as well as recaptured in [14, Prop. 2.17], meet complementation is closely related to dependency preservation. Unfortunately, as noted in Definition 3.2, the explicit characterization of view dependencies is a

very difficult issue in general. To sidestep that issue, at least to some degree, the alternative notion of a constructive join dependency, which involves only states of the main schema $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$, is employed. The idea is to enhance the concept of pairwise definability identified in Definition 3.3 to the case that the constraint set \mathcal{F} contains not only the JD but other constraints as well.

Continuing with the context of Definition 3.3, the join dependency $\bowtie[J] = \bowtie[\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_k]$ is said to be *constructive* (or a *CJD*) for \mathcal{F} if any set of legal databases on the views $\{\Pi_{\mathbf{U}_1}, \Pi_{\mathbf{U}_2}, \dots, \Pi_{\mathbf{U}_k}\}$ which are pairwise consistent define a legal database of the main schema. Formally, $\bowtie[J]$ is a CJD for \mathcal{F} if for any $M_1, \dots, M_k \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ with $\pi_{\mathbf{U}_j \cap \mathbf{U}_{j'}}(M_j) = \pi_{\mathbf{U}_j \cap \mathbf{U}_{j'}}(M_{j'})$ for all $j, j' \in \{1, \dots, k\}$, there is an $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ with $\pi_{\mathbf{U}_i}(M_i) = \pi_{\mathbf{U}_i}(M)$ for $i \in \{1, \dots, k\}$.

Think of $\pi_{\mathbf{U}_i}(M_i)$ as a local database on the schema of $\Pi_{\mathbf{U}_i}$. Thus, the constraints of $\Pi_{\mathbf{U}_i}$ are recaptured implicitly via a database, rather than via an explicit representation. This construction recaptures the essence of dependency preservation by requiring that any collection of view databases which agree on their common columns must arise from a single database of the main schema. If a cover of the dependencies in $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ were not embeddable, there would be view states which would not be so combinable.

It is easy to see that constructibility recaptures the situations of the examples of Sec. 1. In particular, $\bowtie[ABC, CD]$ is a CJD for \mathcal{F}_0 and \mathcal{F}_2 , but not for \mathcal{F}_1 or \mathcal{F}_3 . However, $\bowtie[ABC, BCD]$ is a CJD for \mathcal{F}_1 .

It is important to observe that, unlike pairwise definability, constructibility is a property of the entire set \mathcal{F} of constraints, and not just a JD. Nevertheless, the constructibility condition excludes cyclic JDs, even those induced by FDs [2, Sec. 7].

Definition 3.5 (Type classification of Π -complements). If $J = \{\mathbf{U}_1, \mathbf{U}_2\}$ is an attribute-set collection over \mathbf{U} consisting of exactly two elements, so that $\bowtie[J]$ is a BJD which holds on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$, then it is immediate that $\{\Pi_{\mathbf{U}_1}, \Pi_{\mathbf{U}_2}\}$ forms a complementary pair, called a *Π -complementary pair*, with $\Pi_{\mathbf{U}_1}$ and $\Pi_{\mathbf{U}_2}$ called *Π -complements*. The notions of *minimal* and *optimal Π -complement* are obtained in the natural way by specializing the definitions of Summary 2.2 with $\mathcal{V} = \Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$.

If $\bowtie[J]$ is furthermore a CJD with respect to \mathcal{F} , then $\{\Pi_{\mathbf{U}_1}, \Pi_{\mathbf{U}_2}\}$ necessarily forms a meet-complementary pair. This follows immediately from [14, Prop. 2.14]. Such a meet complementary pair of Π -views is called a *$\wedge\Pi$ -complementary pair*, and Γ_1 and Γ_2 are called *$\wedge\Pi$ -complements* of each other. Minimal and optimal $\wedge\Pi$ -complements for a Π -view Γ are defined by taking \mathcal{V} to be the set $\text{MeetSet}(\Gamma; \Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}))$ of all Π -views of $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ which are also $\wedge\Pi$ -complements of Γ .

Definition 3.6 (The natural order on attribute-set collections). If J consists of more than two elements, then complementary views may still be obtained from $\bowtie[J]$ by partitioning J into two disjoint sets. To show this, first recall that there is a natural order on full attribute-set collections over \mathbf{U} , given by $J_1 \leq J_2$ iff for each $\mathbf{W}_1 \in J_1$ there is a $\mathbf{W}_2 \in J_2$ such that $\mathbf{W}_2 \subseteq \mathbf{W}_1$ [5,

Sec. 3]. This ordering induces a natural implication order on full JDs, which is made explicit in the following.

Lemma 3.7 (Full JD implication via \leq). *Let J_1 and J_2 be full attribute-set collections over \mathbf{U} . Then $\bowtie[J_1] \models \bowtie[J_2]$ iff $J_2 \leq J_1$.*

Proof. The “only if” direction is a special case of a standard inference rule for JDs, called the *covering rule* [6]. The “if” part is obtained easily by applying the specialization of the classical chase inference procedure to full join dependencies [25, Cor. 5.3]. The details are straightforward and left to the reader. \square

Proposition 3.8 (Complementary pairs from JDs). *Let J be a full attribute-set collection over \mathbf{U} with $\mathcal{F} \models \bowtie[J]$, and let $\{J_1, J_2\}$ be a partition of J ; i.e., $J_1 \cup J_2 = J$ and $J_1 \cap J_2 = \emptyset$. Define $\mathbf{P}_i = \bigcup J_i$ for $i \in \{1, 2\}$.*

- (a) $\{\Pi_{\mathbf{P}_1}, \Pi_{\mathbf{P}_2}\}$ forms a complementary pair of views of $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$.
- (b) If $\bowtie[J]$ is furthermore a CJD with respect to \mathcal{F} , then $\{\Pi_{\mathbf{P}_1}, \Pi_{\mathbf{P}_2}\}$ forms a meet-complementary pair of views of $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$.

