

FD Covers and Universal Complements of Simple Projections

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 constant-complement strategy, in which the admissible updates to a given view are those which hold a second complementary view constant, remains one of the most attractive formalisms for identifying suitable translation mechanisms for updates to views of database schemata. However, in general, it suffers from the drawback that the reflections of view updates to the main schema can depend upon the choice of complement in various ways. To overcome this drawback completely, a special kind of complement, called a universal complement, is required. In this paper, sufficient conditions for the existence of such a complement are established for a classical but nevertheless very important setting — views defined by simple projection of a universal relational schema constrained by functional dependencies (FDs). Certain uniqueness properties of covers of these dependencies prove critical in the characterization. The results are extended to quasi-universal complements, which are unique up to exchange of equivalent attributes, thus recapturing certain situations for which unique covers do not exist.

1 Introduction

Although introduced more than thirty years ago, the constant-complement approach to the support of view updates [4] remains an attractive solution formalism. A *complement* of the view Γ is a second view Γ' with the property that the pair $\{\Gamma, \Gamma'\}$ defines a lossless decomposition of the main schema \mathbf{D} . In the constant-complement update strategy on Γ for complement Γ' , an allowable update u to Γ is precisely one which may be realized as an update on \mathbf{D} in such a way that the state of Γ' does not change. Since $\{\Gamma, \Gamma'\}$ forms a lossless decomposition, there can be at most one such *reflection* of u to an update on \mathbf{D} .

This strategy has substantial intuitive appeal because the update to \mathbf{D} is localized — it affects only the part of \mathbf{D} defined by Γ . The “rest” of \mathbf{D} , defined by Γ' , is unaffected by the update. For this reason, among others, it has seen renewed interest in recent years as a fundamental technique [18] [19] [13], as well as a related principle for other approaches requiring view updates, such as those based upon lenses [9], [11].

In the ideal case, the complement Γ' should possess three distinct properties relative to a class \mathcal{V} of views (such as simple projections) to which it belongs:

STATE INVARIANCE: The admissibility of an update u to Γ with constant complement Γ' should be independent of the state of Γ' . This mirrors the idea of a *closed* view, as described in [13], in which the user of the view only needs to know the state of the view itself to determine whether or not an update to it is allowed.

REFLECTION INVARIANCE; If $\Gamma'' \in \mathcal{V}$ is another complement of Γ , and the update u to Γ is also supported by constant complement Γ'' , then the reflection of the view update to the main schema should be the same in each case.

UNIVERSALITY: If $\Gamma'' \in \mathcal{V}$ is another complement of Γ , and the update u to Γ is supported by constant-complement Γ'' , then u should also be supported by constant complement Γ' .

The first two of these topics have already been studied extensively. State invariance is obtained precisely in the case that the congruences (i.e., the equivalence relations defined by the view mappings) of the two views commute. Basic ideas along these lines were discussed already in [3, Def. 14 and Thm. 14], although a specific characterization in terms of the meet is first presented in [12] and refined greatly in [13], where it is also shown that for the case of projections of relational schemata constrained by usual database dependencies, state invariance is obtained precisely in the case that a cover of the dependencies of the main schema embed into the two views [13, Prop. 2.17].

Reflection invariance is also investigated extensively in [13], where it is shown that order structure is critical. In particular, for the case of relational schemata, if the constraints are usual database dependencies and the views are defined as combinations of select, project, join, and rename, then reflection invariance is assured [13, Prop. 2.5 and Thm. 4.3].

Comparable studies of universality are not to be found in the literature. Formally, a universal complement is one which is least within a given class \mathcal{V} of views, since the smaller the complement, the greater the set of updates supported. For \mathcal{V} the set of projections on a relational schema \mathbf{E} , this means that the projection $\Pi_{\mathbf{W}'}^{\mathbf{E}}$ onto attributes \mathbf{W}' is a universal complement of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ if for any other complement $\Pi_{\mathbf{W}''}^{\mathbf{E}}$, $\mathbf{W}' \subseteq \mathbf{W}''$. A study of universality thus amounts to a study of least complements within a given class \mathcal{V} . In [15], conditions which guarantee universality for projections are developed in terms of so-called *governing* join dependencies, but they are difficult to translate into more fundamental properties. The primary goal of this paper is to seek a more concrete characterization of universality, in terms of the basic properties of schemata and views. In contrast to [15], the work here requires state invariance as a precondition for all complements. Thus, \mathcal{V} is restricted to complements which exhibit state invariance – the so-called *meet complements*. The reason is that without it, there are simple counterexamples to universality, such as the following, taken from [15]. Let \mathbf{E}_0 be the relational schema with the single relation symbol $R[ABCD]$, governed by the functional dependencies (FDs) in $\mathcal{F}_0 = \{B \rightarrow D, C \rightarrow D\}$. Observe that this schema is completely symmetric with respect to attributes B and C . Consider the view $\Pi_{ABC}^{\mathbf{E}_0}$, the projection onto attributes ABC . The

projection $\Pi_{CD}^{\mathbf{E}_0}$ onto CD is a minimal complement. Indeed, $ABC \cap CD = C$ is a key for the schema of $\Pi_{CD}^{\mathbf{E}_0}$, and so from the classical result of Rissanen [25, Thm. 1] it follows that the decomposition is lossless. That $\Pi_{CD}^{\mathbf{E}_0}$ is a minimal complement amongst projections is immediate, since neither $\Pi_C^{\mathbf{E}_0}$ nor $\Pi_D^{\mathbf{E}_0}$ can be complements. On the other hand, given the symmetry in B and C , it follows that $\Pi_{BD}^{\mathbf{E}_0}$ is also a minimal complement of $\Pi_{ABC}^{\mathbf{E}_0}$. Thus, there can be no universal complement amongst the projections of \mathbf{E}_0 . This implies that there are some updates to $\Pi_{ABC}^{\mathbf{E}_0}$ which are translatable with constant complement $\Pi_{BD}^{\mathbf{E}_0}$ but not with constant complement $\Pi_{CD}^{\mathbf{E}_0}$, and conversely. Specifically, no update which does not keep $\Pi_B^{\mathbf{E}_0}$ constant can be realized with constant complement $\Pi_{BD}^{\mathbf{E}_0}$, and no update which does not keep $\Pi_C^{\mathbf{E}_0}$ constant can be realized with constant complement $\Pi_{CD}^{\mathbf{E}_0}$. The only updates to $\Pi_{ABC}^{\mathbf{E}_0}$ which keep both complements constant are those in which neither the BD nor the CD projection is allowed to vary. On the other hand, in view of the results of [13], this example does however exhibit reflection invariance — updates to $\Pi_{ABC}^{\mathbf{E}_0}$ which keep both CD and BD constant; i.e., those which change only the values on attribute A , are reflected in the same way with complement $\Pi_{BD}^{\mathbf{E}_0}$ as with $\Pi_{CD}^{\mathbf{E}_0}$.

While it is possible to find examples of situations which exhibit both update-state and reflection invariance, yet deny a universal complement, it is established in this paper that there are reasonable and nontrivial conditions which do guarantee a universal complement. The context is the classical setting of universal relational schemata constrained by FDs, with views defined by simple (i.e., single) projections. The properties which guarantee universality are of covers of the FDs governing the main schema.

The approach is divided into three main phases. In the first, a strong condition is identified which ensures that every projection $\Pi_{\mathbf{W}}^{\mathbf{E}}$ of the universal-relational schema \mathbf{E} has a unique minimal *cover precomplement*; that is, a second projection $\Pi_{\mathbf{W}'}^{\mathbf{E}}$, with the property that a cover of the FDs of \mathbf{E} embed into these views. In the second phase, conditions are established for extension of a such a minimal precomplement to a universal complement. In the third and final phase, it is shown how to relax the conditions of the first two phases to allow for equivalences, and thus include common constructions such as multiple candidate keys. The first two phases are covered in Sec. 3, while the third is presented in 4. Sec. 2 provides background information and Sec. 5 provides conclusions and further directions.

2 Schemata, Views, and Complements

The setting for of this work is the classical relational model, with which it is assumed that the reader is familiar. For any necessary clarification, [21] provides perhaps the best presentation. However, other classical textbooks, such as [24] and [1] should prove more than adequate as well.

The purpose of this section is primarily to present necessary notation and terminology, as well as constructs which build upon that classical model, par-

ticularly on the subject of views. Thus, this section need only be skimmed on a first reading, and then used as a reference as the need arises.

Summary 2.1 (Attributes and relation schemes). Attributes will normally be written using (possibly subscripted and/or primed) capital letters from the beginning of the alphabet. Thus, A , B_1 , and C' denote attributes. Sets of attributes will normally be written using boldface letters from the end of the alphabet; thus \mathbf{U} , \mathbf{W}' , and \mathbf{V}_1 denote sets of attributes. Following the long-standard notation, sets of attributes are represented via concatenation. Thus, ABC is shorthand for $\{A, B, C\}$. In particular, a single attribute name, such as A , represents both attribute itself and the singleton set $\{A\}$. Context will always make clear which is intended.

A *relation scheme* is defined by a name and a nonempty set of attributes. Thus, $R[ABC]$ is a relation scheme. Constraints are not part of a relation scheme.

It is assumed that the reader is familiar with the notions of domain for an attribute and tuple for a relation scheme. These notions will be used only tangentially, and so there is no need to develop a notation carefully.

Summary 2.2 (Functional dependencies). Familiarity with the basic syntax and semantics of functional dependencies (FDs) is assumed. The FD $\mathbf{X} \rightarrow \mathbf{Y}$ is *degenerate* if $\mathbf{X} = \emptyset$. In this work, degenerate FDs will not be considered, so the term *FD* should always be taken to mean *nondegenerate FD*. The FD $\mathbf{X} \rightarrow \mathbf{Y}$ is *trivial* if $\mathbf{Y} \subseteq \mathbf{X}$. Since trivial FDs always hold, they are not particularly interesting, although they may arise in certain constructions.

