Finite Open-World Query Answering with Number Restrictions

ANTOINE AMARILLI, Télécom ParisTech, Institut Mines–Télécom, CNRS LTCI
MICHAEL BENEDIKT, University of Oxford

Open-world query answering is the problem of deciding, given a set of facts, conjunction of constraints, and query, whether the facts and constraints imply the query. This amounts to reasoning over all instances that include the facts and satisfy the constraints. We study *finite open-world query answering* (FQA), which assumes that the underlying world is finite and thus only considers the *finite* completions of the instance. The major known decidable cases of FQA derive from the following: the guarded fragment of first-order logic, which can express referential constraints (data in one place points to data in another) but cannot express number restrictions such as functional dependencies; and the guarded fragment with number restrictions but on a signature of arity only two. In this paper, we give the first decidability results for FQA that combine both referential constraints and number restrictions for arbitrary signatures: we show that, for unary inclusion dependencies and functional dependencies, the finiteness assumption of FQA can be lifted up to taking the finite implication closure of the dependencies. Our result relies on new techniques to construct finite universal models of such constraints, for any bound on the maximal query size.

1. INTRODUCTION

A longstanding goal in computational logic is to design logical languages that are both decidable and expressive. One approach is to distinguish integrity constraints and queries, and have separate languages for them. We would then seek decidability of the *query answering with constraints* problem: given a query q, a conjunction of constraints Σ, and a finite instance I_0, determine which answers to q are certain to hold over any instance I that extends I_0 and satisfies Σ. This problem is often called *open-world query answering*. It is fundamental for deciding query containment under constraints, querying in the presence of ontologies, or reformulating queries with constraints. Thus it has been the subject of intense study within several communities for decades (e.g. [Johnson and Klug 1984; Calì et al. 2003a; Bárány et al. 2010; Pratt-Hartmann 2009; Ibáñez-García et al. 2014]).

In many cases (e.g., in databases) the instances I of interest are the finite ones, and hence we can define *finite open-world query answering* (denoted here as FQA), which restricts the quantification to *finite* extensions I of I_0. In contrast, by *unrestricted open-world query answering* (UQA) we refer to the problem where I can be either finite or infinite. Generally the class of queries is taken to be the conjunctive queries (CQs) — queries built up from relational atoms via existential quantification and conjunction. We will restrict to CQs here, and thus omit explicit mention of the query language, focusing on the constraint language.

A first constraint class known to have tractable open-world query answering problems are *inclusion dependencies* (IDs) — constraints of the form, e.g., $\forall xyz\, R(x,y,z) \rightarrow \exists vw\, S(z,v,w,y)$. The fundamental results of Johnson and Klug [Johnson and Klug 1984] and Rosati [Rosati 2011] show that both FQA and UQA are decidable for ID and that, in fact, they coincide. When this happens, the constraints are said to be *finitely controllable*. These results have been generalized by Bárány et al. [Bárány et al. 2010] to a much richer class of constraints, the guarded fragment of first-order logic.

However, those results do not cover a second important kind of constraints, namely *number restrictions*, which express, e.g., uniqueness. We represent them by the class of *functional dependencies* (FDs) — of the form $\forall \boldsymbol{xy}\ (R(x_1,\ldots,x_n) \wedge R(y_1,\ldots,y_n) \wedge \bigwedge_{i \in L} x_i = y_i) \rightarrow x_r = y_r$. The implication problem (does one FD follow from a set of others) is decidable for FDs, and coincides with implication restricted to finite instances [Abiteboul et al. 1995]. Trivially, the FQA and UQA problems are also decidable for FDs alone, and coincide.

IDs require the model to contain some elements, while FDs restrict the ability to add elements. The interaction is severe enough that trying to combine IDs and FDs makes both UQA and FQA undecidable in general [Calì et al. 2003a]. Some progress has been made on obtaining decidable cases for UQA. UQA is known to be decidable when the FDs and the IDs are *non-conflicting* [Johnson and Klug 1984; Calì et al. 2003a]. Intuitively, this condition guarantees that the FDs can be ignored, as long as they hold on the initial instance I_0, and one can then solve the query answering problem by considering the IDs alone. But the non-conflicting condition only applies to UQA and not to FQA. In fact it is known that even for very simple classes of IDs and FDs, including non-conflicting classes, FQA and UQA do not coincide. Rosati [Rosati 2011] showed that FQA is undecidable for non-conflicting IDs and FDs (indeed, for IDs and keys, which are less rich than FDs).

Thus a broad question is to what extent these classes, FDs and IDs, can be combined while retaining decidable FQA. The only decidable cases impose very severe requirements. For example, the constraint class of "single KDs and FKs" introduced in [Rosati 2011] has decidable FQA, but such constraints cannot model, e.g., FDs which are not keys. Further, in contrast with the general case of FDs and IDs, single KDs and FKs are always finitely controllable, which limits their expressiveness. Indeed, *we know of no tools to deal with FQA for non-finitely-controllable constraints on relations of arbitrary arity.*

A second decidable case is where all relation symbols and all subformulae of the constraints have arity at most two. In this context, results of Pratt-Hartmann [Pratt-Hartmann 2009] imply the decidability of both FQA and UQA for a very rich non-finitely-controllable sublogic of first-order logic. For some fragments of this arity-two logic, the complexity of FQA has recently been isolated by Ibáñez-García et al. [Ibáñez-García et al. 2014]. Yet these results do not apply to arbitrary arity signatures.

The contribution of this paper is to provide the first result about finite query answering for non-finitely-controllable IDs and FDs over relations of arbitrary arity. As the problem is undecidable in general, we must naturally make some restriction. Our choice is to limit to *Unary IDs* (UIDs), which export only one variable: for instance, $\forall xyz\ R(x,y,z) \rightarrow \exists w\ S(w,x)$. UIDs and FDs are an interesting class to study because they are not finitely controllable, and allow the modeling, e.g., of single-attribute foreign keys, a common use case in database systems. In contrast, Johnson and Klug [Johnson and Klug 1984] showed that UIDs in isolation are finitely controllable. The decidability of UQA for UIDs and FDs is known because they are always non-conflicting. In this paper, we show that finite query answering is decidable for UIDs and FDs, and obtain tight bounds on its complexity.

The idea is to *reduce the finite case to the unrestricted case*, but in a more complex way than by finite controllability. We make use of a technique originating in Cosmadakis et al. [Cosmadakis et al. 1990] to study finite implication on UIDs and FDs: the *finite closure* operation which takes a conjunction of UIDs and FDs and determines exactly which additional UIDs and FDs are implied over finite instances. Rosati [Rosati 2008] and Ibáñez-García [Ibáñez-García et al. 2014] make use of the closure operation in their study of constraint classes over schemas of arity two. They show that finite query answering for a query q, instance I_0, and constraints Σ reduces to unrestricted query answering for I_0, q, and the finite closure $\Sigma^{\mathrm{f}*}$ of Σ. In other words, the closure construction which is sound for implication is also sound for query answering.

We show that the same approach applies to arbitrary arity signatures, with constraints being UIDs and FDs. Our main result thus reduces finite query answering to unrestricted query answering, for UIDs and FDs in arbitrary arity:

THEOREM I.1. *For any finite instance I_0, conjunctive query q, and constraints Σ consisting of UIDs and FDs, the finite open-world query answering problem for I_0, q under Σ has the same answer as the unrestricted open-world query answering problem for I_0, q under the finite closure of Σ.*

Using the known results about the complexity of UQA for UIDs, we isolate the precise complexity of finite query answering with respect to UIDs and FDs, showing that it matches that of UQA:

COROLLARY I.2. *The combined complexity of the finite open-world query answering problem for UIDs and FDs is* NP*-complete, and it is PTIME in data complexity (that is, when the constraints and query are fixed).*

Our proof of Theorem I.1 is quite involved, since building finite models that satisfy number restrictions and inclusion dependencies in a signature with arbitrary arity introduces a multitude of new difficulties that do not arise in the arity-two case or in the case of IDs in isolation.

We borrow and adapt a variety of techniques from prior work:

— using k-bounded simulations to preserve small acyclic CQs [Ibáñez-García et al. 2014],
— partitioning UIDs into components that have limited interaction, and satisfying the UIDs component-by-component [Cosmadakis et al. 1990; Ibáñez-García et al. 2014],
— performing a chase that reuses sufficiently similar elements [Rosati 2011],
— taking the product with groups of large girth to blow up cycles [Otto 2002].

However, we must also develop some new infrastructure to deal with number restrictions in an arbitrary arity setting: distinguishing between so-called *dangerous* and *non-dangerous* positions when creating a new element to satisfy some ids, constructing realizations for relations in a *piecewise* manner following the FDs, reusing elements in a *combinatorial* way that shuffles them to avoid violating the higher-arity FDs, and a new notion of *mixed product* to blow cycles up while preserving fact overlaps to avoid violating the higher-arity FDs.

Paper structure. The overall structure of the proof, presented in Section III, is to extend a given instance I_0 to a finite model of UIDs and FDs such that for every conjunctive query of size at most k, the model satisfies it only when it is implied. We call these *k-universal models*. It is easy to show that if a k-universal model exists for an instance and set of constraints, then finite implication and implication of CQs coincide.

We start with only *unary* FDs (UFDs) and *acyclic* CQs (ACQs), and by assuming that the UIDs and UFDs are *reversible*, a condition inspired by the finite closure construction.

As a warm-up, Sections IV and V approximate even further by replacing k-universality by a weaker notion, proving the corresponding result starting with binary signatures and generalizing to arbitrary arity. We extend the result to k-universality in Section VI, maintaining a k-bounded simulation to the chase, and performing *thrifty* chase steps that reuse sufficiently similar elements without violating UFDs. We also rely on a structural observation about the chase under UIDs (Theorem VI.20). Section VII eliminates the assumption that dependencies are reversible, by partitioning the UIDs into classes that are either reversible or trivial, and satisfying successively each class following a certain ordering.

We then generalize our result to higher-arity (non-unary) FDs in Section VIII. This requires us to define a new notion of thrifty chase steps that apply to instances with many ways to reuse elements; the existence of these instances relies on a combinatorial construction of models of FDs with a high number of facts but a small domain (Theorem VIII.11). Last, in Section IX, we apply a cycle blowup process to the result of the previous constructions, to go from acyclic to arbitrary CQs through a product with acyclic groups. The technique is inspired by Otto [Otto 2002] but must be adapted to respect FDs.

II. BACKGROUND

II.1. Instances and Constraints

Instances. We assume an infinite countable set of *elements* (or *values*) $a, b, c, \ldots$ and *variable names* $x, y, z, \ldots$. A *schema* σ consists of *relation names* (e.g., R) with an *arity* (e.g., $|R|$) which we assume is ≥ 1: we write $|\sigma| := \max_{R \in \sigma} |R|$. Following the *unnamed perspective*, the set of *positions* of R is $\mathsf{Pos}(R) := \{R^i \mid 1 \leq i \leq |R|\}$, and we define $\mathsf{Pos}(\sigma) := \bigsqcup_{R \in \sigma} \mathsf{Pos}(R)$. We identify R^i and i when no confusion can result.

A relational *instance* I of σ is a set of *ground facts* of the form $R(\boldsymbol{a})$ where R is a relation name and $\boldsymbol{a}$ an $|R|$-tuple of values. The *size* $|I|$ of a finite instance I is its number of facts. The *active domain* $\mathrm{dom}(I)$ of I is the set of the elements which appear in I. For any position $R^i \in \mathsf{Pos}(\sigma)$, we define the *projection* $\pi_{R^i}(I)$ of I to R^i as the set of the elements of $\mathrm{dom}(I)$ that occur at position R^i in I. For $L \subseteq \mathsf{Pos}(R)$, the projection $\pi_L(I)$ is a set of $|L|$-tuples defined analogously; for convenience, departing from the unnamed perspective, we index those tuples by the positions of L rather than by $\{1,\ldots,|L|\}$. A *superinstance* of I is a (not necessarily finite) instance I' such that $I \subseteq I'$.

A *homomorphism* from an instance I to an instance I' is a mapping $h : \mathrm{dom}(I) \to \mathrm{dom}(I')$ such that, for every fact $F = R(\boldsymbol{a})$ of I, the fact $h(F) := R(h(a_1),\ldots,h(a_{|R|}))$ is in I'.

Constraints. We consider *integrity constraints* (or *dependencies*) which are special sentences of first-order logic without function symbols. We write $I \models \Sigma$ when instance I satisfies constraints Σ, and we then call I a *model* of Σ.

An *inclusion dependency* ID is a sentence of the form $\tau : \forall \boldsymbol{x}\, R(x_1,\ldots,x_n) \to \exists \boldsymbol{y}\, S(z_1,\ldots,z_m)$, where $\boldsymbol{z} \subseteq \boldsymbol{x} \cup \boldsymbol{y}$ and no variable occurs at two different positions of the same fact. The *exported variables* are the variables of $\boldsymbol{x}$ that occur in $\boldsymbol{z}$. This work only studies *unary inclusion dependencies* (UIDs) which are the IDs with exactly one exported variable. We write a UID τ as $R^p \subseteq S^q$, where R^p and S^q are the positions of $R(\boldsymbol{x})$ and $S(\boldsymbol{z})$ where the exported variable occurs. For instance, the UID $\forall xy\, R(x,y) \to \exists z\, S(y,z)$ is written $R^2 \subseteq S^1$. We assume without loss of generality that there are no *trivial* UIDs of the form $R^p \subseteq R^p$.

A *functional dependency* FD is a sentence of the form $\phi : \forall \boldsymbol{x}\boldsymbol{y}\, (R(x_1,\ldots,x_n) \wedge R(y_1,\ldots,y_n) \wedge \bigwedge_{R^l \in L} x_l = y_l) \to x_r = y_r$, where $L \subseteq \mathsf{Pos}(R)$ and $R^r \in \mathsf{Pos}(R)$. For brevity, we write ϕ as $R^L \to R^r$. We call ϕ a *unary functional dependency* UFD if $|L| = 1$; otherwise it is *higher-arity*. For instance, $\forall xx'yy'\, R(x,x') \wedge R(y,y') \wedge x' = y' \to x = y$ is a UFD, and we write it $R^2 \to R^1$. We assume that $|L| > 0$, i.e., we do not allow nonstandard or degenerate FDs. We call ϕ *trivial* if $R^r \in R^L$, in which case ϕ always holds; again we disallow trivial FDs. Two facts $R(\boldsymbol{a})$ and $R(\boldsymbol{b})$ *violate* a non-trivial FD ϕ if $\pi_L(\boldsymbol{a}) = \pi_L(\boldsymbol{b})$ but $a_r \neq b_r$.

For $L, L' \subseteq \mathsf{Pos}(R)$, we write $R^L \to R^{L'}$ the conjunction of FDs $R^L \to R^l$ for $R^l \in L'$. In particular, conjunctions of the form $\kappa : R^L \to R$ (i.e., $L' = \mathsf{Pos}(R)$) are called *key dependencies*. The key κ is *unary* if $|L| = 1$. If κ holds on a relation R, we call L a *key* (or *unary key*) of R.

II.2. Implication and Finite Implication

We say that a conjunction of constraints Σ in a class CL *finitely implies* a constraint ϕ if any *finite* instance that satisfies Σ also satisfies ϕ. We say that Σ *implies* ϕ if the same holds even for infinite instances. The *closure* Σ^* of Σ is the set of constraints of CL which are implied by Σ, and the *finite closure* $\Sigma^{\mathrm{f}*}$ is the set of those which are finitely implied.

An *axiomatization* of implication for CL is a set of deduction rules (which, given dependencies in CL, deduce new dependencies in CL), with the following property: for any conjunction Σ of dependencies in CL, letting Σ' be the result of defining $\Sigma' := \Sigma$ and applying iteratively the deduction rules while possible to inflate Σ', then Σ' is *exactly* Σ^*. An *axiomatization of finite implication* is defined similarly but for $\Sigma^{\mathrm{f}*}$.

Implication for ID. Given a set Σ of IDs, it is known [Casanova et al. 1984] that an ID τ is implied by Σ iff it is finitely implied. Further, when Σ are UIDs, we can easily compute in PTIME the set of implied UIDs (from which we exclude the trivial ones), by closing Σ under the *transitivity rule* [Casanova et al. 1984]: if $R^p \subseteq S^q$ and $S^q \subseteq T^r$ hold in Σ, then so is $R^p \subseteq T^r$ unless it is trivial. We call Σ *transitively closed* if it is thus closed.

Implication for FDs. Again, a set Σ_{FD} of FDs implies an FD ϕ iff it finitely implies it: see, e.g., [Cosmadakis et al. 1990]. The standard axiomatization of FD implication is given in [Armstrong 1974], and includes the *transitivity rule*: for any $R \in \sigma$ and $L, L', L'' \subseteq \mathsf{Pos}(R)$, if $R^L \to R^{L'}$ and $R^{L'} \to R^{L''}$ hold in Σ_{FD}, then so does $R^L \to R^{L''}$.

Implication for UIDs and FDs. It was shown in [Cosmadakis et al. 1990] that implication for conjunctions Σ_{UID} of UIDs and Σ_{FD} of FDs can be axiomatized by the above UID and FD rules in isolation. However, for *finite* implication, we must add a *cycle rule*, which we now define.

Let Σ be a conjunction of dependencies formed of UIDs Σ_{UID} and FDs Σ_{FD}. Define the *reverse* of an UFD $\phi : R^p \to R^q$ as $\phi^{-1} := R^q \to R^p$, and the *reverse* of a UID $\tau : R^p \subseteq S^q$ as $\tau^{-1} := S^q \subseteq R^p$.

A *cycle* in Σ is a sequence of UIDs and UFDs of Σ_{UID} and Σ_{FD} of the following form: $R_1^{p_1} \subseteq R_2^{q_2}$, $R_2^{p_2} \rightarrow R_2^{q_2}$, $R_2^{p_2} \subseteq R_3^{q_3}$, $R_3^{p_3} \rightarrow R_3^{q_3}$, $\ldots$, $R_{n-1}^{p_{n-1}} \subseteq R_n^{q_n}$, $R_n^{p_n} \rightarrow R_n^{q_n}$, $R_n^{p_n} \subseteq R_1^{q_1}$, $R_1^{p_1} \rightarrow R_1^{q_1}$. The *cycle rule*, out of such a cycle, deduces the reverse of each UID and of each UFD in the cycle. We then have:

THEOREM II.1 ([COSMADAKIS ET AL. 1990], THEOREM 4.1). *The UID and FD deduction rules and the cycle rule are an axiomatization of finite implication for UIDs and FDs.*

In terms of complexity, this implies:

COROLLARY II.2 ([COSMADAKIS ET AL. 1990], COROLLARY 4.4). *Given UIDs Σ_{UID} and FDs Σ_{FD}, and a UID or FD τ, we can check in PTIME whether τ is finitely implied by Σ_{UID} and Σ_{FD}.*

II.3. Queries and QA

Queries. An *atom* $A = R(\boldsymbol{t})$ consists of a relation name R and an $|R|$-tuple $\boldsymbol{t}$ of variables or constants. This work studies the *conjunctive queries* CQ, which are existentially quantified conjunctions of atoms, such that each variable in the quantification occurs in some atom. The *size* $|q|$ of a CQ q is its number of atoms. A CQ is *Boolean* if it has no free variables.

A Boolean CQ q *holds* in an instance I exactly when there is a homomorphism h from the atoms of q to I such that h is the identity on the constants of q (we call this a *homomorphism from q to I*). We call such an h a *match* of q in I, and by a slight abuse of terminology we also call the image of h a *match* of q in I.

QA problems. We define the *unrestricted open-world query answering* problem (UQA) as follows: given a finite instance I, a conjunction of constraints Σ, and a Boolean CQ q, decide whether there is a superinstance of I that satisfies Σ and violates q. If there is none, we say that I and Σ *entail* q and write $(I,\Sigma) \models_{\text{unr}} q$. In other words, UQA asks whether the first-order formula $I \wedge \Sigma \wedge \neg q$ has some (possibly infinite) model.

This work focuses on the *finite query answering problem* (FQA), which is the variant of open-world query answering where we require the counterexample superinstance to be *finite*; if no such counterexample exists, we write $(I,\Sigma) \models_{\text{fin}} q$. Of course $(I,\Sigma) \models_{\text{unr}} q$ implies $(I,\Sigma) \models_{\text{fin}} q$.

The *combined complexity* of the UQA and FQA problems, for a fixed class CL of constraints, is the complexity of deciding it when all of I, Σ (in CL) and q are given as input. The *data complexity* is defined by assuming that Σ and q are fixed, and only I is given as input.

Assumptions on queries. Throughout this work, we will make three assumptions about CQs, without loss of generality for UQA and FQA. First, *we assume that CQs are constant-free*. Indeed, for each constant $c \in \text{dom}(I_0)$, we could otherwise do the following: add a fresh relation P_c to the signature, add a fact $P_c(c)$ to I_0, replace c in q by an existentially quantified variable x_c, and add the atom $P_c(x_c)$ to q. It is then clear that UQA with the rewritten instance and query is equivalent to UQA with the original instance and query under any constraints (remember that our constraints are constant-free); the same is true for FQA.

Second, we *assume all CQs to be Boolean*, unless otherwise specified. Indeed, to perform UQA for non-Boolean queries (where the domain of the free variables is that of the base instance I_0), we can always enumerate all possible assignments, and solve our problem by solving polynomially many instances of the UQA problem with Boolean queries. Again, the same is true of FQA.

Third, *we assume all CQs to be connected*. A CQ q is *disconnected* if there is a partition of its atoms in two non-empty sets $\mathcal{A}$ and $\mathcal{A}'$, such that no variable occurs both in an atom of $\mathcal{A}$ and in one atom of $\mathcal{A}'$. In this case, the query $q : \exists \boldsymbol{x}\boldsymbol{y}\, \mathcal{A}(\boldsymbol{x}) \wedge \mathcal{A}(\boldsymbol{y})$ can be rewritten to $q_2 \wedge q_2'$, for two CQs q_2 and q_2' of strictly smaller size. In this paper, we will show that, on finitely closed dependencies, FQA and UQA coincide for *connected* queries. This clearly implies the same for disconnected queries, by considering all their connected subqueries. Hence, we can assume that queries are connected.

Chase. We say that a superinstance I' of an instance I is *universal* for constraints Σ if $I' \models \Sigma$ and if for any CQ q, $I' \models q$ iff $(I, \Sigma) \models_{\text{unr}} q$. We now recall the definition of the *chase* [Abiteboul et al. 1995; Onet 2013], a standard construction of (generally infinite) universal superinstances. We assume that we have fixed an infinite set $\mathcal{N}$ of *nulls* which is disjoint from $\text{dom}(I)$. We only define the chase for transitively closed UIDs, which we call the *UID chase*.

We say that a fact $F_{\text{a}} = R(\boldsymbol{a})$ of an instance I is an *active fact* for a UID $\tau : R^p \subseteq S^q$ if, writing $\tau : \forall \boldsymbol{x}\, R(\boldsymbol{x}) \to \exists \boldsymbol{y}\, S(\boldsymbol{z})$, there is a homomorphism from $R(\boldsymbol{x})$ to F_{a} but no such homomorphism can be extended to a homomorphism from $\{R(\boldsymbol{x}), S(\boldsymbol{z})\}$ to I. In this case we say that we *want* to apply the UID τ to a_p, written $a_p \in \mathsf{Wants}(I, \tau)$. Note that $\mathsf{Wants}(I, \tau) = \pi_{R^p}(I) \backslash \pi_{S^q}(I)$. For a conjunction Σ_{UID} of UIDs, we may also write $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}}(I, S^q)$ if there is $\tau \in \Sigma_{\text{UID}}$ of the form $\tau : U^v \subseteq S^q$ such that $a \in \mathsf{Wants}(I, \tau)$; we drop the subscript when there is no ambiguity.

The result of a *chase step* on the active fact $F_{\text{a}} = R(\boldsymbol{a})$ for $\tau : R^p \subseteq S^q$ in I (we call this *applying* τ to F_{a}) is the superinstance I' of I obtained by adding a new fact $F_{\text{n}} = S(\boldsymbol{b})$ defined as follows: we set $b_q := a_p$, which we call the *exported element* (and S^q the *exported position* of F_{n}), and use fresh nulls from $\mathcal{N}$ to instantiate the existentially quantified variables of τ and complete F_{n}, using a different null at each position; we say the corresponding elements are *introduced* at F_{n}. This ensures that F_{a} is no longer an active fact in I' for τ.

A *chase round* of a conjunction Σ_{UID} of UIDs on I is the result of applying simultaneous chase steps on all active facts for all UIDs of Σ_{UID}, using distinct fresh elements. The *UID chase* $\mathsf{Chase}(I, \Sigma_{\text{UID}})$ of I by Σ_{UID} is the (generally infinite) fixpoint of applying chase rounds. It is a universal superinstance for Σ_{UID} [Fagin et al. 2003].

As we are chasing by transitively closed UIDs, if we perform the *core chase* [Deutsch et al. 2008; Onet 2013] rather than the UID chase that we just defined, we can ensure the following *Unique Witness Property*: for any element $a \in \text{dom}(\mathsf{Chase}(I, \Sigma_{\text{UID}}))$ and position R^p of σ, if two different facts of $\mathsf{Chase}(I, \Sigma_{\text{UID}})$ contain a at position R^p, then they are both facts of I. In our context, however, the core chase matches the UID chase defined above, except at the first round. Thus, modulo the first round, by $\mathsf{Chase}(I, \Sigma_{\text{UID}})$ we refer to the UID chase, which has the Unique Witness Property. See Appendix A for details.

Finite controllability. We say a conjunction of constraints Σ is *finitely controllable* for CQ if FQA and UQA coincide: for every finite instance I and every Boolean CQ q, $(I, \Sigma) \models_{\text{unr}} q$ iff $(I, \Sigma) \models_{\text{fin}} q$.

It was shown in [Rosati 2006; Rosati 2011] that, while conjunctions of IDs are finitely controllable, even conjunctions of UIDs and FDs may not be. It was later shown in [Rosati 2008] that the finite closure process could be used to reduce UQA to FQA for some constraints on relations of arity at most two. Following the same idea, we say that a conjunction of constraints Σ is *finitely controllable up to finite closure* if for every finite instance I, and Boolean CQ q, $(I, \Sigma) \models_{\text{fin}} q$ iff $(I, \Sigma^{\text{f*}}) \models_{\text{unr}} q$, where $\Sigma^{\text{f*}}$ is the finite closure defined by Theorem II.1. If Σ is finitely controllable up to finite closure, then we can reduce FQA to UQA, even if finite controllability does not hold, by computing the finite closure $\Sigma^{\text{f*}}$ of Σ and solving UQA on $\Sigma^{\text{f*}}$.

III. MAIN RESULT AND OVERALL APPROACH

We study open-world query answering for FDs and UIDs. For UQA, the following is already known:

PROPOSITION III.1. *UQA for FDs and UIDs has* AC^0 *data complexity and* NP*-complete combined complexity.*

PROOF. UQA for UIDs in isolation is NP-complete in combined complexity. The lower bound is immediate from query evaluation [Abiteboul et al. 1995], and [Johnson and Klug 1984] showed an NP upper bound for IDs with any fixed bound on the number of exported variables (which they call "width": in their terminology, UIDs are IDs of width 1). For data complexity, the upper bound is from the first-order rewritability of certain answers for arbitrary IDs, from [Calì et al. 2003b].

For UIDs and FDs, clearly the lower bound on combined complexity also applies. The upper bounds are proved by observing that UIDs and FDs are *separable*, namely, for any FDs Σ_{FD} and UIDs Σ_{UID}, for any instance I_0 and CQ q, if $I_0 \models \Sigma_{\text{FD}}$ then we have $(I_0, \Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}) \models_{\text{unr}} q$ iff $(I_0, \Sigma_{\text{UID}}) \models_{\text{unr}} q$. Assuming separability, to decide UQA for Σ_{UID} and Σ_{FD}, we first check whether $I_0 \models \Sigma_{\text{FD}}$, in PTIME combined complexity, and AC^0 data complexity as Σ_{FD} is expressible in first-order logic. If $I_0 \not\models \Sigma_{\text{FD}}$, UQA is vacuously true. Otherwise, we then determine whether $(I_0, \Sigma_{\text{UID}}) \models_{\text{unr}} q$, using the upper bound for UQA for UIDs. By separability, the answer to UQA under Σ_{UID} is the same as the answer to UQA under $\Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}$.

Hence, all that remains to show is that UIDs and FDs are always separable. This follows from the *non-conflicting condition* of [Calì et al. 2003a; Calì et al. 2012] but we give a simpler self-contained argument. Assume that I_0 satisfies Σ_{FD}. It is obvious that $(I_0, \Sigma_{\text{UID}}) \models_{\text{unr}} q$ implies $(I_0, \Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}) \models_{\text{unr}} q$, so let us prove the converse implication. We do it by noticing that the chase $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ satisfies Σ_{FD}. Indeed, assuming to the contrary the existence of F and F' in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ violating an FD of Σ_{FD}, there must exist a position $R^p \in \mathsf{Pos}(\sigma)$ such that $\pi_{R^p}(F) = \pi_{R^p}(F')$. Yet, by the Unique Witness Property, this implies that F and F' are facts of I_0, but we assumed that $I_0 \models \Sigma_{\text{FD}}$, a contradiction.

Hence, $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ satisfies Σ_{FD}, so it is a superinstance of I_0 that satisfies $\Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}$. Hence, $(I_0, \Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}) \models_{\text{unr}} q$ implies that we must have $\mathsf{Chase}(I_0, \Sigma_{\text{UID}}) \models q$. By universality of the chase, this implies $(I_0, \Sigma_{\text{UID}}) \models_{\text{unr}} q$. Hence, the converse implication is proven, so the two UQA problems are equivalent, which implies that Σ_{UID} and Σ_{FD} are separable. □

In the *finite case*, however, even the decidability of FQA for FDs and UIDs was not known. This paper shows that it is decidable, and that the complexity matches that of UQA:

THEOREM III.2. *FQA for FDs and UIDs has* AC^0 *data complexity and* NP-*complete combined complexity.*

This result follows from our Main Theorem, which is proven in the rest of this paper:

THEOREM III.3 (MAIN THEOREM). *Conjunctions of FDs and UIDs are finitely controllable up to finite closure.*

From the Main Theorem, we can prove Theorem III.2, using the closure process of [Cosmadakis et al. 1990]:

PROOF OF THEOREM III.2. Again, the NP-hardness lower bound is immediate from query evaluation [Abiteboul et al. 1995], so we only show the upper bounds. Consider an instance of FQA for FDs and UIDs, consisting of an instance I_0, a conjunction Σ of IDs Σ_{UID} and FDs Σ_{FD}, and a CQ q. Let Σ^*_{FD} be the FDs and Σ^*_{UID} the UIDs of the finite closure $\Sigma^{\text{f}*}$. By our Main Theorem, we have $(I_0, \Sigma) \models_{\text{fin}} q$ iff $(I_0, \Sigma^{\text{f}*}) \models_{\text{unr}} q$. As the computation of $\Sigma^{\text{f}*}$ from Σ is data-independent, the data complexity upper bounds follow from Proposition III.1, so we need only show the combined complexity upper bound.

Materializing $\Sigma^{\text{f}*}$ from the input may take exponential time, which we cannot afford, so we need a more clever approach. Remember from the proof of Proposition III.1 that, as $\Sigma^{\text{f}*}$ consists of UIDs and FDs, it is separable. Hence, to solve UQA for I_0, $\Sigma^{\text{f}*}$ and q, as $\Sigma^{\text{f}*}$ is separable, we need to perform two steps: (1) check whether $I_0 \models \Sigma^*_{\text{FD}}$ (2) if yes, solve UQA for I_0, Σ^*_{UID} and q.

To perform step 1, compute in PTIME the set Σ^*_{UFD} of the UFDs of $\Sigma^{\text{f}*}$, using Corollary II.2. By [Cosmadakis et al. 1990] (remark above Corollary 4.4), all non-unary FDs in Σ^*_{FD} are implied by $\Sigma^*_{\text{UFD}} \wedge \Sigma_{\text{FD}}$ under the axiomatization of FD implication; hence, to check whether $I_0 \models \Sigma^*_{\text{FD}}$, it suffices to check whether $I_0 \models \Sigma^*_{\text{UFD}}$ and $I_0 \models \Sigma_{\text{FD}}$, which we do in PTIME.

To perform step 2, compute Σ^*_{UID} in PTIME by considering each possible UID (there are polynomially many) and determining in PTIME from Σ whether it is in $\Sigma^{\text{f}*}$, using Corollary II.2. Then, solve UQA in NP combined complexity by Proposition III.1. The entire process takes NP combined complexity, and the answer matches that of FQA by our Main Theorem, which proves the NP upper bound. □

In this section, we first explain how we can prove Theorem III.3 from a different statement, namely: we can construct *finite universal models* for finitely closed UIDs and FDs. We conclude this section with the outline of the proof of this result (Theorem III.6) which will be developed in the rest of this paper.