Proof. Part (a) follows immediately from Lemma 3.7. For (b), suppose that $J = \{\mathbf{U}_1, \mathbf{U}_2, \dots, \mathbf{U}_k\}$. Let $M'_1, M'_2 \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ with $\pi_{\mathbf{P}_1 \cap \mathbf{P}_2}(M'_1) = \pi_{\mathbf{P}_1 \cap \mathbf{P}_2}(M'_2)$, and for $i \in \{1, \dots, k\}$, define $M_i = M'_1$ if $\mathbf{U}_i \subseteq \mathbf{P}_1$ and $M_i = M'_2$ if $\mathbf{U}_i \subseteq \mathbf{P}_2$. If $\mathbf{U}_i \subseteq \mathbf{P}_1 \cap \mathbf{P}_2$, then choose M_i arbitrarily as M'_1 or M'_2 . Then, in particular, $\pi_{\mathbf{U}_i \cap \mathbf{U}_j}(M'_1) = \pi_{\mathbf{U}_i \cap \mathbf{U}_j}(M'_2)$ for any $i, j \in \{1, \dots, k\}$. Since $\bowtie[J]$ is assumed to be constructive, there is an $M \in \text{LDB}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ with $\pi_{\mathbf{U}_i}(M_i) = M$ for $i \in \{1, \dots, k\}$. From this it follows that $\pi_{\mathbf{P}_i}(M'_i) = M$ for $i \in \{1, 2\}$, as required. \square

Definition 3.9 (Redundancy and governing JDs). Proposition 3.8 says nothing about the optimality of the complements. For the decomposition identified by Proposition 3.8 to yield optimal Π -complements of a given type i (for $i \in \{0, 1, 2, 3\}$), it is necessary to establish further conditions on the join dependency which governs the decomposition.

There are two distinct ways in which redundancy may arise. First of all, if $\mathbf{U}_i \subsetneq \mathbf{U}_j$ (i.e., if \mathbf{U}_i is a proper subset of \mathbf{U}_j) for (necessarily distinct) i and j , then \mathbf{U}_i may be removed from J without altering the semantics of $\bowtie[J]$. This form is called *trivial redundancy*, and is a strictly set-theoretic property of the set J itself. Call the set J *reduced* if for no distinct i, j is it the case that $\mathbf{U}_i \subsetneq \mathbf{U}_j$. This is tantamount to requiring that J form an *antichain* [8, 1.3] under set inclusion. It is always a trivial matter to require that J be reduced, and this will be done from now on.

There are two far less trivial, semantic forms of (non-)redundancy. Let $\bowtie[J]$ be a full and reduced JD on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ which is entailed by \mathcal{F} .

- (a) Call $\bowtie[J]$ *JD-essential* on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ if it is acyclic and for any full JD φ on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ with $\mathcal{F} \models \varphi$ and $\varphi \models \bowtie[J]$, it must be the case that $\varphi = \bowtie[J]$.
- (b) Call $\bowtie[J]$ *JD-covering* on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ if for any full JD φ on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$, it must be the case that $\bowtie[J] \models \varphi$.

Thus, $\bowtie[J]$ is JD-essential if it is not implied by any stronger JD which holds on $\mathbf{E}_U^{\mathcal{F}}$, while it is JD-governing if it implies all other full JDs which hold on $\mathbf{E}_U^{\mathcal{F}}$. In particular, if $\bowtie[J]$ is JD-covering, then it is also JD-essential.

There are analogous notions for CJDs. Let $\bowtie[J]$ be a full and reduced CJD on $\mathbf{E}_U^{\mathcal{F}}$ which is entailed by \mathcal{F} .

- (a') Call $\bowtie[J]$ *CJD-essential* on $\mathbf{E}_U^{\mathcal{F}}$ if for any full CJD φ on $\mathbf{E}_U^{\mathcal{F}}$ with $\mathcal{F} \models \varphi$ and $\varphi \models \bowtie[J]$, it must be the case that $\varphi = \bowtie[J]$.
- (b') Call $\bowtie[J]$ *CJD-covering* on $\mathbf{E}_U^{\mathcal{F}}$ if for any full CJD φ on $\mathbf{E}_U^{\mathcal{F}}$, it must be the case that $\bowtie[J] \models \varphi$.

Now it is possible to define types of non-redundancy for JDs and CJDs which correspond to the types of optimality for complements which are given in Table 2. These are shown as types of *governing* in Table 3.

n for type- n -governing JD	Form of non-redundancy	
	JD-	CJD-
0	covering	—
1	—	covering
2	essential	covering
3	covering	covering

Table 3. Types of semantic non-redundancy

For example, the JD $\bowtie[J]$ is *type 0 governing* (on $\mathbf{E}_U^{\mathcal{F}}$) if it is JD-covering, and it is *type 2 governing* if it is a CJD which is JD-essential and CJD-covering.

Example 3.10. Consider again the four schemata of Sec. 1. The dependency $\bowtie[ABC, CD]$ is both JD-covering and CJD-covering on \mathbf{E}_0 , and hence it is governing of type i on \mathbf{E}_0 for all $i \in \{0, 1, 2, 3\}$.

This same dependency is not JD-essential on \mathbf{E}_1 , since both $\bowtie[AB, BC, CD]$ and $\bowtie[AB, BD, CD]$ hold on \mathbf{E}_1 and each implies $\bowtie[ABC, CD]$. Indeed both $\bowtie[AB, BC, CD]$ and $\bowtie[AB, BD, CD]$ are JD-essential, so neither can be JD-covering. The JD $\bowtie[ABC, BCD]$ is CJD-optimal, but cannot be JD-essential, since it is implied by $\bowtie[AB, BC, CD]$. Thus $\bowtie[ABC, BCD]$ is governing of type 1 on \mathbf{E}_1 , but for no other type.