$\mathcal{F} \models \varphi$ denotes that the set \mathcal{F} of FDs *entails* or *implies* the FD φ ; i.e., φ holds whenever \mathcal{F} holds. For a second set \mathcal{F}' of FDs, write $\mathcal{F} \models \mathcal{F}'$ if $\mathcal{F} \models \varphi$ for every $\varphi \in \mathcal{F}'$. Write $\mathcal{F} \models \mathcal{F}'$ to denote that both $\mathcal{F} \models \mathcal{F}'$ and $\mathcal{F}' \models \mathcal{F}$ hold, and say that \mathcal{F} and \mathcal{F}' are *equivalent* in this case.

Given an FD $\varphi = \mathbf{X} \rightarrow \mathbf{Y}$, define $\text{LHS}(\varphi) = \mathbf{X}$, $\text{RHS}(\varphi) = \mathbf{Y}$, and $\text{Attr}(\varphi) = \text{LHS}(\varphi) \cup \text{RHS}(\varphi)$. This notation is extended to sets of FDs in a natural way: $\text{LHS}(\mathcal{F}) = \{\text{LHS}(\varphi) \mid \varphi \in \mathcal{F}\}$; $\text{RHS}(\mathcal{F}) = \{\text{RHS}(\varphi) \mid \varphi \in \mathcal{F}\}$; $\text{Attr}(\mathcal{F}) = \{\text{Attr}(\varphi) \mid \varphi \in \mathcal{F}\}$. Furthermore, define $\text{Attr}_{\cup}(\mathcal{F}) = \bigcup \text{Attr}(\mathcal{F})$.

The FD φ is *RHS-simple* if $\text{RHS}(\varphi)$ consists of exactly one attribute, and the set \mathcal{F} of FDs is *RHS-simple* if each of its members has that property. φ is *simple* if both $\text{LHS}(\varphi)$ and $\text{RHS}(\varphi)$ consist of exactly one attribute.

Given a set \mathcal{F} of FDs, the *closure* of \mathcal{F} , denoted \mathcal{F}^+ , is the set of all FDs which are implied by those in \mathcal{F} . The set \mathcal{F} is *nonredundant* if for any $\mathcal{F}' \subseteq \mathcal{F}$, if $\mathcal{F}'^+ = \mathcal{F}^+$, then $\mathcal{F}' = \mathcal{F}$.

An FD $\mathbf{X} \rightarrow \mathbf{Y} \in \mathcal{F}^+$ is *LHS-reduced* or *full* for \mathcal{F} if for any $\mathbf{X}' \subseteq \mathbf{X}$, if $(\mathcal{F} \setminus \{\mathbf{X} \rightarrow \mathbf{Y}\}) \cup \{\mathbf{X}' \rightarrow \mathbf{Y}\}^+ = \mathcal{F}^+$, then $\mathbf{X}' = \mathbf{X}$. The set \mathcal{F} is *LHS-reduced* if each of its members has that property with respect to \mathcal{F} .

If \mathcal{F} is nonredundant, LHS-reduced, and RHS-simple, it is called a *canonical* set of FDs.

A *cover* of \mathcal{F} is any set \mathcal{C} of FDs with the property that $\mathcal{C}^+ = \mathcal{F}^+$; or, equivalently $\mathcal{C} \models \mathcal{F}$. The set \mathcal{F} is *minimum* if each of its covers contains at least as many FDs as \mathcal{F} itself.

A *context* is a pair $\langle \mathbf{U}, \mathcal{F} \rangle$ in which \mathbf{U} is a nonempty set of attributes and \mathcal{F} is set of FDs with $\text{Attr}_{\cup}(\mathcal{F}) \subseteq \mathbf{U}$. Notions such as entailment and closure may be used within a context, provided a little care is taken. Formally, define $\text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle) = \mathcal{F} \cup \{A \rightarrow A \mid A \in \mathbf{U} \setminus \text{Attr}_{\cup}(\mathcal{F})\}$, and $\langle \mathbf{U}, \mathcal{F} \rangle^+ = \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$. Thus, $\text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)$ is \mathcal{F} augmented with trivial FDs which involve the attributes of \mathbf{U} not found in any FD of \mathcal{F} .

A *superkey* for $\langle \mathbf{U}, \mathcal{F} \rangle$ is a $\mathbf{K} \subseteq \mathbf{U}$ with the property $\mathbf{K} \rightarrow A \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$ for every $A \in \mathbf{U}$. A minimal superkey (with respect to subset inclusion) is called a *key*. Note that the property of being a superkey or key depends upon both \mathbf{U} and \mathcal{F} , and not upon \mathcal{F} alone, since there may be attributes in \mathbf{U} which do not occur in \mathcal{F} . However, it is easy to see that \mathbf{K} is a key (resp. superkey) for $\langle \mathbf{U}, \mathcal{F} \rangle$ iff $\mathbf{K} \cap \text{Attr}_{\cup}(\mathcal{F})$ is a key (resp. superkey) for $\langle \text{Attr}_{\cup}(\mathcal{F}), \mathcal{F} \rangle$. In other words, a key (resp. superkey) for $\langle \mathbf{U}, \mathcal{F} \rangle$ may be obtained from one for $\langle \text{Attr}_{\cup}(\mathcal{F}), \mathcal{F} \rangle$ by adding those attributes in \mathbf{U} which do not occur in any FD of \mathcal{F} . Thus, a *key* (resp. *superkey*) for \mathcal{F} may and will be defined unambiguously to be a key for $\langle \text{Attr}_{\cup}(\mathcal{F}), \mathcal{F} \rangle$.

If $\mathbf{X} \rightarrow \mathbf{Y} \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)$ with $\mathbf{X} \subseteq \mathbf{Y}$, it will also be said that \mathbf{X} is a *superkey* for \mathbf{Y} .

$\mathbf{X} \leftrightarrow \mathbf{Y}$ is an abbreviation for $\mathbf{X} \rightarrow \mathbf{Y} \wedge \mathbf{Y} \rightarrow \mathbf{X}$, and the statement $\mathbf{X} \leftrightarrow \mathbf{Y} \in \mathcal{F}^+$ is used as a synonym for $\{\mathbf{X} \rightarrow \mathbf{Y}, \mathbf{Y} \rightarrow \mathbf{X}\} \subseteq \mathcal{F}^+$. In this case, it is said that \mathbf{X} and \mathbf{Y} are *equivalent* in \mathcal{F} . If both \mathbf{X} and \mathbf{Y} are singletons, so that $\mathbf{X} \leftrightarrow \mathbf{Y}$ is of the form $A \leftrightarrow B$, then it is called a *simple equivalence*.

Summary 2.3 (Relational Schemata and Π -Views). A *universal relational schema* is a pair $\mathbf{E} = (R[\mathbf{U}], \text{Constr}(\mathbf{E}))$ in which $R[\mathbf{U}]$ is a relation scheme and $\text{Constr}(\mathbf{E})$ is a set of dependencies on $R[\mathbf{U}]$. $\text{LDB}(\mathbf{E})$ denotes the set of all *legal databases* of \mathbf{E} ; that is, all sets of tuples on $R[\mathbf{U}]$ which satisfy the constraints in $\text{Constr}(\mathbf{E})$. If $\text{Constr}(\mathbf{E})$ is a set of FDs, then \mathbf{E} is called a *universal FD-schema*, with $\langle \mathbf{U}, \mathcal{F} \rangle$ the *context* of \mathbf{E} .

It is worth noting that the universal FD-schema $\mathbf{E} = (R[\mathbf{U}], \mathcal{F})$ and its context $\langle \mathbf{U}, \mathcal{F} \rangle$ determine each other, up to the name of the relation R . However, it nevertheless seems appropriate to maintain this distinction. With $\langle \mathbf{U}, \mathcal{F} \rangle$, the focus is on the dependencies and nothing more, while with \mathbf{E} , notions such as $\text{LDB}(\mathbf{E})$ (which involve the relation symbol R) are central.

A Π -view of the universal relational schema \mathbf{E} is a pair $\Pi_{\mathbf{W}}^{\mathbf{E}} = (\mathbf{E}^{\mathbf{W}}, \pi_{\mathbf{W}}^{\mathbf{E}})$ in which $\mathbf{E}^{\mathbf{W}} = (R_{\mathbf{W}}[\mathbf{W}], \text{Constr}(\mathbf{E}^{\mathbf{W}}))$ is a relational schema and $\pi_{\mathbf{W}}^{\mathbf{E}}$ is the projection morphism which sends a relation on attribute set \mathbf{U} to its projection onto attribute set \mathbf{W} . The set $\text{Constr}(\mathbf{E}^{\mathbf{W}})$ consists of the *implied constraints* [17] on the view schema. The notation $\text{LDB}(\Pi_{\mathbf{W}}^{\mathbf{E}})$ is used as a synonym for $\text{LDB}(\mathbf{E}^{\mathbf{W}})$.

Now assume further that \mathbf{E} is a universal FD-schema. It is unfortunately not the case that $\mathbf{E}^{\mathbf{W}}$ will always have that same property. Indeed $\text{Constr}(\mathbf{E}^{\mathbf{W}})$ may not have a finite cover and thus need not consist only of functional dependencies [16, Lem. 4.1], [14, App. A]. Call $\Pi_{\mathbf{W}}^{\mathbf{E}}$ *closed* if $\text{Constr}(\mathbf{E}^{\mathbf{W}})$ is equivalent to a subset of \mathcal{F}^+ ; that is, if the implied constraints on the view have a basis consisting of FDs.

Fortunately, the views to be considered in this work will always be closed. Specifically, if $\Pi_{\mathbf{W}}^{\mathbf{E}}$ has a meet complement (as defined in Summary 2.4 below), then it is closed [13, Prop. 2.17]. Since the complements constructed in this work will always be meet complements, it follows that all views of universal FD-schemata which are considered in this work will be closed.