III.1. Finite Universal Superinstances

Our Main Theorem claims that a certain class of constraints, namely finitely closed UIDs and FDs, are finitely controllable for the class of conjunctive queries (CQ). To prove this, it will be easier to work with a notion of *k-sound* and *k-universal instances*.

Definition III.4. For $k \in \mathbb{N}$, we say that a superinstance I of an instance I_0 is ***k*-sound** for constraints Σ and CQs (and I_0) if, for every CQ q of size $\leq k$ such that $I \models q$, we have $(I_0, \Sigma) \models_{\mathrm{unr}} q$. We say it is ***k*-universal** if the converse also holds: $I \models q$ whenever $(I_0, \Sigma) \models_{\mathrm{unr}} q$. For a subclass Q of CQs, we call I ***k*-sound** or ***k*-universal** for Σ and Q if the same holds for all queries q of size $\leq k$ that are in Q.

We say that a class CL of constraints *has finite universal superinstances* for a class Q of CQs, if for any constraints Σ of CL, for any $k \in \mathbb{N}$, for any instance I_0, if I_0 has some superinstance that satisfies Σ, then it has a *finite* superinstance that satisfies Σ and is k-sound for Σ and Q (and hence is also k-universal for Σ and Q).

We will thus show that the class of finitely closed UIDs and FDs have finite universal superinstances for CQs. We explain why this implies our Main Theorem:

PROPOSITION III.5. *If constraint class CL has finite universal superinstances for query class Q, then CL is finitely controllable for Q.*

PROOF. Let Σ be constraints in CL, I_0 be a finite instance and q be a query in Q. We show that $(I_0, \Sigma) \models_{\mathrm{unr}} q$ iff $(I_0, \Sigma) \models_{\mathrm{fin}} q$. The forward implication is immediate: if all superinstances of I_0 that satisfy Σ must satisfy q, then so do the finite ones.

For the converse implication, assume that $(I_0, \Sigma) \not\models_{\mathrm{unr}} q$. In particular, this implies that I_0 has some superinstance that satisfies Σ, as otherwise the entailment would be vacuously true. As CL has finite universal superinstances for Q, let I be a finite k-sound superinstance of I_0 that satisfies Σ, where $k := |q|$. As I is k-sound, we have $I \not\models q$, and as $I \models \Sigma$, I witnesses that $(I_0, \Sigma) \not\models_{\mathrm{fin}} q$. This proves the converse direction, so we have established finite controllability. □

So, in this paper, we will actually show the following restatement of the Main Theorem:

THEOREM III.6 (UNIVERSAL MODELS). *The class of finitely closed UIDs and FDs has finite universal models for CQ: for every conjunction Σ of FDs Σ_{FD} and UIDs Σ_{UID} closed under finite implication, for any $k \in \mathbb{N}$, for every finite instance I_0 that satisfies Σ_{FD}, there exists a finite k-sound superinstance I of I_0 that satisfies Σ.*

Indeed, once we have shown this, we can easily deduce the Main Theorem, namely, that any conjunction Σ of FDs and UIDs is finitely controllable up to finite closure. Indeed, for any such Σ, for any instance I_0 and CQ q, we have $(I_0, \Sigma) \models_{\mathrm{fin}} q$ iff $(I_0, \Sigma^{\mathrm{f}*}) \models_{\mathrm{fin}} q$: the forward statement is because any finite model of Σ is a model of $\Sigma^{\mathrm{f}*}$, and the backward statement is tautological. Now, from the Universal Models Theorem and Proposition III.5, we know that $\Sigma^{\mathrm{f}*}$ is finitely controllable, so that $(I_0, \Sigma^{\mathrm{f}*}) \models_{\mathrm{fin}} q$ iff $(I_0, \Sigma^{\mathrm{f}*}) \models_{\mathrm{unr}} q$. We have thus shown that $(I_0, \Sigma) \models_{\mathrm{fin}} q$ iff $(I_0, \Sigma^{\mathrm{f}*}) \models_{\mathrm{unr}} q$, which concludes the proof of the Main Theorem.

Hence, we will show the Universal Models Theorem in the rest of this paper. We proceed in incremental steps, following the plan that we outline next.

III.2. Proof Structure

We first make a simplifying assumption on the signature, without loss of generality, to remove *useless* relations. Given an instance I_0, UIDs Σ_{UID} and FDs Σ_{FD}, it may the be case that the signature σ

Table I: Roadmap of intermediate results.

	Signature	Universality	Constraints	Query
Section IV:	binary	weakly-sound	reversible UIDs, UFDs	ACQ
Section V:	**arbitrary**	weakly-sound	reversible UIDs, UFDs	ACQ
Section VI:	arbitrary	***k*-sound**	reversible UIDs, UFDs	ACQ
Section VII:	arbitrary	*k*-sound	**finitely closed** UIDs, UFDs	ACQ
Section VIII:	arbitrary	*k*-sound	finitely closed UIDs, **FDs**	ACQ
Section IX:	arbitrary	*k*-sound	finitely closed UIDs, FDs	**CQ**

contains a relation R that does not occur in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, namely, it does not occur in I_0 and the existence of an R-fact is not implied by Σ_{UID}. In this case, relation R is *useless*: a CQ q involving R will never be entailed under Σ, neither on unrestricted nor on finite models, unless I_0 has no completion at all satisfying the constraints. In any case, the query q can be replaced by the trivial CQ False, which is only (vacuously) entailed if there are no completions.

Hence, we can always remove useless relations from the signature, up to rewriting the query to the false query. Thus, without loss of generality, *we always assume that the signature contains no useless relations* in this sense: all relations of the signature occur in the chase.

We now present several assumptions that we use to prove weakenings of the Universal Models Theorem. The first one is on queries, which we require to be *acyclic*. The second is on FDs, which we require to be *unary*, i.e., UFDs. The third one is to replace k-soundness by the simpler notion of *weak-soundness*. Then we present two additional assumptions: the first one, reversible, is on the constraints, and requires that they have a certain special form; the second one, binary, is on the constraints and signature, which we require to be binary. In the next section, we show the Universal Models Theorem under all these assumptions, and then we lift the assumptions one by one, in each section. See Table I for a synopsis.

Hence, let us present the assumptions that we will make (and later lift).

Acyclic queries. It will be helpful to focus first on the subset of *acyclic* CQs, denoted ACQ, which are the queries that contain no *Berge cycle*. Formally:

Definition III.7. A *Berge cycle* in a CQ q is a sequence $A_1, x_1, A_2, x_2, \ldots, A_n, x_n$ with $n \geq 2$, where the A_i are pairwise distinct atoms of q, the x_i are pairwise distinct variables of q, and x_i occurs in A_i and A_{i+1} for $1 \leq i \leq n$ (with addition modulo n, so x_n occurs in A_n and A_1). A query q is in ACQ if q has no Berge cycle and if no variable of q occurs more than once in the same atom.

Equivalently, consider the *incidence multigraph* of q, namely, the bipartite undirected multigraph on variables and atoms obtained by putting one edge between variable x and atom A for every time where x occurs in A (possibly multiple times). Then q is in ACQ iff its incidence multigraph is acyclic in the standard sense.

Example III.8. The queries $\exists x\, R(x,x)$, $\exists xy\, R(x,y) \wedge S(x,y)$, and $\exists xyz\, R(x,y) \wedge R(y,z) \wedge R(z,x)$ are not in ACQ: the first one has an atom with two occurrences of the same variable, the other two have a Berge cycle. The following query is in ACQ: $\exists xyzw\, R(x,y,z) \wedge S(x) \wedge T(y,w) \wedge U(w)$.

Intuitively, in the chase, all query matches are acyclic unless they involve some cycle in the initial instance I_0. Hence, only acyclic CQs have matches, except those that match on I_0 or those whose cycles have self-homomorphic matches, so, in a k-sound model, the CQs of size $\leq k$ which hold are usually acyclic. For this reason, we focus only on ACQ queries first. We will ensure in Section IX that cyclic queries of size $\leq k$ have no matches.

Unary FDs. We will first show our result for *unary* FDs (UFDs); recall from Section II that they are the FDs with exactly one determining attribute. We do this because the finite closure construction of [Cosmadakis et al. 1990] is not concerned with higher-arity FDs, except for the UFDs that they imply. Hence, while the UFDs of the finite closure have a special structure that we can rely on, the higher-arity FDs are essentially arbitrary. This is why we deal with them only in Section VIII, using a different approach.

***k*-*soundness and weak-soundness*.** Rather than proving that UIDs and UFDs have finite universal models for ACQ, it will be easier to prove first that they have *1-universal models*. More specifically, we will construct *weakly-sound superinstances*, which satisfy a *sufficient* condition for them to be *1-sound*:

Definition III.9. A superinstance I of an instance I_0 is **weakly-sound** for a set of UIDs Σ_{UID} and for I_0 if the following holds:

— for any $a \in \mathrm{dom}(I_0)$ and $R^p \in \mathsf{Pos}(\sigma)$, if $a \in \pi_{R^p}(I)$, then either $a \in \pi_{R^p}(I_0)$ or $a \in \mathsf{Wants}(I_0, R^p)$;
— for any $a \in \mathrm{dom}(I) \setminus \mathrm{dom}(I_0)$ and $R^p, S^q \in \mathsf{Pos}(\sigma)$, if a $\in \pi_{R^p}(I)$ and $a \in \pi_{S^q}(I)$ then either we have $R^p = S^q$ or $R^p \subseteq S^q$ and $S^q \subseteq R^p$ are in Σ_{UID}.

Thus, we first show that UFDs and UIDs have *finite weakly-universal superinstances* for ACQ, defined analogously to Definition III.4: for any constraints Σ_{U} of UFDs Σ_{UFD} and UIDs Σ_{UID}, for any query q in ACQ, for any instance I_0, if I_0 has a superinstance that satisfies Σ_{U}, then it has a *finite* superinstance that does and is weakly-sound for Σ_{UID} and I_0. This restriction is lifted in Section VI.

***Assumption reversible*.** We will initially make a simplifying assumption on the structure of the UIDs and UFDs, which we call reversible. This assumption is motivated by the finite closure rules of Theorem II.1; intuitively, it amounts to assuming that a certain *constraint graph* defined from the dependencies has a single connected component:

Definition III.10. Let $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ be a set of UIDs and Σ_{UFD} be a set of UFDs. We call $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and Σ_{UFD} **reversible** if:

— $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ is closed under implication, and so is Σ_{UFD};
— All UIDs in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ are reversible (i.e., their reverses are also in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$);
— for any UFD $\phi : R^p \to R^q$ in Σ_{UFD}, if R^p occurs in some UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and R^q also occurs in some UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, then ϕ is reversible, i.e., ϕ^{-1} is also in Σ_{UFD}.

Assumption reversible:. The UIDs Σ_{UID} and UFDs Σ_{UFD} are reversible.

When assumption reversible is made, we will write the UIDs $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ rather than Σ_{UID}. Observe that $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and Σ_{UFD} are then finitely closed: they are closed under UID and UFD implication, and the UIDs and UFDs of any cycle must be reversible. To lift reversible and generalize to the general case, we will follow an SCC decomposition of the constraint graph to manage each SCC separately. See Section VII for details.

***Second assumption*.** We will start our proof in Section IV by introducing important notions in the much simpler case of a binary signature. For this, we will initially make the following assumption binary on the signature and on Σ:

Assumption binary:. Each relation R has arity 2 and the UFDs $R^1 \to R^2$ and $R^2 \to R^1$ hold in Σ. We will lift this assumption in Section V.

***Roadmap*.** Each of the next sections will prove that a certain constraint class has finite universal models for a certain query class in a certain sense, under certain assumptions. Table I summarizes the results that are proved in each section.

The rest of the paper follows this roadmap: each section starts by stating the result that it proves.

IV. WEAK SOUNDNESS ON BINARY SIGNATURES

THEOREM IV.1. *Reversible UIDs and UFDs have finite weakly-universal superinstances for ACQs under assumption binary.*

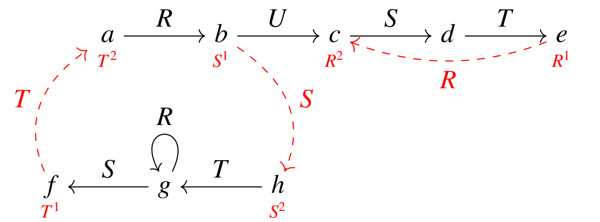

Fig. 1: Connecting balanced instances (see Example IV.4)

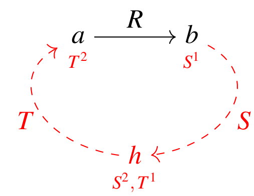

Fig. 2: Using helper elements to balance (see Example IV.5)

We prove this result in this section. Fix an instance I_0 and reversible constraints $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$ formed of UIDs $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and UFDs Σ_{UFD}. Assume that $I_0 \models \Sigma_{\mathrm{UFD}}$ as the question is vacuous otherwise, and make assumption binary.

Our goal is to construct a weakly-sound superinstance I of I_0 that satisfies $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$. We do so by a *completion process* that adds new (binary) facts to connect elements together. As all possible UFDs hold, if we extend I_0 to I by adding a new fact $R(a_1, a_2)$, we must have $a_i \notin \pi_{R^i}(I_0)$ for $i \in \{1,2\}$. Hence, by weak soundness, if $a_i \in \mathrm{dom}(I_0)$ then we must have $a_i \in \mathsf{Wants}(I_0, R^i)$. Our task is this section is thus to complete I_0 to I by adding R-facts, for each relation R, that connect together elements of $\mathsf{Wants}(I_0, R^1)$ and $\mathsf{Wants}(I_0, R^2)$.

IV.1. Completing Balanced Instances

One easy situation to do this is when the instance I_0 is *balanced*: for every relation R, we can construct a bijection between the elements that want to be in R^1 and those that want to be in R^2:

Definition IV.2. We call I_0 **balanced** (for UIDs $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$) if, for every two positions R^p and R^q such that $R^p \to R^q$ and $R^q \to R^p$ are in Σ_{UFD}, we have $\left|\mathsf{Wants}_{\Sigma_{\mathrm{UID}}^{\mathrm{rev}}}(I_0, R^p)\right| = \left|\mathsf{Wants}_{\Sigma_{\mathrm{UID}}^{\mathrm{rev}}}(I_0, R^q)\right|$.

If I_0 is balanced, we can show Theorem IV.1 by constructing I with $\mathrm{dom}(I) = \mathrm{dom}(I_0)$, adding new facts that pair together the existing elements:

PROPOSITION IV.3. *Assuming binary and reversible, any balanced finite instance I_0 satisfying Σ_{UFD} has a finite weakly-sound superinstance I that satisfies $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$, with $\mathrm{dom}(I) = \mathrm{dom}(I_0)$.*

We first exemplify this process:

Example IV.4. Consider four binary relations R, S, T, and U, with the UIDs $R^2 \subseteq S^1$, $S^2 \subseteq T^1$, $T^2 \subseteq R^1$ and their reverses, and the FDs prescribed by assumption binary. Consider $I_0 := \{R(a,b), U(b,c), S(c,d), T(d,e), S(g,f), R(g,g), T(h,g)\}$, as depicted by the black elements and solid black arrows in Figure 1.

We compute, for each element, the set of positions where it wants to be, and write it in red under each element in Figure 1 (in this example, it is a singleton set for each element). For instance, we have $\mathsf{Wants}(I_0, T^1) = \{f\}$. We observe that the instance is balanced: we have $\left|\mathsf{Wants}(I_0, R^1)\right| = \left|\mathsf{Wants}(I_0, R^2)\right|$, and likewise for S, T, and U.

We can construct a weakly-sound superinstance I of I_0 as $I := I_0 \sqcup \{R(e,c), S(b,h), T(f,a\}$: the additional facts are represented as dashed red arrows in Figure 1. Intuitively, we just create new facts that connect together elements to want to occur at the right positions.

We now give the formal proof of the result:

PROOF OF PROPOSITION IV.3. Define a bijection f_R from $\mathsf{Wants}(I_0, R^1)$ to $\mathsf{Wants}(I_0, R^2)$ for every relation R of σ; this is possible because I_0 is balanced.

Consider the superinstance I of I_0, with $\mathrm{dom}(I) = \mathrm{dom}(I_0)$, obtained by adding, for every R of σ, the fact $R(a, f_R(a))$ for every $a \in \mathsf{Wants}(I_0, R^1)$. I is clearly a finite weakly-sound superinstance of I_0, because for every $a \in \text{dom}(I)$, if a occurs at some position R^p in some fact F of I, then either F is a fact of I_0 and $a \in \pi_{R^p}(I_0)$, or F is a new fact in $I \backslash I_0$ and by definition $a \in \mathsf{Wants}(I_0, R^p)$.

Let us show that $I \models \Sigma_{\text{UFD}}$. Assume to the contrary that two facts $F = R(a_1, a_2)$ and $F' = R(a'_1, a'_2)$ in I witness a violation of a UFD $\phi : R^1 \to R^2$ of Σ_{UFD}. As $I_0 \models \Sigma_{\text{UFD}}$, one of F and F', say F, must be a new fact. By definition of the new facts, we have $a_1 \in \mathsf{Wants}(I_0, R^p)$, so that $a_1 \notin \pi_{R^1}(I_0)$. Now, as $\{F, F'\}$ is a violation, we must have $\pi_{R^1}(F) = \pi_{R^1}(F')$, so as $a_1 \notin \pi_{R^1}(I_0)$, F' must also be a new fact. Hence, by definition of the new facts, we have $a_2 = a'_2 = f_R(a_1)$, so $F = F'$, which contradicts the fact that F and F' violate ϕ. For UFDs ϕ of the form $R^2 \to R^1$, the proof is similar, but we have $a_1 = a'_1 = f_R^{-1}(a_2)$.

Let us now show that $I \models \Sigma^{\text{rev}}_{\text{UID}}$. Assume to the contrary that there is an active fact $F = R(a_1, a_2)$ that witnesses the violation of a UID $\tau : R^p \subseteq S^q$. If F is a fact of I_0, we had $a_p \in \mathsf{Wants}(I_0, S^q)$, so F cannot be an active fact in I as this violation was solved in I. So we must have $F \in I \backslash I_0$. Hence, by definition of the new facts, we had $a_p \in \mathsf{Wants}(I_0, R^p)$; so there must be $\tau' : T^r \subseteq R^p$ in $\Sigma^{\text{rev}}_{\text{UID}}$ such that $a_p \in \pi_{T^r}(I_0)$. Hence, because $\Sigma^{\text{rev}}_{\text{UID}}$ is transitively closed, either $T^r = S^q$ or the UID $T^r \subseteq S^q$ is in $\Sigma^{\text{rev}}_{\text{UID}}$. In the first case, as $a_p \in \pi_{T^r}(I_0)$, F cannot be an active fact for τ, a contradiction. In the second case, we had $a_p \in \mathsf{Wants}(I_0, S^q)$, so $a_p \in \pi_{S^q}(I)$ by definition of I, so again F cannot be an active fact for τ.

Hence, I is a finite weakly-sound superinstance of I_0 that satisfies $\Sigma^{\text{rev}}_{\text{U}}$ and with $\text{dom}(I) = \text{dom}(I_0)$, the desired claim. □

IV.2. Adding Helper Elements

If our instance I_0 is not balanced, we cannot use the construction that we just presented. The idea is then to *make* I_0 *balanced*, which we do by adding "helper" elements that we assign to positions. The following example illustrates this:

Example IV.5. We use the same signature and dependencies as in Example IV.4. Consider $I_0 := \{R(a,b)\}$, as depicted in Figure 2. We have $a \in \mathsf{Wants}(I_0, T^2)$ and $b \in \mathsf{Wants}(I_0, S^1)$; however $\mathsf{Wants}(I_0, S^2) = \mathsf{Wants}(I_0, T^1) = \emptyset$, so I_0 is not balanced.

Still, we can construct the weakly-sound superinstance $I := I_0 \sqcup \{S(b,h), T(h,a)\}$ that satisfies the constraints. Intuitively, we have added a "helper" element h and "assigned" it to the positions $\{S^2, T^1\}$, so we could connect b to h with S and h to a with T.

We will formalize this idea of augmenting the domain with *helper elements*, as a *partially-specified superinstance*, namely, an instance that is augmented with helpers assigned to positions. However, we first need to understand at which positions the helpers can appear, without violating weak-soundness:

Definition IV.6. For any two positions R^p and S^q, we write $R^p \sim_{\mathsf{ID}} S^q$ when $R^p = S^q$ or when $R^p \subseteq S^q$ is in $\Sigma^{\text{rev}}_{\text{UID}}$ (and hence $S^q \subseteq R^p$ is in $\Sigma^{\text{rev}}_{\text{UID}}$ by assumption reversible). We write $[R^p]_{\mathsf{ID}}$ the $\sim_{\mathsf{ID}}$-class of R^p.

As $\Sigma^{\text{rev}}_{\text{UID}}$ is transitively closed, $\sim_{\mathsf{ID}}$ is indeed an equivalence relation. Our choice of where to assign the helper elements will be represented as a mapping to an $\sim_{\mathsf{ID}}$-class. We call the result a *partially-specified superinstance*, or *pssinstance*:

Definition IV.7. A **pssinstance** of an instance I is a triple $P = (I, \mathcal{H}, \lambda)$ where $\mathcal{H}$ is a finite set of **helpers** and λ maps each $h \in \mathcal{H}$ to an $\sim_{\mathsf{ID}}$-class $\lambda(h)$.

We define $\mathsf{Wants}(P, R^p) := \mathsf{Wants}(I, R^p) \sqcup \{h \in \mathcal{H} \mid R^p \in \lambda(h)\}$.

In other words, in the pssinstance, elements of I want to appear at the same positions as before, and helper elements want to occur at their $\sim_{\mathsf{ID}}$-class according to λ. A *realization* of a pssinstance P is then a superinstance of its underlying instance I which adds the helper elements, and whose additional facts respect $\mathsf{Wants}(P, R^p)$:

Definition IV.8. A **realization** of $P = (I, \mathcal{H}, \lambda)$ is a superinstance I' of I such that $\text{dom}(I') = \text{dom}(I) \sqcup \mathcal{H}$, and, for any fact $R(\boldsymbol{a})$ of $I' \backslash I$ and $R^p \in \mathsf{Pos}(R)$, we have $a_p \in \mathsf{Wants}(P, R^p)$.

Example IV.9. In Example IV.5, a pssinstance of I_0 is $P := (I_0, \{h\}, \lambda)$ where $\lambda(h) := \{S^2, T^1\}$. Further, it is balanced. For instance, $\mathsf{Wants}(P, S^1) = \{b\}$ and $\mathsf{Wants}(P, S^2) = \{h\}$. The instance I in Example IV.5 is a realization of P.

It is easy to see that realizations of pssinstances are weakly-sound:

LEMMA IV.10 (BINARY REALIZATIONS ARE COMPLETIONS). *If I' is a realization of a pssinstance of I_0 then it is a weakly-sound superinstance of I_0.*

PROOF. Consider $a \in \mathrm{dom}(I')$ and $R^p \in \mathsf{Pos}(\sigma)$ such that $a \in \pi_{R^p}(I')$. As I' is a realization, we know that either $a \in \pi_{R^p}(I)$ or $a \in \mathsf{Wants}(P, R^p)$. By definition of $\mathsf{Wants}(P, R^p)$, and because $\mathcal{H} = \mathrm{dom}(I') \setminus \mathrm{dom}(I)$, this means that either $a \in \mathrm{dom}(I)$ and $a \in \pi_{R^p}(I) \sqcup \mathsf{Wants}(I, R^p)$, or $a \in \mathrm{dom}(I') \setminus \mathrm{dom}(I)$ and $R^p \in \lambda(a)$. Hence, let us check from the definition that I' is weakly-sound, which concludes:

— For any $a \in \mathrm{dom}(I)$ and $R^p \in \mathsf{Pos}(\sigma)$, we have established that $a \in \pi_{R^p}(I')$ implied that either $a \in \pi_{R^p}(I)$ or $a \in \mathsf{Wants}(I, R^p)$.
— For any $a \in \mathrm{dom}(I') \setminus \mathrm{dom}(I)$ and for any $R^p, S^q \in \mathsf{Pos}(\sigma)$, we have established that $a \in \pi_{R^p}(I')$ and $a \in \pi_{S^q}(I')$ implies that $R^p, S^q \in \lambda(a)$, so that $R^p \sim_{\mathsf{ID}} S^q$, hence $R^p = S^q$ or $R^p \subseteq S^q$ is in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$. □

IV.3. Putting it Together

What remains to show to conclude the proof of Theorem IV.1 is that we can construct a *balanced* pssinstance of I_0, even when I_0 itself is not balanced. By a *balanced* pssinstance, we mean the exact analogue of Definition IV.2 for pssinstances:

Definition IV.11. A pssinstance $P = (I, \mathcal{H}, \lambda)$ is **balanced** if for every two positions R^p and R^q such that $R^p \to R^q$ and $R^q \to R^p$ are in Σ_{UFD}, we have $|\mathsf{Wants}(P, R^p)| = |\mathsf{Wants}(P, R^q)|$.

If I_0 is balanced, the empty pssinstance $(I, \emptyset, \lambda)$, with λ the empty function, is a balanced pssinstance of I_0, and we could just complete I_0 as we presented before. We now show that, even if I_0 is not balanced, we can always construct a balanced pssinstance, thanks to the helpers:

LEMMA IV.12 (BALANCING). *Any finite instance I satisfying Σ_{UFD} has a balanced pssinstance.*

In fact, this lemma does not use assumption binary. We will accordingly reuse it in the next section.

PROOF. Let I be a finite instance. For any position R^p, define $o(R^p) := \mathsf{Wants}(I, R^p) \sqcup \pi_{R^p}(I)$, i.e., the elements that either appear at R^p or want to appear there. We show that $o(R^p) = o(S^q)$ whenever $R^p \sim_{\mathsf{ID}} S^q$, which is obvious if $R^p = S^q$, so assume $R^p \neq S^q$. First, we have $\pi_{R^p}(I) \subseteq o(S^q)$: elements in $\pi_{R^p}(I)$ want to appear at S^q unless they already do, and in both cases they are in $o(S^q)$. Second, elements of $\mathsf{Wants}(I, R^p)$ either occur at S^q, or at some other position T^r such that $T^r \subseteq R^p$ is a UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, so that by transitivity $T^r = S^q$ or $T^r \subseteq S^q$ also is, and so they want to be at S^q or they already are. Hence $o(R^p) \subseteq o(S^q)$; and symmetrically $o(S^q) \subseteq o(R^p)$. Thus, the set $o(R^p)$ only depends on the $\sim_{\mathsf{ID}}$-class of R^p.

Let $N := \max_{R^p \in \mathsf{Pos}(\sigma)} |o(R^p)|$, which is finite. We define for each $\sim_{\mathsf{ID}}$-class $[R^p]_{\mathsf{ID}}$ a set $p([R^p]_{\mathsf{ID}})$ of $N - |o(R^p)|$ fresh helpers. We let $\mathcal{H}$ be the disjoint union of the $p([R^p]_{\mathsf{ID}})$ for all classes $[R^p]_{\mathsf{ID}}$, and set λ to map the elements of $p([R^p]_{\mathsf{ID}})$ to $[R^p]_{\mathsf{ID}}$. We have thus defined a pssinstance $P = (I, \mathcal{H}, \lambda)$.

Let us now show that P is balanced. Consider now two positions R^p and R^q such that $\phi : R^p \to R^q$ and $\phi^{-1} : R^q \to R^p$ are in Σ_{UFD}, and show that $|\mathsf{Wants}(P, R^p)| = |\mathsf{Wants}(P, R^q)|$. We have $|\mathsf{Wants}(P, R^p)| = |\mathsf{Wants}(I, R^p)| + |p([R^p]_{\mathsf{ID}})| = |o(R^p)| - |\pi_{R^p}(I)| + N - |o(R^p)|$, which simplifies to $N - |\pi_{R^p}(I)|$. Similarly $|\mathsf{Wants}(P, R^q)| = N - |\pi_{R^q}(I)|$. Since $I \models \Sigma_{\mathrm{UFD}}$ and ϕ and ϕ^{-1} are in Σ_{UFD} we know that $|\pi_{R^p}(I)| = |\pi_{R^q}(I)|$. Hence, P is balanced, as we claimed. □

We had seen in Proposition IV.3 that we could construct a weakly-sound superinstance of a balanced I_0 by pairing together elements. We now generalize this claim to the balanced pssinstances that we constructed, showing that we can build realizations of balanced pssinstances that satisfy $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$, using a similar technique:

LEMMA IV.13 (BINARY REALIZATIONS). *For any balanced pssinstance P of an instance I which satisfies Σ_{UFD}, we can construct a realization of P that satisfies $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$.*

PROOF. As in Proposition IV.3, for every relation R, construct a bijection f_R between $\mathsf{Wants}(P,R^1)$ and $\mathsf{Wants}(P,R^2)$: this is possible, as P is balanced. We then construct our realization I' as in Proposition IV.3: we add to I the fact $R(a, f_R(a))$ for every R of σ and every $a \in \mathsf{Wants}(P,R^1)$.

We prove that I' is a realization as in Proposition IV.3 by observing that whenever we create a fact $R(a, f_R(a))$, then we have $a \in \mathsf{Wants}(P,R^1)$ and $f_R(a) \in \mathsf{Wants}(P,R^2)$. Similarly, we show that $I' \models \Sigma_{\mathrm{UFD}}$ as in Proposition IV.3.

We now show that I' satisfies $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$. Assume to the contrary that there is an active fact $F = R(a_1, a_2)$ that witnesses the violation of a UID $\tau : R^p \subseteq S^q$, so that $a_p \in \mathsf{Wants}(I',R^p)$. If $a_p \in \mathrm{dom}(I)$, then the proof is exactly as for Proposition IV.3. Otherwise, if $a_p \in \mathcal{H}$, clearly by construction of f_R and I' we have $a_p \in \pi_{T^r}(I')$ iff $T^r \in \lambda(a_p)$. Hence, as $a_p \in \pi_{R^p}(I')$ and as τ witnesses by assumption reversible that $R^p \sim_{\mathrm{ID}} S^q$, we have $a_p \in \pi_{S^q}(I')$, contradicting the fact that $a_p \in \mathsf{Wants}(I',S^q)$. □

We now conclude the proof of Theorem IV.1. Given the instance I_0, construct a balanced pssinstance P with the Balancing Lemma, construct a realization I' of P that satisfies $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$ with the Binary Realizations Lemma, and conclude by the "Binary Realizations are Completions" Lemma that I' is a weakly-sound superinstance of I_0.

V. WEAK SOUNDNESS ON ARBITRARY ARITY SIGNATURES

We now lift assumption binary and extend the results to arbitrary arity signatures:

THEOREM V.1. *Reversible UIDs and UFDs have finite weakly-universal models for ACQs.*

A first complication when lifting assumption binary is that realizations cannot be created just by pairing two elements. To satisfy the UIDs we may have to create facts that connect elements on more than two positions, so we may need more than the bijections between two positions that we used before. A much more serious problem is that the positions where we connect together elements may still be only a subset of the positions of the relation, which means that the other positions must be filled somehow.

We address these difficulties by defining first *piecewise realizations*, which create partial facts on positions connected by UFDs, similarly to the previous section. We show that we can get piecewise realizations by generalizing the Binary Realizations Lemma. Second, to find elements to reuse at other positions, we define a notion of *saturation*. We show that, by an initial *saturation process*, we can ensure that there are existing elements that we can reuse at positions where this will not violate UFDs (the *non-dangerous positions*). Third, we define a notion of *thrifty chase step* to solve UID violations one by one. We last explain how to use thrifty chase steps to solve all UID violations on saturated instances, using a piecewise realization as a template; this is how we construct our weakly-sound completion.

V.1. Piecewise Realizations

Without assumption binary, we must define a new equivalence relation to reflect the UFDs, in addition to $\sim_{\mathrm{ID}}$ which reflects the UIDs:

Definition V.2. For any two positions R^p and R^q, we write $R^p \leftrightarrow_{\mathrm{FUN}} R^q$ whenever $R^p = R^q$ or $R^p \rightarrow R^q$ and $R^q \rightarrow R^p$ are both in Σ_{UFD}.

By transitivity of Σ_{UFD}, $\leftrightarrow_{\mathrm{FUN}}$ is indeed an equivalence relation.

The definition of *balanced instances* (Definition IV.2) generalizes as-is to arbitrary arity. We do not change the definition of *pssinstance* (Definition IV.7), and talk of them being balanced in the same way. Further, we know that the Balancing Lemma (Lemma IV.12) holds even without assumption binary.