The JD $\bowtie[AB, BC, CD]$ is both JD-essential and CJD-optimal on \mathbf{E}_2 . However, it is not JD-optimal, since $\bowtie[AB, BD, CD]$ is also JD-essential. Thus, it is governing of types 1 and 2 on \mathbf{E}_2 , but not of types 0 or 3.

Each of the six JDs $\bowtie[AB, BC, CD]$, $\bowtie[AB, BD, CD]$, $\bowtie[AC, BC, BD]$, $\bowtie[AC, BD, CD]$, $\bowtie[AD, BC, CD]$, and $\bowtie[AD, BD, CD]$, is both JD-essential and CJD-essential on \mathbf{E}_3 , and so none can be JD-governing or CJD-governing. This schema has no governing JD or CJD.

Definition 3.11 (The join-reconstruction property). The types of Definition 3.9 provide the proper characterizations via JDs to ensure the corresponding types of optimality, provided one further condition is imposed, that

of *join reconstruction*. Although it is natural to think that when a schema is decomposed into projections, reconstruction must be defined by the join, this need not always be the case. For example, let \mathbf{E}_4 be the universal schema with the single relation schema $R[ABC]$, constrained by two rather unconventional dependencies. The first is $\psi_{41} = (\forall x)(\forall y)(\neg R(x, y, x))$, which states that the A and C values of a tuple must not be the same. The second is a conditional join dependency which takes this restriction into account: $\psi_{42} = (\forall x)(\forall y)(\forall z)((R(x, y, z') \wedge R(x', y, z) \wedge (x \neq z)) \Rightarrow R(x, y, z))$. In other words, the usual join dependency $\bowtie[AB, BC]$ holds conditionally, for all pairs of tuples which would not result in a tuple which would be illegal under constraint ψ_{41} . Nevertheless, the decomposition of \mathbf{E}_4 into its AB and BC projections is lossless. The reconstruction mapping is a conditional join which ignores pairs of view tuples which would result in a tuple t with $t[A] = t[C]$. In particular, it is not the usual join.

It is not the purpose of this paper to investigate this phenomenon in detail. The above example is included to illustrate that it must not simply be assumed that reconstruction of projections must be defined by the join. Suffice to say that it is possible to show that this non-join sort of reconstruction cannot occur within the context of common database dependencies which are defined by Horn sentences [9], [17, Def. 21]. For this paper, this difficulty will simply be assumed away. Say that a universal schema $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ has the *join-reconstruction property* if for every lossless decomposition of that schema into projections, the reconstruction mapping is necessarily the join.

Notation 3.12 (Join-reconstruction property assumed). Unless stated specifically to the contrary, for the rest of this section, assume that $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ has the join-reconstruction property.

Lemma 3.13 (Optimal complementary pairs from JDs). *Let J be a full attribute-set collection over \mathbf{U} , let $i \in \{0, 1, 2, 3\}$, and assume that $\bowtie[J]$ is governing of type i relative to Π -Views $\langle \mathbf{E}_{\mathbf{U}}^{\mathcal{F}} \rangle$. Then for every partition $\{J_1, J_2\}$ of J , $\{\Pi_{\cup J_1}, \Pi_{\cup J_2}\}$ forms an optimal complementary pair of type i relative to Π -Views $\langle \mathbf{E}_{\mathbf{U}}^{\mathcal{F}} \rangle$.*

Proof. The requirement that $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ have the join-reconstruction property limits the reconstruction mapping to the join. From Lemma 3.7 follows directly that smaller optimal complements must arise from JDs and CJDs which are larger in the ordering \leq of Definition 3.6. The optimality then follows directly using the concepts of Definition 3.9 and Table 3. \square

Definition 3.14 (The Π -basis of a view relative to J). It is easy to extend the result of Lemma 3.13 to projections which are not defined by the union of some of the elements of J , thus obtaining a theory of optimal Π - and $\wedge \Pi$ -complements for all Π -views on a suitable schema. Given an attribute-set collection J over \mathbf{U} and any $\mathbf{U}' \subseteq \mathbf{U}$, define the elements of J covered by \mathbf{U}' to be $\text{Covered}_{\langle J, \Pi \rangle} \langle \mathbf{U}' \rangle = \{\mathbf{W} \in J \mid \mathbf{W} \subseteq \mathbf{U}'\}$. The Π -complementary cover to \mathbf{U}' relative to J is given by $\overline{\text{Covered}}_{\langle J, \Pi \rangle} \langle \mathbf{U}' \rangle = J \setminus \text{Covered}_{\langle J, \Pi \rangle} \langle \mathbf{U}' \rangle$.

Theorem 3.15 (Governing implies optimal complements). *Let J be a full attribute-set collection over \mathbf{U} , let $i \in \{0, 1, 2, 3\}$, and assume that $\bowtie[J]$ is governing of type i relative to $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$. Then for every $\mathbf{W} \subseteq \mathbf{U}$, the view $\Pi_{\mathbf{W}}$ has an optimal complement of type i relative to $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$, given explicitly by $\Pi_{\mathbf{W}'}$, with $\mathbf{W}' = \bigcup \overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{W})$.*

Proof. With $\mathbf{W}' = \bigcup \overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{W})$, the view $\Pi_{\mathbf{W}'}$ is a complement of $\Pi_{\mathbf{W}}$ of the appropriate type, since $\bigcup \overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{W}) \subseteq \mathbf{W}$, and by Lemma 3.13, $\Pi_{\overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{W})}$ and $\Pi_{\mathbf{W}'}$ are each (optimal) complements of one another of type i . On the other hand, an application of Lemma 3.7 shows that any Π -complement of $\Pi_{\mathbf{W}}$ must include at least the attributes included in the members of $\overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{W})$, whence the characterization of the complement holds. \square

Example 3.16. Continue with the discussion of Example 3.10 of the four schemata of Sec. 1. The optimality properties which have already been presented directly in Definition 2.5 may now be verified formally using Lemma 3.13.