There is a natural order¹ on Π -views of \mathbf{E} given by $\Pi_{\mathbf{W}_1}^{\mathbf{E}} \preceq \Pi_{\mathbf{W}_2}^{\mathbf{E}}$ iff $\mathbf{W}_1 \subseteq \mathbf{W}_2$. Similarly, $\Pi_{\mathbf{W}_1}^{\mathbf{E}} \prec \Pi_{\mathbf{W}_2}^{\mathbf{E}}$ iff $\mathbf{W}_1 \subsetneq \mathbf{W}_2$; i.e., $\mathbf{W}_1 \subseteq \mathbf{W}_2$ but $\mathbf{W}_1 \neq \mathbf{W}_2$.

Summary 2.4 (Complementary and covered pairs of views). For a comprehensive presentation of complementary views based upon congruences, see [13] and [14], and for a general characterization of view ordering, also based upon congruences, see [15, Summary 2.1]. In this paragraph, only the ideas necessary to the special setting of this paper are presented, and in some cases, the definitions themselves are recast to capture this special setting only.

Let $\mathbf{E} = (R[\mathbf{U}], \mathcal{F})$ be a universal FD-schema, and let $\Pi_{\mathbf{W}_i}^{\mathbf{E}} = (\mathbf{E}^{\mathbf{W}_i}, \pi_{\mathbf{W}_i}^{\mathbf{E}})$ be Π -views of \mathbf{E} for $i \in \{1, 2\}$. The pair $P = \{\Pi_{\mathbf{W}_1}^{\mathbf{E}}, \Pi_{\mathbf{W}_2}^{\mathbf{E}}\}$ is *complementary* if it is possible to recover the state of \mathbf{E} from the combined states of the two views. In light of the classical result [25, Thm. 1], this condition may be recaptured by defining P to be a *complementary pair* if $\mathbf{W}_1 \cup \mathbf{W}_2 = \mathbf{U}$ and at least one of the FDs $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_1$, $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_2$ is in \mathcal{F}^+ . In this paper, this classical characterization will be used as the definition of a complementary pair.

The members of a complementary pair are called *Π -complements* of each other. As only Π -views are considered in this work, the term *complement* should always be interpreted as *Π -complement*, unless stated explicitly to the contrary.

The set \mathcal{F} *embeds into* P if every $\varphi \in \mathcal{F}$ embeds into at least one of the views in P (i.e., $\text{Attr}(\varphi) \subseteq \mathbf{W}_i$ for $i = 1$ or $i = 2$.) An *embedded cover* of \mathcal{F} for P is a set \mathcal{C} of FDs which is a cover for \mathcal{F} and each element of \mathcal{C} embeds in P . Call P a *covered pair* if $\mathbf{W}_1 \cup \mathbf{W}_2 = \mathbf{U}$ and there is an embedded cover of \mathcal{F} for P . A covered pair need not be complementary. For example, if $\mathcal{F} = \emptyset$, then P is trivially a covered pair, but it is not a complementary pair unless $\mathbf{W}_1 = \mathbf{U}$ or $\mathbf{W}_2 = \mathbf{U}$.

Call P a *meet-complementary pair* if it is a complementary and covered pair. In this case, the Π -view $\Pi_{\mathbf{W}_1 \cap \mathbf{W}_2}^{\mathbf{E}}$ defined by the attributes common to both $\Pi_{\mathbf{W}_1}^{\mathbf{E}}$ and $\Pi_{\mathbf{W}_2}^{\mathbf{E}}$ is called the *meet* of P . Meet complementary pairs have the following important *independence* property: If $N_1 \in \text{LDB}(\Pi_{\mathbf{W}_1}^{\mathbf{E}})$ and $N_2 \in \text{LDB}(\Pi_{\mathbf{W}_2}^{\mathbf{E}})$ have the property that they agree on the meet $\Pi_{\mathbf{W}_1 \cap \mathbf{W}_2}^{\mathbf{E}}$, then there is a (necessarily unique) $M \in \text{LDB}(\mathbf{E})$ with $\pi_{\mathbf{W}_i}^{\mathbf{E}}(M) = N_i$ for $i \in \{1, 2\}$ [13, Thm. 2.14 and Prop. 2.17].

If P is a complementary pair (resp. covered pair, resp. meet-complementary pair), then $\Pi_{\mathbf{W}_1}^{\mathbf{E}}$ and $\Pi_{\mathbf{W}_2}^{\mathbf{E}}$ are called *complements* (resp. *cover precomplements*, resp. *meet complements*) of each other. Minimality and optimality are defined with respect to the ordering \preceq [10, 1.23]. Specifically, $\Pi_{\mathbf{W}_2}^{\mathbf{E}}$ is a *minimal Π -*

¹ The definition given here is specialized to projections only. For a general definition based upon congruences, see [15, Summary 2.1].

complement (resp. cover Π -precomplement, resp. meet Π -complement) of $\Pi_{\mathbf{W}_1}^{\mathbf{E}}$ if it is a minimal element with respect to the ordering \preceq when restricted to all complements (resp. cover precomplements, resp. meet complements) of $\Pi_{\mathbf{W}_1}^{\mathbf{E}}$. Similarly, $\Pi_{\mathbf{W}_2}^{\mathbf{E}}$ is an *optimal* Π -complement (resp. cover Π -precomplement, resp. meet Π -complement) of $\Pi_{\mathbf{W}_2}^{\mathbf{E}}$ if it is a least element with respect to the ordering \preceq when restricted to all complements (resp. cover precomplements, resp. meet complements) of $\Pi_{\mathbf{W}_1}^{\mathbf{E}}$. See [15, Summary 2.2] for an elaboration of this idea.

A minimal (resp. optimal) meet Π -complement is also called a *fully minimal* (resp. *fully optimal*) Π -complement. A fully optimal Π -complement satisfies all three conditions identified in the introduction: state-invariance, reflection invariance, and universality, while a fully minimal Π -complement need only satisfy the first two. A fully optimal Π -complement is also called a *universal Π -complement*. Since the smaller the complement, the more updates which are supported via constant complement, a universal Π -complement supports all view updates which any other complement does. Thus, this formalization recaptures the notion of universality identified in Sec. 1, when restricted to meet complements.

These relationships need not be symmetric. For example, let \mathbf{E}_1 have a $R[ABCD]$ constrained by the FDs in $\mathcal{F}_1 = \{A \rightarrow B, B \rightarrow C, C \rightarrow D\}$. The universal Π -complement of $\Pi_{ABC}^{\mathbf{E}_1}$ is $\Pi_{BCD}^{\mathbf{E}_1}$, while the universal Π -complement of $\Pi_{BCD}^{\mathbf{E}_1}$ is $\Pi_{AB}^{\mathbf{E}_1}$.

3 Universal Π -Complements and Unique Covers

Notation 3.1 (Notational convention). Unless specifically stated to the contrary, for the rest of this paper, take $\mathbf{E} = (R[\mathbf{U}], \mathcal{F})$ to be a universal FD-schema. In particular, \mathbf{U} is a nonempty set of attributes and \mathcal{F} is a set of FDs over \mathbf{U} .

Definition 3.2 (Cover precomplements with respect to fixed set of FDs). Let $\mathbf{W} \subseteq \mathbf{U}$. Suppose that the goal is to identify a cover precomplement $\Pi_{\mathbf{W}'}^{\mathbf{E}}$ of $\Pi_{\mathbf{W}}^{\mathbf{E}}$. The first requirement is that $\mathbf{W} \cup \mathbf{W}' = \mathbf{U}$, and this implies in particular that every attribute of \mathbf{U} which is not in \mathbf{W} must lie in \mathbf{W}' . The set of all such attributes is called the *base complement* of \mathbf{W} with respect to \mathbf{U} , and it occurs often enough to warrant its own notation.

$$(a) \text{BaseCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle = \mathbf{U} \setminus \mathbf{W}.$$

The second requirement for a precomplement is that a cover of \mathcal{F} embed into the pair $\{\Pi_{\mathbf{W}}^{\mathbf{E}}, \Pi_{\mathbf{W}'}^{\mathbf{E}}\}$. As a first step, consider the simpler problem of identifying the least \mathbf{W}' with the property that each element of \mathcal{F} embed into the pair. In other words, \mathcal{F} is to be used as its own embedded cover. Although the end result of a solution to that problem is not a goal in itself, the associated construction will prove to be central in the more general case. In this case, all FDs in \mathcal{F} which do not embed into the view $\Pi_{\mathbf{W}}^{\mathbf{E}}$ must embed into $\Pi_{\mathbf{W}'}^{\mathbf{E}}$. This implies that the

attributes in the *dependency-preserving extension* for \mathbf{W} with respect to $\langle \mathbf{U}, \mathcal{F} \rangle$, defined as follows, must be a subset of \mathbf{W}' .

$$(b) \text{DPExt}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle = \bigcup \{ \mathbf{V} \in \text{Attr}(\mathcal{F}) \mid \mathbf{V} \not\subseteq \mathbf{W} \}.$$

Putting these two together, the formula for the attributes of \mathbf{W}' is the *pre-complement set* of \mathbf{W} with respect to $\langle \mathbf{U}, \mathcal{F} \rangle$, defined as follows.

$$(c) \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle = \text{BaseCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle \cup \text{DPExt}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle.$$

Owing to the importance of this construction, it is recorded as a formal result.