Our general scheme is the same: construct a balanced pssinstance of I_0, and use it to construct the completion I. What we need is to change the notion of *realization*. We replace it by *piecewise realizations*, which are defined on $\leftrightarrow_{\text{FUN}}$-classes. We number the $\leftrightarrow_{\text{FUN}}$-classes of $\mathsf{Pos}(\sigma)$ as $\Pi_1,\ldots,\Pi_n$ and define *piecewise instances* by their projections to the Π_i:

Definition V.3. A **piecewise instance** is an n-tuple $PI = (K_1,\ldots,K_n)$, where each K_i is a set of $|\Pi_i|$-tuples, indexed by Π_i for convenience. The **domain** of PI is $\text{dom}(PI) := \bigcup_i \text{dom}(K_i)$. For $1 \le i \le n$ and $R^p \in \Pi_i$, we define $\pi_{R^p}(PI) := \pi_{R^p}(K_i)$.

We will realize a pssinstance P, not as an instance as in the previous section, but as a piecewise instance. The tuples in each K_i will be defined from P, and will connect elements that want to occur at the corresponding position in Π_i, generalizing the ordered pairs constructed with bijections in the proof of the Binary Realizations Lemma. Let us define accordingly the notion of a *piecewise realization* of a pssinstance as a piecewise instance:

Definition V.4. A piecewise instance $PI = (K_1,\ldots,K_n)$ is a **piecewise realization** of the pssinstance $P = (I,\mathcal{H},\lambda)$ if:

— $\pi_{\Pi_i}(I) \subseteq K_i$ for all $1 \le i \le n$,
— $\text{dom}(PI) = \text{dom}(I) \sqcup \mathcal{H}$,
— for all $1 \le i \le n$, for all $R^p \in \Pi_i$, for every tuple $\boldsymbol{a} \in K_i \backslash \pi_{\Pi_i}(I)$, we have $a_p \in \mathsf{Wants}(P,R^p)$.

Notice that the definition is similar to the conditions imposed on realizations (Definition IV.8), although piecewise realizations are piecewise instances, not actual instances; so we will need one extra step to make real instances out of them: this is done in Section V.4.

We must now generalize the Binary Realizations Lemma (Lemma IV.13) to construct these piecewise realizations out of balanced pssinstances. For this, we need to define what it means for a piecewise instance PI to "satisfy" $\Sigma_{\text{U}}^{\text{rev}}$. For Σ_{UFD}, we require that PI respects the UFDs within each $\leftrightarrow_{\text{FUN}}$-class. For $\Sigma_{\text{UID}}^{\text{rev}}$, we define it directly from the projections of PI.

Definition V.5. A piecewise instance PI is Σ_{UFD}**-compliant** if, for all $1 \le i \le n$, there are no two tuples $\boldsymbol{a} \neq \boldsymbol{b}$ in K_i such that $a_p = b_p$ for some $R^p \in \Pi_i$.
PI is $\Sigma_{\text{UID}}^{\text{rev}}$**-compliant** if $\mathsf{Wants}(PI,\tau) := \pi_{R^p}(PI) \backslash \pi_{S^q}(PI)$ is empty for all $\tau : R^p \subseteq S^q$ in $\Sigma_{\text{UID}}^{\text{rev}}$.
PI is $\Sigma_{\text{U}}^{\text{rev}}$**-compliant** if it is Σ_{UFD}- and $\Sigma_{\text{UID}}^{\text{rev}}$-compliant.

We can then state and prove the generalization of the Binary Realizations Lemma:

LEMMA V.6 (REALIZATIONS). *For any balanced pssinstance P of an instance I that satisfies Σ_{UFD}, we can construct a piecewise realization of P which is $\Sigma_{\text{U}}^{\text{rev}}$-compliant.*

Before we prove the Realizations Lemma, we show a simple example:

Example V.7. Consider a 4-ary relation R and the UIDs $\tau : R^1 \subseteq R^2$, $\tau' : R^3 \subseteq R^4$ and their reverses, and the UFDs $\phi : R^1 \to R^2$, $\phi' : R^3 \to R^4$ and their reverses. We have $\Pi_1 = \{R^1,R^2\}$ and $\Pi_2 = \{R^3,R^4\}$. Consider $I_0 := \{R(a,b,c,d)\}$, which is balanced, and the trivial balanced pssinstance $P := (I_0,\emptyset,\lambda)$, where λ is the empty function. A $\Sigma_{\text{U}}^{\text{rev}}$-compliant piecewise realization of P is $PI := (\{(a,b),(b,a)\},\{(c,d),(d,c)\})$.

We conclude the subsection with the proof of the Realizations Lemma:

PROOF OF LEMMA V.6. Let $P = (I,\mathcal{H},\lambda)$ be the balanced pssinstance. Recall that the $\leftrightarrow_{\text{FUN}}$-classes of σ are numbered $\Pi_1,\ldots,\Pi_n$. By definition of P being balanced (Definition IV.2 applied to arbitrary arity), for any $\leftrightarrow_{\text{FUN}}$-class Π_i, for any two positions $R^p, R^q \in \Pi_i$, we have $|\mathsf{Wants}(P,R^p)| = |\mathsf{Wants}(P,R^q)|$. Hence, for all $1 \le i \le n$, we can define s_i as the value of $|\mathsf{Wants}(P,R^p)|$ for any $R^p \in \Pi_i$.

For $1 \leq i \leq n$, we let m_i be $|\Pi_i|$, and number the positions of Π_i as $R^{p^i_1}, \ldots, R^{p^i_{m_i}}$. We choose for each $1 \leq i \leq n$ and $1 \leq j \leq m_i$ an arbitrary bijection ϕ^i_j from $\{1, \ldots, s_i\}$ to $\mathsf{Wants}(P, R^{p^i_j})$. We construct the piecewise realization $PI = (K_1, \ldots, K_n)$ by setting each K_i for $1 \leq i \leq n$ to be $\pi_{\Pi_i}(I)$ plus the tuples $(\phi^i_1(l), \ldots, \phi^i_{m_i}(l))$ for $1 \leq l \leq s_i$.

It is clear that PI is a piecewise realization. Indeed, the first two conditions are immediate. Further, whenever we create a tuple $\boldsymbol{a} \in \Pi_i$ for any $1 \leq i \leq n$, then, for any $R^p \in \Pi_i$, we have $a_p \in \mathsf{Wants}(P, R^p)$.

Let us then show that PI is Σ_{UFD}-compliant. Assume by contradiction that there is $1 \leq i \leq n$ and $\boldsymbol{a}, \boldsymbol{b} \in K_i$ such that $a_l = b_l$ but $a_r \neq b_r$ for some $R^l, R^r \in \Pi_i$. As I satisfies Σ_{UFD}, we assume without loss of generality that $\boldsymbol{a} \in K_i \backslash \pi_{\Pi_i}(I)$. Now either $\boldsymbol{b} \in \pi_{\Pi_i}(I)$ or $\boldsymbol{b} \in K_i \backslash \pi_{\Pi_i}(I)$.

- If $\boldsymbol{b} \in \pi_{\Pi_i}(I)$, then $b_l \in \pi_{R^l}(I)$. Yet, we know by construction that, as $\boldsymbol{a} \in K_i \backslash \pi_{\Pi_i}(I)$, we have $a_l \in \mathsf{Wants}(P, R^l)$, so that by definition of $\mathsf{Wants}(P, R^l)$ we have $a_l \in \mathsf{Wants}(I, R^l)$. But we have $a_l = b_l$, so we have a contradiction.
- If $\boldsymbol{b} \in K_i \backslash \pi_{\Pi_i}(I)$, then, writing $R^l = R^{p^i_j}$ and $R^r = R^{p^i_{j'}}$, the fact that $a_l = b_l$ but $a_r \neq b_r$ contradicts the fact that $\phi^i_j \circ (\phi^i_{j'})^{-1}$ is injective.

Hence, PI is Σ_{UFD}-compliant.

Let us now show that PI is $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$-compliant. We must show that, for every UID $\tau : R^p \subseteq S^q$ of $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, we have $\mathsf{Wants}(PI, \tau) = \emptyset$, which means that we have $\pi_{R^p}(PI) \subseteq \pi_{S^q}(PI)$. Let Π_i be the $\leftrightarrow_{\mathrm{FUN}}$-class of R^p, and assume to the contrary the existence of a tuple $\boldsymbol{a}$ of K_i such that $a_p \notin \pi_{S^q}(PI)$. Either we have $a_p \in \mathrm{dom}(I)$, or we have $a_p \in \mathcal{H}$.

- If $a_p \in \mathrm{dom}(I)$, as $a_p \notin \pi_{S^q}(PI)$, in particular $a_p \notin \pi_{S^q}(I)$, and as $a_p \in \pi_{R^p}(I)$, τ witnesses that $a_p \in \mathsf{Wants}(I, S^q)$. By construction of PI, then, letting $\Pi_{i'}$ be the $\leftrightarrow_{\mathrm{FUN}}$-class of S^q and letting $S^q = S^{p^{i'}_j}$, as $\phi^{i'}_j$ is surjective, we must have $a_p \in \pi_{S^q}(K_{i'})$, that is, $a_p \in \pi_{S^q}(PI)$, a contradiction.
- If $a_p \in \mathcal{H}$, clearly by construction we have $a_p \in \pi_{T^r}(PI)$ iff $T^r \in \lambda(a_p)$, so that, given that τ witnesses $R^p \sim_{\mathsf{ID}} S^q$, if $a_p \in \pi_{R^p}(PI)$ then $a_p \in \pi_{S^q}(PI)$, a contradiction.

We conclude that PI is indeed a $\Sigma^{\mathrm{rev}}_{\mathrm{U}}$-compliant piecewise realization of P. □

V.2. Relation-Saturation

The Realizations Lemma gives us a $\Sigma^{\mathrm{rev}}_{\mathrm{U}}$-compliant piecewise realization which is a piecewise instance. To construct an actual superinstance from it, we will have to expand each tuple $\boldsymbol{t}$ of each K_i, defined on the $\leftrightarrow_{\mathrm{FUN}}$-class Π_i, to an entire fact $F_{\boldsymbol{t}}$ of the corresponding relation.

However, to fill the other positions of $F_{\boldsymbol{t}}$, we will need to reuse existing elements of I_0. To do this, it is easier to assume that I_0 contains some R-fact for every relation R of the signature.

Definition V.8. A superinstance I of I_0 is **relation-saturated** if for every $R \in \sigma$ there is an R-fact in I.

We illustrate why it is easier to work with relation-saturated instances:

Example V.9. Suppose our schema has two binary relations R and T and a unary relation S, the UIDs $\tau : S^1 \subseteq R^1$, $\tau' : R^2 \subseteq T^1$ and their reverses, and no UFDs. Consider the non-relation-saturated instance $I_0 := \{S(a)\}$. It is balanced, so $P := (I_0, \emptyset, \lambda)$, with λ the empty function, is a pssinstance of I.

Now, a $\Sigma^{\mathrm{rev}}_{\mathrm{U}}$-compliant piecewise realization of P is $PI = (K_1, \ldots, K_5)$ with $K_2 = K_4 = K_5 = \emptyset$ and $K_1 = K_3 = \{a\}$, where Π_1 and Π_3 are the $\leftrightarrow_{\mathrm{FUN}}$-classes of R^1 and S^1. However, we cannot easily complete PI to an actual superinstance of I_0 satisfying τ and τ'. Indeed, to create the fact $R(a, \bullet)$, as indicated by K_1, we need to fill position R^2. Using an existing element would violate weak-soundness, and using a fresh element would introduce a violation of τ', which P and PI would not tell us how to solve.

Consider instead the relation-saturated instance $I_1 := I_0 \sqcup \{S(c), R(c,d), T(d,e)\}$. We can complete I_1 to a weakly-sound superinstance that satisfies τ and τ', by adding the fact $R(a,d)$. Observe how we reused d to fill position R^2: this does not violate weak-soundness or introduce new UID violations.

Relation-saturation can clearly be ensured by initial chasing, which does not violate weak-soundness. We call this a *saturation process* to ensure relation-saturation:

LEMMA V.10 (RELATION-SATURATED SOLUTIONS). *For any reversible UIDs $\Sigma_{\text{UID}}^{\text{rev}}$, UFDs Σ_{UFD}, and instance I_0 satisfying Σ_{UFD}, the result of performing sufficiently many chase rounds on I_0 by $\Sigma_{\text{UID}}^{\text{rev}}$ is a weakly-sound relation-saturated superinstance of I_0 that satisfies Σ_{UFD}.*

This allows us to assume that I_0 was preprocessed with initial chasing if needed, so we can assume it to be relation-saturated. To show the lemma, and also for further use, we make a simple observation on weak-soundness:

LEMMA V.11 (WEAK-SOUNDNESS TRANSITIVITY). *If I' is a weakly-sound superinstance of I, and I is a weakly-sound superinstance of I_0, then I' is a weakly-sound superinstance of I_0.*

PROOF. Let $a \in \text{dom}(I')$, and let us show that it does not witness a violation of the weak-soundness of I' for I_0. We distinguish three cases:

— If $a \in \text{dom}(I_0)$, then in particular $a \in \text{dom}(I)$. Hence, letting S^q be any position such that $a \in \pi_{S^q}(I')$, as I' is a weakly-sound superinstance of I, either $a \in \pi_{S^q}(I)$ or we have $a \in \mathsf{Wants}(I, S^q)$. Let R^p be a position such that $a \in \pi_{R^p}(I)$, and such that $R^p = S^q$ (in the first case) or $R^p \subseteq S^q$ holds in $\Sigma_{\text{UID}}^{\text{rev}}$ (in the second case). As I is a weakly-sound superinstance of I_0, either $a \in \pi_{R^p}(I_0)$ or $a \in \mathsf{Wants}(I_0, R^p)$. As $\Sigma_{\text{UID}}^{\text{rev}}$ is transitively closed, we conclude that $a \in \mathsf{Wants}(I_0, S^q)$ or $a \in \pi_{S^q}(I_0)$. Hence, the fact that a occurs at position S^q in I' does not cause a violation of weak-soundness in I' for I_0.
— If $a \in \text{dom}(I) \setminus \text{dom}(I_0)$, we must show that for any two positions R^p, S^q where a occurs in I', we have $R^p \sim_{\mathsf{ID}} S^q$. Let us fix two such positions, i.e., we have $a \in \pi_{R^p}(I')$ and $a \in \pi_{S^q}(I')$. As I' is a weakly-sound superinstance of I, we have either $a \in \pi_{R^p}(I)$ or $a \in \mathsf{Wants}(I, R^p)$, and we have either $a \in \pi_{S^q}(I)$ or $a \in \mathsf{Wants}(I, S^q)$. As in the previous case, let T^v and U^w be positions such that $a \in \pi_{T^v}(I)$ and $a \in \pi_{U^w}(I)$, and $T^v = R^p$ or the UID $\tau : T^v \subseteq R^p$ holds in $\Sigma_{\text{UID}}^{\text{rev}}$, and $U^w = S^q$ or the UID $\tau' : U^w \subseteq S^q$ holds in $\Sigma_{\text{UID}}^{\text{rev}}$. As I is a weakly-sound superinstance of I_0, and $a \notin \text{dom}(I_0)$, we know that $T^v \sim_{\mathsf{ID}} U^w$. By assumption reversible and as $\Sigma_{\text{UID}}^{\text{rev}}$ is transitively closed, we deduce (using τ and τ' if necessary) that $R^p \sim_{\mathsf{ID}} S^q$, which is what we wanted to show. Hence, the fact that a occurs at positions R^p and S^q in I' does not cause a violation of weak-soundness in I' for I_0.
— If $a \in \text{dom}(I') \setminus \text{dom}(I)$, then from the fact that I' is a weakly-sound superinstance of I, we deduce immediately about a what is needed to show that it does not witness a violation of the weak-soundness of I' for I_0.

So we conclude that I' is a weakly-sound instance of I_0, as desired. □

We conclude the subsection by proving the Relation-Saturated Solutions Lemma:

PROOF OF LEMMA V.10. Remember that the signature σ was assumed without loss of generality not to contain any useless relation. Hence, for every relation $R \in \sigma$, there is an R-fact in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}}^{\text{rev}})$, which was generated at the n_R-th round of the chase, for some $n_R \in \mathbb{N}$. Let $n := \max_{R \in \sigma} n_R$, which is finite because the number of relations in σ is finite. We take I to be the result of applying n chase rounds to I_0.

It is clear that I is relation-saturated. The fact that I is weakly-sound is by the Weak-Soundness Transitivity Lemma, because each chase step clearly preserves weak-soundness: the exported element occurs at a position where it wants to occur, so we can use assumption reversible, and new elements only occur at one position. □

V.3. Thrifty Chase Steps

We have explained why I_0 can be assumed to be relation-saturated, and we know we can build a $\Sigma_{\mathrm{U}}^{\mathrm{rev}}$-compliant piecewise realization PI of a balanced pssinstance. Our goal is now to satisfy the UIDs using PI. We will do so by a *completion process* that fixes each violation one by one, following PI. This subsection presents the tool that we use for this, and the next subsection describes the actual process.

Our tool is a form of chase step, a *thrifty chase step*, which adds a new fact F_{n} to satisfy a UID violation. For some of the positions, the elements of F_{n} will be defined from the realization PI, using one of its tuples. For each of these elements, either F_{n} makes them occur at a position that they want to be (thus satisfying another violation) or these elements are helpers that did not occur already in the domain. At any other position S^r of F_{n}, we may either reuse an existing element (by relation saturation, one can always reuse an element that already occurs in that position) or create a fresh element (arguing that no UID will be violated on that element). This depends on whether S^r is *dangerous* or *non-dangerous*:

Definition V.12. We say a position $S^r \in \mathsf{Pos}(\sigma)$ is **dangerous** for the position $S^q \neq S^r$ if $S^r \rightarrow S^q$ is in Σ_{UFD}, and write $S^r \in \mathsf{Dng}(S^q)$. Otherwise, still assuming $S^q \neq S^r$ S^r is **non-dangerous** for S^q, written $S^r \in \mathsf{NDng}(S^q)$. Note that $\{S^q\} \sqcup \mathsf{Dng}(S^q) \sqcup \mathsf{NDng}(S^q) = \mathsf{Pos}(S)$.

We can now define *thrifty chase steps*. The details of the definition are designed for the completion process defined in the next subsection (Proposition V.17), and for the specialized notions that we will introduce later in this subsection as well as in the following sections.

Definition V.13. Let I be a superinstance of I_0, let $\tau : R^p \subseteq S^q$ be a UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, and let $F_{\mathrm{a}} = R(\boldsymbol{a})$ be an active fact for τ in I. We call S^q the **exported position**, and write Π_i for its $\leftrightarrow_{\mathrm{FUN}}$-class.

Applying a **thrifty chase step** to F_{a} (or a) in I by τ yields a superinstance I' of I_0 which is I plus a single new fact $F_{\mathrm{n}} = S(\boldsymbol{b})$. We require the following on b_r for all $S^r \in \mathsf{Pos}(S)$:

— For $S^r = S^q$, we require $b_q = a_p$ and $b_q \in \mathsf{Wants}(I, \tau)$;
— For $S^r \in \Pi_i \backslash \{S^q\}$, we require that one of the following holds:
  — $b_r \in \mathsf{Wants}(I, S^r)$;
  — $b_r \notin \mathrm{dom}(I)$ and for all $S^s \in \Pi_i$, such that $b_r = b_s$, we have $S^r \sim_{\mathrm{ID}} S^s$;
— For $S^r \in \mathsf{Dng}(S^q) \backslash \Pi_i$, we require b_r to be fresh and occur only at that position;
— For $S^r \in \mathsf{NDng}(S^q)$, we require that $b_r \in \pi_{S^r}(I)$.

Thrifty chase steps eliminate UID violations on the element at the exported position S^q of the new fact (which is why we call them "chase steps"), and also eliminate violations on positions in the same $\leftrightarrow_{\mathrm{FUN}}$-class as S^q, unless a fresh element is used there. The completion process that we will define in the next subsection will *only* apply thrifty chase steps (namely, *relation-thrifty steps*, which we will define shortly), and indeed this will be true of *all* completion processes used in this paper.

For now, we can observe that thrifty chase steps cannot break weak-soundness:

LEMMA V.14 (THRIFTY PRESERVES WEAK-SOUNDNESS). *For any weakly-sound superinstance I of an instance I_0, letting I' be the result of applying a thrifty chase step on I, I' is a weakly-sound superinstance of I_0.*

PROOF. By the Weak-Soundness Transitivity Lemma, it suffices to show that I' is a weakly-sound superinstance of I. It suffices to check this for the elements occurring in the one fact $F_{\mathrm{n}} = S(\boldsymbol{b})$ of $I' \backslash I$, as the other elements occur at the same positions as before. Let us show for each b_r for $S^r \in \mathsf{Pos}(S)$ that b_r does not cause a violation of weak-soundness:

— For $S^r = S^q$, we have $b_r \in \mathsf{Wants}(I, S^r)$, so b_r does not violate weak-soundness;
— For $S^r \in \Pi_i \backslash \{S^q\}$, there are two possible cases:

  — $b_r \in \mathsf{Wants}(I, S^r)$, so b_r does not violate weak-soundness;
  — $b_r \notin \mathrm{dom}(I)$ and b_r occurs only at positions related by $\sim_{\mathrm{ID}}$, so b_r does not violate weak-soundness;

— For $S^r \in \mathsf{Dng}(S^q) \backslash \Pi_i$, b_r is fresh and occurs at a single position in I', so b_r does not violate weak-soundness;
— For $S^r \in \mathsf{NDng}(S^q)$, as $b_r \in \pi_{S^r}(I)$, b_r does not violate weak-soundness. □

Thrifty chase steps may introduce UFD violations. For this reason, we introduce the special case of *relation-thrifty* chase steps, which can not introduce such violations. (Relation-thrifty chase steps may still introduce FD violations; we will deal with this in Section VIII.)

Definition V.15 (*Relation-thrifty*). A **relation-thrifty chase step** is a thrifty chase step where we choose one fact $F_{\mathrm{r}} = S(\boldsymbol{c})$ of I, and use F_{r} to define $b_r := c_r$ for all $S^r \in \mathsf{NDng}(S^q)$.

Remember that relation-saturation ensures that such a fact $S(\boldsymbol{c})$ can always be found, so clearly any UID violation can be solved on a relation-saturated instance by applying some relation-thrifty chase step. Further, we can show that relation-thrifty chase steps, unlike thrifty chase steps, preserve UFDs:

LEMMA V.16 (RELATION-THRIFTY PRESERVATION). *For any superinstance I of an instance I_0 such that I satisfies Σ_{UFD}, letting I' be the result of applying a relation-thrifty chase step on I, then I' satisfies Σ_{UFD}. Further, if I is relation-saturated, then I' is relation-saturated.*

PROOF. Assume to the contrary the existence of two facts $F = S(\boldsymbol{a})$ and $F' = S(\boldsymbol{b})$ in I' that witness a violation of some UFD $\phi : S^r \to S^p$ of Σ_{UFD}. As $I \models \Sigma_{\mathrm{UFD}}$, we may assume without loss of generality that F' is $F_{\mathrm{n}} = S(\boldsymbol{b})$, the unique fact of $I' \backslash I$. Write $\tau : R^p \subseteq S^q$ the UID of $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ applied in the relation-thrifty chase step.

We first note that we must have S^r in $\mathsf{NDng}(S^q)$. Indeed, assuming to the contrary that $S^r = S^q$ or $S^r \in \mathsf{Dng}(S^q)$, the definition of thrifty chase steps requires that either $b_r \notin \mathrm{dom}(I)$ or $b_r \in \mathsf{Wants}(I, S^r)$, so that in either case $b_r \notin \pi_{S^r}(I)$. Yet, as $a_r = b_r$, F witnesses that $b_r \in \pi_{S^r}(I)$, a contradiction. Thus, $S^r \in \mathsf{NDng}(S^q)$.

Now, because ϕ holds in Σ_{UFD} and Σ_{UFD} is closed under the transitivity rule, unwinding the definitions we can see that $S^p \in \mathsf{NDng}(S^q)$ as well. Now, let $F_{\mathrm{r}} = S(\boldsymbol{c})$ be the chosen fact for the relation-thrifty chase step. Observe that we must have $F \neq F_{\mathrm{r}}$: this follows because we have $\pi_{S^r}(F_{\mathrm{r}}) = c_q = b_q$ but $\pi_{S^r}(F) = a_q$ and $a_q \neq b_q$ by definition of a UFD violation. Hence, as $b_q = c_q$ and $b_r = c_r$ by definition of F_{n} from F_{r}, as $F \neq F_{\mathrm{r}}$, F and F_{r} are also a violation of ϕ, which is in I, contradicting that $I \models \Sigma_{\mathrm{UFD}}$.

The second part of the claim is immediate. □

To summarize: we have defined the general tool used in our completion process, *thrifty chase steps*, along with a special case that preserves UFDs, *relation-thrifty chase steps*, which applies to relation-saturated instances. We now move to the last part of this section, where we use this tool to satisfy UID violations, also using the tools previously defined in this section.

V.4. Relation-Thrifty Completions

To prove Theorem V.1, let us start by taking our initial finite instance I_0, which satisfies Σ_{UFD}, and use the Relation-Saturated Solutions Lemma to obtain a finite weakly-sound superinstance I'_0 which is relation-saturated and still satisfies Σ_{UFD}. We now obtain our weakly-sound superinstance from I'_0 by performing a *completion process* by relation-thrifty chase steps, which we phrase as follows:

PROPOSITION V.17 (REVERSIBLE RELATION-THRIFTY COMPLETION). *For any reversible Σ_{UFD} and $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, for any finite relation-saturated instance I'_0 that satisfies Σ_{UFD}, we can use relation-thrifty chase steps to construct a finite weakly-sound superinstance I_{f} of I'_0 that satisfies $\Sigma^{\mathrm{rev}}_{\mathrm{U}} = \Sigma^{\mathrm{rev}}_{\mathrm{UID}} \cup \Sigma_{\mathrm{UFD}}$.*

Indeed, once this result is proven, we can immediately conclude the proof of Theorem V.1 with it, by applying it to I'_0 and obtaining I_f which is a weakly-sound superinstance of I'_0, hence of I_0 by the Weak-Soundness Transitivity Lemma. So we conclude the section with the proof of this proposition.

Recall that we number $\Pi_1, \ldots, \Pi_n$ the $\leftrightarrow_{\text{FUN}}$-classes of $\mathsf{Pos}(\sigma)$. Let us write $\Pi_i \to \Pi_j$ to mean that all corresponding UFDs hold in Σ_{UFD} for positions in Π_i and Π_j. That is, equivalently, if one of them does (by definition of a $\leftrightarrow_{\text{FUN}}$-class and the fact that Σ_{UFD} is transitively closed). We first define the *inner* classes, where creating elements may cause UID violations, and the *outer* classes, where this cannot happen because no position of the class occurs in any UID:

Definition V.18. We say that Π_j is an **inner** $\leftrightarrow_{\text{FUN}}$-class if it contains a position occurring in $\Sigma^{\text{rev}}_{\text{UID}}$; otherwise, it is an **outer** $\leftrightarrow_{\text{FUN}}$-class.

The fundamental property is:

LEMMA V.19. *For any $1 \leq i, j \leq n$ with $i \neq j$, if Π_i is inner and $\Pi_j \to \Pi_i$ then Π_j is outer.*

PROOF. Assume to the contrary that Π_j is inner. This means that it contains a position R^q that occurs in $\Sigma^{\text{rev}}_{\text{UID}}$. As Π_i is inner, pick any $R^p \in \Pi_i$ that occurs in $\Sigma^{\text{rev}}_{\text{UID}}$. As $\Pi_j \to \Pi_i$, $\phi : R^q \to R^p$ holds in Σ_{UFD}. Hence, by assumption reversible, ϕ^{-1} also does. But then we have $R^p \leftrightarrow_{\text{FUN}} R^q$, contradicting the maximality of $\leftrightarrow_{\text{FUN}}$-classes Π_i and Π_j. □

Let us now start the actual proof of Proposition V.17, and fix the finite relation-saturated instance I'_0 that satisfies Σ_{UFD}. We start by constructing a balanced pssinstance P of I'_0 using the Balancing Lemma (Lemma IV.12), and a finite $\Sigma^{\text{rev}}_{\text{U}}$-compliant piecewise realization $PI = (K_1, \ldots, K_n)$ of P by the Realizations Lemma (Lemma V.6). Let $\mathcal{F}$ be an infinite set of fresh elements (not in $\text{dom}(P)$) from which we will take the (finitely many) fresh elements that we will introduce (only at dangerous positions, in outer classes) during the relation-thrifty chase steps.

We will use PI to construct a weakly-sound superinstance I_f by relation-thrifty chase steps. We maintain the following invariant when doing so:

Definition V.20. A superinstance I of the instance I'_0 **follows** the piecewise realization $PI = (K_1, \ldots, K_n)$ if for every inner $\leftrightarrow_{\text{FUN}}$-class Π_i, we have $\pi_{\Pi_i}(I) \subseteq K_i$.

We prove the Reversible Relation-Thrifty Completion Proposition by satisfying UID violations in I'_0 with relation-thrifty chase steps using the piecewise realization PI. We call I the current state of our superinstance, starting at $I := I'_0$, and we perform relation-thrifty chase steps on I to satisfy UID violations, until we reach a finite weakly-sound superinstance I_f of I'_0 such that I_f satisfies $\Sigma^{\text{rev}}_{\text{UID}}$ and I_f follows PI. This I_f will be the final result of the Reversible Relation-Thrifty Completion Proposition.

Chasing by relation-thrifty chase steps preserves the following invariants:

sub:. $I'_0 \subseteq I$ (this is clearly monotone);
wsnd:. I is weakly-sound (by Lemma V.14);
fun:. $I \models \Sigma_{\text{UFD}}$ (by Lemma V.16);
rsat:. I is relation-saturated (by Lemma V.16).

Further, we maintain the following invariants:

fw:. I follows PI;
help:. For any position R^p of an outer class, $\pi_{R^p}(I)$ and $\mathcal{H}$ are disjoint.

Let us show that any UID violation in I at any stage of the construction can be solved by applying a relation-thrifty chase step that preserves these invariants. To show this, let $a \in \mathsf{Wants}(I, \tau)$ be an element to which some UID $\tau : R^p \subseteq S^q$ of $\Sigma^{\text{rev}}_{\text{UID}}$ is applicable. Let $F_a = R(\boldsymbol{a})$ be the active fact, with $a = a_p$. Let Π_i denote the $\leftrightarrow_{\text{FUN}}$-classes of R^p and S^q respectively. The UID τ witnesses that Π_i is inner, so by invariant fw we have $a \in \pi_{R^p}(PI)$. As PI is $\Sigma^{\text{rev}}_{\text{UID}}$-compliant, we must have $a \in \pi_{S^q}(PI)$, and there is a $|\Pi_i|$-tuple $\boldsymbol{t} \in K_i$ such that $t_q = a$; in fact, by Σ_{UFD}-compliance, there is exactly one such tuple.

Let $F_\mathrm{r} = S(\boldsymbol{c})$ be an S-fact of I'_0, which is possible by invariant rsat. We create a new fact $F_\mathrm{n} = S(\boldsymbol{b})$ with the relation-thrifty chase step defined as follows:

— For the exported position S^q, we set $b_q := a_p$.
— For any $S^r \in \Pi_i$, we set $b_r := t_r$.
— For any position $S^r \in \mathsf{Dng}(S^q) \backslash \Pi_i$, we take b_r to be a fresh element f_r from $\mathcal{F}$.
— For any position $S^r \in \mathsf{NDng}(S^q)$, we set $b_r := c_r$.

We first verify that this satisfies the conditions of thrifty chase steps. We have set $b_q = a$, and by definition of F_r it is immediate that $b_r \in \pi_{S^r}(I)$ for $S^r \in \mathsf{NDng}(S^q)$. For $S^r \in \mathsf{Dng}(S^q) \backslash \Pi_i$, we use a fresh element f_r from $\mathcal{F}$ which occurs only at position S^r, as we should.

The last case to check is for $S^r \in \Pi_i \backslash \{S^q\}$. The first case is if $b_r \notin \mathrm{dom}(I)$, in which case we must show that all positions at which b_r occurs are $\sim_{\mathsf{ID}}$-equivalent. Assume that b_r occurs at some other position $S^s \in \Pi_i$. Now as b_r is in $\pi_{S^s}(PI)$, by definition of PI being a piecewise realization of P, we have $b_r \in \mathsf{Wants}(P, S^s)$. Now, as $b_r \notin \mathrm{dom}(I)$, by invariant sub we also have $b_r \notin \mathrm{dom}(I'_0)$. But as $b_r \in \mathrm{dom}(PI)$, we must have $b_r \in \mathcal{H}$. So by definition of a pssinstance we have $S^s \in \lambda(b_r)$. As $b_r \in \mathsf{Wants}(P, S^r)$ also, we have $S^r \in \lambda(b_r)$. By definition of $\lambda(b_r)$ being an $\sim_{\mathsf{ID}}$-class, this means that $S^r \sim_{\mathsf{ID}} S^s$, as required.