Because the dependency $\bowtie[ABC, CD]$ is governing of all four types on \mathbf{E}_0 , $\{\Pi_{ABC}^{\mathbf{E}_0}, \Pi_{CD}^{\mathbf{E}_0}\}$ forms an optimal complementary pair of type i for all four types relative to $\Pi\text{-Views}(\mathbf{E}_0)$. It is optimal regardless of whether or not dependency preservation is required.

The dependency $\bowtie[ABC, BCD]$ is CJD-optimal on \mathbf{E}_1 , but is not even JD-essential, so $\{\Pi_{ABC}^{\mathbf{E}_1}, \Pi_{BCD}^{\mathbf{E}_1}\}$ forms an optimal complementary pair of type 1, but for no other types. In other words, it is optimal for dependency preservation, but not generally.

On \mathbf{E}_2 , $\bowtie[ABC, CD]$ is governing of types 1 and 2, so $\{\Pi_{ABC}^{\mathbf{E}_2}, \Pi_{CD}^{\mathbf{E}_2}\}$ forms an optimal complementary pair of types 1 and 2, but not for types 0 or 3. In other words, it is optimal for dependency preservation, and, additionally, it is essential (i.e., minimal) amongst full JDs on \mathbf{E}_2 .

On \mathbf{E}_3 , $\bowtie[AB, BC, CD]$ is both JD- and CJD-essential, but not governing for any of the four types, and so $\{\Pi_{ABC}^{\mathbf{E}_3}, \Pi_{CD}^{\mathbf{E}_3}\}$ does not form an optimal complementary pair in any of these senses. Indeed, $\{\Pi_{ABC}^{\mathbf{E}_3}, \Pi_{BD}^{\mathbf{E}_3}\}$ and $\{\Pi_{ABD}^{\mathbf{E}_3}, \Pi_{CD}^{\mathbf{E}_3}\}$ form alternate minimal complementary pairs relative to $\Pi\text{-Views}(\mathbf{E}_3)$. As shown in Example 3.10, it has six distinct JDs, each of which is both JD-essential and CJD-essential.

Example 3.17. Let \mathbf{E}_5 be the universal relational schema whose only relation symbol is $R[ABCDE]$, constrained by the FDs $\mathcal{F}_5 = \{B \rightarrow C, C \rightarrow D, D \rightarrow E\}$. Arguing in a fashion similar to that for \mathbf{E}_2 , it is not difficult to see that $\bowtie[AB, BC, CD, DE]$ is governing of types 1 and 2, but not of types 0 or 3, on the schema \mathbf{E}_3 .

Let $J_1 = \{AB, DE\}$ and $\mathbf{P}_1 = \bigcup J_1 = ABDE$. In the notation of Proposition 3.8 and Lemma 3.13, J_2 must be $\{BC, CD\}$, and so $\mathbf{P}_2 = \bigcup J_2 = BCD$. Thus, the $\wedge\Pi$ -complement of $\Pi_{ABDE}^{\mathbf{E}_5}$ is $\Pi_{BCD}^{\mathbf{E}_5}$, and vice-versa. Note that the embedded join dependency $\bowtie[AB, DE]$ is not satisfied in the schema of $\Pi_{ABDE}^{\mathbf{E}_5}$. There is no requirement that J_1 and J_2 in Proposition 3.8 and Theorem 3.15 satisfy any embedded join constraints.

The use of this schema to illustrate key points continues in Example 3.21 below.

Theorem 3.15 establishes that if $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ is governed by a JD of type i , then all projections have optimal complements of that type. It is also possible to establish a sort of converse, which states that if all projections have optimal complements of a given type, then the schema must be governed by a JD of that type. The only qualification is that the set of all JDs which are generated by the projections and their complements must be equivalent to a single JD; that is, it must be BJD representable. The details are sketched below.

Discussion 3.18 (Further properties of representation via BJDs). Recall from Definition 3.3 that one of the “desirable” properties of a JD $\bowtie[J]$ is that it be representable by a set of BJDs. There is a further refinement of this idea which will prove useful. Let $\varphi_1 = \bowtie[\mathbf{U}_{11}, \mathbf{U}_{12}]$ and $\varphi_2 = \bowtie[\mathbf{U}_{21}, \mathbf{U}_{22}]$ be BJDs. Say that $\{\varphi_1, \varphi_2\}$ has the *subset property* if at least one of the following four conditions holds:

$$\begin{array}{ll} \mathbf{U}_{11} \subseteq \mathbf{U}_{21} \text{ and } \mathbf{U}_{22} \subseteq \mathbf{U}_{12} & \mathbf{U}_{11} \subseteq \mathbf{U}_{22} \text{ and } \mathbf{U}_{21} \subseteq \mathbf{U}_{12} \\ \mathbf{U}_{21} \subseteq \mathbf{U}_{11} \text{ and } \mathbf{U}_{12} \subseteq \mathbf{U}_{22} & \mathbf{U}_{22} \subseteq \mathbf{U}_{11} \text{ and } \mathbf{U}_{12} \subseteq \mathbf{U}_{21} \end{array}$$

In other words, φ_1 and φ_2 have an asymmetry relationship in that one element of φ_1 is at least as large as one of the elements of φ_2 , and the other element of φ_2 is at least as large as the remaining element of φ_1 . The set S of BJDs on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ has the *subset property* if every pair of elements from S has the property. The key result is that if $\bowtie[J]$ is equivalent to a set of BJDs, then it is equivalent to a set of BJDs with the subset property [13, Thm. 6].

For $i \in \{0, 1, 2, 3\}$, call a BJD $\bowtie[\mathbf{W}_1, \mathbf{W}_2]$ *optimal of type i* relative to $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$ on $\mathbf{E}_{\mathbf{U}}^{\mathcal{F}}$ if $\{\Pi_{\mathbf{W}_1}, \Pi_{\mathbf{W}_2}\}$ is an optimal pair of type i relative to $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$.