Proposition 3.3. *For any $\mathbf{W} \subseteq \mathbf{U}$, $\Pi_{\mathbf{X}}^{\mathbf{E}}$ with $\mathbf{X} = \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle$ is the least Π -view of E for which \mathcal{F} embeds into $\{\Pi_{\mathbf{W}}^{\mathbf{E}}, \Pi_{\mathbf{X}}^{\mathbf{E}}\}$. It is thus a Π -precomplement of $\Pi_{\mathbf{W}}^{\mathbf{E}}$. \square*

Example 3.4. Let $\mathbf{E}_1 = (R[\mathbf{U}_1], \mathcal{F}_1)$ with $\mathbf{U}_1 = ABCDEF$ and $\mathcal{F}_1 = \{A \rightarrow B, B \rightarrow C, CD \rightarrow E\}$. Let $\mathbf{W} = ABCE$, so the view whose precomplement is sought is $\Pi_{\mathbf{W}}^{\mathbf{E}_1} = \Pi_{ABCE}^{\mathbf{E}_1}$. Then $\text{BaseCompl}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle} \langle ABCE \rangle = ABCDEF \setminus ABCE = DF$ and $\text{DPExt}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle} \langle ABCE \rangle = CDE$, and so the least precomplement of $\Pi_{ABCE}^{\mathbf{E}_1}$ which preserves \mathcal{F}'_1 is $\Pi_{CDE \cup DF}^{\mathbf{E}_1} = \Pi_{CDEF}^{\mathbf{E}_1}$.

Discussion 3.5 (The inadequacy of canonical and minimal covers).

Proposition 3.3 only identifies a precomplement which is optimal with respect to a fixed set \mathcal{F} of FDs. In general, such a set of FDs will have many covers, with each giving rise to a possibly distinct precomplement. An optimal cover precomplement must be least when taken over all possible covers of \mathcal{F} . A potential strategy is to work with a cover of \mathcal{F} which is normalized in some way which guarantees that the precomplement which it induces is no larger (in terms of \preceq) than any other.

There are a number of results which suggest that such an approach might be fruitful. In one of the earliest papers on the theory of FDs [5], it is shown that that any two nonredundant covers \mathcal{C}_1 and \mathcal{C}_2 of \mathcal{F} have the remarkable property that for any FD $\varphi_1 \in \mathcal{C}_1$, there is an FD $\varphi_2 \in \mathcal{C}_2$ with the property that $\text{LHS}(\varphi_1) \leftrightarrow \text{LHS}(\varphi_2)$. Thus, up to FD equivalence, the left-hand sides of the FDs in any two nonredundant covers are the same. In [20], this result is strengthened: for any two minimum sets \mathcal{C}_1 and \mathcal{C}_2 of FDs, there is a bijective correspondence between the equivalent left-hand sides of each. In other words, the number of occurrences of an equivalent left-hand side in each set is the same.

Armed with these results, it would appear to be possible to work with minimum covers and then show that any two minimal precomplements differ only by an equivalence of attribute sets. Unfortunately, there are serious obstacles in the path of this potential tactic. The following examples illustrate the problems.

Examples 3.6 (The need for unique covers). Let \mathbf{E}_2 be the schema whose single relation symbol is $R[\mathbf{U}_2]$ with $\mathbf{U}_2 = ABCDE$, governed by the FDs in $\mathcal{F}_2 = \{A \rightarrow BCDE, CE \rightarrow D, CD \rightarrow E\}$. Then each of $\mathcal{F}_{21} = \{A \rightarrow BCD, CE \rightarrow$

$D, CD \rightarrow E\}$ and $\mathcal{F}_{22} = \{A \rightarrow BCE, CE \rightarrow D, CD \rightarrow E\}$ is a minimum cover of \mathcal{F} . Rewriting these in RHS-simple format makes this clearer:

$$\begin{aligned}\mathcal{F}'_{21} &= \{A \rightarrow B, A \rightarrow C, A \rightarrow D, CE \rightarrow D, CD \rightarrow E\} \\ \mathcal{F}'_{22} &= \{A \rightarrow B, A \rightarrow C, A \rightarrow E, CE \rightarrow D, CD \rightarrow E\}\end{aligned}$$

Indeed, $A \rightarrow E$ may be derived from $A \rightarrow CD \rightarrow E$ and $A \rightarrow D$ may be derived from $A \rightarrow CE \rightarrow D$. Thus, either $A \rightarrow D$ or else $A \rightarrow E$, but not both, may be removed from \mathcal{F}_2 while retaining the same closure. From this observation and Proposition 3.3 it follows that $\Pi_{BCDE}^{\mathbf{E}_2}$ has two distinct minimal Π -precomplements, $\Pi_{ABCD}^{\mathbf{E}_2}$ and $\Pi_{ABCE}^{\mathbf{E}_2}$.

Note that the left-hand sides of the \mathcal{F}_{2i} 's are exactly the same in each case. Indeed, it is not difficult to see that $\text{LHS}(\mathcal{C}_1) = \text{LHS}(\mathcal{C}_2)$ for any two nonredundant minimum covers \mathcal{C}_1 and \mathcal{C}_2 of \mathcal{F}_2 . It is the right-hand sides and only the right-hand sides which differ, and which lead to alternative minimal precomplements. Furthermore, $\mathcal{F}_2 \not\models D \leftrightarrow E$, so there can be no equivalence of the right-hand sides. Indeed, the two cover precomplements $\Pi_{ABCD}^{\mathbf{E}_2}$ and $\Pi_{ABCE}^{\mathbf{E}_2}$ are not equivalent in any simple sense.

This example does not rule out the possibility of identifying situations in which the right-hand sides of the FDs remain constant while the left-hand sides vary amongst equivalent sets. However, the following example suggests that this approach might be problematic in a more general setting.

Let \mathbf{E}_3 be the schema whose single relation symbol is $R[\mathbf{U}_3]$, with $\mathbf{U}_3 = A_1A_2B_1B_2CD$, governed by the FDs in $\mathcal{F}_3 = \{A_1B_1 \leftrightarrow A_2B_2, A_1B_1 \rightarrow C, A_2B_2 \rightarrow C, B_2 \rightarrow D\}$. A minimum cover is obtained by removing one of the FDs $A_1B_1 \rightarrow C$ or $A_2B_2 \rightarrow C$, but not both. Thus, the following two minimum covers are obtained.

$$\begin{aligned}\mathcal{F}'_{31} &= \{A_1B_1 \leftrightarrow A_2B_2, A_1B_1 \rightarrow C, B_2 \rightarrow D\} \\ \mathcal{F}'_{32} &= \{A_1B_1 \leftrightarrow A_2B_2, A_2B_2 \rightarrow C, B_2 \rightarrow D\}\end{aligned}$$

Let $\mathbf{W}_3 = A_1B_1A_2B_2$, and consider the problem of finding an optimal precomplement of $\Pi_{\mathbf{W}_3}^{\mathbf{E}_3}$. Using the construction of Definition 3.2 and Proposition 3.3, the two precomplements $\Pi_{\mathbf{W}_{31}}^{\mathbf{E}_3}$ and $\Pi_{\mathbf{W}_{32}}^{\mathbf{E}_3}$, with respect to $\langle \mathbf{U}_3, \mathcal{F}'_{31} \rangle$ and $\langle \mathbf{U}_3, \mathcal{F}'_{32} \rangle$ respectively, have $\mathbf{W}_{31} = A_1B_1B_2CD$ and $\mathbf{W}_{32} = A_2B_2CD$. While these two sets are equivalent for \mathcal{F}_3 , the relationship is not simple. The problem is that while $\mathcal{F} \models A_1B_1 \leftrightarrow A_2B_2$, A_1B_1 and A_2B_2 do not have the same substructure. If A_2B_2 is used, then B_2 may serve double duty as the left-hand side of $B_2 \rightarrow D$. On the other hand, if A_1B_1 is used, then B_2 must be included to cover $B_2 \rightarrow D$. Thus there is a sense in which A_1B_1 and A_2B_2 are not fully equivalent, and so it seems a bit of a stretch to regard $\Pi_{\mathbf{W}_{31}}^{\mathbf{E}_3}$ and $\Pi_{\mathbf{W}_{32}}^{\mathbf{E}_3}$ as similar enough to be the ‘‘same’’ precomplement under some form of isomorphism. For this reason, this sort of equivalence will not be allowed in the first instance. Rather, the approach is to first rule out allowing distinct precomplements which are equivalent via FD associations, and to reintroduce a limited form of equivalence subsequently as a special construction.

Discussion 3.7 (Indecomposable FDs). Call the RHS-simple FD $\varphi \in \mathcal{F}^+$ *indecomposable* with respect to \mathcal{F} if for every RHS-simple cover \mathcal{C} of \mathcal{F} , $\varphi \in \mathcal{C}$. Put another way, if $\varphi = \mathbf{X} \rightarrow A$, then for every cover \mathcal{C} of \mathcal{F} , there is an FD $\mathbf{X} \rightarrow \mathbf{Y} \in \mathcal{C}$ with $A \in \mathbf{Y}$. Define $\text{Indecomp}(\mathcal{F})$ to be the set of all RHS-simple members of \mathcal{F}^+ which are indecomposable with respect to that set.

The following is immediate, invoking Proposition 3.3 for (b).

Proposition 3.8 (Indecomposable FDs and unique covers).

- (a) \mathcal{F} admits a unique canonical cover iff $\text{Indecomp}(\mathcal{F})$ is such a cover.
- (b) If \mathcal{F} admits a unique canonical cover \mathcal{C} , then for any $\mathbf{W} \subseteq \mathbf{U}$, $\Pi_{\mathbf{W}}^{\mathbf{E}}$ is an optimal cover Π -precomplement of $\Pi_{\mathbf{W}'}^{\mathbf{E}}$ for $\mathbf{W}' = \text{PreCompl}_{\langle \mathbf{U}, \mathcal{C} \rangle}(\mathbf{W})$. \square

Example 3.9 (Additional attributes necessary for meet complements).

The example schema \mathbf{E}_1 of Definition 3.2 provides a setting with unique canonical covers. Indeed, \mathcal{F}_1 is its own unique canonical cover. Thus, in particular, $\Pi_{CDEF}^{\mathbf{E}_1}$ is the optimal cover Π -precomplement of $\Pi_{ABCE}^{\mathbf{E}_1}$.