The second case is $b_r \in \mathrm{dom}(I)$. We will show that we have $b_r \in \mathsf{Wants}(I, S^r)$. Observe first that $b_r \notin \pi_{S^r}(I)$. Indeed, assuming to the contrary that $b_r \in \pi_{S^r}(I)$, let $F = S(\boldsymbol{d})$ be a witnessing fact in I. As Π_i is inner, by invariant fw, we deduce that $\pi_{\Pi_i}(\boldsymbol{d}) \in \pi_{\Pi_i}(PI)$. Now, as $d_r = t_r$ and PI is Σ_{UFD}-compliant, we deduce that $\boldsymbol{d} = \boldsymbol{t}$, so that F witnesses that d_q is in $\pi_{S^q}(I)$. As we have $d_q = t_q = a$, this contradicts the applicability of τ to a. Hence, we have $b_r \notin \pi_{S^r}(I)$.

Second, observe that we have $t_r \in \mathsf{Wants}(P, S^r)$. Indeed, we have $b_r = t_r$ which is in $\pi_{S^r}(PI)$, and we cannot have $\boldsymbol{t} \in \pi_{\Pi_i}(I)$, as otherwise this would contradict the applicability of τ to a, as we showed; so in particular, by invariant sub, we cannot have $\boldsymbol{t} \in \pi_{\Pi_i}(I'_0)$. Thus, by definition of a piecewise realization, we have $t_r \in \mathsf{Wants}(P, S^r)$.

Now, as $t_r \in \mathsf{Wants}(P, S^r)$, by definition of $\mathsf{Wants}(P, S^r)$, there are two cases:

— We have $t_r \in \mathrm{dom}(I'_0)$ and $t_r \in \mathsf{Wants}(I'_0, S^r)$. In this case, as we have shown that $t_r \notin \pi_{S^r}(I)$, we conclude immediately that $\mathrm{t}_r \in \mathsf{Wants}(I, S^r)$.
— We have $t_r \in \mathcal{H}$ and $S^r \in \lambda(t_r)$. In this case, consider a fact F' of I witnessing $t_r \in \mathrm{dom}(I)$, where t_r occurs at a position T^l; let $\Pi_{i'}$ be the $\leftrightarrow_{\mathrm{FUN}}$-class of T^l. As $t_r \in \mathcal{H}$, by invariant help, $\Pi_{i'}$ is inner, so by invariant fw there is a tuple $\boldsymbol{t}'$ of $K_{i'}$ such that $t'_l = t_r$. Now, as $t_r \in \mathcal{H}$, by definition of piecewise realizations, we have $T^l \in \lambda(t_r)$. Hence, either the UID $\tau' : T^l \subseteq S^r$ is in $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ or we have $T^l = S^r$. As $t_r \in \pi_{T^l}(I)$ and we have shown earlier that $t_r \notin \pi_{S^r}(I)$, we know that $T^l \neq S^r$, so τ' is in $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$. Hence, as F' witnesses that $t_r \in \pi_{T^l}(I)$, and as $t_r \notin \pi_{S^r}(I)$, we conclude that $t_r \in \mathsf{Wants}(I, S^r)$.

Hence, in either case we have $t_r \in \mathsf{Wants}(I, S^r)$, as claimed. This concludes the proof of the fact that we have indeed defined a thrifty chase step. Further, the step is clearly relation-thrifty by construction. The last thing to do is to check that invariants fw and help are preserved by the relation-thrifty chase step:

— For invariant fw, τ witnesses that the class Π_i of S^q is inner. Hence, for any $S^r \in \mathsf{Dng}(S^q) \backslash \Pi_i$, by Lemma V.19, the $\leftrightarrow_{\mathrm{FUN}}$-class of S^r is outer. Thus, to show that fw is preserved, it suffices to show it for the $\leftrightarrow_{\mathrm{FUN}}$-class Π_i and on the $\leftrightarrow_{\mathrm{FUN}}$-classes included in $\mathsf{NDng}(S^q)$ (clearly no $\leftrightarrow_{\mathrm{FUN}}$-class includes both a position of $\mathsf{Dng}(S^q)$ and a position of $\mathsf{NDng}(S^q)$). For Π_i, the new fact F_n is defined following $\boldsymbol{t}$; for the classes in $\mathsf{NDng}(S^q)$, it is defined following an existing fact of I. Hence, invariant fw is preserved.
— Invariant help is preserved because the only new elements of F_n that may be in $\mathcal{H}$ are those used at positions of Π_i, which is inner.

Let I_{f} be the result of the process that we have described. It satisfies Σ_{UID} by definition, and it is a finite weakly-sound superinstance of I'_0 that satisfies Σ_{UID}, by invariants wsnd, sub, and fun. Further, it follows *PI* by invariant fw, and *PI* is finite. This implies that I_{f} is finite, because we apply chase steps by $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, so each chase step makes an element of $\mathrm{dom}(PI)$ occur at a new position, so we only applied finitely many chase steps. This concludes the proof of the Reversible Relation-Thrifty Completion Proposition, and concludes the section.

VI. ENSURING k-UNIVERSALITY

We build on the constructions of the previous section to replace weak-soundness by k-soundness for acyclic queries in ACQ, for some $k > 0$ fixed in this section. That is, we aim to prove:

THEOREM VI.1. *Reversible UIDs and UFDs have finite k-universal models for ACQs.*

We first introduce the concept of *aligned superinstances*, which give us an invariant that ensures k-soundness. We then give the *fact-saturation* process that generalizes relation-saturation, and a related notion of *fact-thrifty chase step*. We then define *essentiality*, which must additionally be ensured for us to be able to reuse the weakly-sound completions of the previous section. We conclude by the construction of a generalized completion process that uses these chase steps to repair UID violations in the instance while preserving k-soundness.

In this section, we still make assumption reversible on $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ and Σ_{UFD}. However, we will also be considering a superset Σ_{UID} of $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, which we assume to be transitively closed, but which may not satisfy assumption reversible. To prove Theorem VI.1, it suffices to define $\Sigma_{\mathrm{UID}} := \Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, so *the distinction can be safely ignored on first reading*. The reason for the distinction will become apparent in the next section.

VI.1. Aligned Superinstances

In this subsection, we only work with the superset Σ_{UID}, and we do not use assumption reversible. We ensure k-soundness relative to Σ_{UID} by maintaining a *k-bounded simulation* from our superinstance of I_0 to the chase $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$.

Definition VI.2. For I, I' two instances, $a \in \mathrm{dom}(I)$, $b \in \mathrm{dom}(I')$, and $n \in \mathbb{N}$, we write $(I,a) \leq_n (I',b)$ if, for any fact $R(\boldsymbol{a})$ of I with $a_p = a$ for some $R^p \in \mathsf{Pos}(R)$, there exists a fact $R(\boldsymbol{b})$ of I' such that $b_p = b$, and $(I,a_q) \leq_{n-1} (I',b_q)$ for all $R^q \in \mathsf{Pos}(R)$ (note that this is tautological for $R^q = R^p$). The base case $(I,a) \leq_0 (I',b)$ always holds.

An ***n*-bounded simulation** from I to I' is a mapping sim such that for all $a \in \mathrm{dom}(I)$, we have $(I,a) \leq_n (I', \mathsf{sim}(a))$.

We write $a \simeq_n b$ for $a,b \in \mathrm{dom}(I)$ if both $(I,a) \leq_n (I,b)$ and $(I,b) \leq_n (I,a)$; this is an equivalence relation on $\mathrm{dom}(I)$.

Example VI.3. We illustrate in Figure 3 some examples of 2-bounded simulations from one instance to another, on a binary signature. For any element a in a left instance I and image a' of a in the right instance I' by the 2-bounded simulation (represented by the dashed red arrows), we have $(I,a) \leq_2 (I',a')$. This means that, for any element b in I which is adjacent to a by some relation R, there must be an element b' in I' which is adjacent to a' by R and satisfies $(I,b) \leq_1 (I',b')$; however, note that b' need not be the image of b by the bounded simulation.

Figure 3a illustrates how a homomorphism is a special case of a 2-bounded simulation (indeed, it is an n-bounded simulation for any $n \in \mathbb{N}$).

Figure 3b illustrates how a 2-bounded simulation from I to I' does not guarantee that any ACQ satisfied by I is also true in I': for this example, consider the query $\exists xyzuvw\ R(x,y) \wedge S(y,z) \wedge T(z,u) \wedge U(u,v) \wedge V(v,w)$. However, we will soon see that n-bounded simulations preserve ACQ of size $\leq n$ (Lemma VI.4).

Figure 3c shows that a 2-bounded simulation does not preserve CQs that are not ACQs, as witnessed by $\exists xyz\, R(x,y) \wedge S(y,z) \wedge T(z,x)$.

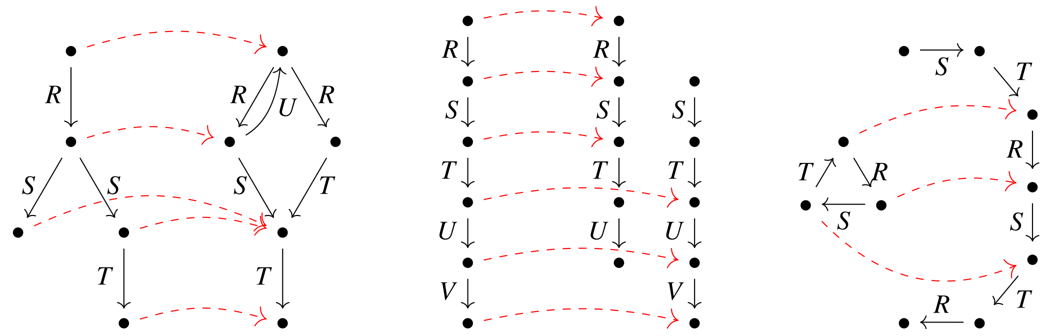

(a) Homomorphisms are bounded simulations

(b) Bounded simulations do not preserve large ACQs

(c) Bounded simulations do not preserve CQs

Fig. 3: Examples of 2-bounded simulations (represented as dashed red lines): see Example VI.3

The point of bounded simulations is that they preserve acyclic queries of size smaller than the bound:

LEMMA VI.4 (ACQ PRESERVATION). *For any instance I and ACQ q of size $\leq n$ such that $I \models q$, if there is an n-bounded simulation from I to I', then $I' \models q$.*

To show this lemma, we introduce a different way to write queries in ACQ. Consider the following alternate query language:

Definition VI.5. We inductively define a special kind of query with at most one free variable, a **pointed query**. The base case is that of a tautological query with no atoms. Inductively, pointed queries include all queries of the form:

$$q(x) : \bigwedge_i \left(\exists \boldsymbol{y^i} \left(A^i(x, \boldsymbol{y^i}) \wedge \bigwedge_{y^i_j \in \boldsymbol{y^i}} q^i_j(y^i_j) \right) \right)$$

where the $\boldsymbol{y^i}$ are vectors of pairwise distinct variables (also distinct from x), A^i are atoms with free variables as indicated and with no repeated variables (each free variable occurs at exactly one position), and the q^i_j are pointed queries.

The **size** $|q|$ of a pointed query q is the total number of atoms in q, including its subqueries.

It is easily seen that, for any pointed query q', the query $q : \exists x\, q'(x)$ is an ACQ. Conversely, we can show:

LEMMA VI.6. *For any (Boolean) ACQ q and variable x of q, we can rewrite q as $\exists x\, q'(x)$ with q' a pointed query such that $|q| = |q'|$.*

PROOF. We show the claim by induction on the size of q. It is clearly true for the empty query.

Otherwise, let $\mathcal{A} = A_1, \ldots, A_m$ be the atoms of q where x occurs. Because q is an ACQ, x occurs exactly once in each of them, and each variable y occurring in one of the A_i occurs exactly once in them overall: y cannot occur twice in the same atom, nor can occur in two different atoms A_p and A_q (as in this case A_p, y, A_q, x would be a Berge cycle of q). Let $\mathcal{Y}$ be the set of the variables occurring in $\mathcal{A}$, not including x.

Consider the incidence multigraph G of q (Definition III.7). Remember that we assume queries to be connected, so q is connected, and G is connected. Let $\mathcal{Z}$ be the variables of q which are not in $\mathcal{Y} \cup \{x\}$. For each $z \in \mathcal{Z}$, there must be a path p_z from x to z in G, written $x = w^z_1, \ldots, w^z_{n_z} = z$. Observe that, by definition of $\mathcal{Y}$, we must have $w^z_2 \in \mathcal{Y}$ for any such path. Further, for each $z \in \mathcal{Z}$, we claim that there is a *single* $y_z \in \mathcal{Y}$ such that $w_2^z = y_z$ for any such path. Indeed, assuming to the contrary that there are $y_z \neq y'_z$ in $\mathcal{Y}$, a path p_z whose second element is y_z, and a path p'_z whose second element is y'_z, we deduce from p_z and p'_z a Berge cycle in q.

Thus we can partition $\mathcal{Z}$ into sets of variables $\mathcal{Z}_y$ for $y \in \mathcal{Y}$, where $\mathcal{Z}_y$ contains all variables z of $\mathcal{Z}$ such that y is the variable used to reach z from x. Let $\mathcal{A}_y$ for $y \in \mathcal{Y}$ be the atoms of q whose variables are a subset of $\mathcal{Z}_y \cup \{y\}$. It is clear that $\mathcal{A}$ and the $\mathcal{A}_y$ are a partition of the atoms of q: no atom A can include a variable z from $\mathcal{Z}_y$ and a variable z' from $\mathcal{Z}_{y'}$ for $y \neq y'$ in $\mathcal{Y}$, as otherwise a path from x to z and a path from x to z', together with A, imply that q has a Berge cycle.

Now, we form for each $y \in \mathcal{Y}$ a query q_y as the set of atoms $\mathcal{Z}_y$, with all variables existentially quantified except for y. As the queries $\exists y\, q_y(y)$ are connected queries in ACQ which are strictly smaller than q, by induction we can rewrite q_y to a pointed query of the same size. Hence, we have shown that q can be rewritten as a pointed query built from the A_i and, for each i, the q_y for $y \in \mathcal{Y}$. □

We use this normal form to prove the ACQ Preservation Lemma:

PROOF OF LEMMA VI.4. Fix the instances I and I', and the ACQ q. We show, by induction on $n \in \mathbb{N}$, the following claim: for any $n \in \mathbb{N}$, for any *pointed query* q such that $|q| \leq n$, for any $a \in \mathrm{dom}(I)$, if $I \models q(a)$, then for any $a' \in \mathrm{dom}(I')$ such that $(I,a) \leq_n (I',a')$, we have $I' \models q(a')$. Clearly this claim implies the statement of the Lemma, as by Lemma VI.6 any ACQ query can be written as $\exists x\, q(x)$ with q a pointed query. The case of the trivial query is immediate.

For the induction step, consider a pointed query $q(x)$ of size $n := |q|$, $n > 0$, written in the form of Definition VI.5, and fix $a \in \mathrm{dom}(I)$. Consider a match h of $q(a)$ on I, which must map x to a. Let $a' \in \mathrm{dom}(I')$ be such that $(I,a) \leq_n (I',a')$. We show that $I' \models q'(a)$.

Using notation from Definition VI.5, write $\boldsymbol{y}$ the (disjoint) union of the $\boldsymbol{y^i}$, write $\mathcal{A} = A^1, \ldots, A^n$, and write $q^i_j(y^i_j)$ the subqueries. Let $b^i_j := h(y^i_j)$ for all $y^i_j \in \boldsymbol{y}$. We show that there is a match $h_{\mathcal{A}}$ of $\mathcal{A}$ on I' that maps x to a' and such that every $y^i_j \in \boldsymbol{y}$ is mapped to some element $(b^i_j)'$ of I' such that $(I, b^i_j) \leq_{n-1} (I', (b^i_j)')$. Indeed, start by fixing $h_{\mathcal{A}}(x) := a'$. Now, for each atom $A^i = R(x, \boldsymbol{y^i})$ of $\mathcal{A}$, x occurs at some position, say R^p, and $h(A^i) = R(\boldsymbol{b^i})$ is a fact of I where $h(x) = a$ occurs at position R^p. As each variable in $\boldsymbol{y^i}$ occurs at precisely one position of A^i, we index each of these variables by the one position in A^i where it occurs. Now, as $(I,a) \leq_n (I',a')$, there is a fact $(A^i)' = R((\boldsymbol{b^i})')$ of I' such that $(b^i_p)' = a'$ and, for all $1 \leq j \leq |R|$ with $p \neq j$, we have $(I, b^i_j) \leq_{n-1} (I', (b^i_j)')$. We define $h_{\mathcal{A}}(y^i_j) := (b^i_j)'$ for all i and j. As each variable of $\boldsymbol{y}$ occurs exactly once in $\mathcal{A}$ overall, there definitions cannot conflict, so this correctly defines a function $h_{\mathcal{A}}$ which is clearly indeed a match of $\mathcal{A}$ on I' with the claimed properties.

Now, each of the q^i_j is a pointed query which is strictly smaller than q. Further, the restriction of h to the variables of q^i_j is a match of q^i_j on I that maps each y^i_j (indexing the variables of y^i in the same way as before) to $b^i_j \in \mathrm{dom}(I)$. As we have $(I, b^i_j) \leq_{n-1} (I', (b^i_j)')$, then we can apply the induction hypothesis to show that each of the q^i_j has a match $h_{i,j}$ in I' that maps y^i_j to $(b^i_j)'$. As these queries have disjoint sets of variables, the range of the $h_{i,j}$ is disjoint, and the range of each $h_{i,j}$ overlaps with $h_{\mathcal{A}}$ only on $\{y^i_j\}$, where we have $h_{\mathcal{A}}(y^i_j) = h_{i,j}(y^i_j) = (b^i_j)'$. Thus, we can combine the $h_{i,j}$ and the previously defined $h_{\mathcal{A}}$ to obtain an overall match of q in I' that matches x to a'. This concludes the proof of the induction step, and proves our claim on pointed queries. □

This implies that any superinstance of I_0 that has a k-bounded simulation to $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ must be k-sound for Σ_{U} (no matter whether it satisfies Σ_{U} or not). Indeed, the chase is a universal model for Σ_{UID}, and it satisfies Σ_{UFD} (by the Unique Witness Property, and because I_0 does). Hence, the chase is in particular k-universal for Σ_{U}. Hence, by the ACQ preservation lemma, any superinstance with a k-bounded simulation to the chase is k-sound.

We give a name to such superinstances. For convenience, we also require them to be finite and satisfy Σ_{UFD}. For technical reasons we require that the simulation is the identity on I_0, that it does not map other elements to I_0, and that elements occur in the superinstance at least at the position where their sim-image was introduced in the chase (the *directionality condition*):

Definition VI.7. An **aligned superinstance** $J = (I, \mathsf{sim})$ of I_0 (for Σ_{UFD} and Σ_{UID}) is a finite superinstance I of I_0 that satisfies Σ_{UFD}, and a k-bounded simulation sim from I to $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ such that $\mathsf{sim}_{|I_0}$ is the identity and $\mathsf{sim}_{|(I \setminus I_0)}$ maps to $\mathsf{Chase}(I_0, \Sigma_{\text{UID}}) \setminus I_0$.

Further, for any $a \in \text{dom}(I) \setminus \text{dom}(I_0)$, letting R^p be the position where $\mathsf{sim}(a)$ was introduced in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, we require that $a \in \pi_{R^p}(I)$. We call this the **directionality condition**.

We write $\text{dom}(J)$ to mean $\text{dom}(I)$, and extend other existing notation in the same manner when relevant, e.g., $\mathsf{Wants}(J, \tau)$ means $\mathsf{Wants}(I, \tau)$.

VI.2. Fact-Saturation

Before we perform the *completion process* that allows us to satisfy the UIDs $\Sigma_{\text{UID}}^{\text{rev}}$, we need to perform a *saturation process*. Like aligned superinstances, this process is defined with respect to the superset Σ_{UID}, and does not depend on assumption reversible. The process generalizes relation-saturation from the previous section: instead of achieving all relations, we want the aligned superinstance to achieve all *fact classes*:

Definition VI.8. A **fact class** is a pair $(R^p, \boldsymbol{C})$ of a position $R^p \in \mathsf{Pos}(\sigma)$ and a $|R|$-tuple of $\simeq_k$-classes of elements of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, with $\simeq_k$ as in Definition VI.2.

The **fact class** of a fact $F = R(\boldsymbol{a})$ of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}}) \setminus I_0$ is $(R^p, \boldsymbol{C})$, where a_p is the exported element of F and C_i is the $\simeq_k$-class of a_i in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ for all $R^i \in \mathsf{Pos}(R)$.

A fact class $(R^p, \boldsymbol{C})$ is **achieved** in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ if $\mathsf{NDng}(R^p) \neq \emptyset$ and if it is the fact class of some fact of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}}) \setminus I_0$. Such a fact is an *achiever* of the fact class. We write AFactCl for the set of all achieved fact classes.

For brevity, the dependence on I_0, Σ_{UID}, and k is omitted from this notation.

The requirement that $\mathsf{NDng}(R^p)$ is non-empty is a technicality that will prove useful in Section VIII. The following is easy to see:

LEMMA VI.9. *For any initial instance I_0, set Σ_{UID} of UIDs, and $k \in \mathbb{N}$, AFactCl is finite.*

PROOF. We first show that $\simeq_k$ has only a finite number of equivalence classes on $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$. Indeed, for any element $a \in \text{dom}(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}))$, by the Unique Witness Property, the number of facts in which a occurs is bounded by a constant depending only on I_0 and Σ_{UID}. Hence, there is a constant M depending only on I_0, Σ_{UID}, and k, such that, for any element $a \in \text{dom}(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}))$, the number of elements of $\text{dom}(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}))$ which are relevant to determine the $\simeq_k$-class of a (that is, the elements whose distance to d in the Gaifman graph of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ is $\leq k$) is bounded by M.

This clearly implies that AFactCl is finite, because the number of m-tuples of equivalence classes of $\simeq_k$ that occur in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ is then finite for any $m \leq \max_{R \in \sigma} |R|$, and $\mathsf{Pos}(\sigma)$ is finite. □

We define *fact-saturated* superinstances, which achieve all fact classes in AFactCl:

Definition VI.10. An aligned superinstance $J = (I, \mathsf{sim})$ of I_0 is **fact-saturated** if, for any achieved fact class $D = (R^p, \boldsymbol{C})$ in AFactCl, there is a fact $F_D = R(\boldsymbol{a})$ of $I \setminus I_0$ such that $\mathsf{sim}(a_i) \in C_i$ for all $R^i \in \mathsf{Pos}(R)$. We say that F_D **achieves** D in J.

Note that this definition does not depend on the position R^p of the fact class.

The point of fact-saturation is that, when we perform thrifty chase steps, we can reuse elements from a suitable achiever at the non-dangerous positions. With relation-saturation, the facts were of the right relation; with fact-saturation, they further achieve the right *fact class*, which will be important to maintain the bounded simulation sim.

The fact-saturation completion process, which replaces the relation-saturation process of the previous section, works in the same way.

LEMMA VI.11 (FACT-SATURATED SOLUTIONS). *For any UIDs Σ_{UID}, UFDs Σ_{UFD}, and instance I_0, the result I of performing sufficiently many chase rounds on I_0 is such that $J_0 = (I, \text{id})$ is a fact-saturated aligned superinstance of I_0.*

PROOF. For every $D \in \mathsf{AFactCl}$, let $n_D \in \mathbb{N}$ be such that D is achieved by a fact of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ created at round n_D. As $\mathsf{AFactCl}$ is finite, $n := \max_{D \in \mathsf{AFactCl}} n_D$ is finite. Hence, all classes of $\mathsf{AFactCl}$ are achieved after n chase rounds on I_0.

Consider now I_0' obtained from the aligned superinstance I_0 by n rounds of the UID chase, and $J_0 = (I_0', \text{id})$. It is clear that for any $D \in \mathsf{AFactCl}$, considering an achiever of D in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, the corresponding fact in J_0 is an achiever of D in J_0. Hence, J_0 is indeed fact-saturated. □

We thus obtain a fact-saturated aligned superinstance J_0 of our initial instance I_0, which we now want to complete to one that satisfies the UIDs we are interested in, namely $\Sigma_{\text{UID}}^{\text{rev}}$.

VI.3. Fact-Thrifty Steps

In the previous section, we defined relation-thrifty chase steps, which reused non-dangerous elements from any fact of the correct relation, assuming relation-saturation. We now define *fact-thrifty* steps, which are thrifty steps that reuse elements from a fact achieving the right fact class, thanks to fact-saturation. To do so, however, we must first refine the notion of *thrifty* chase step, to make them apply to aligned superinstances. We will *always* apply them to aligned superinstances for Σ_{UID} and Σ_{UFD}; however, we will always chase by the UIDs of $\Sigma_{\text{UID}}^{\text{rev}}$.

Definition VI.12 (*Thrifty chase steps*). Let $J = (I, \mathsf{sim})$ be an aligned superinstance of I_0 for Σ_{UID} and Σ_{UFD}, let $R^p \subseteq S^q$ be a UID of $\Sigma_{\text{UID}}^{\text{rev}}$, and let $a \in \mathsf{Wants}(J, \tau)$. The result of applying a **thrifty** chase step to a in J by τ is a pair (I', sim') where:

— The instance I' is the result of applying some thrifty step to a in I by τ, as in Definition V.13 (note that this *only* depends on $\Sigma_{\text{UID}}^{\text{rev}}$ and Σ_{UFD}, *not* on Σ_{UID}).

— The mapping sim' extends sim to elements of $\text{dom}(I') \setminus \text{dom}(I)$ as follows. Because sim is a k-bounded simulation and $k > 0$, it is in particular a 1-bounded simulation, so we have $\mathsf{sim}(a) \in \pi_{R^p}(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}))$. Hence, because $\tau \in \Sigma_{\text{UID}}^{\text{rev}} \subseteq \Sigma_{\text{UID}}$, there is a fact $F_{\text{w}} = S(\boldsymbol{b}')$ in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ with $b_q' = \mathsf{sim}(a)$. We call F_{w} the **chase witness**. For any $b \in \text{dom}(I') \setminus \text{dom}(I)$, letting S^r be some position of the new fact F_{n} where b appears (so $b = b_r$), we define $\mathsf{sim}'(b_r) := b_r'$.

We do not know yet whether the result (I', sim') of a thrifty chase step on an aligned superinstance (I, sim) is still an aligned superinstance; we will investigate this later.

Now that we have defined thrifty chase steps on aligned superinstances, we can clarify the role of the directionality condition. Its goal is to ensure, intuitively, that as chase steps go “downwards” in the original chase, thrifty chase steps on aligned superinstances makes the sim mapping go “downwards” in the chase as well. Formally:

LEMMA VI.13 (DIRECTIONALITY). *Let J be an aligned superinstance of I_0 for Σ_{UID} and Σ_{UFD}, and consider the application of a thrifty chase step for a UID $\tau : R^p \subseteq S^q$. Consider the chase witness $F_{\text{w}} = S(\boldsymbol{b}')$. Then b_q' is the exported element of F_{w}.*

PROOF. Let $F_{\text{a}} = R(\boldsymbol{a})$ be the active fact in J, let $F_{\text{n}} = S(\boldsymbol{b})$ be the new fact of J', and let $\tau : R^p \subseteq S^q$ be the UID, so $a_p = b_q$ is the exported element of this chase step. Let $F_{\text{w}} = S(\boldsymbol{b}')$ be the chase witness in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$. Assume by way of contradiction that b_q' was not the exported element in F_{w}, so that it was introduced in F_{w}. In this case, as $\mathsf{sim}(a_p) = \mathsf{sim}(b_q) = b_q'$, by the directionality condition in the definition of aligned superinstances, we have $a_p \in \pi_{S^q}(J)$, which contradicts the fact that $a_p \in \mathsf{Wants}(J, \tau)$. Hence, we have proved by contradiction that b_q' was the exported element in F_{w}. □

This observation will be important to connect fact-saturation to the *fact-thrifty chase steps* that we now define:

Definition VI.14. We define a **fact-thrifty chase step**, using the notation of Definition V.13, as follows: if $\mathsf{NDng}(S^q)$ is non-empty, choose one fact $F_\mathrm{r} = S(\boldsymbol{c})$ of $I \backslash I_0$ that achieves the fact class of $F_\mathrm{w} = S(\boldsymbol{b}')$ (that is, $\mathsf{sim}(c_i) \simeq_k b'_i$ for all i), and use F_r to define $b_r := c_r$ for all $S^r \in \mathsf{NDng}(S^q)$.
We also call a fact-thrifty chase step **fresh** if for all $S^r \in \mathsf{Dng}(S^q)$, we take b_r to be a fresh element only occurring at that position (and extend sim' accordingly).

We first show that, on *fact-saturated* instances, any UID violation can be repaired by a fact-thrifty chase step; this uses Lemma VI.13. More specifically, we show that, for any relation-thrifty chase step that we could want to apply, we could apply a fact-thrifty chase step instead.

LEMMA VI.15 (FACT-THRIFTY APPLICABILITY). *For any fact-saturated superinstance I of an instance I_0, for any UID $\tau : R^p \subseteq S^q$ of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, for any element $a \in \mathit{Wants}(I, \tau)$, we can apply a fact-thrifty chase step on a with τ to satisfy this violation. Further, for any new fact $S(\boldsymbol{e})$ that we can create by chasing on a with τ with a relation-thrifty chase step, we can instead apply a fact-thrifty chase step on a with τ to create a fact $S(\boldsymbol{b})$ with $b_r = e_r$ for all $S^r \in \mathit{Pos}(S) \backslash \mathit{NDng}(S^r)$.*

PROOF. We prove the first part of the statement by justifying the existence of the fact F_r, which only needs to be done if $\mathsf{NDng}(S^q)$ is non-empty. In this case, considering the fact $F_\mathrm{w} = S(\boldsymbol{b}')$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$, we know by Lemma VI.13 that b'_q is the exported element in F_w. Hence, letting D be the fact class of F_w, we have $D = (S^q, \boldsymbol{C})$ for some $\boldsymbol{C}$, and D is in $\mathsf{AFactCl}$ because $\mathsf{NDng}(S^q)$ is non-empty. Hence, by definition of fact-saturation, there is a fact $F_\mathrm{r} = S(\boldsymbol{c})$ in J such that, for all $S^r \in \mathsf{Pos}(R)$, we have $\mathsf{sim}(c_i) \in C_i$, i.e., $\mathsf{sim}(c_i) \simeq_k b'_i$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$. This proves the first part of the claim.
For the second part of the claim, observe that the definition of fact-thrifty chase steps only imposes conditions on the non-dangerous positions, so considering any new fact $S(\boldsymbol{e})$ created by a relation-thrifty chase step, changing its non-dangerous positions to follow the definition of fact-thrifty chase steps, we can create it with a fact-thrifty chase step. □

We now look at which properties are preserved on the result (I', sim') of fact-thrifty chase steps. First note that fact-thrifty chase steps are in particular relation-thrifty, so I' is still weakly-sound and still satisfies Σ_{UFD} (by Lemmas V.14 and V.16). However, we do not know yet whether (I', sim') is an aligned superinstance for Σ_{UFD} and Σ_{UID}.
For now, we show that it is the case for *fresh* fact-thrifty chase steps:

LEMMA VI.16 (FRESH FACT-THRIFTY PRESERVATION). *For any fact-saturated aligned superinstance J of I_0 (for Σ_{UFD} and Σ_{UID}), the result J' of a fresh fact-thrifty chase step on J is still a fact-saturated aligned superinstance of I_0.*

We prove this result in the rest of the subsection. For *non-fresh* fact-thrifty chase steps, the analogous claim is not true in general: it requires us to introduce *essentiality*, the focus of the next subsection, and relies on the assumption reversible that we made on $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and Σ_{UFD}.
To prove the Fresh Fact-Thrifty Preservation Lemma, we first make a general claim about how we can extend a superinstance by adding a fact, and preserve bounded simulations.

LEMMA VI.17. *Let $n \in \mathbb{N}$. Let I_1 and I be instances and sim be a n-bounded simulation from I_1 to I. Let I_2 be a superinstance of I_1 defined by adding one fact $F_\mathrm{n} = R(\boldsymbol{a})$ to I_1, and let sim' be a mapping from I_2 to I such that $\mathsf{sim}'_{|I_1} = \mathsf{sim}$. Assume there is a fact $F_\mathrm{w} = R(\boldsymbol{b})$ in I such that, for all $R^i \in \mathit{Pos}(R)$, $\mathsf{sim}'(a_i) \simeq_n b_i$. Then sim' is an n-bounded simulation from I_2 to I.*

PROOF. We prove the claim by induction on n. The base case of $n = 0$ is immediate.
Let $n > 0$, assume that the claim holds for $n-1$, and show that it holds for n. As sim is an n-bounded simulation, it is an $(n-1)$-bounded simulation, so we know by the induction hypothesis that sim' is an $(n-1)$-bounded simulation.