Call the BJD $\bowtie[\mathbf{W}_1, \mathbf{W}_2]$ *primary* relative to $\bowtie[J]$ if there is a partition $\{J_1, J_2\}$ of J with the property that for $i \in \{1, 2\}$, $\mathbf{W}_i = \bigcup J_i$ and for each connected block B of J , the subset of B which is part of J_i is itself connected. Put another way, this last condition means that the embedded join dependency $\bowtie[J_i]$ holds on $\Pi_{\mathbf{W}_i}$.

Call a binary basis B of $\bowtie[J]$ *primary* if it consists of BJDs with the subset property and with the additional property that each $\varphi \in B$ is primary relative to $\bowtie[J]$.

Lemma 3.19 (Optimal complements and BJDs). *Let J be a full attribute-set collection over \mathbf{U} , and assume that $\bowtie[J]$ is acyclic and reduced.*

- (a) $\bowtie[J]$ has a primary basis relative to $\bowtie[J]$.
- (b) If $\bowtie[J]$ is JD-essential (resp. JD-governing), then it has a primary basis relative to $\bowtie[J]$ with the further property that for each $\bowtie[\mathbf{W}_1, \mathbf{W}_2]$ in that basis, $\{\Pi_{\mathbf{W}_1}, \Pi_{\mathbf{W}_2}\}$ is a minimal (resp. optimal) complementary pair relative to $\Pi\text{-Views}(\mathbf{E}_{\mathbf{U}}^{\mathcal{F}})$.
- (c) If $\bowtie[J]$ is constructive, then every full JD φ with $\bowtie[J] \models \varphi$ is constructive as well.

- (d) For $i \in \{0, 1, 2, 3\}$, if $\bowtie[J]$ is type i governing on $\mathbf{E}_U^{\mathcal{F}}$, then it has a primary basis consisting of optimal complementary pairs of type i relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$.
- (e) If S is a primary basis of $\bowtie[J]$ with the property that every $\varphi \in S$ is constructive, then $\bowtie[J]$ is constructive as well.
- (f) If S is a binary basis of $\bowtie[J]$ consisting of optimal complementary pairs of type i relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$, then $\bowtie[J]$ is governing of type i on $\mathbf{E}_U^{\mathcal{F}}$.

Proof outline. To show (a), let B_1, \dots, B_ℓ denote the connected blocks of $\bowtie[J]$. For each such block B_i , let $\langle \mathbf{U}_{i1}, \dots, \mathbf{U}_{i\ell_i} \rangle$ denote an ordering of its elements with the property that each element except the first is connected to at least one element which precedes it in the ordering. Let $\langle n_1, \dots, n_\ell \rangle$ denote a sequence of ℓ numbers, with $0 \leq n_i \leq \ell_i$ for each i and at least one of the numbers not zero. Define $K_{\langle n_1, \dots, n_\ell \rangle}$ to be the subset of J consisting of the first n_i elements of B_i for each i , and define $K'_{\langle n_1, \dots, n_\ell \rangle} = J \setminus K_{\langle n_1, \dots, n_\ell \rangle}$. It is straightforward to verify that the set of all BJDs of the form $\bowtie[\bigcup K_{\langle n_1, \dots, n_\ell \rangle}, \bigcup K'_{\langle n_1, \dots, n_\ell \rangle}]$ is a primary basis for $\bowtie[J]$.

Keeping Lemma 3.7 in mind, the proof of (b) is a simple verification.

Again keeping Lemma 3.7 in mind, part (c) follows directly from the definition of constructive.

Part (d) is a consequence of (b) and (c).

The details of the proof of (e) depend upon the algorithm of [13]; the idea is as follows. Let $\bowtie[\mathbf{W}_1, \mathbf{W}_2] \in S$, and let $M_1, M_2 \in \text{LDB}(\mathbf{E}_U^{\mathcal{F}})$ with $\pi_{\mathbf{W}_1 \cap \mathbf{W}_2}(M_1) = \pi_{\mathbf{W}_1 \cap \mathbf{W}_2}(M_2)$. By the definition of constructive, there is an $M \in \text{LDB}(\mathbf{D})$ with $\pi_{\mathbf{W}_i}(M_i) = \pi_{\mathbf{W}_i}(M)$ for $i \in \{1, 2\}$. Now, owing to the subset condition, applying a second $\bowtie[\mathbf{W}_3, \mathbf{W}_4] \in S$ will result in a split of only one of \mathbf{W}_1 and \mathbf{W}_2 . Assume, say, that $\mathbf{W}_1 \subseteq \mathbf{W}_3$ and $\mathbf{W}_4 \subseteq \mathbf{W}_2$. The consequence of combining these two JDs is then $\bowtie[\mathbf{W}_1, \mathbf{W}_2 \cap \mathbf{W}_3, \mathbf{W}_4]$; i.e., \mathbf{W}_2 is split into $\mathbf{W}_2 \cap \mathbf{W}_3$ and \mathbf{W}_4 , while \mathbf{W}_1 is not split. Now, letting \mathbf{W}_{23} denote $\mathbf{W}_2 \cap \mathbf{W}_3$, if $M_1, M_{23}, M_4 \in \text{LDB}(\mathbf{E}_U^{\mathcal{F}})$ with the property that $\pi_{\mathbf{W}_j \cap \mathbf{W}_{j'}}(M_j) = \pi_{\mathbf{W}_j \cap \mathbf{W}_{j'}}(M_{j'})$ for $j, j' \in \{1, 23, 4\}$, then an $M \in \text{LDB}(\mathbf{E}_U^{\mathcal{F}})$ with the property that $\pi_{\mathbf{W}_i}(M_i) = \pi_{\mathbf{W}_i}(M)$ for $i \in \{1, 23, 4\}$ may be constructed in steps. First, construct an $M_2 \in \text{LDB}(\mathbf{E}_U^{\mathcal{F}})$ with the property that $\pi_{\mathbf{W}_i}(M_i) = \pi_{\mathbf{W}_i}(M_2)$ for $i \in \{23, 4\}$, and then combine that M_2 with M_1 to obtain an M which agrees with each of M_1, M_{23} , and M_4 on the associated projections. Continue on with further primary BJDs in an inductive fashion. The details of this construction are left to the reader.