However, $\Pi_{CDEF}^{\mathbf{E}_1}$ is not a complement of $\Pi_{ABCE}^{\mathbf{E}_1}$, since $ABCE \cap CDEF = CE$ is not a key for either $ABCE$ or $CDEF$ (see Summary 2.4). To remedy this situation, a set \mathbf{X} of attributes must be added to $CDEF$ to render the attributes in $ABCE \cap (CDEF \cup \mathbf{X})$ a superkey for $ABCE$. The unique minimal solution is $\mathbf{X} = A$, so the optimal meet complement of $\Pi_{ABCE}^{\mathbf{E}_1}$ is $\Pi_{ACDEF}^{\mathbf{E}_1}$.

Definition 3.10 (Extension to optimal meet complements). To formalize this process, a few definitions are helpful. For these definitions, return to the general setting of Notation 3.1.

- (a) $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) = \mathbf{W} \cap \text{DPExt}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$.
- (b) $\text{DPMutRHS}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) = \{A \in \mathbf{W} \mid \text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \rightarrow A \in \mathcal{F}^+\}$.
- (c) $\text{DPMut}\overline{\text{RHS}}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) = \mathbf{W} \setminus \text{DPMutRHS}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$.

Thus, $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is the set of attribute which are mutual (i.e., common) to $\Pi_{\mathbf{W}}^{\mathbf{E}_1}$ and its optimal cover precomplement for \mathcal{F} . $\text{DPMutRHS}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is the set of attributes in \mathbf{W} which are already functionally determined by $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$, and $\text{DPMut}\overline{\text{RHS}}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is the set of attributes in \mathbf{W} which are not functionally determined by $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$.

Example 3.11 (Continuation of Example 3.9). In Example 3.9, the definitions of Definition 3.10 yield $\text{DPMut}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle}(ABCE) = CE$, $\text{DPMutRHS}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle}(ABCE) = CE$, and $\text{DPMut}\overline{\text{RHS}}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle}(ABCE) = ABCE \setminus CE = AB$. Only a subset of $\text{DPMut}\overline{\text{RHS}}_{\langle \mathbf{U}_1, \mathcal{F}_1 \rangle}(ABCE)$ which generates all of its attributes via FD closure which must be added to the mutual elements in order to obtain a meet complement. In this case, only A need be added to the mutual elements, since $A \rightarrow B$. Hence the optimal meet complement of $\Pi_{ABCE}^{\mathbf{E}_1}$ is $\Pi_{A \cup CDEF}^{\mathbf{E}_1} = \Pi_{ACDEF}^{\mathbf{E}_1}$. Unfortunately, such a unique minimal set of generating elements need not exist, even in the case that the canonical cover is unique, as shown by the example of the next paragraph.

Example 3.12 (Non-uniqueness of extension from optimal cover pre-complement to meet complement). Let $\mathbf{E}_4 = (R[\mathbf{U}_4], \mathcal{F}_4)$ with $\mathbf{U}_4 = ABCD$ and $\mathcal{F}_4 = \{AB \rightarrow C, C \rightarrow A, D \rightarrow A\}$. Then \mathcal{F}_4 is its own unique canonical cover, and so the construction of Proposition 3.3 delivers a unique cover precomplement for each subset of \mathbf{U} . However, it is not necessarily the case that such a precomplement has a unique extension to a meet complement. The optimal cover precomplement of $\Pi_{ABC}^{\mathbf{E}_4}$ is $\Pi_{AD}^{\mathbf{E}_4}$, but $\text{DPMut}_{\langle \mathbf{U}_4, \mathcal{F}_4 \rangle} \langle ABC \rangle = A$ is not a key for either ABC or BD . Here $\text{DPMutRHS}_{\langle \mathbf{U}_4, \mathcal{F}_4 \rangle} \langle ABC \rangle = A$ and $\text{DPMutRHS}_{\langle \mathbf{U}_4, \mathcal{F}_4 \rangle} \langle ABC \rangle = BC$, but only one of B or C need be added to $\text{DPMut}_{\langle \mathbf{U}_4, \mathcal{F}_4 \rangle} \langle ABC \rangle$ to make it a key for $\Pi_{ABC}^{\mathbf{E}_4}$. Thus each of $\Pi_{ABD}^{\mathbf{E}_4}$ and $\Pi_{ACD}^{\mathbf{E}_4}$ is a fully minimal Π -complement of $\Pi_{ABC}^{\mathbf{E}_4}$, so it has no optimal meet Π -complement. Hence, requiring unique canonical covers is not sufficient to guarantee universal (i.e., fully optimal) Π -complements. The simplest remedy is to require that the set of FDs admit a unique key, as developed below.

In [6, Cor. 3], it is shown that a general dependency-preserving decomposition of a universal schema \mathbf{E} (into possibly more than just two Π -views) is lossless if and only if the attributes of at least one of those views forms a superkey for \mathbf{E} . For a decomposition into two components, this result can be strengthened and related to the classical result [25, Thm. 2] as follows.

Lemma 3.13 (Meet complements and superkeys). *Let $\{\Pi_{\mathbf{W}_1}^{\mathbf{E}}, \Pi_{\mathbf{W}_2}^{\mathbf{E}}\}$ be a covered pair of views of \mathbf{E} . Then for $i \in \{1, 2\}$, $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_i \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$ iff \mathbf{W}_{3-i} is a superkey; i.e., $\mathbf{W}_{3-i} \rightarrow \mathbf{U} \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$.*

Proof. Let $i \in \{1, 2\}$. If $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_i \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$, then it is immediate that \mathbf{W}_{3-i} is a superkey. Conversely, suppose that \mathbf{W}_{3-i} is a superkey. Let $M_1, M_2 \in \text{LDB}(\mathbf{E})$ have the property that they agree on the attributes of $\mathbf{W}_1 \cap \mathbf{W}_2$; i.e., $\pi_{\mathbf{W}_1 \cap \mathbf{W}_2}(M_1) = \pi_{\mathbf{W}_1 \cap \mathbf{W}_2}(M_2)$. Let M be any relation on \mathbf{U} with $\pi_{\mathbf{W}_i}(M) = \pi_{\mathbf{W}_i}(M_i)$ and $\pi_{\mathbf{W}_{3-i}}(M) = \pi_{\mathbf{W}_{3-i}}(M_{3-i})$. Such an M always exists; indeed, the natural join $\pi_{\mathbf{W}_1}^{\mathbf{E}}(M_i) \bowtie \pi_{\mathbf{W}_2}^{\mathbf{E}}(M_{3-i})$ is one such relation. Since such a relation satisfies all FDs in \mathcal{F} by construction, $M \in \text{LDB}(\mathbf{E})$. Since \mathbf{W}_{3-i} is a superkey, $M = M_{3-i}$, whence $\pi_{\mathbf{W}_i}(M_{3-i}) = \pi_{\mathbf{W}_i}(M)$, and since $\pi_{\mathbf{W}_i}(M) = \pi_{\mathbf{W}_i}(M_i)$ by construction, it follows that $\pi_{\mathbf{W}_i}(M_{3-i}) = \pi_{\mathbf{W}_i}(M_i)$; i.e., $\pi_{\mathbf{W}_i}(M_1) = \pi_{\mathbf{W}_i}(M_2)$. Since the only initial assumption about M_1 and M_2 is that they agree on $\mathbf{W}_1 \cap \mathbf{W}_2$, it follows that $\mathbf{W}_1 \cap \mathbf{W}_2 \rightarrow \mathbf{W}_i \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$, as required. \square

Definition 3.14 (The strong cover property). Say that $\langle \mathbf{U}, \mathcal{F} \rangle$ has the *strong cover property* if \mathcal{F} admits a unique canonical cover and $\langle \mathbf{U}, \mathcal{F} \rangle$ has just one key. In view of the discussion in Summary 2.2, this latter condition may be stated equivalently as requiring that \mathcal{F} admit a unique key. In other words, $\langle \mathbf{U}, \mathcal{F} \rangle$ has the strong cover property if \mathcal{F} admits both a unique canonical cover and a unique key, so this property does not depend upon attributes in \mathbf{U} which do not occur in \mathcal{F} . However, the key itself for $\langle \text{Attr}_{\cup}(\mathcal{F}), \mathcal{F} \rangle$ will in fact differ from that of $\langle \mathbf{U}, \mathcal{F} \rangle$ if $\mathbf{U} \neq \text{Attr}_{\cup}(\mathcal{F})$. In the case that $\langle \mathbf{U}, \mathcal{F} \rangle$ admits a unique

key, $\text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle}$ will be used to denote that key. For further studies of schemata which admit unique keys, see [7] and [2, Thm. 3.4].

The main result on optimal meet Π -complements may now be presented. Define $\text{MeetCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ to be:

$$\begin{aligned} \text{(mc-i)} \quad & \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) && \text{if } \text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \rightarrow \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \\ & && \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+ \\ \text{(mc-ii)} \quad & \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \cup \text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle} && \text{otherwise} \end{aligned}$$

Theorem 3.15 (Universal Π -complements). *If $\langle \mathbf{U}, \mathcal{F} \rangle$ has the strong cover property, then for any $\mathbf{W} \subseteq \mathbf{U}$, $\Pi_{\mathbf{W}}^{\mathbf{E}}$ admits a universal Π -complement $\Pi_{\mathbf{X}}^{\mathbf{E}}$, with $\mathbf{X} = \text{MeetCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$. In words, the universal Π -complement of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ is just the optimal cover Π -precomplement if $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is a superkey for $\text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$, and it is the projection defined by augmenting the attributes of $\text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ with the unique key $\text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle}$ otherwise.*

Proof. If $\text{DPMut}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \rightarrow \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \in \text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle)^+$, the optimal meet Π -precomplement is already a meet complement (see Summary 2.4). Otherwise, for $\Pi_{\mathbf{X}}^{\mathbf{E}}$ to be a meet complement of $\Pi_{\mathbf{W}}^{\mathbf{E}}$, it is necessary that $\mathbf{W} \cap \mathbf{X}$ be a key for \mathbf{W} . In view of Lemma 3.13, this holds iff \mathbf{X} is a superkey for $\langle \mathbf{U}, \mathcal{F} \rangle$. To make \mathbf{X} such a superkey, the least set attributes which must be added to $\text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is precisely those elements of $\text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle}$ which are not already in $\text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$; exactly as (mc-ii) dictates.