Let us now show that it is an n-bounded simulation. Let $a \in \mathrm{dom}(I_2)$ be an element and show that $(I_2, a) \leq_n (I, \mathsf{sim}'(a))$. Hence, for any $F = S(\boldsymbol{a})$ a fact of I_2 with $a_p = a$ for some p, we must show that there exists a fact $F' = S(\boldsymbol{a}')$ of I with $a'_p = \mathsf{sim}'(a_p)$ and $(I_2, a_q) \leq_{n-1} (I, a'_q)$ for all $S^q \in \mathsf{Pos}(S)$.

The first possibility is that F is the new fact $F_\mathrm{n} = R(\boldsymbol{a})$. In this case, as we have $(I, b_p) \leq_n (I, \mathsf{sim}'(a_p))$, considering F_w, we deduce the existence of a fact $F'_\mathrm{w} = R(\boldsymbol{c})$ in I such that $c_p = \mathsf{sim}'(a_p)$ and $(I, b_q) \leq_{n-1} (I, c_q)$ for all $1 \leq q \leq |R|$. We take $F' = F'_\mathrm{w}$ as our witness fact for F. By construction we have $c_p = \mathsf{sim}'(a_p)$. Fixing $1 \leq q \leq |R|$, to show that $(I_2, a_q) \leq_{n-1} (I, c_q)$, we use the fact that sim' is an $(n-1)$-bounded simulation to deduce that $(I_2, a_q) \leq_{n-1} (I, \mathsf{sim}'(a_q))$. Now, we have $(I, \mathsf{sim}'(a_q)) \leq_{n-1} (I, b_q)$, and as we explained we have $(I, b_q) \leq_{n-1} (I, c_q)$, so we conclude by transitivity.

If F is another fact, then it is a fact of I_1, so its elements are in $\mathrm{dom}(I_1)$, and as sim' coincides with sim on such elements, we conclude because sim is a n-bounded simulation. □

We now prove the Fresh Fact-Thrifty Preservation Lemma, which concludes the subsection:

PROOF OF LEMMA VI.16. It is immediate that, letting $J' = (I', \mathsf{sim}')$ be the result of the fact-thrifty chase step, I' is still a finite superinstance of I_0, and it still satisfies Σ_{UFD}, because fact-thrifty chase steps are relation-thrifty chase steps, so we can still apply Lemma V.16.

To show that sim' is still a k-bounded simulation, we apply Lemma VI.17 with $F_\mathrm{n} = S(\boldsymbol{b})$ and $F_\mathrm{w} = S(\boldsymbol{b}')$. Indeed, letting $\tau : R^p \subseteq S^q$ be the applied UID in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, we have $\mathsf{sim}'(b_q) = b'_q$ by definition, and have set $\mathsf{sim}'(b_r) := b'_r$ for all $S^r \in \mathsf{Dng}(S^q)$ (note that each such b_r occurs at only one position). For $S^r \in \mathsf{NDng}(S^q)$, we have $\mathsf{sim}'(b_r) \simeq_k b'_r$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ by definition of a fact-thrifty chase step. Hence, by Lemma VI.17, sim' is still a k-bounded simulation from I' to $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$.

We now check the directionality condition on elements of $\mathrm{dom}(I') \setminus \mathrm{dom}(I)$, namely, we show: for $S^r \neq S^q$, if $b_r \in \mathrm{dom}(I') \setminus \mathrm{dom}(I)$, then b_r occurs in J' at the position where $\mathsf{sim}'(b_r)$ was introduced in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$. By the Directionality Lemma (Lemma VI.13) we know that b'_q was the exported element of F_w. Hence, as $\mathsf{sim}'(b_r) := b'_r$, we know that b'_r was introduced at position S^r in F_w in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$, so the condition is respected.

Last, the preservation of fact-saturation is immediate, and the fact that sim' is the identity on I_0 is immediate because $\mathsf{sim}'_{|I_0} = \mathsf{sim}_{|I_0}$. We show that $\mathsf{sim}'_{|I' \setminus I_0}$ maps to $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \setminus I_0$, using the directionality condition. Indeed, for all elements $b_r \in \mathrm{dom}(I') \setminus \mathrm{dom}(I)$ (with $S^r \neq S^q$), which are clearly not in I_0, we have fixed $\mathsf{sim}'(b_r) := b'_r$, and as we explained b'_r is introduced in F_w in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ so it cannot be an element of I_0; hence b'_r is indeed an element of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \setminus I_0$. This is the last point we had to verify. □

VI.4. Essentiality

The problem of non-fresh fact-thrifty chase steps is that, while they try to preserve k-soundness on the non-dangerous positions, they may not preserve it overall:

Example VI.18. Consider the instance $I_0 = \{U(a,u), R(a,b), V(v,b)\}$ depicted as the solid black elements and edges in Figure 4. Consider the UID $\tau : R^1 \subseteq R^2$, and the UFD $\phi : R^1 \to R^2$. We define $\Sigma_{\mathrm{UID}}^{\mathrm{rev}} = \Sigma_{\mathrm{UID}} = \{\tau, \tau^{-1}\}$ and $\Sigma_{\mathrm{UFD}} = \{\phi, \phi^{-1}\}$, so that Σ_{UFD} and $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ are reversible. We have $a \in \mathsf{Wants}(I, \tau^{-1})$ and $b \in \mathsf{Wants}(I, \tau)$. To satisfy these violations, we can apply a fact-thrifty chase step by τ^{-1} on a and create $F = R(b,a)$, noting that there are no non-dangerous positions. However, the superinstance $I_0 \sqcup \{F\}$ is not a k-sound superinstance of I_0 for $k \geq 3$. For instance, it makes the following ACQ true, which is not true in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$: $\exists xyzw\, V(x,y), R(y,z), U(z,w)$.

However, for any value of k, this problem can be avoided in the following way. First, apply k fresh fact-thrifty chase steps by τ to create the chain $R(b,b_1), R(b_1,b_2), \ldots, R(b_{k-1},b_k)$. Then apply k fresh fact-thrifty chase steps by τ^{-1} to create $R(a_1,a), R(a_2,a_1), \ldots, R(a_k,a_{k-1})$. Now we can apply a non-fresh fact-thrifty chase step by τ^{-1} on a and create $R(b_k,a_k)$, and this does not make any new

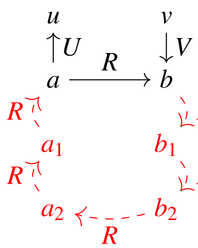

Fig. 4: Soundness and fact-thrifty chase steps (see Example VI.18)

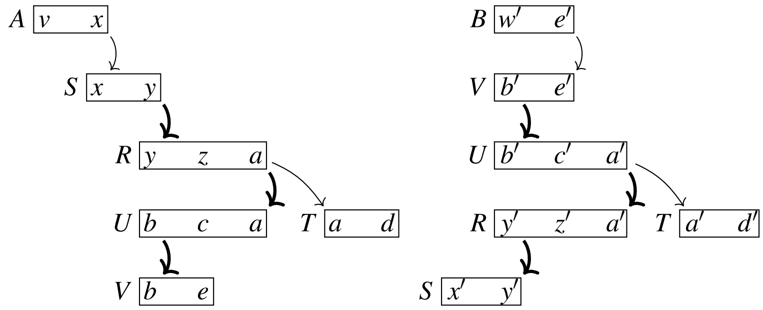

Fig. 5: UID Chase Similarity Theorem; see Example VI.29.

ACQ of size $\leq k$ true. This process is illustrated with red elements and red dashed edges in Figure 4 for $k = 2$.

The intuition behind this example is that non-fresh fact-thrifty chase steps may connect together elements at the dangerous positions, but their image by sim may be too dissimilar, so the bounded simulation does not extend. This implies that, in general, the result of a fact-thrifty chase step may not be an aligned superinstance. As the example shows, however, we can avoid that problem if we chase for sufficiently long, so that the "history" of elements no longer contains anything specific about them.

We first formalize this notion for elements of the chase $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$, which we call *essentiality*. We will then define it for aligned superinstances using the sim mapping.

Definition VI.19. We define a forest structure on the facts of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$: the facts of I_0 are the roots, and the parent of a fact F not in I_0 is the fact F' that was the active fact for which F was created, so that F' and F share the exported element of F.

For $a \in \mathrm{dom}(\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}))$, if a was introduced at position S^r of an S-fact $F = S(\boldsymbol{a})$ created by applying the UID $\tau : R^p \subseteq S^q$ (with $S^q \neq S^r$) to its parent fact F', we call τ the **last UID** of a. The **last two UIDs** of a are (τ, τ') where τ' is the last UID of the exported element a_q of F (which was introduced in F'). For $n \in \mathbb{N}$, we define the **last** n **UIDs** in the same way, for elements of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ introduced after sufficiently many rounds.

We say that a is n**-essential** if its last n UIDs are reversible in Σ_{UID}. This is in particular the case if these last UIDs are in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$: indeed, $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ satisfies assumption reversible, so for any $\tau \in \Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, we have $\tau^{-1} \in \Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, so that $\tau^{-1} \in \Sigma_{\mathrm{UID}}$.

The point of this definition is the following result, which we state without proof for now. We will prove it in Section VI.5:

THEOREM VI.20 (UID CHASE SIMILARITY THEOREM). *For any instance I_0, transitively closed set of UIDs Σ_{UID}, and $n \in \mathbb{N}$, for any two elements a and b respectively introduced at positions R^p and S^q in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$, if a and b are n-essential, and if $R^p \subseteq S^q$ and $S^q \subseteq R^p$ hold in Σ_{UID}, then $a \simeq_n b$.*

In other words, in the chase, when your last n UIDs were reversible, then your $\simeq_n$-class only depends on the position where you were introduced.

We use this to define a corresponding notion on aligned superinstances: an aligned superinstance is *n-essential* if, for all elements that witness a violation of the UIDs $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ that we wish to solve, their sim image is an n-essential element of the chase, introduced at a suitable position. In fact, we introduce a more general definition, which does not require the superinstance to be aligned, i.e., does not require that sim is a k-bounded simulation.

Definition VI.21. Let $J = (I, \mathsf{sim})$ be a pair of a superinstance I of I_0 and a mapping sim from I to $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$. Let $k \in \mathbb{N}$. We call $a \in \text{dom}(I)$ ***n*-essential** in J for $\Sigma_{\text{UID}}^{\text{rev}}$ if for any position $S^q \in \mathsf{Pos}(\sigma)$ such that $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(I, S^q)$, then:

— $\mathsf{sim}(a)$ is an n-essential element of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$;
— the position T^v where $\mathsf{sim}(a)$ was introduced in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ is such that $T^v \subseteq S^q$ and $S^q \subseteq T^v$ hold in $\Sigma_{\text{UID}}^{\text{rev}}$, which we write $T^v \sim_{\text{ID}} S^q$ as in the previous sections.

Note that if there is no UID of $\Sigma_{\text{UID}}^{\text{rev}}$ applicable to a, then a is vacuously n-essential. We call J ***n*-essential** for $\Sigma_{\text{UID}}^{\text{rev}}$ if, for all $a \in \text{dom}(J)$, a is n-essential in J for $\Sigma_{\text{UID}}^{\text{rev}}$.

We now show that, as we assumed the UIDs of $\Sigma_{\text{UID}}^{\text{rev}}$ to be reversible (assumption reversible), fresh fact-thrifty chase steps by $\Sigma_{\text{UID}}^{\text{rev}}$ never make essentiality decrease, and even make it *increase* on new elements:

LEMMA VI.22 (THRIFTY STEPS AND ESSENTIALITY). *For any n-essential aligned superinstance J, letting $J' = (I', \mathsf{sim}')$ be the result of a thrifty chase step on J by a UID of $\Sigma_{\text{UID}}^{\text{rev}}$, J' is still n-essential. Further, all elements of $\text{dom}(J') \setminus \text{dom}(J)$ are $(n+1)$-essential in J'.*

PROOF. Fix J and J'; note that J' may not be an aligned superinstance. Consider first the elements of $\text{dom}(J)$ in J'. For any $a \in \text{dom}(J)$, by definition of thrifty chase steps, we know that for any $T^v \in \mathsf{Pos}(\sigma)$ such that $a \in \pi_{T^v}(J')$, we have either $a \in \pi_{T^v}(J)$ or $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(J, T^v)$. Hence, as $\Sigma_{\text{UID}}^{\text{rev}}$ is transitively closed, for any $U^w \in \mathsf{Pos}(\sigma)$ such that $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(J', U^w)$, we have also $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(J, U^w)$, and as J is n-essential, we conclude that a is n-essential in J'. Hence, it suffices to show that any element in $\text{dom}(J') \setminus \text{dom}(J)$ is $(n+1)$-essential in J'.

To do this, write $\tau : R^p \subseteq S^q$ the UID applied in the chase step, and let $F_{\text{n}} = S(\boldsymbol{b})$ be the new fact. By definition of thrifty chase steps, we had $b_q \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(J, S^q)$, so that b_q was n-essential because J was; hence, $\mathsf{sim}(b_q)$ is n-essential in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$. By the Directionality Lemma (Lemma VI.13), $b'_q := \mathsf{sim}(b_q)$ is also the exported element of the chase witness $F_{\text{w}} = S(\boldsymbol{b}')$. Now, the UID τ applied in the thrifty chase step is also the one applied to create F_{w} in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, and as $\tau \in \Sigma_{\text{UID}}^{\text{rev}}$, by assumption reversible, we have $\tau^{-1} \in \Sigma_{\text{UID}}^{\text{rev}}$, hence $\tau^{-1} \in \Sigma_{\text{UID}}$. Hence, for all $S^r \in \mathsf{Pos}(S) \setminus \{S^q\}$, b'_r is $(n+1)$-essential in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, and is introduced at position S^r.

Now, let $a \in \text{dom}(J') \setminus \text{dom}(J)$, and let T^v such that $a \in \mathsf{Wants}_{\Sigma_{\text{UID}}^{\text{rev}}}(J', T^v)$. Let $U^w \in \mathsf{Pos}(\sigma)$ and $\tau' : U^w \subseteq T^v$ that witness this, i.e., $\tau' \in \Sigma_{\text{UID}}^{\text{rev}}$ and $a \in \pi_{U^w}(J')$. By assumption reversible, we have $U^w \sim_{\text{ID}} T^v$ in $\Sigma_{\text{UID}}^{\text{rev}}$. Now, by definition of thrifty chase steps on aligned superinstances, we know that we defined $\mathsf{sim}'(a) := b'_r$ for some S^r where a occurred in F_{n}. Further, by definition of thrifty chase steps, we know that all positions in which a occurs in F_{n}, and thus all positions where it occurs in J', are $\sim_{\text{ID}}$-equivalent in $\Sigma_{\text{UID}}^{\text{rev}}$; in particular $S^r \sim_{\text{ID}} U^w$, hence by transitivity $S^r \sim_{\text{ID}} T^v$. By the previous paragraph, $\mathsf{sim}(a) = b'_r$ is an $(n+1)$-reversible element introduced in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ at position S^r, and we have $S^r \sim_{\text{ID}} T^v$. This shows that a is $(n+1)$-reversible in J'.

Hence, J' is indeed n-reversible, and the elements of $\text{dom}(J') \setminus \text{dom}(J)$ are indeed $(n+1)$-reversible, which concludes the proof. □

In conjunction with the Fresh Fact-Thrifty Preservation Lemma, this implies that applying sufficiently many fresh fact-thrifty chase rounds yields an n-essential aligned superinstance:

LEMMA VI.23 (ENSURING ESSENTIALITY). *For any $n \in \mathbb{N}$, applying $n+1$ fresh fact-thrifty chase rounds on a fact-saturated aligned superinstance J by the UIDs of $\Sigma_{\text{UID}}^{\text{rev}}$ yields an n-essential aligned superinstance J'.*

PROOF. Fix the aligned superinstance $J = (I, \mathsf{sim})$. We use the Fresh Fact-Thrifty Preservation Lemma to show that the property of being aligned is preserved, so we only show that the result is n-essential. We prove this claim by induction on n.

For the base case, we must show that the result $J' = (I', \mathsf{sim}')$ of a fresh fact-thrifty chase round on J by $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ is 0-essential. Let $S^q \in \mathsf{Pos}(\sigma)$ and $a \in \mathsf{Wants}_{\Sigma_{\mathrm{UID}}^{\mathrm{rev}}}(J', S^q)$. As $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ is transitively closed, by definition of a chase round, we have $a \in \mathrm{dom}(J') \setminus \mathrm{dom}(J)$, because UID violations on elements of $\mathrm{dom}(J)$ must have been solved in J'; hence, a was created by a fact-thrifty chase step on J. By similar reasoning as in the proof of Lemma VI.22, considering the chase witness F_{w} for this chase step, we conclude that $\mathsf{sim}(a)$ was introduced at a position T^v in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ such that $T^v \sim_{\mathsf{ID}} S^q$ in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$. Further, $\mathsf{sim}(a)$ is vacuously 0-essential. Hence, J' is indeed 0-essential.

For the induction step, let J' be the result of $n+1$ fresh fact-thrifty chase rounds on J, and show that it is n-essential. By induction hypothesis, the result $J'' = (I'', \mathsf{sim}'')$ of n fresh fact-thrifty chase rounds is $(n-1)$-essential. Now, again by definition of a chase round, for any position $S^q \in \mathsf{Pos}(\sigma)$ and $a \in \mathsf{Wants}_{\Sigma_{\mathrm{UID}}^{\mathrm{rev}}}(J'', S^q)$, we must have $a \in \mathrm{dom}(J') \setminus \mathrm{dom}(J'')$, so that a was created by applying a fact-thrifty chase step on an element a'' in J'' which witnessed a violation of a UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$. As J'' is $(n-1)$-essential, a' was $(n-1)$-essential in J'', so we conclude by Lemma VI.22 that a is n-essential in J'. Hence, we conclude that J' is indeed n-essential. □

Hence, we can ensure k-essentiality. The point of essentiality is to guarantee that the result of *non-fresh* fact-thrifty chase steps on a k-essential aligned superinstance is also an aligned superinstance.

LEMMA VI.24 (FACT-THRIFTY PRESERVATION). *For any fact-saturated k-essential aligned superinstance J for Σ_{UID} and Σ_{UFD}, the result J' of* any *fact-thrifty chase step on J by a UID of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ is still a fact-saturated and k-essential aligned superinstance.*

PROOF. Fix $J' = (I', \mathsf{sim}')$, the UID $\tau : R^p \subseteq S^q$ of $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, which is reversible by assumption reversible, and the element $a \in \mathrm{dom}(J)$ to which it is applied.

The fact that k-essentiality is preserved is by Lemma VI.22, and fact-saturation is clearly preserved, so we must only show that J' is still an aligned superinstance. The fact that J' is a finite superinstance of I_0 is immediate, and it still satisfies Σ_{UFD} by Lemma V.16 because fact-thrifty chase steps are relation-thrifty chase steps. The directionality condition is clearly respected because any new element in $\mathrm{dom}(J') \setminus \mathrm{dom}(J)$ occurs at least at the position at which its sim'-image was introduced in the chase (namely, the position where it occurs in F_{w}), and the additional conditions on $\mathsf{sim}'_{|I_0}$ and $\mathsf{sim}'_{|J' \setminus I_0}$ are still verified.

The only thing to show is that sim' is still a k-bounded simulation. Let $F_{\mathrm{n}} = S(\boldsymbol{b})$ be the new fact and $F_{\mathrm{w}} = S(\boldsymbol{b}')$ be the chase witness. Now, as in the proof of the Thrifty Steps And Essentiality Lemma, and using the Directionality Lemma, all elements of F_{w} are n-essential (and, except for b'_q, they were introduced at their position of F_{w}).

Now, to show that sim' is a k-bounded simulation, we use Lemma VI.17, so it suffices to show that we have $\mathsf{sim}(b_r) \simeq_k b'_r$ for all r. This is the case whenever we have $\mathsf{sim}(b_r) = b'_r$, which is guaranteed by definition for $S^r = S^q$ and for elements in $\mathsf{Dng}(S^q)$ such that $S^r \leftrightarrow_{\mathrm{FUN}} S^q$ does not hold. For non-dangerous elements, the fact that $\mathsf{sim}(b_r) \simeq_k b'_r$ is by definition of fact-thrifty chase steps. For the other positions, there are two cases:

— $b_r \in \mathrm{dom}(I)$, in which case $b_r \in \mathsf{Wants}_{\Sigma_{\mathrm{UID}}^{\mathrm{rev}}}(I, S^r)$. As J is n-essential, $\mathsf{sim}(b_r)$ is an n-essential element of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ introduced at a position T^v such that $T^v \sim_{\mathsf{ID}} S^r$ holds in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$. Now, b'_r is an n-essential element of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ introduced at position S^r. By the UID Chase Similarity Theorem, we then have $\mathsf{sim}'(b_r) \simeq_k b'_r$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$

— $b_r \notin \mathrm{dom}(I)$, in which case the claim is immediate unless it occurs at multiple positions. However, by definition of thrifty chase steps, all positions at which it occurs are related by $\sim_{\mathsf{ID}}$ in $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$, so the corresponding elements of F_{w} are also $\simeq_k$-equivalent by the UID Chase Similarity Theorem: hence we have $\mathsf{sim}'(b_r) \simeq_k b'_r$.

We conclude by Lemma VI.17 that J' is indeed an aligned superinstance, which concludes the proof. □

We can now conclude the proof of Theorem VI.1. Let I_0 be the initial instance, and consider $J_0 = (I_0, \mathrm{id})$ which is trivially an aligned superinstance of I_0. Apply the Fact-Saturated Solutions Lemma to obtain a fact-saturated aligned superinstance $J'_0 = (I'_0, \mathsf{sim}')$. We must now show that we can complete J'_0 to a superinstance that satisfies $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ as well, which we do with the following variant of the Reversible Relation-Thrifty Completion Proposition (see Proposition V.17):

PROPOSITION VI.25 (REVERSIBLE FACT-THRIFTY COMPLETION). *For any reversible Σ_{UFD} and $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, for any transitively closed UIDs $\Sigma_{\mathrm{UID}} \supseteq \Sigma^{\mathrm{rev}}_{\mathrm{UID}}$, for any fact-saturated aligned superinstance J'_0 of I_0 (for Σ_{UFD} and Σ_{UID}), we can use fact-thrifty chase steps by UIDs of $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ to construct an aligned fact-saturated superinstance J_{f} of I_0 (for Σ_{UFD} and Σ_{UID}) that satisfies $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$.*

To prove this lemma, we first apply the Ensuring Essentiality Lemma with the UIDs of $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$ to make J'_0 k-essential. By the Fresh Fact-Thrifty Preservation Lemma, the result $J_1 = (I_1, \mathsf{sim}_1)$ is then a fact-saturated k-essential aligned superinstance of I_0 (for Σ_{UID} and Σ_{UFD}).

To prove the Reversible Fact-Thrifty Completion Proposition, we will now use the Reversible Relation-Thrifty Completion Proposition (Proposition V.17) on J_1; but we must refine it to a stronger claim. We do so using the following definition:

Definition VI.26. A **thrifty sequence** on an instance I for UIDs Σ_{UID} and UFDs Σ_{UFD} is a sequence L defined inductively as follows, with an **output** $L(I)$ which is a superinstance of I that we also define inductively:

— The empty sequence $L = ()$ is a thrifty sequence, with $L(I) = I$
— Let L' be a thrifty sequence, let $I' = L'(I)$ be the output of L', and let $t = (a, \tau, \boldsymbol{b})$ be a triple formed of an element $a \in \mathrm{dom}(I')$, a UID $\tau : R^p \subseteq S^q$ of Σ_{UID}, and an $|S|$-tuple $\boldsymbol{b}$. We require that the fact $S(\boldsymbol{b})$ can be created in I' by applying a thrifty chase step to a in $L'(I)$ by τ (Definition V.13). Then the concatenation L of L' and t is a thrifty sequence, and its output $L(I)$ is the result of performing this chase step on $L'(I)$, namely, $L(I) := L'(I) \sqcup \{S(\boldsymbol{b})\}$.

The **length** of L is written $|L|$ and the elements of L are indexed by $L_1, \ldots, L_{|L|}$. We define a **relation-thrifty sequence** in the same way with relation-thrifty steps, and likewise define a **fact-thrifty sequence**.

With this definition, the Reversible Relation-Thrifty Completion Proposition (Proposition V.17) implies that there is a relation-thrifty sequence L such that $L(I'_0)$ is a finite weakly-sound superinstance I_{f} of I_0 that satisfies Σ_{UFD} and $\Sigma^{\mathrm{rev}}_{\mathrm{UID}}$. Our goal to prove the Reversible Fact-Thrifty Completion Proposition is to rewrite L to a fact-thrifty sequence. To do this, we first need to show that any two thrifty sequences that coincide on non-dangerous positions have the same effect in terms of UID violations:

Definition VI.27. Let Σ_{UID} be UIDs and Σ_{UFD} be UFDs, let I_0 be an instance, and L and L' be thrifty sequences on I_0. We say that L and L' **non-dangerously match** if $|L| = |L'|$ and that for all $1 \leq i \leq |L|$, writing $L_i = (a, \tau, \boldsymbol{b})$ and $L'_i = (a', \tau', \boldsymbol{b}')$, we have $a = a'$, $\tau = \tau'$, and, writing $\tau : R^p \subseteq S^q$, we have $b_r = b'_r$ for all $S^r \in \mathsf{Pos}(S) \backslash \mathsf{NDng}(S^q)$.

LEMMA VI.28 (THRIFTY SEQUENCE REWRITING). *Let Σ_{UID} be UIDs and Σ_{UFD} be UFDs, let I_0 be an instance, and let L and L' be thrifty sequences on I_0 that non-dangerously match. Then $L(I_0)$ satisfies Σ_{UID} iff $L'(I_0)$ satisfies Σ_{UID}.*

PROOF. We prove by induction on the common length of L and L' that, if L and L' non-dangerously match, then, for all $U^v \in \mathsf{Pos}(\sigma)$, we have $\pi_{U^v}(L(I_0)) = \pi_{U^v}(L'(I_0))$. If both L and L' have length 0, the claim is trivial. For the induction step, write $I := L(I_0)$ and $I' := L'(I_0)$. Write L as the concatenation of L_2 and its last tuple $t = (a, \tau, \boldsymbol{b})$, and write similarly L' as the concatenation of L'_2 and the last tuple $t' = (a', \tau', \boldsymbol{b}')$. Let $U^v \in \mathsf{Pos}(\sigma)$ and show that $\pi_{U^v}(L(I_0)) = \pi_{U^v}(L'(I_0))$. Clearly L_2 and L'_2 non-dangerously match and are strictly shorter than L and L', respectively, so by the induction hypothesis, writing $I_2 := L_2(I_0)$ and $I'_2 := L'_2(I_0)$, we have $\pi_{U^v}(I_2) = \pi_{U^v}(I'_2)$. Further, we have $\tau = \tau'$; write them as $R^p \subseteq S^q$. We then have $I = I_2 \sqcup S(\boldsymbol{b})$, and $I' = I_2' \sqcup S(\boldsymbol{b}')$. As we must have $b_r = b_r'$ if $U^v \notin \mathsf{NDng}(S^q)$, there is nothing to show unless we have $U^v \in \mathsf{NDng}(S^q)$. However, in this case, writing U^v as S^r, then, by definition of thrifty chase steps, we have $b_r \in \pi_{S^r}(I_2)$, so that $\pi_{S^r}(I) = \pi_{S^r}(I_2)$. Likewise, $\pi_{S^r}(I') = \pi_{S^r}(I_2')$, hence $\pi_{S^r}(I) = \pi_{S^r}(I')$. This concludes the induction proof.

We now prove the lemma. Fix $\tau : R^p \subseteq S^q$ in Σ_{UID}. We have $L(I_0) \models \tau$ iff $\pi_{R^p}(L(I_0)) \setminus \pi_{S^q}(L(I_0)) = \emptyset$, and likewise for $L'(I_0)$. By the result proved in the paragraph above, these conditions are equivalent, and thus we have $L(I_0) \models \tau$ iff $L'(I_0) \models \tau$. □

Hence, considering our fact-saturated aligned superinstance $J_1 = (I_1, \mathsf{sim}_1)$ (for Σ_{UID} and Σ_{UFD}), use the rephrasing of the Reversible Relation-Thrifty Completion Proposition to obtain a relation-thrifty sequence L such that $L(I_1)$ satisfies Σ_{UID}. We modify L inductively to obtain a *fact-thrifty sequence* L' that non-dangerously matches L, in the following manner. Whenever L applies a relation-thrifty step $t = (a, \tau, \boldsymbol{b})$ to the previous instance $L_2(I_1)$, then observe that $L_2(I_1)$ is fact-saturated, because I_1 was fact-saturated and fact-thrifty chase steps preserve fact-saturation, by the Fact-Thrifty Preservation Lemma. Hence, by the Fact-Thrifty Applicability Lemma, instead of applying the relation-thrifty step described by t, we can choose to apply a fact-thrifty step on a with τ, defining the new fact using $\boldsymbol{b}$ except on the non-dangerous positions. By Lemma VI.28, the resulting L' also ensures that $L'(I_1)$ satisfies Σ_{UID}.

Considering now the fact-thrifty sequence L', as J_1 is a fact-saturated k-essential aligned superinstance of I_0 (for Σ_{UID} and Σ_{UFD}), letting $I_{\text{f}} := L'(I_1)$, we can use the Fact-Thrifty Preservation Lemma to define an aligned fact-saturated superinstance $J_{\text{f}} = (I_{\text{f}}, \mathsf{sim}_{\text{f}})$ (for Σ_{UID} and Σ_{UFD}), following each fact-thrifty step, and we have shown that I_{f} satisfies Σ_{UID}. Hence, we have proven the Reversible Fact-Thrifty Completion Proposition.

To prove Theorem VI.1, we can simply apply the proposition with $\Sigma_{\text{UID}} = \Sigma_{\text{UID}}^{\text{rev}}$, and the resulting aligned superinstance $J_{\text{f}} = (I_{\text{f}}, \mathsf{sim}_{\text{f}})$ of I_0 satisfies Σ_{UID} and is k-sound for Σ_{U} and ACQ. Further, it satisfies Σ_{UFD} and is finite, by definition of being an aligned superinstance. Hence, I_{f} is the desired k-universal model, which proves Theorem VI.1.

VI.5. UID Chase Similarity Theorem

We conclude the section by proving the UID Chase Similarity Theorem:

THEOREM VI.20 (UID CHASE SIMILARITY THEOREM). *For any instance I_0, transitively closed set of UIDs Σ_{UID}, and $n \in \mathbb{N}$, for any two elements a and b respectively introduced at positions R^p and S^q in $\mathit{Chase}(I_0, \Sigma_{\text{UID}})$, if a and b are n-essential, and if $R^p \subseteq S^q$ and $S^q \subseteq R^p$ hold in Σ_{UID}, then $a \simeq_n b$.*

Note that this result does not involve FDs, and applies to any arbitrary transitively closed set of UIDs, not relying on any finite closure properties, or on assumption reversible. It only assumes that the last n dependencies used to create a and b were reversible.

Example VI.29. Consider Figure 5 on page 29, which illustrates the neighborhood of two elements, a and a', in the UID chase by some UIDs. Each rectangle represents a higher-arity fact, and edges represent the UIDs used in the chase, with thick edges representing reversible UIDs.

The last UID applied to create a was $S^2 \subseteq R^1$, and the last UID for a' is $V^1 \subseteq U^1$; they are reversible. Further, a is introduced at position R^3 and a' at position U^3, and $R^3 \subseteq U^3$ holds and is reversible. The theorem claims that a that a' are 1-bounded-bisimilar, which is easily verified; in fact, they are 2-bounded-bisimilar. This is intuitively because all child facts of the R-fact at the left must occur at the right by definition of the UID chase, and the parent fact must occur as well because of the reverse of the last UID for a; a similar argument ensures that the facts at the right must be reflected at the left.

However, note that a and a' are not 3-bounded-bisimilar: the A-fact at the left is not reflected at the right, and vice-versa for the B-fact, because these UIDs are not reversible, To prove the theorem, fix the instance I_0 and the set Σ_{UID} of UIDs. We first show the following easy lemma:

LEMMA VI.30. *For any $n > 0$ and position R^p, for any two elements a, b of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ introduced at position R^p in two facts F_a and F_b, letting a' and b' be the exported elements of F_a and F_b, if $a' \simeq_{n-1} b'$, then $a \simeq_n b$.*

PROOF. We proceed by induction on n. By symmetry, it suffices to show that $(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}), a) \leq_n (\mathsf{Chase}(I_0, \Sigma_{\text{UID}}), b)$.