Finally, (f) follows from the previous assertions. \square

The converse of Theorem 3.15, which follows from Lemma 3.19, as the following.

Theorem 3.20 (Optimal complements imply governing). *Let $i \in \{0, 1, 2, 3\}$, and assume that for every $\mathbf{W} \subseteq \mathbf{U}$, $\Pi_{\mathbf{W}}$ has an optimal complement of type i on $\mathbf{E}_U^{\mathcal{F}}$ relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$. Assume further that the set of all BJDs of the form $\bowtie[\mathbf{W}_1, \mathbf{W}_2]$ for which $\{\Pi_{\mathbf{W}_1}, \Pi_{\mathbf{W}_2}\}$ is an optimal com-*

plementary pair of type i relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$ is equivalent to a single join dependency. Then that JD is type i governing on $\mathbf{E}_U^{\mathcal{F}}$. \square

Observe that in Theorem 3.20 it is necessary to assume that the BJDs generate a single join dependency. This corresponds to assuming that there is a governing JD which is acyclic. Of course, there are cases which lie outside of this framework, but it is unclear whether they ever occur in real modelling situations.

To conclude the investigation, the idea of working with views which are sets of projections, rather than single projections, is examined briefly. First, a motivating example is discussed.

Example 3.21. Continue with the schema \mathbf{E}_5 of Example 3.17; i.e., with $J = \{AB, BC, CD, DE\}$. Let let $\mathbf{U}' = ABE$. Then $\text{Covered}_{\langle J, \Pi \rangle}(\mathbf{U}') = \{AB\}$, and so $\overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{U}') = BC \cup CD \cup DE = BCDE$, whence the optimal complement of $\Pi_{ABE}^{\mathbf{E}_5}$ of types 1 and 2 relative to $\Pi\text{-Views}(\mathbf{E}_5)$ must be $\Pi_{BCDE}^{\mathbf{E}_5}$. On the other hand, if $\mathbf{U}' = BCE$, then $\text{Covered}_{\langle J, \Pi \rangle}(\mathbf{U}') = \{BC\}$, and so $\overline{\text{Covered}}_{\langle J, \Pi \rangle}(\mathbf{U}') = AB \cup CD \cup DE = ABCDE = \mathbf{U}$, i.e., the optimal complements of both types relative to $\Pi\text{-Views}(\mathbf{E}_5)$ is Π_{ABCDE} , which is the identity view. With that complement, no updates at all are possible under the constant-complement strategy [3] [14]. This drawback may be remedied, at least to some degree, by working with views which are defined by sets of projections instead of single projections. A brief presentation of this extension follows.

Definition 3.22 ($\vee\Pi$ -views). A *joined Π -view*, or $\vee\Pi$ -view, is defined by a set of projections, rather than by just a single projection. Continue with the context of a full attribute-set collection J on $\mathbf{E}_U^{\mathcal{F}}$, and let $J' \subseteq J$. The $\vee\Pi$ -view defined by J' , denoted $\Pi_{J'} = (\mathbf{E}_{J'}, \pi_{J'})$, has in its schema one relation $R_{\mathbf{W}}[\mathbf{W}]$ for each $\mathbf{W} \in J'$. The view morphism $\pi_{J'}$ sends an instance of the universal relation $R[\mathbf{U}]$ to each of its projections defined by J' . More formally, this view is the product of the set $\{\Pi_{\mathbf{W}}^{\mathbf{E}_U} \mid \mathbf{W} \in J'\}$ of views, as elaborated in [16, 3.4]. The set of all $\vee\Pi$ -views on $\mathbf{E}_U^{\mathcal{F}}$ is denoted $\vee\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$.

For $i \in \{0, 1, 2, 3\}$, the notions of optimal complement of type i , as well as optimal complementary pairs of type i , relative to $\vee\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$ are defined as the natural extensions of the corresponding notions relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$.

Since the notation to describe such schemata fully becomes cumbersome, the idea will instead be illustrated by example, using \mathbf{E}_5 of Example 3.17 and Example 3.21 as the main schema. For $J' = \{AB, CD\}$, $\Pi_{J'}^{\mathbf{E}_5} = \Pi_{\{AB, CD\}}^{\mathbf{E}_5}$ has two relation symbols $R_{AB}[AB]$ and $R_{CD}[CD]$. For $M \in \text{LDB}(\mathbf{E}_5)$, the view mapping $\pi_{J'}^{\mathbf{E}_5} = \pi_{\{AB, CD\}}^{\mathbf{E}_5} : M \mapsto (\pi_{AB}^{\mathbf{E}_5}(M), \pi_{CD}^{\mathbf{E}_5}(M))$. This view is strictly weaker than $\Pi_{ABCD}^{\mathbf{E}_5}$; i.e., $\Pi_{\{AB, CD\}}^{\mathbf{E}_5} \prec_{\mathbf{E}_5} \Pi_{ABCD}^{\mathbf{E}_5}$, since it does not preserve information on which AB -projections are associated with which CD -projections.

On the other hand, $\Pi_{\{AB, BC\}}^{\mathbf{E}_5}$ is equivalent to $\Pi_{ABC}^{\mathbf{E}_5}$, since the embedded join dependency $\bowtie[AB, BC]$ holds on \mathbf{E}_5 . (This is thus an example, as promised in Summary 2.1, of the same congruence defining two distinct views. Clearly, these views are equivalent in any reasonable sense.)