Note that if $\text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle} \subseteq \text{PreCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ already, then adding $\text{Key}_{\langle \mathbf{U}, \mathcal{F} \rangle}$ has no effect, but specifying that it be added in all cases does no harm. \square

4 Quasi-Universal Π -Complements

The conditions imposed in Theorem 3.15 in order to obtain optimal meet complements for all Π -views of a given schema are quite strong. In particular, they rule out most situations with multiple (candidate) keys. Since such multiple keys are quite common in practice, it is important to seek extensions which relax, at least to some degree, the strict uniqueness properties imposed by the strong cover property. In this section, one such extension is pursued, in which individual attributes may be equivalent to one another. An example will help to identify the main ideas.

Example 4.1. Let $\mathbf{E}_5 = (R[\mathbf{U}_5], \mathcal{F}_5)$ with $\mathbf{U}_5 = AB_1B_2B_3C$ and $\mathcal{F}_5 = \{B_1 \leftrightarrow B_2, B_2 \leftrightarrow B_3, B_1 \rightarrow C, B_2 \rightarrow C, B_3 \rightarrow C\}$. Since \mathcal{F}_5 contains (simple) equivalences, it admits neither a unique canonical cover nor a unique key. Indeed, each of

$$\begin{aligned} \mathcal{F}'_{51} &= \{B_1 \rightarrow B_2, B_2 \rightarrow B_3, B_3 \rightarrow B_1, B_1 \rightarrow C\} \\ \mathcal{F}'_{52} &= \{B_1 \rightarrow B_2, B_2 \rightarrow B_3, B_3 \rightarrow B_1, B_2 \rightarrow C\} \\ \mathcal{F}'_{53} &= \{B_1 \rightarrow B_2, B_2 \rightarrow B_3, B_3 \rightarrow B_1, B_3 \rightarrow C\} \end{aligned}$$

is a minimum and canonical cover, and the view $\Pi_{AB_1B_2B_3}^{\mathbf{E}_5}$ has three distinct minimal meet Π -complements, $\Pi_{B_1C}^{\mathbf{E}_5}$, $\Pi_{B_2C}^{\mathbf{E}_5}$, and $\Pi_{B_3C}^{\mathbf{E}_5}$. Furthermore, each of AB_1 , AB_2 , and AB_3 is a key for $\langle \mathbf{U}_5, \mathcal{F}_5 \rangle$. Nevertheless, these three are closely related in the sense that each one is obtained from the other by replacing each attribute by an equivalent one. This is also the case for the three distinct canonical covers above, as well as the three distinct keys. Indeed, if the three attributes in $B_1B_2B_3$ are collapsed into one, say B_1 , with the resulting schema $\tilde{\mathbf{E}}_5 = (R[\tilde{\mathbf{U}}_5], \tilde{\mathcal{F}}_5)$ with $\tilde{\mathcal{F}}_5 = \{B_1 \rightarrow C\}$, then $\langle \tilde{\mathbf{U}}_5, \tilde{\mathcal{F}}_5 \rangle$ has the strong cover property and the universal Π -complement of $\Pi_{AB_1}^{\tilde{\mathbf{E}}_5}$ is $\Pi_{B_1C}^{\tilde{\mathbf{E}}_5}$.

The fully minimal Π -complements may be recovered easily from this solution. For example, if the view to be complemented is $\Pi_{AB_1B_2B_3}^{\mathbf{E}_5}$, then first compute the complement of $\Pi_{AB_1}^{\tilde{\mathbf{E}}_5}$, which is $\Pi_{B_1C}^{\tilde{\mathbf{E}}_5}$. Then, restore the individual values associated with B_1 , placing as few as possible in the complement. In general, there will be several ways to do this. In the example, one of $B_1B_2B_3$ is placed in the complement, with the other two in the main view. This leads exactly to the three complements identified above.

A formalization of these ideas makes up the rest of this section.

Definition 4.2 (Simple equivalence of attributes and quasi-universal Π -complements). Recall that two attributes A and B are *simply equivalent* in \mathcal{F} if $\mathcal{F} \models A \leftrightarrow B$. For convenience, let $\text{SimpEquiv}(\mathcal{F})$ denote the set of all simple equivalences implied by \mathcal{F} ; thus $\text{SimpEquiv}(\mathcal{F}) = \{A \leftrightarrow B \mid \mathcal{F} \models A \leftrightarrow B\}$. Define $\equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1$ to be the equivalence relation on \mathbf{U} given by $A \equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1 B$ iff $\text{FDAug}(\langle \mathbf{U}, \mathcal{F} \rangle) \models A \leftrightarrow B$.

Given $\mathbf{W}, \mathbf{W}' \subseteq \mathbf{U}$, $\Pi_{\mathbf{W}'}^{\mathbf{E}}$ is called a *quasi-universal Π -complement* of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ with respect to $\equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1$ if for any other minimal meet Π -complement $\Pi_{\mathbf{W}''}^{\mathbf{E}}$, there is a bijection $\beta : \mathbf{W}' \rightarrow \mathbf{W}''$ with the property that for each $A \in \mathbf{W}'$, $A \equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1 \beta(A)$. In other words, a minimal meet complement is quasi-universal with respect to $\equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1$ if all other meet complements differ only by simple equivalences on the attributes. A universal complement is a special case in which the only such bijection β is the identity on \mathbf{W}' . Since all quasi-universal complements in this paper will be with respect to the simple equivalence of the underlying schema, that qualification will be dropped, and it will simply be said that one view is a *quasi-universal complement* of another.

Definition 4.3 (Aliases and representation functions). Any $B \in \mathbf{U}$ with $A \equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1 B$ is called an *alias* for A in $\langle \mathbf{U}, \mathcal{F} \rangle$. Call A *aliased* in $\langle \mathbf{U}, \mathcal{F} \rangle$ if it has an alias other than itself. Otherwise, call A *unaliased* in $\langle \mathbf{U}, \mathcal{F} \rangle$. The set of all aliases of A in $\langle \mathbf{U}, \mathcal{F} \rangle$ is denoted $\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$. $\tilde{\mathbf{U}} = \{\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A) \mid A \in \mathbf{U}\}$; thus, $\tilde{\mathbf{U}}$ is just the set of all equivalence classes of $\equiv_{\langle \mathbf{U}, \mathcal{F} \rangle}^1$. For $\mathbf{W} \subseteq \mathbf{U}$, define $\tilde{\mathbf{W}} = \{\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A) \mid A \in \mathbf{W}\}$. It is useful to have a more compact representation for elements of the form $\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$. To this end, define a *representation* for $\tilde{\mathbf{U}}$ to be a function $\tau : \mathbf{U} \rightarrow \mathbf{U}$ with the properties that for each $A \in \mathbf{U}$, $\tau(A) \in \text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$ and $\tau(A) = \tau(B)$ iff $B \in \text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$. Thus, a representation

maps all elements of a given class $\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$ to the same element, which must itself be a member of $\text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A)$. In the above example, the representation used is τ_5 defined by $\tau_5(A) = A$, $\tau_5(C) = C$, and $\tau_5(B_1) = \tau_5(B_2) = \tau_5(B_3) = B_1$.

Fix a representation function τ for \mathcal{F} , and define $\tilde{\mathcal{F}}$ to be the set of FDs obtained from \mathcal{F} by replacing each $A \in \text{Attr}(\mathcal{F})$ with $\tau(A)$, and then removing all trivial FDs. In Example 4.1, $\tilde{\mathcal{F}}_5$ is obtained in this way from \mathcal{F}_5 . Choosing a different representation function will change the names of the symbols and nothing else, so it is justified to write \mathcal{F} without any reference to τ .

Finally, define $\tilde{\mathbf{E}} = (R[\tilde{\mathbf{U}}], \tilde{\mathcal{F}})$, the *simple-alias reduction* of \mathbf{E} .

It is perhaps worth noting that a representation as defined above is a special case of a *system of distinct representatives (SDR)*, also called a *transversal* [22], [8, p. 192], [23, Def. 1]. In a general SDR, the sets to be represented may overlap, and so the question of existence is nontrivial; the solution being given by *Hall's Marriage Theorem* (see the references identified above). In the special case considered here, the sets do not overlap and so the existence of an SDR is trivial.

Example 4.4 (An issue surrounding simple-alias reduction). The idea is to represent the original $\langle \mathbf{U}, \mathcal{F} \rangle$ using $\tilde{\mathcal{F}}$ together with a set \mathcal{E} of simple FDs which defines all of the simple equivalences of \mathcal{F} . Then, the different minimal complements may be obtained by using the “unique” canonical cover $\tilde{\mathcal{F}}$ as the representation function τ varies. This almost works; however, there is one minor complication, which is best illustrated by example. Let $\mathbf{E}_6 = (R[A_1A_2BC], \mathcal{F}_6)$ be the relational schema with $\mathcal{F}_6 = \{A_1 \leftrightarrow A_2, A_2 \leftrightarrow BC\}$. Then, with $\tau_6(A_1) = \tau_6(A_2) = A_2$, $\tilde{\mathcal{F}}_6 = \{A_2 \leftrightarrow BC\}$ which is its own unique canonical cover. The view $\Pi_{A_2BC}^{\mathbf{E}_6}$ has two minimal Π -complements, $\Pi_{A_1A_2}^{\mathbf{E}_6}$ and $\Pi_{A_1BC}^{\mathbf{E}_6}$ because, in addition to the alias A_2 , BC is also formally equivalent to A_1 . This property depends upon the isolation of BC from other FDs. For example, let $\mathbf{E}'_6 = (R[A_1A_2BCD], \mathcal{F}'_6)$ be the relational schema with $\mathcal{F}'_6 = \mathcal{F}_6 \cup \{BC \rightarrow D\}$. With this simple addition, $\tilde{\mathcal{F}}'_6 = \{A_2 \leftrightarrow BC, BC \rightarrow D\}$ no longer has the strong cover property. Nevertheless, the case illustrated by \mathbf{E}_6 must be ruled out.