For the base case $n = 1$, observe that, for every fact F of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ where a occurs at some position S^q, there are two cases. Either $F = F_a$, so we can pick F_b as the representative fact, or the UID $R^p \subseteq S^q$ is in Σ_{UID} so we can pick a corresponding fact for b by definition of the chase. Hence, the claim is shown for $n = 1$.

For the induction step, we proceed in the same way. If $F = F_a$, we pick F_b as representative fact, and use either the hypothesis on a' and b' or the induction hypothesis (for other elements of F_a and F_b) to justify that F_b is suitable. Otherwise, we pick the corresponding fact for b which must exist by definition of the chase, and apply the induction hypothesis to the other elements of the fact to conclude. □

We now prove the UID Chase Similarity Theorem. Throughout the proof, we write $R^p \sim_{\mathsf{ID}} S^q$ as shorthand to mean that $R^p \subseteq S^q$ and $S^q \subseteq R^p$ are in Σ_{UID}: it is still the case that $\sim_{\mathsf{ID}}$ is an equivalence relation, even without assumption reversible.

We prove the main claim by induction on n: for any positions R^p and S^q such that $R^p \sim_{\mathsf{ID}} S^q$, for any two n-essential elements a and b respectively introduced at positions R^p and S^q, we have $a \simeq_n b$. By symmetry it suffices to show that $a \leq_n b$ in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, formally, $(\mathsf{Chase}(I_0, \Sigma_{\text{UID}}), a) \leq_n (\mathsf{Chase}(I_0, \Sigma_{\text{UID}}), b)$.

The base case of $n = 0$ is immediate.

For the induction step, fix $n > 0$, and assume that the result holds for $n-1$. Fix R^p and S^q such that $R^p \sim_{\mathsf{ID}} S^q$, and let a, b be two n-essential elements introduced respectively at R^p and S^q in facts F_a and F_b. Note that by the induction hypothesis we already know that $a \leq_{n-1} b$ in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$; we must show that this holds for n.

First, observe that, as a and b are n-essential with $n > 0$, they are not elements of I_0. Hence, by definition of the chase, for each one of them, the following is true: for each fact of the chase where the element occurs, it only occurs at one position, and all other elements co-occurring with it in a fact of the chase occur only at one position and in exactly one of these facts. Thus, to prove the claim, it suffices to construct a mapping ϕ from the set $N_1(a)$ of the facts of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ where a occurs, to the set $N_1(b)$ of the facts where b occurs, such that the following holds: for every fact $F = T(\boldsymbol{a})$ of $N_1(a)$, letting T^c be the one position of F such that $a_c = a$, the element b occurs at position T^c in $\phi(F) = T(\boldsymbol{b})$, and for every i, we have $a_i \leq_{n-1} b_i$.

By construction of the chase (using the Unique Witness Property), $N_1(a)$ consists of exactly the following facts:

— The fact $F_a = R(\boldsymbol{a})$, where $a_d = a'$ is the exported element (for a certain $R^d \neq R^p$) and $a_p = a$ was introduced at R^p in F_a. Further, for $i \notin \{p, d\}$, the element a_i was introduced at R^i in F_a.
— For every UID $\tau : R^p \subseteq V^g$ of Σ_{UID}, a V-fact F_a^τ where the element at position V^g is a. Further, for $i \neq g$, the element at position V^i in F_a^τ was introduced at this position in that fact.

A similar characterization holds for b: we write the corresponding facts F_b and F_b^τ. We construct the mapping ϕ as follows:

— If $R^p = S^q$ then set $\phi(F_a) := F_b$; otherwise, as $\tau : S^q \subseteq R^p$ is in Σ_{UID}, set $\phi(F_a) := F_b^\tau$.
— For every UID $\tau : R^p \subseteq V^g$ of Σ_{UID}, as $R^p \sim_{\mathsf{ID}} S^q$, by transitivity, either $S^q = V^g$ or the UID $\tau' : S^q \subseteq V^g$ is in Σ_{UID}. In the first case, set $\phi(F_a^\tau) := F_b$. In the second case, set $\phi(F_a^\tau) := F_b^{\tau'}$.

We must now show that this mapping ϕ from $N_1(a)$ to $N_1(b)$ satisfies the required conditions. Verify that indeed, by construction, whenever a occurs at position T^c in F, then $\phi(F)$ is a T-fact where b occurs at position T^c. So we must show that for any $F \in N_1(a)$, writing $F = T(\boldsymbol{a})$ and $\phi(F) = T(\boldsymbol{b})$, with $a_c = a$ and $b_c = b$ for some c, we have indeed $a_i \leq_{n-1} b_i$ for all $T^i \in \mathsf{Pos}(T)$. If $n = 1$ there is nothing to show and we are done, so we assume $n \geq 2$. If $i = c$ then the claim is immediate by the induction hypothesis; otherwise, we distinguish two cases:

(1) $F = F_a$ (so that $T = R$ and $c = p$), or $F = F_a^\tau$ such that the UID $\tau : R^p \subseteq T^c$ is reversible, meaning that $\tau^{-1} \in \Sigma_{\text{UID}}$. In this case, by construction, either $\phi(F) = F_b$ or $\phi(F) = F_b^{\tau'}$ for $\tau' : S^q \subseteq T^c$; τ' is then reversible, because $R^p \sim_{\mathsf{ID}} S^q$ and $R^p \sim_{\mathsf{ID}} T^c$.
We show that for all $1 \leq i \leq |T|$ such that $i \neq c$, the element a_i is $(n-1)$-essential and was introduced in $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$ at a position in the $\sim_{\mathsf{ID}}$-class of T^i. Once we have proved this, we can show the same for all b_i in a symmetric way, so that we can conclude that $a_i \leq_{n-1} b_i$ by induction hypothesis. To see why the claim holds, we distinguish two subcases. Either a_i was introduced in F, or we have $F = F_a$, $i = d$ and a_i is the exported element for a.
In the first subcase, a_i was created by applying the reversible UID τ and the exported element was a, which is n-essential, so a_i is $(n-1)$-essential (in fact it is even $(n+1)$-essential), and a_i is introduced at position T^i.
In the second subcase, a_i is the exported element used to create a, which is n-essential, so a_i is $(n-1)$-essential; and as $n \geq 2$, the last dependency applied to create a_i is reversible, so that a_i was introduced at a position in the same $\sim_{\mathsf{ID}}$-class as T^i.
Hence, we have proved the desired claim for the first case.

(2) $F = F_a^\tau$ such that $\tau : R^p \subseteq T^c$ is not reversible. In this case, we cannot have $T^c = S^q$ (because we have $R^p \sim_{\mathsf{ID}} S^q$), so we must have $\phi(F) = F_b^{\tau'}$ with $\tau' : S^q \subseteq T^c$. Now, all a_i for $i \neq c$ were introduced in F at position T^i, and likewise for the b_i in $\phi(F)$. Using Lemma VI.30, as $a \simeq_{n-1} b$, we conclude that $a_i \simeq_n b_i$, hence $a_i \leq_{n-1} b$.

This concludes the proof of the UID Chase Similarity Theorem, thus completing the proof of Theorem VI.1.

VII. DECOMPOSING THE CONSTRAINTS

In this section, we lift assumption reversible, proving:

THEOREM VII.1. *Finitely-closed UIDs and UFDs have finite k-universal models for ACQs.*

VII.1. Partitioning the UIDs

We write the UFDs as Σ_{UFD} and the UIDs as Σ_{UID}. We will proceed by partitioning Σ_{UID} into subsets of UIDs which are either reversible or are much simpler to deal with.

Our desired notion of partition respects an order on UID, which we now define. As we will show (Lemma VII.8), the order is also respected by thrifty chase steps.

Definition VII.2. For any $\tau, \tau' \in \Sigma_{\text{UID}}$, we write $\tau \rightarrowtail \tau'$ when we can write $\tau = R^p \subseteq S^q$ and $\tau' = S^r \subseteq T^v$ with $S^q \neq S^r$, and the UFD $S^r \rightarrow S^q$ is in Σ_{UFD}. An **ordered partition** $(P_1, \ldots, P_n)$ of Σ_{UID} is a partition of Σ_{UID} (i.e., $\Sigma_{\text{UID}} = \bigsqcup_i P_i$) such that for any $\tau \in P_i$, $\tau' \in P_j$, if $\tau \rightarrowtail \tau'$ then $i \leq j$.

The point of partitioning Σ_{UID} is to be able to control the structure of the UIDs in each class:

Definition VII.3. We call $P \subseteq \Sigma_{\text{UID}}$ **reversible** if P and Σ_{UFD} are reversible (Definition III.10). We say $P \subseteq \Sigma_{\text{UID}}$ is **trivial** if we have $P = \{\tau\}$ for some $\tau \in \Sigma_{\text{UID}}$ such that $\tau \not\rightarrowtail \tau$. A partition is **manageable** if it is ordered and all of its classes are either reversible or trivial.

As we will show in Section VII.3, we can always construct a manageable partition of Σ_{UID}:

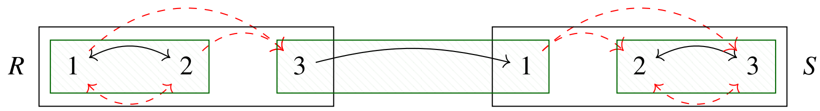

Fig. 6: Manageable partition (see Example VII.5)

PROPOSITION VII.4. *Any conjunction Σ_{UID} of UIDs closed under finite implication has a manageable partition.*

Example VII.5. Consider two ternary relations R and S. Consider the UIDs $\tau_R : R^1 \subseteq R^2$, $\tau_S : S^2 \subseteq S^3$, $\tau_{RS} : R^3 \subseteq S^1$, and the UFDs $\phi_R : R^1 \to R^2$, $\phi_S : S^2 \to S^3$, $\phi'_R : R^3 \to R^1$, and $\phi'_S : S^3 \to S^1$. The UIDs τ_R^{-1} and τ_S^{-1}, and the UFDs ϕ_R^{-1}, ϕ_S^{-1}, and $R^3 \to R^2$, $S^2 \to S^1$, are finitely implied. The two relations R and S are illustrated in Figure 6, where UIDs as drawn as solid black edges, and UFDs as dashed red edges that are *reversed* (this follows Definition VII.10).

A manageable partition of the UIDs of the finite closure is $(\{\tau_R, \tau_R^{-1}\}, \{\tau_{RS}\}, \{\tau_S, \tau_S^{-1}\})$, where the first and third classes are reversible and the second is trivial. The classes of the partition are drawn as green hatched rectangles in Figure 6; they are intuitively related to a topological sort of the graph of the black and red edges (see Definition VII.12).

VII.2. Using Manageable Partitions

Fix the instance I_0 and the finitely closed constraints Σ_{U} formed of UIDs Σ_{UID} and UFDs Σ_{UFD}. To prove Theorem VII.1, starting with the initial aligned superinstance $J_0 = (I_0, \mathrm{id})$ of I_0 (for Σ_{UID} and Σ_{UFD}), we first note that the Fact-Saturated Solutions Lemma (Lemma VI.11) does not use assumption reversible. Hence, we apply it (with Σ_{UID}) to obtain from I_0 an aligned fact-saturated superinstance J_1 of I_0 (for Σ_{UFD} and Σ_{UID}). This is the **saturation process**.

The goal is now to apply a **completion process** to satisfy Σ_{UID}, which we formalize as the following proposition. Recall the definition of thrifty sequences (Definition VI.26). We refine the definition below.

Definition VII.6. We define a **preserving fact-thrifty sequence** L (for UIDs Σ_{UID} and UFDs Σ_{UFD}) on any *fact-saturated aligned superinstance* J of I_0 in the following inductive way, with its **output** $L(J)$ also being a fact-saturated aligned superinstance:

— The empty list $L = ()$ is a preserving fact-thrifty sequence, with output $L(J) := J$.
— Let L be the concatenation of a preserving fact-thrifty sequence L' and a triple $t = (a, \tau, \boldsymbol{b})$. Let $J' := L'(J)$ be the output of L'. We call L **preserving**, if one of the following holds:
  — t is a fresh fact-thrifty chase step. In this case, by the Fresh Fact-Thrifty Preservation Lemma, J' is indeed a fact-saturated aligned superinstance of I_0.
  — J' is k-essential for some subset $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ of Σ_{UID} such that $\tau \in \Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and $\Sigma_{\mathrm{UID}}^{\mathrm{rev}}$ and Σ_{UFD} are reversible. In this case, by the Fact-Thrifty Preservation Lemma, J is a fact-saturated aligned superinstance of I_0 (which is also k-essential for the same subset).

  In either case, the output of L is the aligned superinstance obtained as the result of applying the fact-thrifty chase step represented by t on J'.

We can now state our intended result, which implies Theorem VII.1:

PROPOSITION VII.7 (FACT-THRIFTY COMPLETION). *Let $\Sigma_{\mathrm{U}} = \Sigma_{\mathrm{UFD}} \sqcup \Sigma_{\mathrm{UID}}$ be finitely closed UFDs and UIDs, and let I_0 be an instance that satisfies UFDs. For any fact-saturated aligned superinstance J of I_0 for Σ_{U}, there is a preserving fact-thrifty sequence L such that $L(J)$ satisfies Σ_{UID}.*

We prove Proposition VII.7, and from it Theorem VII.1, in the rest of the subsection. We construct a manageable partition $\boldsymbol{P} = (P_1, \ldots, P_n)$ of Σ_{UID} using Proposition VII.4. Now, for $1 \leq i \leq n$, we use fact-thrifty chase steps by UIDs of P_i to extend the fact-saturated aligned superinstance J_i to a larger one J_{i+1} that satisfies P_i.

The crucial point is that we can apply fact-thrifty chase steps to satisfy P_i without creating any new violations of P_j for $j < i$, and hence we can make progress following the partition. The reason for this is the following easy fact about thrifty chase steps:

LEMMA VII.8. *Let J be an aligned superinstance of I_0 and J' be the result of applying a thrifty chase step on J for a UID τ of Σ_{UID}. Assume that a UID τ' of Σ_{UID} was satisfied by J but is not satisfied by J'. Then $\tau \rightarrowtail \tau'$.*

PROOF. Fix J, J', $\tau : R^p \subseteq S^q$ and τ'. As chase steps add a single fact, the only new UID violations in J' relative to I are on elements in the newly created fact $F_{\text{n}} = S(\boldsymbol{b})$, As Σ_{UID} is transitively closed, F_{n} can introduce no new violation on the exported element b_q. Now, as thrifty chase steps always reuse existing elements at non-dangerous positions, we know that if $S^r \in \mathsf{NDng}(S^q)$ then no new UID can be applicable to b_r. Hence, if a new UID is applicable to b_r for $S^r \in \mathsf{Pos}(S)$, then necessarily $S^r \in \mathsf{Dng}(S^q)$. By definition of dangerous positions, the UFD $S^r \rightarrow S^q$ is then in Σ_{UFD}, and we have $S^r \neq S^q$. Hence, writing $\tau' : S^r \subseteq T^r$, we see that $\tau \rightarrowtail \tau'$. □

The lemma justifies our definition of ordered partition, since it will allow us to do an inductive argument to prove Proposition VII.7. Using the fact that $\boldsymbol{P}$ is ordered ensures that we can indeed apply fact-thrifty chase steps to satisfy each P_i individually, dealing with them in the order of the partition.

Thus, to prove Proposition VII.7, consider each class P_i in order. As $\boldsymbol{P}$ is manageable, there are two cases: either P_i is trivial or it is reversible.

First, if P_i is trivial, it can simply be satisfied by a preserving fact-thrifty sequence L_i of fresh fact-thrifty chase steps using the one UID of P_i, as follows immediately from Lemma VII.8.

LEMMA VII.9. *For any trivial class $\{\tau\}$, performing one chase round on an aligned fact-saturated superinstance J of I_0 by fresh fact-thrifty chase steps for τ yields an aligned superinstance J' of I_0 that satisfies τ.*

PROOF. Fix J, J' and τ. All violations of τ in J have been satisfied in J' by definition of J', so we only have to show that no new violations of τ were introduced in J'. But by Lemma VII.8, as $\tau \not\rightarrowtail \tau$, each fresh fact-thrifty chase step cannot introduce such a violation, hence there is no new violation of τ in J'. Hence, $J' \models \tau$. □

Second, returning to the proof of Proposition VII.7, the interesting case is that of a *reversible* P_i, for which we have done the work of the last three sections. We satisfy a reversible P_i by a preserving fact-thrifty sequence L_i obtained using the Reversible Fact-Thrifty Completion Proposition (Proposition VI.25). Indeed, J_i is a fact-saturated aligned superinstance of I_0 for Σ_{UFD} and Σ_{UID}, and by definition of P_i being reversible, letting $\Sigma_{\text{UID}}^{\text{rev}} := P_i$, the constraints Σ_{UFD} and $\Sigma_{\text{UID}}^{\text{rev}}$ are reversible. By the Reversible Fact-Thrifty Completion Proposition, we can thus construct a fact-thrifty sequence L_i (by UIDs of $\Sigma_{\text{UID}}^{\text{rev}}$) such that $J_{i+1} := L_i(J_i)$ is a fact-saturated aligned superinstance of I_0 for Σ_{UFD} and Σ_{UID} that satisfies P_i. Further, from the proof, it is clear that L_i is preserving.

Hence, in either of the two cases, we construct a preserving fact-thrifty sequence L_i and $J_{i+1} := L_i(J_i)$ satisfies P_i. Further, as L_i only performs fact-thrifty chase steps by UIDs of P_i, J_{i+1} actually satisfies $\bigcup_{j \leq i} P_j$, thanks to Lemma VII.8.

The concatenation of the preserving fact-thrifty sequences L_i for each P_i is thus a preserving fact-thrifty sequence L whose final result $L(J) = J_{n+1}$ is thus an aligned superinstance of I_0 that satisfies Σ_{UID}, which proves the Fact-Thrifty Completion Proposition. As an aligned superinstance, J_{n+1} is also finite, satisfies Σ_{UFD}, and is k-sound for ACQ; so it is k-universal for Σ_{U} and ACQ. This concludes the proof of Theorem VII.1.

VII.3. Building Manageable Partitions

The only missing part is to show how manageable partitions are constructed (Proposition VII.4), which we show in this subsection. We will construct the manageable partition using a *constraint graph* defined from the dependencies, inspired by the multigraph used in the proof of Theorem II.1 in [Cosmadakis et al. 1990].

Definition VII.10. Given a set Σ_{U} of finitely closed UIDs and UFDs on signature σ, the **constraint graph** $G(\Sigma_{\mathrm{U}})$ is the directed graph with vertex set $\mathsf{Pos}(\sigma)$ and with the following edges:

— For each UID $R^p \subseteq S^q$ in Σ_{U}, an edge from R^p to S^q
— For each UFD $R^a \rightarrow R^b$ in Σ_{U}, an edge from R^b to R^a.

As we forbid trivial UIDs and UFDs, $G(\Sigma_{\mathrm{U}})$ has no self-loop, but it may contain both the edge (R^p, S^q) and (S^q, R^p). However, we do not represent multiple edges in $G(\Sigma_{\mathrm{U}})$: for instance, if the UID $R^a \subseteq R^b$ and the UFD $R^b \rightarrow R^a$ hold in $G(\Sigma_{\mathrm{U}})$, we only create a single copy of the edge (R^a, R^b).

Hence, fix the finitely closed UIDs and UFDs $\Sigma_{\mathrm{U}} := \Sigma_{\mathrm{UID}} \wedge \Sigma_{\mathrm{UFD}}$, and construct the graph $G(\Sigma_{\mathrm{U}})$. As observed by [Cosmadakis et al. 1990], the graph $G(\Sigma_{\mathrm{U}})$ has the following property, which will be needed to show that classes are reversible:

LEMMA VII.11. *For any edge e occurring in a cycle in $G(\Sigma_{\mathrm{U}})$, for any dependency τ which caused the creation of e, the reverse τ^{-1} of τ is in Σ_{U}.*

PROOF. Let e_1 be the edge, and $e_1, \ldots, e_n$ be the cycle (the first vertex of e_1 is the second vertex of e_n), and let τ be the dependency. Consider a cycle of dependencies $\tau_1, \ldots, \tau_n$, with $\tau_1 = \tau$, such that each τ_i caused the creation of edge e_i in $G(\Sigma_{\mathrm{U}})$. We must show that the reverse τ^{-1} of τ is in Σ_{U}.

If all the τ_i are UIDs, then, as Σ_{UID} is closed under the transitivity rule, we apply it to $\tau_2, \ldots, \tau_n$ and deduce that τ_1^{-1} is in Σ_{UID}. Likewise, if all the τ_i are UFDs, then we proceed in the same way because Σ_{UFD} is closed under the transitivity rule.

If the τ_i are of alternating types (alternatively UIDs and UFDs), then, recalling that Σ_{U} is closed under the *cycle rule* (see Section II.2) we deduce that τ_i^{-1} is in Σ_{U} for all i.

In the general case, consider the maximal subsequence $\tau_j, \ldots, \tau_n, \tau_1, \ldots, \tau_i$ ($i < j$) of consecutive dependencies in the cycle that includes τ and where all dependencies are of the same type. Let τ_{m} be the result of combining these dependencies by the transitivity rule, and consider the cycle $\tau_{\mathrm{m}}, \tau_{i+1}, \ldots, \tau_n, \tau_1, \ldots, \tau_{j-1}$. Collapsing all other consecutive sequences of dependencies to a single dependency using the transitivity rule, and applying the cycle rule as in the previous case, we deduce that τ_{m}^{-1} is in Σ_{U}. Hence, the cycle $\tau_j, \ldots, \tau_n, \tau_1, \ldots, \tau_i, \tau_{\mathrm{m}}^{-1}$ is a cycle of dependencies of the same type as τ, and it includes τ, so we conclude as in the first two cases that τ^{-1} is in Σ_{U}.

Hence, in all cases τ^{-1} is in Σ_{U}. This concludes the proof. $\square$

Compute the strongly connected components of $G(\Sigma_{\mathrm{U}})$, ordered following a topological sort: we label them $V_1, \ldots, V_n$. The order of the V_i guarantees that there are no edges in $G(\Sigma_{\mathrm{U}})$ from V_i to V_j unless $i \leq j$.

We will build each class of the manageable partition, either as the set of UIDs within the positions of an SCC (a *reversible* class), or as a singleton UID going from a class V_i to a class V_j with $j > i$ (a *trivial* class). Formally:

Definition VII.12. The topological sort of the SCCs of $G(\Sigma_{\mathrm{U}})$, written $V_1, \ldots, V_n$, defines a partition $\boldsymbol{P}$ of the UIDs of Σ_{UID}, in the following manner. For each V_i, if there are any non-trivial UIDs of the form $R^p \subseteq S^q$ with $R^p, S^q \in V_i$, create a class of UIDs (the *main* class) containing all of them. Then, for each UID of the form $R^p \subseteq S^q$ with $R^p \in V_i$ and $S^q \in V_j$ with $j > i$, create a singleton class of UIDs containing exactly that UID (a *satellite* class). The partition $\boldsymbol{P}$ is obtained by taking the concatenation, for i from 1 to n, of the main class of V_i (if it exists) and then all satellite classes of V_i (if any) in an arbitrary order.

Remember that, while the constraint graph reflects both the UIDs and the UFDs, the partition $\boldsymbol{P}$ that we define is a partition of Σ_{UID}, that is, a partition of UIDs, and does not contain UFDs. We first show that $\boldsymbol{P}$ is indeed a partition, and then that it is an ordered partition.

LEMMA VII.13. *$\boldsymbol{P}$ is indeed a partition of Σ_{UID}.*

PROOF. As the SCCs of $G(\Sigma_{\mathrm{U}})$ partition the vertex set of $G(\Sigma_{\mathrm{U}})$, it is clear by construction that any UID occurs in at most a single class of the partition, which must be a class for the SCC of its first position, and either a satellite class or the main class depending on the SCC of its second position.

Conversely, each UID τ is reflected in some class of the partition, for the SCC V_i of its first position: either the second position of τ is also in V_i, so τ is in the main class for V_i; or the second position of τ is in an SCC V_j with $i \neq j$, in which case $i < j$ by definition of a topological sort, and τ is in some satellite class for V_i. Hence, $\boldsymbol{P}$ is indeed a partition of Σ_{UID}. □

LEMMA VII.14. *$\boldsymbol{P}$ is an ordered partition.*

PROOF. Assume by way of contradiction that there are two classes P_i and P_j and $\tau \in P_i$, $\tau' \in P_j$, such that $\tau \rightarrowtail \tau'$ but $i > j$. Let V_p and V_q be the SCCs in which P_i and P_j were created. We must have $p \geq q$, so there are two possibilities.

First, if $p = q$, then the first positions of τ and τ' must both be in $V_p = V_q$, and as P_i is not the first class created for $V_p = V_q$, it must be a satellite class. Hence, the second position of τ is in another SCC, say V_r, with $r > p$. Now, as $\tau \rightarrowtail \tau'$, there is a UFD from the first position of τ' to the second position of τ, which implies that there is an edge from V_r to V_p in $G(\Sigma_{\mathrm{U}})$. As $r > p$, this contradicts the fact that the SCCs are ordered following a topological sort.

Second, if we have $p > q$, then again the first position of τ must be in V_p, and the first position of τ' is in V_q. Let V_r be the SCC of the second position of τ. As $\tau \rightarrowtail \tau'$, the UFD from the first position of τ' to the second position of τ witnesses that there is an edge in $G(\Sigma_{\mathrm{U}})$ from V_r to V_q. Hence, we must have $r \leq q$. But τ witnesses that there is an edge from p to r in $G(\Sigma_{\mathrm{U}})$, so that we must have $p \leq r$. Hence, $p \leq q$, but we had assumed $p > q$, a contradiction. □

We now show that $\boldsymbol{P}$ is manageable, by considering each class and showing that it is either trivial or that it is reversible:

LEMMA VII.15. *Each satellite class in $\boldsymbol{P}$ is trivial.*

PROOF. Each satellite class consists by construction of a singleton dependency $\tau = R^p \subseteq S^q$, implying the existence of an edge in the constraint graph $G(\Sigma_{\mathrm{U}})$ from R^p to S^q. Assume by way of contradiction that $\tau \rightarrowtail \tau$. This implies that $R^p \rightarrow S^q$ holds in Σ_{UFD}, so there is an edge in $G(\Sigma_{\mathrm{U}})$ from S^q to R^p. Hence, $\{R^p, S^q\}$ is strongly connected, so R^p and S^q belong to the same SCC, which contradicts the definition of a satellite class. □

LEMMA VII.16. *Each main class in the partition is reversible.*

PROOF. Let P_i be the class and V_i be the corresponding SCC. We first show that P_i is transitively closed. Consider two UIDs τ and τ' of P_i that would be combined by the transitivity rule to the UID τ''. As Σ_{UID} is transitively closed, we have $\tau'' \in \Sigma_{\mathrm{UID}}$. Now, if both τ and τ' have both positions in V_i, then so does τ'', so we also have $\tau'' \in P_i$.

Second, to see that every UID τ in P_i is reversible, consider a UID $\tau : R^p \subseteq S^q$ of P_i, with $R^p, S^q \in V_i$. We forbid trivial UIDs, so $R^p \neq S^q$. As V_i is strongly connected, consider a directed path π of edges of $G(\Sigma_{\mathrm{U}})$ from S^q to R^p. Combining π with the edge created in $G(\Sigma_{\mathrm{U}})$ for the UID τ, we deduce the existence of a cycle in $G(\Sigma_{\mathrm{U}})$. Hence, by Lemma VII.11, the UID τ^{-1} holds in Σ_{UID}, and it also has both positions in V_i, so τ^{-1} is in P_i.

Third, we prove the claim about UFDs. Consider a UFD $\phi : R^p \rightarrow R^q$ of Σ_{UFD}, with $R^p \neq R^q$. Assume that R^p and R^q occur in a UID of P_i; this implies that $R^p, R^q \in V_i$. By the same reasoning as before, we find a cycle in $G(\Sigma_{\mathrm{U}})$ that includes the edge that corresponds to ϕ, and deduce that ϕ^{-1} holds in Σ_{UFD}. □

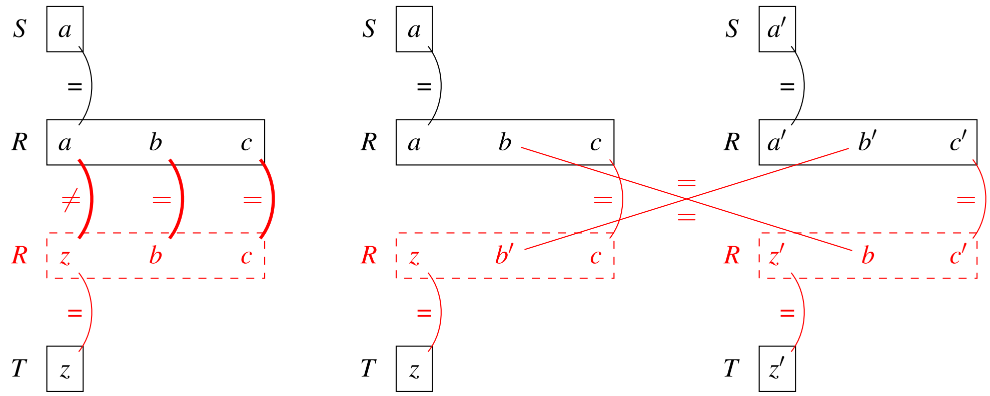

(a) Example VIII.2: an FD violation (b) Example VIII.15: a way to avoid the FD violation

Fig. 7: Example of a higher-arity FD violation in our process, and the proposed solution

Hence, $\boldsymbol{P}$ is an ordered partition of Σ_{UID} where each class is either reversible or trivial, i.e., it is a manageable partition. This concludes the proof of Proposition VII.4.

VIII. HIGHER-ARITY FDs

The goal of this section is to generalize our results to functional dependencies of arbitrary arity:

THEOREM VIII.1. *Finitely-closed UIDs and FDs have finite universal models for ACQs.*

We fix the finitely-closed constraints $\Sigma := \Sigma_{\text{FD}} \wedge \Sigma_{\text{UID}}$, consisting of arbitrary-arity FDs Σ_{FD} and UIDs Σ_{UID}. We denote by Σ_{UFD} the *unary* FDs among Σ_{FD}, and write $\Sigma_{\text{U}} := \Sigma_{\text{UFD}} \wedge \Sigma_{\text{UID}}$. From the definition of the finite closure (Section II.2), it is clear that Σ_{U} is finitely closed as well, so the construction of the previous sections applies to Σ_{U}.

The problem to address in this section is that our completion process to satisfy Σ_{UID} was defined with fact-thrifty chase steps. These chase steps may reuse elements from the same facts at the same positions multiple times. This may violate Σ_{FD}, and it is in fact the only point where we do so in the construction.

Example VIII.2. For simplicity, we work with instances rather than aligned superinstances. Consider $I_0 := \{S(a), T(z)\}$, the UIDs $\tau : S^1 \subseteq R^1$ and $\tau' : T^1 \subseteq R^1$ for a 3-ary relation R, and the FD $\phi : R^2R^3 \to R^1$. Consider $I := I_0 \sqcup \{R(a,b,c)\}$ obtained by one chase step of τ on $S(a)$. Figure 7a represents I in solid black, using edges to highlight equalities between elements.

We can perform a fact-thrifty chase step of τ' on z to create $R(z,b,c)$, reusing (b,c) at $\mathsf{NDng}(R^1) = \{R^2, R^3\}$; this is illustrated in dashed red in Figure 7a. However, the two R-facts would then be a violation of ϕ, as shown by the patterns of equalities and inequalities illustrated as thick red edges.

The goal of this section is to define a new version of thrifty chase steps that preserves Σ_{FD} rather than just Σ_{UFD}; we call them *envelope-thrifty chase steps*. We first describe the new saturation process designed for them, which is much more complex because we need to saturate *sufficiently* with respect to the completion process that we do next. To saturate, we use a separate combinatorial result, of possible independent interest: Theorem VIII.11, proved in Section VIII.3. Second, we redefine the completion process of the previous section for this new notion of chase step, and use this new completion process to prove Theorem VIII.1.

VIII.1. Envelopes and Envelope-Saturation

We start by defining our new notion of saturated instances. Recall the notions of fact classes (Definition VI.8) and thrifty chase steps (Definition V.13). When a fact-thrifty chase step creates a fact F_{n} whose chase witness F_{w} has fact class $(R^p, \boldsymbol{C})$, we need elements to reuse in F_{n} at positions of $\mathsf{NDng}(R^p)$, which need to already occur at the positions where we reuse them. Further, the reused elements must have sim-images of the right class.