The key advantage of \sqrt{II} -views over II -views is that since they form a larger class, the optimal complements will be smaller in general, thus allowing larger update sets under the constant-complement strategy. For example, continuing with \mathbf{E}_5 , as illustrated in Example 3.21 above, the optimal II -complement of $\Pi_{BC}^{\mathbf{E}_5}$ is the identity view. On the other hand, it is easy to see that $\Pi_{\{AB,CD,DE\}}^{\mathbf{E}_5}$ is a \sqrt{II} -complement of $\Pi_{BC}^{\mathbf{E}_5}$ which is strictly weaker than the identity view. Indeed, with $M = \{R(a_1, b_1, c_1, d_1, e_1), R(a_2, b_2, c_2, d_2, e_2)\}$, the update from $\{R_{BC}(b_1, c_1), R_{BC}(b_2, c_2)\}$ to $\{R_{BC}(b_1, c_2), R_{BC}(b_2, c_1)\}$ on the state of $\Pi_{BC}^{\mathbf{E}_5}$ is possible with \sqrt{II} -complement $\Pi_{\{AB,CD,DE\}}^{\mathbf{E}_5}$, constant, but not with the optimal II -complement $\Pi_{ABCDE}^{\mathbf{E}_5}$ constant.

The formalization and results for the various flavors of optimal II -complements extend easily to the \sqrt{II} -framework. Return to the general framework of $\mathbf{E}_U^{\mathcal{F}}$ and an attribute-set collection J over \mathbf{U} . For J' any attribute-set collection over \mathbf{U} , define $\text{Covered}_{\langle J; \sqrt{II} \rangle} \langle J' \rangle = \{\mathbf{W} \in J \mid (\exists \mathbf{W}' \in J') (\mathbf{W} \subseteq \mathbf{W}')\}$.

The \sqrt{II} -complementary cover to J' relative to J is given by $\overline{\text{Covered}}_{\langle J; \sqrt{II} \rangle} \langle J' \rangle = J \setminus \text{Covered}_{\langle J; \sqrt{II} \rangle} \langle J' \rangle$.

The proof of the following theorem is almost identical to that of Theorem 3.15, and is left to the reader.

Theorem 3.23 (Optimal complements for \sqrt{II} -views). *Let $i \in \{0, 1, 2, 3\}$. If the schema $\mathbf{E}_U^{\mathcal{F}}$ has the join-reconstruction property and the $JD \bowtie [J]$ is governing of type i , then every \sqrt{II} -view of $\mathbf{E}_U^{\mathcal{F}}$ has an optimal complement of type i relative to $II\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$. More specifically, given $J' \subseteq J$, in each case the optimal complement of $\Pi_{J'}$ of type i relative to $\sqrt{II}\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$ is given by $\Pi_{J''}$, with $J'' = \overline{\text{Covered}}_{\langle J; \sqrt{II} \rangle} \langle J' \rangle$. \square*

Example 3.24 (Optimal \sqrt{II} -complements). Continuing with \mathbf{E}_5 and the discussion at the end of Definition 3.22, the optimal complements of $\Pi_{\{BC,CE\}}^{\mathbf{E}_5}$ (which is equivalent to $\Pi_{BCE}^{\mathbf{E}_5}$), of all types, relative to $\sqrt{II}\text{-Views}(\mathbf{E}_5)$, is identical to those of $\Pi_{BC}^{\mathbf{E}_5}$. In each case, that complement is $\Pi_{\{AB,CD,DE\}}^{\mathbf{E}_5}$ (which is equivalent to $\Pi_{\{AB,CDE\}}^{\mathbf{E}_5}$).

Also, $\Pi_{\{AB,CD\}}^{\mathbf{E}_5}$ and $\Pi_{\{BC,DE\}}^{\mathbf{E}_5}$ are each optimal complements of the other for type i for $i \in \{0, 1, 2, 3\}$, relative to $\sqrt{II}\text{-Views}(\mathbf{E}_5)$. These views are each strictly smaller than their II -counterparts $\Pi_{ABCD}^{\mathbf{E}_5}$ and $\Pi_{BCDE}^{\mathbf{E}_5}$.

4 Conclusions and Further Directions

A characterization of optimal complements for views defined by projections on a universal-relational schema has been developed. To cover a variety of situations, four distinct notions of optimality were identified. The characterization of optimality for each notion is rooted in notions of governing dependencies. Specifically, optimal complements of a given type exist precisely in the situation that a governing join dependency of a corresponding type exists, with “governing” meaning that all other join dependencies of that type are implied by it.

There are several key areas for further work on this subject.

Individual complements: The theory developed in Sec. 3 addresses the situation in which every Π -view or $\sqrt{\Pi}$ -view of $\mathbf{E}_U^{\mathcal{F}}$ has an optimal complement of a given type. However, there are cases in which some views may have optimal complements, while others do not. For example, in the context of \mathbf{E}_2 of Sec. 1, the view $\Pi_{AB}^{\mathbf{E}_1}$ has $\Pi_{BCD}^{\mathbf{E}_1}$ as an optimal complement of all four types relative to $\Pi\text{-Views}(\mathbf{E}_U^{\mathcal{F}})$, even though $\Pi_{ABC}^{\mathbf{E}_1}$ has an optimal complement only of type 1. The results of Sec. 3 should be expanded to identify such individual optimal complements.

Complements in the context of normalization: The theory developed in Sec. 3 addresses only the situation of a universal relational schema as the main schema. However, in practice, such schemata exist only as part of the design process; the schemata which are used in practice are multi-relational, and often obtained from a universal schema after normalization via decomposition. If a lossless and dependency preserving normalization is employed, then the results of this paper apply equally well to the result schema. For example, if the single relation $R[ABCD]$ of \mathbf{E}_0 of Sec. 1 is decomposed into R_{ABC} and R_{CD} , and the two-way inclusion dependency $R_{ABC}[C] = R_{CD}[C]$ is enforced as well, then the normalized schema is isomorphic to the original one, and all of the results of the theory apply equally well to it. However, it is usually the case that the two-way inclusion dependency is replaced with a one-way *foreign-key dependency*, in this situation $R_{ABC}[C] \subseteq R_{CD}[C]$. Now, the normalized schema is no longer isomorphic to the original one, and so which views identify optimal complements of others may change. This is an important practical question which will be addressed in a forthcoming paper.