Definition 4.5 (The quasi-strong cover property). For $\mathbf{X} \subseteq \mathbf{U}$ with $\text{Card}(\mathbf{X}) > 1$ (i.e., \mathbf{X} contains at least two distinct elements) and $A, B \in \mathbf{U}$, call $\{A, B, \mathbf{X}\}$ a *complex triple* for \mathcal{F} if both $A \leftrightarrow B, B \leftrightarrow \mathbf{X} \in \mathcal{F}^+$, and for no proper subset $\mathbf{X}' \subsetneq \mathbf{X}$ is it the case that $\mathbf{X}' \leftrightarrow B \in \mathcal{F}^+$. In other words, $A \leftrightarrow B \leftrightarrow \mathbf{X}$ holds with \mathbf{X} consisting of at least two attributes, and no element of \mathbf{X} may be removed while retaining that property. For both \mathcal{F}_6 and \mathcal{F}'_6 , $\{A_1, A_2, BC\}$ is a complex triple. Say that \mathcal{F} is *free of complex triples* if \mathcal{F}^+ contains no such triples. Define \mathcal{F} to have the *quasi-strong cover property* if it is free of complex triples and has the strong equivalence-cover property. It is worth emphasizing that requiring that \mathcal{F} be free of complex triples does not impose substantial further limits beyond those already implied by the strong cover property on $\tilde{\mathcal{F}}$, since as illustrated by the schema \mathbf{E}'_6 , if the attributes of a complex triple are connected to other attributes via additional FDs, then the strong cover property is lost anyway.

Lemma 4.6 (Derivation of simple equivalences). *If \mathcal{F} is free of complex triples and \mathcal{C} is a canonical cover for \mathcal{F} , then for every simple equivalence $A \leftrightarrow B \in \mathcal{F}^+$ with $A \neq B$, there are $A_1, \dots, A_k \in \mathbf{U}$ with $A_1 = A$, $A_k = B$, and $A_i \rightarrow A_{i+1} \in \mathcal{C}$ for $i \in 1, \dots, k-1$.*

Proof outline. Without loss of generality, assume that \mathcal{F} is canonical. The proof is based upon the claim that there must be a sequence $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_k$ of subsets of \mathbf{U} with $\mathbf{X}_1 = A$, $\mathbf{X}_k = B$, and for each $i \in \{1, \dots, k-1\}$, $\mathbf{X}_i \rightarrow C \in \mathcal{F}$ for each $C \in \mathbf{X}_{i+1}$. Then, using the property that \mathcal{F} be free of complex triples, it follows that each \mathbf{X}_i must consist of single attribute, whence the result.

To place this argument on rigorous grounds, a deductive formalism for FDs is required. The most appropriate one in this case is the *derivation directed acyclic graphs* or *DDAGs* [20], [21, Sec. 4.5]. In this case, there must be a DDAG with single initial vertex labelled with A and a final vertex labelled B . Furthermore, there must be such a graph with exactly one outgoing edge from each vertex other than that labelled with B ; otherwise, it is possible to deduce a complex cycle. The full elaboration is straightforward but tedious; space limitation preclude expanding it further. \square

Proposition 4.7 (Representation of quasi-strong covers). *If \mathcal{F} has the quasi-strong cover property, then every canonical cover of \mathcal{F} is of the form $\mathcal{C}_{\text{Str}} \cup \mathcal{C}_{\text{Eq}}$, with the following properties.*

- (a) \mathcal{C}_{Str} is the unique canonical cover for $\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle$ for some representation function.
- (b) \mathcal{C}_{Eq} is a canonical cover for $\text{SimpEquiv}(\mathcal{F})$.

Proof. In view of Lemma 4.6, every canonical cover of \mathcal{F} must contain a subset consisting of simple equivalences which form a cover for $\text{SimpEquiv}(\mathcal{F})$. Let \mathcal{C}_{Eq} be any such set. Since \mathcal{F} is assumed to have the quasi-strong cover property, the remainder of the canonical cover must be \mathcal{F}_{Str} . \square

Definition 4.8 (Constructions for quasi-universal Π -complements).

The precise construction of the quasi-universal Π -complements of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ may now be specified. Although the definitions are straightforward, it may prove useful to follow them in parallel with Example 4.10, which illustrates their application to a fixed schema.

Let $\mathbf{W} \subseteq \mathbf{U}$. The process of working with $\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle$ rather than $\langle \mathbf{U}, \mathcal{F} \rangle$ collapses all simply equivalent elements into one. In order to recover the views corresponding to $\langle \mathbf{U}, \mathcal{F} \rangle$, it is necessary to “uncollapse” these equivalent elements in the appropriate fashion. To begin, the attributes in \mathbf{W} which have an alias in the base complement of \mathbf{W} are identified.

- (a) $\text{DPBoth}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle = \{A \in \mathbf{W} \mid \text{BaseComp}_{\langle \mathbf{U}, \mathcal{F} \rangle} \langle \mathbf{W} \rangle \cap \text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A) \neq \emptyset\}$

In order to maintain an embedded cover of \mathcal{C}_{Eq} as identified in Proposition 4.7 above, if an attribute A_1 occurs in the view $\Pi_{\mathbf{W}}^{\mathbf{E}}$ to be complemented while an equivalent attribute A_2 occurs in a potential complement $\Pi_{\mathbf{X}}^{\mathbf{E}}$ but not in $\Pi_{\mathbf{W}}^{\mathbf{E}}$, then either one of these, or else a third attribute equivalent to both, must

appear in $\mathbf{W} \cap \mathbf{X}$. This alias must be an element of \mathbf{W} , since the meet is a subset of \mathbf{W} . Additionally, any equivalence class which has a representative in $\text{DPMut}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}})$ must also have at least one representative in the meet of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ and any of its minimal complements. The possible elements for the meet are defined as follows.

- (b) $\text{DPMeetSets}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) = \{ \mathbf{W} \cap \text{Alias}_{\langle \mathbf{U}, \mathcal{F} \rangle}(A) \mid (A \in \text{DPBoth}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})) \text{ or } (\mathfrak{r}(A) \in \text{DPMut}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}})) \}$
- (c) $\text{DPMeetAlts}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ is all subsets of \mathbf{W} formed by selecting exactly one attribute from each member of $\text{DPMeetSets}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$.

Some set of attributes in $\text{DPMeetAlts}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$ must be contained in the meet of any quasi-universal Π -complement. Observe in particular that if an attribute has no alias other than itself, then it will be in the meet of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ and its minimal complement iff it occurs in the meet of $\tilde{\Pi}_{\tilde{\mathbf{W}}}^{\tilde{\mathbf{E}}}$ and its optimal complement.

Finally, the alternatives for the quasi-universal Π -complements may be specified by adding to the base complement a set of suitable elements for the meet.

- (d) $\text{MeetComplAlts}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) = \{ \mathbf{X} \cup \text{BaseCompl}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \mid \mathbf{X} \in \text{DPMeetAlts}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W}) \}$

This is summarized in the following theorem. The proof is a straightforward argument based upon the above constructions, and is omitted due to space considerations.

Theorem 4.9. *Let $\mathbf{W} \subseteq \mathbf{U}$, and assume that $\langle \mathbf{U}, \mathcal{F} \rangle$ has the quasi-strong cover property. Then $\Pi_{\mathbf{X}}^{\mathbf{E}}$ is a quasi-universal Π -complement of $\Pi_{\mathbf{W}}^{\mathbf{E}}$ iff $\mathbf{X} \in \text{MeetComplAlts}_{\langle \mathbf{U}, \mathcal{F} \rangle}(\mathbf{W})$. \square*

Example 4.10. Define the schema

$$\mathbf{E}_7 = (R[A_1 A_2 A_3 B C_1 C_2 C_3 D E_1 E_2 E_3 F_1 F_2 F_3 G H I J_1 J_2 J_3 K_1 K_2 K_3], \mathcal{F}_7)$$

with $\mathcal{F}_7 = \mathcal{C}_{\text{Str}7} \cup \mathcal{C}_{\text{Eq}7}$ and these latter two sets given by

$$\begin{aligned} \mathcal{C}_{\text{Str}7} &= \{ A_1 B \rightarrow C_1 D, C_1 D \rightarrow E_1 F_1 G, E_1 F_1 \rightarrow H, H \rightarrow I J_1 K_1 \} \\ \mathcal{C}_{\text{Eq}7} &= \{ A_1 \rightarrow A_2, A_2 \rightarrow A_3, A_3 \rightarrow A_1, C_1 \rightarrow C_2, C_2 \rightarrow C_3, C_3 \rightarrow C_1, \\ &E_1 \rightarrow E_2, E_2 \rightarrow E_3, E_3 \rightarrow E_1, F_1 \rightarrow F_2, F_2 \rightarrow F_3, F_3 \rightarrow F_1, \\ &J_1 \rightarrow J_2, J_2 \rightarrow J_3, J_3 \rightarrow J_1, K_1 \rightarrow K_2, K_2 \rightarrow K_3, K_3 \rightarrow K_1 \} \end{aligned}$$

Upon RHS-simplification, $\mathcal{C}_{\text{Str}7}$ becomes a unique canonical cover for $\tilde{\mathcal{F}}_7$, while $\mathcal{C}_{\text{Eq}7}$ is a canonical cover for $\text{SimpEquiv}(\mathcal{F})$, with $\alpha_i \equiv_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}^1 \alpha_j$ for $\alpha \in \{A, C, E, F, J, K\}$ and $i \in \{1, 2, 3\}$. The representation used is $\mathfrak{r}_7 : \alpha_i \mapsto \alpha_1$ for $\alpha \in \{A, C, E, F, J, K\}$ and $\mathfrak{r}_7 : \alpha \mapsto \alpha$ for $\alpha \in \{B, D, G, H, I\}$.