Fact-thrifty chase steps reuse a tuple of elements from one fact F_{r}, and thus apply to *fact-saturated instances*, where each fact class D which is achieved in the chase is also achieved by some fact (recall Definitions VI.8 and VI.10). Our new notion of envelope-thrifty chase steps will consider *multiple* tuples that achieve each class D, that we call an *envelope* for D; with the difference, however, that not all *tuples* need to actually occur in an achiever fact in the instance, though each *individual* element needs to occur in some achiever fact. Formally:

Definition VIII.3. Consider $D = (R^p, \boldsymbol{C})$ in $\mathsf{AFactCl}$, and write $O := \mathsf{NDng}(R^p)$. Remember that O is then non-empty. An **envelope** E for D and for an aligned superinstance $J = (I, \mathsf{sim})$ of I_0 is a non-empty set of $|O|$-tuples indexed by O, with domain $\mathrm{dom}(I)$, such that:

(1) for every FD $\phi : R^L \to R^r$ of Σ_{FD} with $R^L \subseteq O$ and $R^r \in O$, for any $\boldsymbol{t}, \boldsymbol{t}' \in E$, $\pi_{R^L}(\boldsymbol{t}) = \pi_{R^L}(\boldsymbol{t}')$ implies $t_r = t'_r$;
(2) for every FD $\phi : R^L \to R^r$ of Σ_{FD} with $R^L \subseteq O$ and $R^r \notin O$, for all $\boldsymbol{t}, \boldsymbol{t}' \in E$, $\pi_{R^L}(\boldsymbol{t}) = \pi_{R^L}(\boldsymbol{t}')$ implies $\boldsymbol{t} = \boldsymbol{t}'$;
(3) for every $a \in \mathrm{dom}(E)$, there is exactly one position $R^q \in O$ such that $a \in \pi_{R^q}(E)$, and then we also have $a \in \pi_{R^q}(J)$;
(4) for any fact $F = R(\boldsymbol{a})$ of J and $R^q \in O$, if $a_q \in \pi_{R^q}(E)$, then F achieves D in J and $\pi_O(\boldsymbol{a}) \in E$.

Intuitively, the tuples in the envelope E satisfy the FDs of Σ_{FD} within $\mathsf{NDng}(R^p)$ (condition 1), and never overlap on positions that determine a position out of $\mathsf{NDng}(R^p)$ (condition 2). Further, their elements already occur at the position where they will be reused, and we require for simplicity that there is exactly one such position (condition 3). Last, the elements have the right sim-image for the fact class D, and for simplicity, whenever a fact reuses an envelope element, we require that it reuses a whole envelope tuple (condition 4).

We then extend this definition across all achieved fact classes in the natural way:

Definition VIII.4. A **global envelope** $\mathcal{E}$ for an aligned superinstance $J = (I, \mathsf{sim})$ of I_0 is a mapping from each $D \in \mathsf{AFactCl}$ to an envelope $\mathcal{E}(D)$ for D and J, such that the envelopes have pairwise disjoint domains.

It is not difficult to see that an aligned superinstance with a global envelope must be fact-saturated, as for each $D \in \mathsf{AFactCl}$, the envelope $\mathcal{E}(D)$ is a non-empty set of non-empty tuples, and any element of this tuple must occur in a fact that achieves D, by conditions 3 and 4. However, the point of envelopes is that they can contain more than a single tuple, so we have multiple choices of elements to reuse.

For some fact classes $(R^p, \boldsymbol{C})$ it is not useful for envelopes to contain more than one tuple. This is the case if the position R^p is *safe*, meaning that no FD from positions in $O := \mathsf{NDng}(R^p)$ determines a position outside of O. (Notice that by definition of $\mathsf{NDng}(R^p)$, such an FD could never be a UFD.) Formally:

Definition VIII.5. We call $R^p \in \mathsf{Pos}(\sigma)$ **safe** if there is no FD $R^L \to R^r$ in Σ_{FD} with $R^L \subseteq \mathsf{NDng}(R^p)$ and $R^r \notin \mathsf{NDng}(R^p)$. Otherwise, R^p is **unsafe**.

We accordingly call a fact class $(R^p, \boldsymbol{C}) \in \mathsf{AFactCl}$ **safe** or **unsafe** depending on R^p. Observe that the second condition of Definition VIII.3 is trivial for envelopes on safe fact classes.

It is not hard to see that when we apply a fact-thrifty (or even relation-thrifty) chase step, and the exported position of the new fact is safe, then the problem illustrated by Example VIII.2 cannot arise. In fact, one could show that fact-thrifty or relation-thrifty chase steps cannot introduce FD violations in this case. Because of this, in envelopes for *safe* fact classes, we do not need more than one tuple, which we can reuse as we did with fact-thrifty chase steps.

For unsafe fact classes, however, it will be important to have more tuples, and to *never reuse the same tuple twice*. This motivates our definition of the *remaining tuples* of an envelope, depending on whether the fact class is safe or not; and the definition of *envelope-saturation*, which depends on the number of remaining tuples:

Definition VIII.6. Letting E be an envelope for $(R^p, \boldsymbol{C}) \in \mathsf{AFactCl}$ and J be an aligned superinstance, the **remaining tuples** of E are $E \backslash \pi_{\mathsf{NDng}(R^p)}(J)$ if $(R^p, \boldsymbol{C})$ is unsafe, and just E if it is safe.

We call J n**-envelope-saturated** if it has a global envelope $\mathcal{E}$ such that $\mathcal{E}(D)$ has $\geq n$ remaining tuples for all unsafe $D \in \mathsf{AFactCl}$. J is **envelope-saturated** if it is n-envelope-saturated for some $n > 0$.

In the rest of the subsection, inspired by the fact-saturation lemma, we will show that we can construct envelope-saturated solutions. However, there are some complications when doing so. First, we must show that we can construct *sufficiently* envelope-saturated solutions, i.e., instances with sufficiently many remaining tuples. To do this, we will need multiple copies of the chase, which explains the technical switch from I_0 to I'_0 in the statement of the next result. Second, for reasons that will become clear later in this section, we need to ensure that the envelopes are large *relative to the resulting instance size*. This makes the result substantially harder to show.

PROPOSITION VIII.7 (SUFFICIENTLY ENVELOPE-SATURATED SOLUTIONS). *For any $K \in \mathbb{N}$ and instance I_0, we can construct an instance I'_0 formed of disjoint copies of I_0, and an aligned superinstance J of I'_0 that satisfies Σ_{FD} and is $(K \cdot |J|)$-envelope-saturated.*

We prove the proposition in the rest of the subsection. It is not hard to see that I'_0 and J can be constructed separately for each fact class in $\mathsf{AFactCl}$, and that this is difficult only for unsafe classes. In other words, the crux of the matter is to prove the following:

LEMMA VIII.8 (SINGLE ENVELOPE). *For any unsafe class D in $\mathsf{AFactCl}$, instance I_0 and constant factor $K \in \mathbb{N}$, there exists $N_0 \in \mathbb{N}$ such that, for any $N \geq N_0$, we can construct an instance I'_0 formed of disjoint copies of I_0, and an aligned superinstance $J = (I, \mathsf{sim})$ of I'_0 that satisfies Σ_{FD}, with an envelope E for D of size $\geq K \cdot N$, such that $|J| \leq N$.*

Indeed, let us prove Proposition VIII.7 with this lemma, and we will prove the lemma afterwards:

PROOF OF PROPOSITION VIII.7. Fix the constant $K \in \mathbb{N}$ and the initial instance I_0, and let us build I'_0 and the aligned superinstance $J = (I, \mathsf{sim})$ of I'_0 that has a global envelope $\mathcal{E}$. As $\mathsf{AFactCl}$ is finite, we build one J_D per $D \in \mathsf{AFactCl}$ with an envelope E_D for the class D, and we will define $J := \bigsqcup_{D \in \mathsf{AFactCl}} J_D$ and define $\mathcal{E}$ by $\mathcal{E}(D) := E_D$ for all $D \in \mathsf{AFactCl}$. When $D = (R^p, \boldsymbol{C})$ is safe, we proceed as in the proof of the Fact-Saturated Solutions Lemma: take a single copy J_D of the truncated chase to achieve the class D, and take as the only fact of the envelope E_D the projection to $\mathsf{NDng}(R^p)$ of an achiever of D in J_D. When D is unsafe, we use the Single Envelope Lemma to obtain J_D and the envelope E_D. As $\mathsf{AFactCl}$ is finite and its size does not depend on I_0, we can ensure that that $|E_D| \geq (K+1) \cdot |J|$ for all unsafe $D \in \mathsf{AFactCl}$ by using the Single Envelope Lemma with $K' := (K+1) \cdot |\mathsf{AFactCl}|$, and taking $N \in \mathbb{N}$ which is larger than the largest N_0 of that lemma across all $D \in \mathsf{AFactCl}$. Indeed, the resulting model J then ensures that $|J| \leq |\mathsf{AFactCl}| \cdot N$ and $|E_D| \geq (K+1) \cdot |\mathsf{AFactCl}| \cdot N$.

We now check that the resulting J and E satisfy the conditions. Each J_D is an aligned superinstance of an instance $(I'_0)_D$ which is formed of disjoint copies of I_0 (for unsafe classes) or which is exactly I_0 (for safe classes), so J is an aligned superinstance of $I'_0 := \bigsqcup_{D \in \mathsf{AFactCl}} (I'_0)_D$, so I'_0 is also a union of disjoint copies of I_0. There are no violations of Σ_{FD} in J because there are none in any of the J_D. The disjointness of domains of envelopes in the global envelope $\mathcal{E}$ is because the J_D are disjoint. It is easy to see that J is $(K \cdot |I|)$-envelope-saturated, because $|\mathcal{E}(D)| \geq (K+1) \cdot |I|$ for all unsafe $D \in \mathsf{AFactCl}$, so the number of remaining facts of each envelope for an unsafe class is $\geq K \cdot |I|$ because every fact of I eliminates at most one fact in each envelope. Hence, the proposition is proven. □

So the only thing left to do is to prove the Single Envelope Lemma. Let us accordingly fix the unsafe class $D = (R^p, \boldsymbol{C})$ in $\mathsf{AFactCl}$. We will need to study more precisely the FDs implied by the definition of an envelope for D (Definition VIII.3). We first introduce notation for them:

Definition VIII.9. Given a set Σ_{FD} of FDs on a relation R and $O \subseteq \mathsf{Pos}(R)$, the **FD projection** Σ_{FD}^{O} of Σ_{FD} to O consists of the following FDs, which we close under implication:

(1) the FDs $R^L \to R^r$ of Σ_{FD} such that $R^L \subseteq O$ and $R^r \in O$;
(2) for every FD $R^L \to R^r$ of Σ_{FD} where $R^L \subseteq O$ and $R^r \notin O$, the key dependency $R^L \to O$.

We will need to show that, as R^p is unsafe, Σ_{FD}^{O} cannot have a *unary key* in O, namely, there cannot be $R^q \in O$ such that, for every $R^r \in O$, either $R^q = R^r$ or the UFD $R^q \to R^r$ holds in Σ_{FD}^{O}. We show the contrapositive of this statement:

LEMMA VIII.10. *For any $R^p \in \mathsf{Pos}(\sigma)$, letting $O := \mathsf{NDng}(R^p)$, if O has a unary key in Σ_{FD}^{O}, then R^p is safe.*

PROOF. Fix $R^p \in \mathsf{Pos}(\sigma)$ and let $O := \mathsf{NDng}(R^p)$. We first show that if O has a unary key $R^s \in O$ in the original FDs Σ_{FD}, then R^p is safe. Indeed, assume the existence of such an $R^s \in O$. Assume by way of contradiction that R^p is not safe, so there is an FD $R^L \to R^r$ in Σ_{FD} with $R^L \subseteq O$ and $R^r \notin O$. Then, as Σ_{FD} is closed under the transitivity rule, the UFD $\phi : R^s \to R^r$ is in Σ_{UFD}. Now, as $R^r \notin O$, either $R^r = R^p$ or $R^r \in \mathsf{Dng}(R^p)$; in both cases, ϕ witnesses that $R^s \in \mathsf{Dng}(R^p)$, but we had $R^s \in O$, a contradiction.

We must now show that if O has a unary key in O according to Σ_{FD}^{O}, then O has a unary key in O according to Σ_{FD}. It suffices to show that for any two positions $R^q, R^s \in O$, if the UFD $\phi' : R^q \to R^s$ holds in Σ_{FD}^{O} then it also does in Σ_{FD}. Hence, fix R^q in O, and consider the set S of positions in O that R^q determines according to Σ_{FD}^{O}. Let Φ be the FDs in the list given in Definition VIII.9, so that Σ_{FD}^{O} is the result of closing Φ under FD implication. We can compute S using the well-known "attribute closure algorithm" [Abiteboul et al. 1995], which starts at $S = \{R^q\}$ and iterates the following operation: whenever there is $\phi : R^L \to R^r$ such that $R^L \subseteq S$, add R^r to S.

Assume now that there is a position R^s in S such that $\phi' : R^q \to R^s$ does not hold in Σ_{FD}. This implies that, when computing S, we must have used some FD $R^L \to R^t$ from a key dependency κ in Φ, as they are the only FDs of Φ which are not in Σ_{FD}. The first time we did this, we had derived, using only FDs from Σ_{FD}, that $R^L \subseteq S$, so that, in Σ_{FD}, the key dependency $R^q \to R^L$ holds. Now, κ witnesses that there is an FD $R^L \to R^r$ in Σ_{FD} with $R^r \notin O$, so that, as Σ_{FD} is closed under implication, we deduce that $R^q \to R^r$ holds in Σ_{FD} with R^q in O and $R^r \notin O$. As before, this contradicts the definition of $O := \mathsf{NDng}(R^p)$. So indeed, there is no such R^s in S.

Hence, if O has a unary key in O according to Σ_{FD}^{O}, then it also does according to Σ_{FD}, and then, by the reasoning of the first paragraph, R^p is safe, which is the desired claim. $\square$

We now know that O has no unary key in Σ_{FD}^{O}. This allows us to introduce the crucial tool needed to prove the Sufficiently Envelope-Saturated Solutions Proposition. It is the following independent result, which is proved separately in Section VIII.3 using a combinatorial construction.

THEOREM VIII.11 (DENSE INTERPRETATIONS). *For any set Σ_{FD} of FDs over a relation R with no unary key, for all $K \in \mathbb{N}$, there exists $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$, we can construct a non-empty instance I of R that satisfies Σ_{FD} and such that $|\mathrm{dom}(I)| \leq N$ and $|I| \geq K \cdot N$.*

Further, we can impose a ***disjointness condition*** *on the result I: we can ensure that for all $a \in \mathrm{dom}(I)$, there exists exactly one $R^p \in \mathsf{Pos}(R)$ such that $a \in \pi_{R^p}(I)$.*

We can now prove the Single Envelope Lemma and conclude the subsection. Choose a fact $F_{\mathrm{ach}} = R(\boldsymbol{b})$ of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \backslash I_0$ that achieves the fact class D, and let I_1 be obtained from I_0 by applying UID chase steps on I_0 to obtain a finite truncation of $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ that includes F_{ach} but no child fact of F_{ach}. Consider the aligned superinstance $J_1 = (I_1, \mathsf{sim}_1)$ of I_0, where sim_1 is the identity.

Remember that we wrote $D = (R^p, \boldsymbol{C})$, and $O = \mathsf{NDng}(R^p)$, which is non-empty. Define a $|O|$-ary relation $R_{|O}$ (with positions indexed by O for convenience), define Σ_{FD}^{O} as in Definition VIII.9, and consider Σ_{FD}^{O} as FDs on $R_{|O}$. Because D is unsafe, by Lemma VIII.10, $R_{|O}$ has no unary key in Σ_{FD}^{O}. Letting $K \in \mathbb{N}$ be our target constant for the Single Envelope Lemma, apply the Dense Interpretations Theorem (Theorem VIII.11) to $R_{|O}$ and Σ_{FD}^{O}, taking $K' := 2 \cdot K \cdot |J_1|$ as the constant. Define $N_0 \in \mathbb{N}$ for the Single Envelope Lemma as $2 \cdot \max(|J_1|, 1) \cdot \max(N_0', 1)$ where N_0' is obtained from the Dense Interpretations Theorem for K'. Letting $N' \in \mathbb{N}$ be our target size for the Single Envelope Lemma, using $N := \lfloor N' / |J_1| \rfloor$ as the target size for the Dense Interpretations Theorem (which is $\geq N_0'$), we can build an instance I_{dense} of $R_{|O}$ that satisfies Σ_{FD}^{O} and such that $|I_{\mathrm{dense}}| \geq N \cdot K'$ and $|\mathrm{dom}(I_{\mathrm{dense}})| \leq N$.

Let $I'_{\mathrm{dense}} \subseteq I_{\mathrm{dense}}$ be a subinstance of size exactly N of I_{dense} such that we have $\mathrm{dom}(I'_{\mathrm{dense}}) = \mathrm{dom}(I_{\mathrm{dense}})$, that is, such that each element of $\mathrm{dom}(I_{\mathrm{dense}})$ occurs in some fact of I'_{dense}: we can clearly construct I'_{dense} by picking, for each element of $\mathrm{dom}(I_{\mathrm{dense}})$, one fact of I_{dense} where it occurs, removing duplicate facts, and completing with other arbitrary facts of I_{dense} so the number of facts is exactly N. Number the facts of I'_{dense} as $F'_1, \ldots, F'_N$.

Let us now create $N-1$ disjoint copies of J_1, numbered J_2 to J_N. Let I_{pre} be the disjoint union of the underlying instances of the J_i, let I'_0 be formed of the N disjoint copies of I_0 in I_{pre}, and define a mapping $\mathsf{sim}_{\mathrm{pre}}$ from $\mathrm{dom}(I_{\mathrm{pre}})$ to $\mathsf{Chase}(I'_0, \Sigma_{\mathrm{UID}})$ following the sim_i in the expected way. It is clear that J_{pre} is an aligned superinstance of I'_0. For $1 \leq i \leq N$, we call $F_i = R(\boldsymbol{a^i})$ *the fact of I_i that corresponds to the achiever F_{ach} in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$*. In particular, for all $1 \leq i \leq N$, we have that $\mathsf{sim}(a^i_j) = b_j$ for all j, and a^i_p is the only element of F_i that also occurs in other facts of J_i, as J_i does not contain any descendent fact of F_i.

Intuitively, we will now identify elements in J_{pre} so that the restriction of the F_i to O is exactly the F'_i, and this will allow us to use the instance I_{dense} to define the envelope. Formally, as the a^i_j are pairwise distinct, we can define the function f that maps each a^i_j, for $1 \leq i \leq N$ and $R^j \in O$, to $\pi_{R^j}(F'_i)$. In other words, f is a surjective (but generally not injective) mapping, the domain of f is the projection to O of the F_i in I_i, the range of f is $\mathrm{dom}(I'_{\mathrm{dense}})$, and f maps each element of the projection to the corresponding element in F'_i. Extend f to a mapping f' with domain $\mathrm{dom}(I_{\mathrm{pre}})$ by setting $f'(a) := f(a)$ when a is in the domain of f, and $f'(a) := a$ otherwise. Now, let $I := f'(I_{\mathrm{pre}})$. In other words, I is I_{pre} except that elements in the projection to O of the facts F_i are renamed, and some are identified, so that the projection to O of $\{f'(F_i) \mid 1 \leq i \leq N\}$, seen as an instance of $R_{|O}$-facts, is exactly I'_{dense}. Because a^i_j occurs only in F_i for all $R^j \neq R^p$, nd $R^p \notin O$, this means that the elements identified by f' only occurred in the F_i in I_{pre}.

We now build $J = (I, \mathsf{sim})$ obtained by defining sim from $\mathsf{sim}_{\mathrm{pre}}$ as follows: if a is in the domain of f, then $\mathsf{sim}(a) := \mathsf{sim}_{\mathrm{pre}}(a')$ for any preimage of a' by f' (as we will see, the choice of preimage does not matter), and if a is not in the domain of f, then $\mathsf{sim}(a) := \mathsf{sim}_{\mathrm{pre}}(a)$ because a is then the only preimage of a by f'. We have now defined the instance I'_0 formed of disjoint copies of I_0 and the final J, and we define $E := I_{\mathrm{dense}}$. We must now show that J is indeed an aligned superinstance of I'_0, and that E is an envelope for I and D, and that they satisfy the required conditions.

We note that it is immediate that $J = (I, \mathsf{sim})$ is a superinstance of I'_0. Indeed, we have $I := f(I_{\mathrm{pre}})$, and I_{pre} was a superinstance of I'_0, so it suffices to note that $\mathrm{dom}(I'_0)$ is not in the domain of f: this is because the achiever F_{ach} is not a fact of I_0, so the domain of f, namely, the projection of the F_i on O, does not intersect $\mathrm{dom}(I'_0)$. Further, it is clear that J is finite and has $N \cdot |J_1|$ facts, because this is the case of J_{pre} by definition, and f' cannot have caused any facts of J_{pre} to be identified in J, because we have $R^p \notin O$, so the projection of each F_i to R^p is a different element which is mapped to itself by f'. Hence, we have $|J| = N \cdot |J_1| \leq N'$. Further, we have $|E| = |I_{\mathrm{dense}}| \geq N \cdot K' \geq \lfloor N' / |J_1| \rfloor \cdot 2 \cdot K \cdot |J_1|$, and as $N' \geq N_0 \geq 2 \cdot |J_1|$ we have $\lfloor N' / |J_1| \rfloor \geq (1/2) \cdot (N' / |J_1|)$. Hence, $|E| \geq K \cdot N'$, so we have achieved the required size bound.

We now show that J is indeed an aligned superinstance of I'_0. The technical conditions on sim are clearly respected, because they were respected on J_{pre}, because f' only identifies elements in

$I_{\text{pre}} \backslash I_0$, and because the identified elements occur at the same positions as their preimages so the directionality condition is respected.

We show that sim is a k-bounded simulation from J to $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ by showing the stronger claim that it is actually a k'-bounded simulation for all $k' \in \mathbb{N}$, which we show by induction on k'. The case of $k' = 0$ is trivial. The induction case is trivial for all facts except for the $f'(F_i)$, because the a^i_j only occurred in I_{pre} in the facts F_i, by our assumption that the F_i have no children in the I_i, and because the exported position of F_{ach} is $R^p \notin O$. Hence, consider a fact $F' = R(\boldsymbol{c})$ of I which is the image by f' of some fact F_i. Choose $1 \leq p \leq |R|$. We wish to show that there exists a fact $F'' = R(\boldsymbol{d})$ of $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ such that $\mathsf{sim}(c_p) = d_p$ and for all $1 \leq q \leq |R|$ we have $(I, c_q) \leq_{k'-1} (\mathsf{Chase}(I'_0, \Sigma_{\text{UID}}), d_q)$. Let $a^{i_0}_{j_0}$ be the preimage of c_p used to define $\mathsf{sim}(c_p)$; by the disjointness condition of the Dense Interpretations Theorem, we must have $j_0 = p$. Observe that $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ is formed of disjoint copies of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, so, recalling the definition of J'_{i_0}, consider the fact $F'' = R(\boldsymbol{d})$ of $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ corresponding to F_{i_0} in I. By definition, $\mathsf{sim}(c_p) = \mathsf{sim}(a^{i_0}_{j_0}) = d_p$.

We now show that for all $1 \leq q \leq |R|$ we have $(I, c_q) \leq_{k'-1} (\mathsf{Chase}(I'_0, \Sigma_{\text{UID}}), d_q)$. Fix $1 \leq q \leq |R|$. It suffices to show that $\mathsf{sim}(c_q) \simeq_{k'} d_q$, as we can then use the induction hypothesis to know that $(I, c_q) \leq_{k'-1} (\mathsf{Chase}(I'_0, \Sigma_{\text{UID}}), \mathsf{sim}(c_q))$, so that by transitivity $(I, c_q) \leq_{k'-1} (\mathsf{Chase}(I'_0, \Sigma_{\text{UID}}), d_q)$. Hence, we show that $\mathsf{sim}(c_q) \simeq_{k'} d_q$. Let $a^{i'_0}_{j'_0}$ be the preimage of c_q used to define $\mathsf{sim}(c_q)$. Again we must have $j'_0 = q$ by the disjointness condition, and, considering the fact $F''' = R(\boldsymbol{e})$ of $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ corresponding to $F_{i'_0}$ in I, we have $\mathsf{sim}(c_q) = e_q$. But as both F''' and F'' are copies in $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ of the same fact F_{ach} of $\mathsf{Chase}(I_0, \Sigma_{\text{UID}})$, it is indeed the case that $d_q \simeq_{k'} e_q$. Hence, $\mathsf{sim}(c_q) \simeq_{k'} d_q$, from which we conclude that F'' is a suitable witness fact for F'. By induction, we have shown that sim is indeed a k'-bounded simulation from J to $\mathsf{Chase}(I'_0, \Sigma_{\text{UID}})$ for any $k' \in \mathbb{N}$, so that it is in particular a k-bounded simulation.

We now show that J satisfies Σ_{FD}. For this, it will be convenient to define the *overlap* of two facts:

Definition VIII.12. The **overlap** $\mathsf{OVL}(F, F')$ between two facts $F = R(\boldsymbol{a})$ and $F' = R(\boldsymbol{b})$ of the same relation R in an instance I is the subset O of $\mathsf{Pos}(R)$ such that $a_s = b_s$ iff $R^s \in O$. If $|O| > 0$, we say that F and F' **overlap**.

As I_{pre} satisfies Σ_{FD} by the Unique Witness Property of the UID chase, any new violation of Σ_{FD} in I relative to I_{pre} must include some fact $F = f'(F'_{i_0})$, and some fact $F' \neq F$ that overlaps with F, so necessarily $F' = f'(F'_{i_1})$ for some i_1 by construction of I, and $\mathsf{OVL}(F, F') \subseteq O$. If $\mathsf{OVL}(F, F') = O$, then, by our definition of f and of the F'_i, this implies that $F'_{i_0} = F'_{i_1}$, a contradiction because $F \neq F'$. So the only case to consider is when $\mathsf{OVL}(F, F') \subsetneq O$, but we can also exclude this case:

LEMMA VIII.13. *Let I be an instance, Σ_{FD} be a conjunction of FDs, and $F \neq F'$ be two facts of I. Assume there is a position $R^p \in Pos(\sigma)$ such that, writing $O := NDng(R^p)$, we have $OVL(F, F') \subsetneq O$, and that $\{\pi_O(F), \pi_O(F')\}$ is not a violation of Σ^O_{FD}. Then $\{F, F'\}$ is not a violation of Σ_{FD}.*

PROOF. Assume by way of contradiction that F and F' violate an FD $\phi : R^L \to R^r$ of Σ_{FD}, which implies that $R^L \subseteq \mathsf{OVL}(F, F') \subseteq O$ and $R^r \notin \mathsf{OVL}(F, F')$. Now, if $R^r \in O$, then ϕ is in Σ^O_{FD}, so that $\pi_O(F)$ and $\pi_O(F')$ violate Σ^O_{FD}, a contradiction. Hence, $R^r \in \mathsf{Pos}(R) \backslash O$, and the key dependency $\kappa : R^L \to O$ is in Σ^O_{FD}, so that $\pi_O(F)$ and $\pi_O(F')$ must satisfy κ. Thus, because $R^L \subseteq \mathsf{OVL}(F, F')$, we must have $\mathsf{OVL}(F, F') = O$, which is a contradiction because we assumed $\mathsf{OVL}(F, F') \subsetneq O$. □

Now, by by definition of I'_{dense}, we know that I'_{dense} satisfies Σ^O_{FD}, so that $\{\pi_O(F), \pi_O(F')\}$ is not a violation of Σ^O_{FD}. Thus, we can conclude with Lemma VIII.13 that $\{F, F'\}$ is not a violation of Σ_{FD}, so that J satisfies Σ_{FD}. We have thus shown that J is an aligned superinstance of I'_0.

Last, we check that E is indeed an envelope for D and for J. Indeed, E satisfies Σ_{FD}^{O} by construction, so conditions 1 and 2 are respected. The first part of condition 3 is ensured by the disjointness condition, and its second part follows from our definition of I'_{dense} that ensures that any element in $\mathrm{dom}(E)$ occurs in a fact F'_i of I'_{dense}, hence occurs in $f'(F_i)$ in J. Last, condition 4 is true because the elements of $\mathrm{dom}(E)$ are only used in the $f'(F_i)$, and the sim-images of the $f'(F_i)$ are copies in $\mathsf{Chase}(I'_0, \Sigma_{\mathrm{UID}})$ of the same fact F_{ach} in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ that achieves D, so the F_i are all achievers of D; further, by definition, their projection to O is a tuple of E because it is a fact of I'_{dense}.

Hence, J is indeed an aligned superinstance of a disjoint union I'_0 of copies of I_0, J satisfies Σ_{FD}, $|J| \leq N'$, and J has an envelope E of size $K \cdot N'$ for D. This concludes the proof of the Single Envelope Lemma, and hence of the Sufficiently Envelope-Thrifty Solutions Proposition.

VIII.2. Envelope-Thrifty Chase Steps

We have shown that we can construct sufficiently envelope-saturated superinstances of the input instance. The point of this notion is to introduce *envelope-thrifty chase steps*, namely, thrifty chase steps that use remaining tuples from the envelope to fill the non-dangerous positions:

Definition VIII.14. **Envelope-thrifty chase steps** are thrifty chase steps (Definition V.13) which apply to envelope-saturated aligned superinstances. Following Definitions V.13 and VI.14, we write S^q for the exported position of the new fact F_{n}, we write $F_{\mathrm{w}} = S(\boldsymbol{b}')$ for the chase witness, and we let $D = (S^q, \boldsymbol{C}) \in \mathsf{AFactCl}$ be the fact class of F_{w}. Analogously to Definition V.13, we define an **envelope-thrifty** chase step as follows: if $\mathsf{NDng}(S^q)$ is non-empty, choose one remaining tuple $\boldsymbol{t}$ of $\mathcal{E}(D)$, and set $b_r := t_r$ for all $S^r \in \mathsf{NDng}(S^q)$.

We define a **fresh envelope-thrifty step** in the same way as a fresh fact-thrifty step: all elements at dangerous positions are fresh elements only occurring at that position.

Example VIII.15. Recall I_0, τ, τ' and ϕ from Example VIII.2. Now, consider $I'_0 := \{S(a), T(z), S(a'), S(z')\}$ formed of two copies of I_0, and $I' := I'_0 \sqcup \{R(a,b,c), R(a',b',c')\}$ obtained by two chase steps: this is illustrated in solid black in Figure 7b on page 40. The two facts $R(a,b,c)$ and $R(a',b',c')$ would achieve the same fact class D, so we can define $E(D) := \{(b,c), (b',c'), (b',c), (b,c')\}$.

We can now satisfy Σ_{UID} on I' without violating ϕ, with two envelope-thrifty chase steps that reuse the remaining tuples (b',c) and (b,c') of $E(D)$: the new facts and the pattern of equalities between them is illustrated in red in Figure 7b.

Recall that fact-thrifty chase steps apply to fact-saturated aligned superinstances (Lemma VI.15). Similarly, envelope-thrifty chase steps apply to envelope-saturated aligned superinstances:

LEMMA VIII.16 (ENVELOPE-THRIFTY APPLICABILITY). *For any envelope-saturated superinstance I of an instance I_0, UID $\tau : R^p \subseteq S^q$ and element $a \in \mathsf{Wants}(I,\tau)$, we can apply an envelope-thrifty chase step on a with τ to satisfy this violation.*

Further, for any new fact $S(\boldsymbol{e})$ that we can create by chasing on a with τ with a fact-thrifty chase step, we can instead apply an envelope-thrifty chase step on a with τ to create a fact $S(\boldsymbol{b})$ with $b_r = e_r$ for all $S^r \in \mathsf{Pos}(S) \backslash \mathsf{NDng}(S^r)$.

PROOF. For the first part of the claim, as in the proof of the Fact-Thrifty Applicability Lemma (Lemma VI.15), there is nothing to show unless $\mathsf{NDng}(S^q)$ is non-empty, and the fact class $D = (S^q, \boldsymbol{C})$ is then in $\mathsf{AFactCl}$, where $\boldsymbol{C}$ is the tuple of the $\simeq_k$-equivalence classes of the elements of the chase witness F_{w}. Hence, as J is envelope-saturated, it has some remaining tuple for the class D that we can use to define the non-dangerous positions of the new fact.