Weaker equivalence of non-isomorphic complements: Consider again the situation surrounding \mathbf{E}_3 of Sec. 1, which does not have optimal complements of any of the four types. Indeed, in Example 3.10 it is shown that this schema has six incomparable JDs, all of which are both JD- and CJD-essential. However, the attributes in $\{B, C, D\}$ are equivalent in the sense that $B \rightarrow C \rightarrow D \rightarrow B$. It furthermore turns out that the different minimal complements of a given schema differ only in a permutation of these elements. Thus, the complements are equivalent in a sense weaker than true isomorphism but nevertheless meaningful. The theory which characterizes such equivalence will be presented in a forthcoming paper.

A general algebraic theory of optimal complements: Although the results of Sec. 3 are focused upon universal relational schemata and views defined by projection, the underlying framework of Sec. 2 is much more general, suggesting that the kind of results developed for projections on a universal schema should in fact apply in a more abstract setting. Of particular importance are the additional kinds of decompositions identified in the normalization question above, as well as decompositions on relational schemata which involve both projection and selection, such as those of [18]. An important future direction is the formalization of such a general framework and its application to other forms of decomposition.

Rapprochement with computational approaches: In recent years, there has been renewed interest in using view complements to manage data warehouses, and a number of approaches to computing such complements has arisen, the most comprehensive of which is [21]. Although the goal of that work is much different than that of this paper, in that it develops ways to compute good complements via the manipulation of expressions in the relational algebra, there are nevertheless points of similarity which deserve further investigation.

References

1. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
2. A. V. Aho, C. Beeri, and J. D. Ullman. The theory of joins in relational databases. *ACM TODS*, 4(3):297–314, 1979.
3. F. Bancilhon and N. Spyrtos. Update semantics of relational views. *ACM Trans. Database Systems*, 6:557–575, 1981.
4. C. Beeri, R. Fagin, D. Maier, and M. Yannakakis. On the desirability of acyclic database schemes. *JACM*, 30(3):479–513, 1983.
5. C. Beeri, A. O. Mendelzon, Y. Sagiv, and J. D. Ullman. Equivalence of relational database schemes. *SIAM J. Computing*, 10(2):352–370, 1981.
6. C. Beeri and M. Vardi. On the properties of join dependencies. In H. Gallaire, J. Minker, and J. M. Nicolas, editors, *Advances in Data Base Theory*, volume 1, pages 25–71. Plenum Press, 1981.
7. A. Bohannon, B. C. Pierce, and J. A. Vaughan. Relational lenses: a language for updatable views. In *Proceedings of the Twenty-Fifth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 26-28, 2006, Chicago, Illinois, USA*, pages 338–347, 2006.
8. B. A. Davey and H. A. Priestly. *Introduction to Lattices and Order*. Cambridge University Press, second edition, 2002.
9. R. Fagin. Horn clauses and database dependencies. *J. Assoc. Comp. Mach.*, 29(4):952–985, 1982.
10. R. Fagin. Degrees of acyclicity for hypergraphs and relational database schemes. *J. Assoc. Comp. Mach.*, 30(3):514–550, 1983.
11. W. Fan, F. Geerts, X. Jia, and A. Kementsietsidis. Conditional functional dependencies for capturing data inconsistencies. *ACM Trans. Database Systems*, 33(2), 2008.
12. J. N. Foster, B. C. Pierce, and S. Zdancewic. Updatable security views. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium, CSF 2009, Port Jefferson, New York, USA, July 8-10, 2009*, pages 60–74, 2009.
13. N. Goodman and Y. C. Tay. A characterization of multivalued dependencies equivalent to a join dependency. *Inf. Process. Lett.*, 18(5):261–266, 1984.
14. S. J. Hegner. An order-based theory of updates for closed database views. *Ann. Math. Art. Intell.*, 40:63–125, 2004.
15. S. J. Hegner. The complexity of embedded axiomatization for a class of closed database views. *Ann. Math. Art. Intell.*, 46:38–97, 2006.
16. S. J. Hegner. Semantic bijectivity and the uniqueness of constant-complement updates in the relational context. In K.-D. Schewe and B. Thalheim, editors, *International Workshop on Semantics in Data and Knowledge Bases, SDKB 2008, Nantes, France, March 29, 2008, Proceedings*, volume 4925 of *Lecture Notes in Computer Science*, pages 172–191. Springer-Verlag, 2008.

17. S. J. Hegner. Internal representation of database views. *J. Universal Comp. Sci.*, 16(20):2956–2985, 2010.
18. S. J. Hegner. A model of independence and overlap for transactions on database schemata. In B. Catania, M. Ivanović, and B. Thalheim, editors, *Advances in Databases and Information Systems, 14th East European Conference, ADBIS 2010, Novi Sad, Serbia, September 20-24, 2010, Proceedings*, volume 6295 of *Lecture Notes in Computer Science*, pages 209–223. Springer-Verlag, 2010.
19. R. Hull. Finitely specifiable implicational dependency families. *J. Assoc. Comp. Mach.*, 31(2):210–226, 1984.
20. J. Lechtenbörger. The impact of the constant complement approach towards view updating. In *Proceedings of the Twenty-Second ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, San Diego, California, June 09-11, 2003*, pages 49–55, 2003.
21. J. Lechtenbörger and G. Vossen. On the computation of relational view components. *ACM Trans. Database Systems*, 28:175–208, 2003.
22. O. Ore. Theory of equivalence relations. *Duke Math. J.*, 9:573–627, 1942.
23. J. Paredaens, P. De Bra, M. Gyssens, and D. Van Gucht. *The Structure of the Relational Database Model*. Springer-Verlag, 1989.
24. J. Rissanen. Independent components of relations. *ACM Trans. Database Systems*, 2(4):317–325, 1977.
25. E. Sciore. A complete axiomatization of full join dependencies. *J. Assoc. Comp. Mach.*, 29(2):373–393, 1982.