Let $\mathbf{W} = E_1 E_2 E_3 F_1 F_2 G H I J_1 K_1 K_2 K_3$ with $\Pi_{\mathbf{W}}^{\mathbf{E}_7}$ the view to be complemented. First of all, for $\tilde{\mathbf{E}}_7 = (R[A_1 B C_1 D E_1 F_1 G H I J_1 K_1], \mathcal{C}_{\text{Str}7})$, using the constructions for the strong cover property, $\tilde{\mathbf{W}} = E_1 F_1 G H I J_1 K_1$,

$\text{BaseCompl}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) = A_1BC_1, D$, $\text{DPExt}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) = A_1BC_1DE_1F_1G$, and $\text{DPMut}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) = E_1F_1G$. It is easy to see that A_1B is a key for $\tilde{\mathbf{E}}_7$, so by Theorem 3.15 the optimal complement of $\Pi_{\tilde{\mathbf{W}}}^{\tilde{\mathbf{E}}_7}$ is $\Pi_{\mathbf{X}}^{\tilde{\mathbf{E}}_7}$ with $\mathbf{X} = \text{PreCompl}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) = \text{BaseCompl}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) \cup \text{DPExt}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}}) = A_1BC_1DE_1F_1G$. The task is to uncollapse this solution into one for \mathbf{E}_7 and the view $\Pi_{\mathbf{W}}^{\mathbf{E}_7}$.

First of all, $\text{DPBoth}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W}) = F_1F_2J_1$, the elements of \mathbf{W} which have an alias in $\text{BaseCompl}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W}) = A_1A_2A_3BC_1C_2C_3DF_3J_2J_3$. Next, $\text{DPMeetSets}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W}) = \mathbf{W} \cap (F_1F_2F_3J_1J_2J_3 \cup E_1E_2E_3F_1F_2F_3G) = E_1E_2E_3F_1F_2GJ_1$ consists of all attributes in \mathbf{W} with an alias in either $\text{DPBoth}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W})$ or $\text{DPMut}_{\langle \tilde{\mathbf{U}}, \tilde{\mathcal{F}} \rangle}(\tilde{\mathbf{W}})$. These are the possible elements for inclusion in the meet. The actual meet selects one element from $\text{DPMeetSets}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W})$ for each equivalence class of $\equiv_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}^1$ which is represented in that set. Thus the set of meet alternatives is given by $\text{DPMeetAlts}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W}) = \{E_iF_jGJ_1 \mid i \in \{1, 2, 3\} \text{ and } j \in \{1, 2\}\}$, and so $\text{MeetComplAlts}_{\langle \mathbf{U}_7, \mathcal{F}_7 \rangle}(\mathbf{W}) = A_1A_2A_3BC_1C_2C_3DF_3GJ_1J_2J_3 \cup \{E_iF_j \mid i \in \{1, 2, 3\} \text{ and } j \in \{1, 2\}\}$. Hence, there are six alternatives for a quasi-universal Π -complement, all related by the replacement of simple attributes by equivalent ones. Note, however, that not any equivalent attribute will do. F_1 may only be replaced with F_2 , not F_3 , and J_1 may not be replaced by either J_2 or J_3 .

5 Conclusions and Further Directions

A characterization of universal complements in the setting of projections of a universal relational schema which is constrained by functional dependencies has been presented. Although the conditions are rather strict — existence of a unique canonical cover and a unique key — they reveal the nature of requirements which must be imposed in order to realize such complements. Conditions which allow a modest amount of flexibility — by replacing single attributes by equivalent ones to yield quasi-universal complements, have also been developed. They give useful solutions in certain important settings, such as when there are multiple candidate keys.

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

EXTENSION OF THE CONDITIONS: The presentations of Sections 3 and 4 were designed to illustrate the fundamental ideas without becoming bogged down in small details. As such, they can clearly be strengthened. The strong cover property identified in Sec. 3 guarantees fully optimal Π -complements for all possible views of the main schema \mathbf{E} . It is possible to provide a localized strong cover property which only guarantees fully optimal complements for a fixed view $\Pi_{\mathbf{W}}^{\mathbf{E}}$, or for some set of views. The details are a straightforward extension to the results presented, but nevertheless should be elaborated in a future paper. Similarly, an extension of the results of Sec. 4 on quasi-universal Π -complements to equivalences amongst sets of attributes is also a worthwhile topic, although

the details of exactly what such an extension would look like remain to be developed.

EXTENSION OF THE SCOPE TO \sqrt{II} -VIEWS: In [15], the \sqrt{II} -views of a relational schema, which are defined by sets of projections rather than single projections, are considered as the fundamental sets of views. Since the individual projections which comprise such a view may not join losslessly, this is a genuine extension which often yields smaller complements than the II -framework. An extension of the results of this paper to the \sqrt{II} -setting would provide useful new results.

GENERAL CHARACTERIZATION OF FULLY OPTIMAL COMPLEMENTS: Although this work is an important first step towards the understanding of universality and fully optimal complements, it is clearly limited in comparison to the results on characterization of state invariance and reflection invariance, which hold in very general settings not limited to any particular data model. The next major step for this research program is to seek such a generalized extension, which would proceed in two steps. In the first, the results of this paper would be integrated with the *governing join dependencies* of [15], thus lifting the results from schemata constrained by just FDs to those with more general constraints. In the second step, a more general characterization, independent of any particular data model, would be developed. The exact nature of this characterization remains to be identified.

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley (1995)
2. Atzeni, P., Antonellis, V.D.: Relational Database Theory. Benjamin/Cummings (1993)
3. Bancilhon, F., Spyratos, N.: Independent components of databases. In: Proceedings of the Seventh International Conference on Very Large Data Bases. pp. 398–408 (1981)
4. Bancilhon, F., Spyratos, N.: Update semantics of relational views. ACM Trans. Database Systems 6, 557–575 (1981)
5. Bernstein, P.A.: Synthesizing third normal form relations from functional dependencies. ACM Trans. Database Syst. 1(4), 277–298 (1976)
6. Biskup, J., Dayal, U., Bernstein, P.A.: Synthesizing independent database schemas. In: Proceedings of the ACM-SIGMOD 1979 Conference on Management of Data. pp. 143–151 (1979)
7. Biskup, J., Demetrovics, J., Libkin, L., Muchnik, I.B.: On relational database schemes having unique minimal key. Elektronische Informationsverarbeitung und Kybernetik 27(4), 217–225 (1991)
8. Bogart, K.P.: Introductory Combinatorics. Pitman (1983)
9. Bohannon, A., Pierce, B.C., Vaughan, J.A.: 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. pp. 338–347 (2006)

10. Davey, B.A., Priestly, H.A.: Introduction to Lattices and Order. Cambridge University Press, second edn. (2002)
11. Foster, J.N., Greenwald, M.B., Moore, J.T., Pierce, B.C., Schmitt, A.: Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM Trans. Programming Languages and Systems* 29(3) (2007)
12. Hegner, S.J.: Foundations of canonical update support for closed database views. In: Abiteboul, S., Kanellakis, P.C. (eds.) *ICDT'90, Third International Conference on Database Theory, Paris, France, December 1990*. pp. 422–436. Springer-Verlag (1990)
13. Hegner, S.J.: An order-based theory of updates for closed database views. *Ann. Math. Art. Intell.* 40, 63–125 (2004)
14. Hegner, S.J.: The complexity of embedded axiomatization for a class of closed database views. *Ann. Math. Art. Intell.* 46, 38–97 (2006)
15. Hegner, S.J.: Characterization of optimal complements of database views defined by projection. In: Schewe, K.D., Thalheim, B. (eds.) *4th International Workshop, SDKB 2010, Bordeaux, France, July 5, 2010, Revised Selected Papers*. Lecture Notes in Computer Science, vol. 6384, pp. 73–95. Springer-Verlag (2011)
16. Hull, R.: Finitely specifiable implicational dependency families. *J. Assoc. Comp. Mach.* 31(2), 210–226 (1984)
17. Jacobs, B.E., Aronson, A.R., Klug, A.C.: On interpretations of relational languages and solutions to the implied constraint problem. *ACM Trans. Database Systems* 7(2), 291–315 (1982)
18. Lechtenbörger, J.: 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*. pp. 49–55 (2003)
19. Lechtenbörger, J., Vossen, G.: On the computation of relational view complements. *ACM Trans. Database Systems* 28, 175–208 (2003)
20. Maier, D.: Minimum covers in the relational database model. *J. ACM* 27(4), 664–674 (1980)
21. Maier, D.: *The Theory of Relational Databases*. Computer Science Press (1983)
22. Mann, H.B., Ryser, H.J.: Systems of distinct representatives. *American Math. Monthly* 60(6), 297–401 (1953)
23. Marcus, M.: *A Survey of Finite Mathematics*. Dover (1993)
24. Paredaens, J., De Bra, P., Gyssens, M., Van Gucht, D.: *The Structure of the Relational Database Model*. Springer-Verlag (1989)
25. Rissanen, J.: Independent components of relations. *ACM Trans. Database Systems* 2(4), 317–325 (1977)