For the second part, again as in the proof of the Fact-Thrifty Applicability Lemma, observe that the definition of envelope-thrifty chase steps only poses additional conditions (relative to thrifty chase steps) on $\mathsf{NDng}(S^q)$, so that, for any fact that we would create with a fact-thrifty chase step, we can change the elements at $\mathsf{NDng}(S^q)$ to perform an envelope-thrifty chase step, using the fact that I is envelope-saturated. □

Further, recall that we showed that relation-thrifty chase steps never violate Σ_{UFD} (Lemma V.16). We now show that envelope-thrifty chase steps never violate Σ_{FD}, which is their intended purpose:

LEMMA VIII.17 (ENVELOPE-THRIFTY FD PRESERVATION). *For any n-envelope-saturated aligned superinstance J that satisfies Σ_{FD}, the result of an envelope-thrifty chase step on J satisfies Σ_{FD}.*

PROOF. Fix J and its global envelope $\mathcal{E}$. Let $F_{\mathrm{n}} = S(\boldsymbol{b})$ be the fact created by the envelope-thrifty step, let $\tau : R^p \subseteq S^q$ be the UID, let $J' = (I', \mathsf{sim}')$ be the result of the chase step, let F_{w} be the chase witness, and let D be the fact class of F_{w}. Write $O := \mathsf{NDng}(S^q)$. Assume by contradiction that $I' \not\models \Sigma_{\mathrm{FD}}$; as $I \models \Sigma_{\mathrm{FD}}$, any violation of Σ_{FD} in I' must be between the new fact F_{n} and an existing fact $F = S(\boldsymbol{c})$ of I. Recalling the definition of overlaps (Definition VIII.12), note that we only have $b_r \in \pi_{S^r}(I)$ for $S^r \in O$ by definition of thrifty chase steps, so we must have $\mathsf{OVL}(F_{\mathrm{n}}, F) \subseteq O$. Now, as $\pi_O(F_{\mathrm{n}})$ was defined using elements of $\mathrm{dom}(\mathcal{E}(D))$, taking any $S^r \in \mathsf{OVL}(F_{\mathrm{n}}, F) \subseteq O$ (which is non-empty by definition of an FD violation), we have $c_r = b_r \in \pi_{S^r}(\mathcal{E}(D))$, so that, by condition 4 of the definition of the envelope $\mathcal{E}(D)$, we know that $\pi_O(\boldsymbol{c})$ is a tuple $\boldsymbol{t}'$ of $\mathcal{E}(D)$. Now, either $\mathsf{OVL}(F_{\mathrm{n}}, F) \subsetneq O$ or $\mathsf{OVL}(F_{\mathrm{n}}, F) = O$.

In the first case, we observe that, by conditions 1 and 2 of the definition of the envelope $\mathcal{E}(D)$, we know that $\{\pi_O(\boldsymbol{c}), \pi_O(\boldsymbol{b})\}$ is not a violation of Σ_{FD}^{O}. Together with the fact that $\mathsf{OVL}(F_{\mathrm{n}}, F) \subsetneq O$, this allows us to apply Lemma VIII.13 and deduce that $\{F, F_{\mathrm{n}}\}$ actually does not violate Σ_{FD}, a contradiction.

In the second case, where $\mathsf{OVL}(F_{\mathrm{n}}, F) = O$, we have $\boldsymbol{t} = \boldsymbol{t}'$. Now, either D is safe or D is unsafe. If D is unsafe, we have a contradiction because F witnesses that $\boldsymbol{t}$ was not a remaining tuple of $\mathcal{E}(D)$, so we cannot have used $\boldsymbol{t}$ to define F_{n}. If D is safe, then by definition there is no FD $R^L \to R^r$ of Σ_{FD} with $R^L \subseteq O$ and $R^r \notin O$. Now, as $\mathsf{OVL}(F_{\mathrm{n}}, F) = O$, it is clear that F and F_{n} cannot violate any FD of Σ_{FD}, a contradiction again. □

Last, recall that we showed that fresh fact-thrifty steps preserve the property of being aligned (Lemma VI.16) and that non-fresh fact-thrifty steps also do when we additionally assume k-essentiality, which they also preserve (Lemma VI.24). We now prove the analogous claims for envelope-thrifty steps assuming envelope-saturation. The only difference is that envelope-thrifty chase steps make envelope-saturation *decrease*, unlike fact-thrifty steps which always preserved fact-saturation:

LEMMA VIII.18 (ENVELOPE-THRIFTY PRESERVATION). *For any $n \in \mathbb{N}$, for any n-envelope-saturated aligned superinstance J of I_0, the result J' of a fresh envelope-thrifty chase step on J is an $(n-1)$-envelope-saturated aligned superinstance of I_0. Further, if J is k-essential, the claim holds even for non-fresh envelope-thrifty chase steps, and the result J' is additionally k-essential.*

PROOF. We reuse notation from Lemma VIII.17: considering an application of an envelope-thrifty chase step: let $J = (I, \mathsf{sim})$ be the aligned superinstance of I_0, let $\tau : R^p \subseteq S^q$ be the UID, write $O := \mathsf{NDng}(S^q)$, let $F_{\mathrm{w}} = S(\boldsymbol{b}')$ the chase witness, let $D = (S^q, \boldsymbol{C})$ be the fact class, let $F_{\mathrm{n}} = S(\boldsymbol{b})$ be the new fact to be created, and let $\boldsymbol{t}$ be the remaining tuple of $\mathcal{E}(D)$ used to define F_{n}, and let $J' = (I', \mathsf{sim}')$ be the result.

We now prove that J' is still an aligned superinstance. We first adapt the Fresh Fact-Thrifty Preservation Lemma (Lemma VI.16) to work with envelope-thrifty chase steps. We can no longer use Lemma V.16 to prove that $J' \models \Sigma_{\mathrm{UFD}}$, but we have shown already that $J' \models \Sigma_{\mathrm{FD}}$ in Lemma VIII.17, so this point is already covered. The only other point specific to fact-thriftiness is proving that sim' is still a k-bounded simulation, but it actually only relies on the fact that $\mathsf{sim}'(b_r) \simeq_k b'_r$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ for all $S^r \in \mathsf{NDng}(S^q)$, which is still ensured by envelope-thrifty chase steps: by conditions 3 and 4 of the definition of envelopes, we know that, for any $S^r \in \mathsf{NDng}(S^q)$, the element t_r already occurs at position S^r in a fact of I that achieves D, so that $\mathsf{sim}(t_r) \simeq_k b'_r$.

Second, we adapt the Fact-Thrifty Preservation Lemma (Lemma VI.24) to envelope-thrifty chase steps. Again, the only condition of fact-thrifty chase steps used when proving that lemma is that

$\mathsf{sim}'(b_r) \simeq_k b'_r$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ for all $S^r \in \mathsf{NDng}(S^q)$, which is still true. Hence, having adapted these two lemmas, we conclude that J' has the required properties.

We now prove that $\mathcal{E}$ is still a global envelope of J' after performing an envelope-thrifty chase step. The condition on the disjointness of the envelope domains only concerns $\mathcal{E}$ itself, which is unchanged. Hence, we need only show that, for any $D' \in \mathsf{AFactCl}$, $\mathcal{E}(D')$ is still an envelope. All conditions of the definition of envelopes except condition 4 are clearly true, because they were true in J, and they only depend only on $\mathcal{E}(D')$ or they are preserved when creating more facts. We now check condition 4, which only needs to be verified on the new fact F_{n}.

Consider $S^u \in \mathsf{Pos}(S)$ and $S^t \in \mathsf{NDng}(S^u)$, and assume that $b_t \in \pi_{S^t}(\mathcal{E}(D'))$. As $\mathcal{E}(D)$ is an envelope for J, by condition 3 of the definition, we have $b_t \in \pi_{S^t}(I)$ as well, so that, by definition of thrifty chase steps, we must have $S^t \in O$. Now, as the envelopes of $\mathcal{E}$ are pairwise disjoint, and as the b_r for $S^r \in O$ are all in $\mathrm{dom}(\mathcal{E}(D))$, we must have $D = D'$, and $\boldsymbol{t}$ witnesses that $\pi_O(\boldsymbol{b}) \in \mathcal{E}(D)$. Hence $\mathcal{E}$ is still a global envelope of J'.

Last, to see that the resulting J' is $(n-1)$-envelope-saturated, it suffices to observe that, for each unsafe class $D \in \mathsf{AFactCl}$, the remaining tuples of $\mathcal{E}(D)$ for J' are those of $\mathcal{E}(D)$ for J minus at most one tuple (namely, some projection of F_{n}). This concludes the proof. □

Hence, we know that envelope-thrifty chase steps preserve being aligned and also preserves Σ_{FD} (rather than Σ_{UFD} for fact-thrifty chase steps). Our goal is then to modify the Fact-Thrifty Completion Proposition of the previous section (Proposition VII.7) to use envelope-thrifty rather than fact-thrifty chase steps, relying on the previous lemmas to preserve all invariants. The problem is that unlike fact-saturation, envelope-saturation "runs out"; whenever we use a remaining tuple $\boldsymbol{t}$ in a chase step to create F_{n} and obtain a new aligned superinstance J', then we can no longer use the same $\boldsymbol{t}$ in J'. This is why the result of an envelope-thrifty chase step is less saturated than its input, and it is why we made sure in the Sufficiently Envelope-Saturated Solutions Proposition that we could construct arbitrarily saturated superinstances.

For this reason, before we modify the Fact-Thrifty Completion Proposition, we need to account for the number of chase steps that the proposition performs. We show that it is linear in the size of the input instance.

LEMMA VIII.19 (ACCOUNTING). *There exists $B \in \mathbb{N}$ depending only on σ, k, and Σ_{U}, such that, for any aligned superinstance $J = (I, \mathsf{sim})$ of I_0, letting L be the preserving fact-thrifty sequence constructed in the Fact-Thrifty Completion Proposition, we have $|L| < B \cdot |I|$.*

PROOF. It suffices to show that $|L(J)| < B \cdot |I|$, because, as each chase step creates one fact, we have $|L| \leq |L(J)|$.

Remember that the fact-thrifty completion process starts by constructing an ordered partition $\boldsymbol{P} = (P_1, \ldots, P_n)$ of Σ_{UID} (Definition VII.2). This $\boldsymbol{P}$ does not depend on I. Hence, as we satisfy the UIDs of each P_i in turn, if we can show that the instance size only increases by a multiplicative constant for each class, then the blow-up for the entire process is by a multiplicative constant (obtained as the product of the constants for each P_i).

For trivial classes, we apply one chase round by fresh fact-thrifty chase steps (Lemma VII.9), It is easy to see that applying a chase round by any form of thrifty chase step on an aligned superinstance $J_1 = (I_1, \mathsf{sim}_1)$ yield a result whose size has only increased relative to J_1 by a multiplicative constant. This is because $|\mathrm{dom}(I_1)| \leq |\sigma| \cdot |I_1|$, and the number of facts created per element of I_1 in a chase round is at most $|\mathsf{Pos}(\sigma)|$. Hence, for trivial classes, we only incur a blowup by a constant multiplicative factor.

For non-trivial classes, we apply the Reversible Fact-Thrifty Completion Proposition (Proposition VI.25). Remember that this lemma first ensures k-essentiality by applying $k+1$ fact-thrifty chase rounds (Lemma VI.23) and then makes the result satisfy Σ_{UID} using the sequence constructed by the Reversible Relation-Thrifty Completion Proposition (Proposition V.17). Ensuring k-essentiality only implies a blow-up by a multiplicative constant, because it is performed by applying $k+1$

fact-thrifty chase rounds, so we can use the same reasoning as for trivial classes. Hence, we focus on the Reversible Relation-Thrifty Completion Proposition, and show that it also causes only a blow-up by a multiplicative constant.

When we apply the Reversible Relation-Thrifty Completion Proposition to an instance I, we start by constructing a balanced pssinstance P using the Balancing Lemma (Lemma IV.12), and a Σ_{U}-compliant piecewise realization PI of P by the Realizations Lemma (Lemma V.6), and we then apply fact-thrifty chase steps to satisfy Σ_{UID} following PI. We know that, whenever we apply a fact-thrifty chase step to an element a, the element a occurs after the chase step at a new position where it did not occur before. Hence, it suffices to show that $|\mathrm{dom}(P)|$ is within a constant factor of $|I|$, because then we know that the final number of facts created by the sequence of the Reversible Relation-Thrifty Completion Proposition will be $\leq |\mathrm{dom}(P)| \cdot |\mathsf{Pos}(\sigma)|$.

To show this, remember that $\mathrm{dom}(P) = \mathrm{dom}(I) \sqcup \mathcal{H}$, where $\mathcal{H}$ is the helper set. Hence, we only need to show that $|\mathcal{H}|$ is within a multiplicative constant factor of $|I|$. From the proof of the Balancing Lemma, we know that $\mathcal{H}$ is a disjoint union of $\leq |\mathsf{Pos}(\sigma)|$ sets whose size is linear in $|\mathrm{dom}(I)|$ which is itself $\leq |\sigma| \cdot |I|$. Hence, the Reversible Relation-Thrifty Completion Proposition only causes a blowup by a constant factor. As we justified, this implies the same about the entire completion process, and concludes the proof. □

This allows us to deduce the minimal level of envelope-saturation required to adapt the Fact-Thrifty Completion Proposition:

PROPOSITION VIII.20 (ENVELOPE-THRIFTY COMPLETION). *Let $\Sigma = \Sigma_{\mathrm{FD}} \wedge \Sigma_{\mathrm{UID}}$ be finitely closed FDs and UIDs, let $B \in \mathbb{N}$ be as in the Accounting Lemma, and let I_0 be an instance that satisfies Σ_{FD}. For any $(B \cdot |J|)$-envelope-saturated aligned superinstance J of I_0 that satisfies Σ_{FD}, we can obtain by envelope-thrifty chase steps an aligned superinstance J_{f} of I_0 that satisfies Σ.*

PROOF. We define envelope-thrifty sequences, and preserving envelope-thrifty sequences, analogously to (preserving) fact-thrifty sequences (Definition VI.26 and Definition VII.6) in the expected manner, but further requiring that all intermediate aligned superinstances remain envelope-saturated. This definition makes sense thanks to the Envelope-Thrifty Preservation Lemma.

By the Fact-Thrifty Completion Proposition, there exists a preserving fact-thrifty sequence L such that $L(J)$ satisfies Σ_{UID}, and $|L| < B \cdot |I|$. Construct from L an envelope-thrifty sequence L' that non-dangerously matches L, by changing each fact-thrifty chase step to an envelope-thrifty chase step, which we can do at each individual step thanks to the Envelope-Thrifty Applicability Lemma. It is clear that this is a preserving envelope-thrifty sequence, thanks to the Envelope-Thrifty Preservation Lemma, and thanks to the fact that the Ensuring Essentiality Lemma (Lemma VI.23) clearly adapts from fact-thrifty chase steps to envelope-thrifty chase steps: again, it only relies on the fact that, letting $F_{\mathrm{n}} = S(\boldsymbol{b})$ be the new fact, $F_{\mathrm{w}} = S(\boldsymbol{b}')$ the chase witness, and $\tau : R^p \subseteq S^q$ the UID, we have $\mathsf{sim}'(b_r) \simeq_k b'_r$ in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ for all $S^r \in \mathsf{NDng}(S^q)$. This also uses the fact that, by the Accounting Lemma, we have $|L| \leq B \cdot |I|$, so by the Envelope-Thrifty Preservation Lemma, all intermediate aligned superinstances remain envelope-saturated.

Hence, $J_{\mathrm{f}} := L'(J)$ is an aligned superinstance of I_0. Further, by the Thrifty Sequence Rewriting Lemma (Lemma VI.28), as $L(J) \models \Sigma_{\mathrm{UID}}$, so does J_{f}. Last, as $J \models \Sigma_{\mathrm{FD}}$, by the Envelope-Thrifty FD Preservation Lemma, so does J_{f}. This concludes the proof. □

We can now conclude the proof of Theorem VIII.1. Start by applying the saturation process of the Sufficiently Envelope-Saturated Solutions Proposition to obtain an aligned superinstance $J = (I, \mathsf{sim})$ of a disjoint union I'_0 of copies of I_0, such that J satisfies Σ_{FD} and is $(B \cdot |I|)$-envelope-saturated. Now, apply the Envelope-Thrifty Completion Proposition to obtain an aligned superinstance $J_{\mathrm{f}} = (I_{\mathrm{f}}, \mathsf{sim}_{\mathrm{f}})$ of I'_0 that satisfies Σ. We know that I_{f} satisfies Σ and is a k-sound superinstance of I'_0 for ACQ, but clearly it is also a k-sound superinstance of I_0, as is observed by the k-bounded simulation from I' to $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ obtained by composing sim' with the obvious homomorphism from $\mathsf{Chase}(I'_0, \Sigma_{\mathrm{UID}})$ to $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$. This concludes the proof.

VIII.3. Constructing Dense Interpretations

All that remains is to show the Dense Interpretations Theorem:

THEOREM VIII.11 (DENSE INTERPRETATIONS). *For any set Σ_{FD} of FDs over a relation R with no unary key, for all $K \in \mathbb{N}$, there exists $N_0 \in \mathbb{N}$ such that for all $N \geq N_0$, we can construct a non-empty instance I of R that satisfies Σ_{FD} and such that $|\text{dom}(I)| \leq N$ and $|I| \geq K \cdot N$.*

*Further, we can impose a **disjointness condition** on the result I: we can ensure that for all $a \in \text{dom}(I)$, there exists exactly one $R^p \in \mathsf{Pos}(R)$ such that $a \in \pi_{R^p}(I)$.*

Fix the relation R, and let Σ_{FD} be an arbitrary set of FDs which we assume is closed under FD implication. Let Σ_{UFD} be the UFDs implied by Σ_{FD}; it is also closed under FD implication. Recall the definition of OVL (Definition VIII.12). We introduce a notion of *tame overlaps* for Σ_{UFD}, which depends only on Σ_{UFD} but is a sufficient condition to satisfy Σ_{FD}, as we will show.

Definition VIII.21. We say a subset $O \subseteq \mathsf{Pos}(R)$ is **tame** for Σ_{UFD} if O is empty or for every $R^p \in \mathsf{Pos}(R) \setminus O$, there exists $R^q \in \mathsf{Pos}(R)$ such that:

— for all $R^s \in O$, the UFD $R^q \to R^s$ holds in Σ_{UFD},
— the UFD $R^q \to R^p$ does not hold in Σ_{UFD}.

We say that an instance I has the **tame overlaps** property (for Σ_{UFD}) if for every $F \neq F'$ of I, $\mathsf{OVL}(F,F')$ is tame.

In particular, if an instance ensures that non-empty overlaps between pairs of facts always have a unary key that determines precisely the overlap, then it has tame overlaps; we will show a refinement of this as Lemma VIII.25. We now claim the following lemma, and its immediate corollary:

LEMMA VIII.22. *If $O \subseteq \mathsf{Pos}(R)$ is tame for Σ_{UFD} then there is no FD $\phi : R^L \to R^r$ in Σ_{FD} such that $R^L \subseteq O$ but $R^r \notin O$.*

PROOF. If O is empty the claim is immediate. Otherwise, assume to the contrary the existence of such an FD ϕ. As $R^r \notin O$ and O is tame, there is $R^q \in \mathsf{Pos}(R)$ such that $R^q \to R^s$ holds in Σ_{UFD} for all $R^s \in O$, but $\phi' : R^q \to R^r$ does not hold in Σ_{UFD}. Now, as $R^L \subseteq O$, we know that $R^q \to R^s$ holds in Σ_{UFD} for all $R^s \in R^L$, so that, by transitivity with ϕ, as Σ_{FD} is closed by implication, ϕ' holds in Σ_{FD}. As ϕ' is a UFD, by definition of Σ_{UFD}, ϕ' holds in Σ_{UFD}, a contradiction. □

COROLLARY VIII.23. *For any instance I, if I has the tame overlaps property for Σ_{UFD}, then I satisfies Σ_{FD}.*

PROOF. Considering any two facts F and F' in I, as $O := \mathsf{OVL}(F,F')$ is tame, we know that for any FD $\phi : R^L \to R^r$ in Σ_{FD}, we cannot have $R^L \subseteq O$ but $R^r \notin O$. Hence, F and F' cannot be a violation of ϕ. □

We forget for now the disjointness condition in the Dense Interpretations Theorem, which we will prove at the very end of the subsection (Corollary VIII.26), and focus only on the first part. We claim the following generalization of the result:

THEOREM VIII.24. *Let R be a relation and Σ_{UFD} be a set of UFDs over R. Let D be the smallest possible cardinality of a **key** K of R (i.e., $K \subseteq \mathsf{Pos}(R)$ and for all $R^q \in \mathsf{Pos}(R)$, there is $R^p \in K$ such that $R^p \to R^q$ holds in Σ_{UFD}). Let x be $\frac{D}{D-1}$ if $D > 1$ and 1 otherwise.*

For every $N \in \mathbb{N}$, there exists a finite instance I of R such that $|\text{dom}(I)|$ is $O(N)$, $|I|$ is $\Omega(N^x)$, and I has the tame overlaps property for Σ_{UFD}.

Observe that, thanks to the use of the tame overlaps, the result does not mention higher-arity FDs, only UFDs; intuitively, tame overlaps ensures that the construction works for any FDs that have the same consequences as UFDs.

It is clear that this theorem implies the first part of the Dense Interpretations Theorem, because if R has no unary key for Σ_{FD} then $D > 1$ and thus $x > 1$, which implies that, for any K, by taking a sufficiently large N_0, we can obtain for all $N \geq N_0$ an instance I for R with $\leq N$ elements and $\geq K \cdot N$ facts that has the tame overlaps property for Σ_{UFD}; now, by Lemma VIII.23, this implies that I satisfies Σ_{FD}.

In the rest of this subsection, we prove Theorem VIII.24, until the very end where we additionally show that we can enforce the disjointness condition for the Dense Interpretations Theorem. Fix the relation R and set of UFDs Σ_{UFD}. The case of $D = 1$ is vacuous and can be eliminated directly (consider the instance $\{R(a_i, \ldots, a_i) \mid 1 \leq i \leq N\}$). Hence, assume that $D > 1$, and let $x := \frac{D}{D-1}$.

We first show the claim on a specific relation R_{full} and set $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$ of UFDs. We will then generalize the construction to arbitrary relations and UFDs. Let $T := \{1, \ldots, D\}$, and consider a bijection $\nu : \{1, \ldots, 2^D - 1\} \to \mathfrak{P}(T) \backslash \{\emptyset\}$, where $\mathfrak{P}(T)$ denotes the powerset of T. Let R_{full} be a $(2^D - 1)$-ary relation, and take $\Sigma_{\mathrm{UFD}}^{\mathsf{full}} := \{R^i \to R^j \mid \nu(i) \subseteq \nu(j)\}$. Note that $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$ is clearly closed under implication of UFDs. Fix $N \in \mathbb{N}$, and let us construct an instance I_{full} with $O(N)$ elements and $\Omega(N^x)$ facts.

Fix $n := \lfloor N^{1/(D-1)} \rfloor$. Let $\mathcal{F}$ be the set of partial functions from T to $\{1, \ldots, n\}$, and write $\mathcal{F} = \mathcal{F}_{\mathrm{t}} \sqcup \mathcal{F}_{\mathrm{p}}$, where $\mathcal{F}_{\mathrm{t}}$ and $\mathcal{F}_{\mathrm{p}}$ are respectively the total and the strictly partial functions. We take I_{full} to consist of one fact F_f for each $f \in \mathcal{F}_{\mathrm{t}}$, where $F_f = R_{\mathsf{full}}(\boldsymbol{a^f})$ is defined as follows: for $1 \leq i \leq 2^D - 1$, $a_i^f := f_{|T \backslash \nu(i)}$. In particular:

— $a_{\nu^{-1}(T)}^f$, the element of F_f at the position mapped by ν to $T \in \mathfrak{P}(T) \backslash \{\emptyset\}$, is the strictly partial function that is nowhere defined; note that $R_{\mathsf{full}}^{\nu^{-1}(T)}$ is determined by *all* positions in $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$.
— $a_{\nu^{-1}(\{i\})}^f$, the element of F_f at the position mapped by ν to $\{i\} \in \mathfrak{P}(T) \backslash \{\emptyset\}$, is the strictly partial function equal to f except that it is undefined on i; note that $R_{\mathsf{full}}^{\nu^{-1}(\{i\})}$ is determined by no other position of R_{full} in $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$.

Hence, $\mathrm{dom}(I_{\mathsf{full}}) = \mathcal{F}_{\mathrm{p}}$ (because $\emptyset$ is not in the image of ν), so that $|\mathrm{dom}(I_{\mathsf{full}})| = \sum_{0 \leq i < D} \binom{D}{i} n^i$. Remembering that D is a constant, this implies that $|\mathrm{dom}(I_{\mathsf{full}})|$ is $O(n^{D-1})$, so it is $O(N)$ by definition of n. Further, we claim that $|I_{\mathsf{full}}| = |\mathcal{F}_{\mathrm{t}}| = n^D = N^x$. To show this, consider two facts F_f and F_g. We show that $F_f = F_g$ implies $f = g$, so there are indeed $|\mathcal{F}_{\mathrm{t}}|$ different facts in I_{full}. As $\pi_{\nu^{-1}(\{1\})}(F_f) = \pi_{\nu^{-1}(\{1\})}(F_g)$, we have $f(t) = g(t)$ for all $t \in T \backslash \{1\}$, and as $D \geq 2$, we can look at $\pi_{\nu^{-1}(\{2\})}(F_f)$ and $\pi_{\nu^{-1}(\{2\})}(F_g)$ to conclude that $f(1) = g(1)$, hence $f = g$ as claimed. Hence, the cardinalities of I_{full} and of its domain are suitable.

We must now show that I_{full} has the tame overlaps property for $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$. For this we first make the following general observation:

LEMMA VIII.25. *Let Σ_{UFD} be any conjunction of UFDs and I be an instance such that $I \models \Sigma_{\mathrm{UFD}}$. Assume that, for any pair of facts $F \neq F'$ of I that overlap, there exists $R^p \in \mathsf{OVL}(F, F')$ which is a unary key for $\mathsf{OVL}(F, F')$. Then I has the tame overlaps property for Σ_{UFD}.*

PROOF. Consider $F, F' \in I$ and $O := \mathsf{OVL}(F, F')$. If $F = F'$, then $O = \mathsf{Pos}(R)$, and O is vacuously tame. Otherwise, if $F \neq F'$, let $R^p \in \mathsf{Pos}(R) \backslash O$. We take $R^q \in O$ to be the unary key of O. We know that $R^q \to R^s$ holds in Σ_{UFD} for all $R^s \in O$, so to show that O is tame it suffices to show that $\phi : R^q \to R^p$ does not hold in Σ_{UFD}. However, if it did, then as $R^q \in O$ and $R^p \notin O$, F and F' would witness a violation of ϕ, contradicting the fact that I satisfies Σ_{UFD}. □

So we show that I_{full} satisfies $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$ and that every non-empty overlap between facts of I_{full} has a unary key, so we can conclude by Lemma VIII.25 that I_{full} has tame overlaps.

First, to show that I_{full} satisfies $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$, observe that (*) whenever $\phi : R_{\mathsf{full}}^i \to R_{\mathsf{full}}^j$ holds in $\Sigma_{\mathrm{UFD}}^{\mathsf{full}}$, then $\nu(i) \subseteq \nu(j)$, so that, for any fact F of I_{full}, for any $1 \leq t \leq D$, whenever $(\pi_j(F))(t)$ is defined, so is $(\pi_i(F))(t)$, and we have $(\pi_j(F))(t) = (\pi_i(F))(t)$. Further, by our construction, we easily see that (**) for any fact F of I_{full}, for any $1 \leq i \leq 2^D - 1$ and $1 \leq t \leq D$, the fact that $(\pi_i(F))(t)$ is defined or not only depends on i and t, not on F. Hence, consider a UFD $\phi : R^i_{\mathsf{full}} \to R^j_{\mathsf{full}}$ in $\Sigma^{\mathsf{full}}_{\mathrm{UFD}}$, let F and F' be two facts of I_{full} such that $\pi_i(F) = \pi_i(F')$, and show that $\pi_j(F) = \pi_j(F')$. Take $1 \leq t \leq D$ and show that either $(\pi_j(F))(t)$ and $(\pi_j(F'))(t)$ are both undefined, or they are both defined and equal. By (**), either both are undefined or both are defined, so it suffices to show that if they are defined then they are equal. But then, if both are defined, by (*), we have $(\pi_j(F'))(t) = (\pi_i(F'))(t) = (\pi_i(F))(t) = (\pi_j(F))(t)$. So we conclude indeed that $\pi_j(F) = \pi_j(F')$, so that F and F' cannot witness a violation of ϕ. Hence, $I_{\mathsf{full}} \models \Sigma^{\mathsf{full}}_{\mathrm{UFD}}$.

Second, to show that non-empty overlaps in I_{full} have unary keys, consider two facts $F_f = R_{\mathsf{full}}(\boldsymbol{a^f})$ and $F_g = R_{\mathsf{full}}(\boldsymbol{a^g})$, with $f \neq g$ so that $F_f \neq F_g$. Assume that $\mathsf{OVL}(F_f, F_g)$ is non-empty, and let us show that it has a unary key. Let $O := \{t \in T \mid f(t) = g(t)\}$, and let $X = T \backslash O$; we have $X \neq \emptyset$, because otherwise $f = g$, so we can define $p := \nu^{-1}(X)$. We will show that

$$\mathsf{OVL}(F_f, F_g) = \{R^i \in \mathsf{Pos}(R_{\mathsf{full}}) \mid X \subseteq \nu(i)\}$$

This implies that $R^p \in \mathsf{OVL}(F_f, F_g)$ and that R^p is a unary key of $\mathsf{OVL}(F_f, F_g)$, because, for all $R^q \in \mathsf{OVL}(F_f, F_g)$, $X \subseteq \nu(q)$, so that $R^p \to R^q$ holds in $\Sigma^{\mathsf{full}}_{\mathrm{UFD}}$.

To show the equality above, consider R^i such that $X \subseteq \nu(i)$. Then $T \backslash \nu(i) \subseteq T \backslash X$. Because $a^f_i = f_{|T \backslash \nu(I)}$ and $a^g_i = g_{|T \backslash \nu(I)}$, we have $a^f_i = a^g_i$ by definition of $O = T \backslash X$. Thus $R^i \in \mathsf{OVL}(F_f, F_g)$. Conversely, if $R^i \in \mathsf{OVL}(F_f, F_g)$, then we have $a^f_i = a^g_i$, so by definition of O we must have $T \backslash \nu(i) \subseteq O = T \backslash X$, which implies $X \subseteq \nu(i)$.

Hence, I_{full} is a finite instance of $\Sigma^{\mathsf{full}}_{\mathrm{UFD}}$ which satisfies the tame overlaps property and contains $O(N)$ elements and $\Omega(N^x)$ facts. This concludes the proof of Theorem VIII.24 for the specific case of R_{full} and $\Sigma^{\mathsf{full}}_{\mathrm{UFD}}$.

Let us now show Theorem VIII.24 for an arbitrary relation R and set Σ_{UFD} of UFDs. Let K be a key of R of minimal cardinality, so that $|K| = D$. Let λ be a bijection from K to T. Extend λ to a function μ such that, for all $R^p \in \mathsf{Pos}(R)$, we set $\mu(R^p) := \{\lambda(R^k) \mid R^k \in K \text{ such that } R^k = R^p \text{ or } R^k \to R^p \text{ holds in } \Sigma_{\mathrm{UFD}}\}$; note that this set is never empty.

Consider the instance I_{full} for relation R_{full} that we defined previously, and create an instance I of R that contains, for every fact $R_{\mathsf{full}}(\boldsymbol{a})$ of I_{full}, a fact $F = R(\boldsymbol{b})$ in I, with $b_i = a_{\nu^{-1}(\mu(R^i))}$ for all $1 \leq i \leq |R|$.

We first show that $|\mathrm{dom}(I)| = O(N)$ and $|I| = \Omega(N^x)$. Indeed, for the first point, we have $\mathrm{dom}(I) \subseteq \mathrm{dom}(I_{\mathsf{full}})$, and as we had $|\mathrm{dom}(I_{\mathsf{full}})| = O(N)$, we deduce the same of $\mathrm{dom}(I)$. For the second point, it suffices to show that we never create the same fact twice in I for two different facts of I_{full}. Assume that there are two facts $F_f = R_{\mathsf{full}}(\boldsymbol{a})$ and $F_g = R_{\mathsf{full}}(\boldsymbol{a}')$ in I_{full} for which we created the same fact $F = R(\boldsymbol{b})$ in I, and let us show that we then have $f = g$ so that $F_f = F_g$. As $|K| \geq 2$, consider $R^{k_1} \neq R^{k_2}$ in K. We have $\mu(R^{k_1}) = \{\lambda(R^{k_1})\}$ and $\mu(R^{k_2}) = \{\lambda(R^{k_2})\}$. Hence, let $i_j := \lambda(R^{k_j})$ for $j \in \{1,2\}$; as λ is bijective, we deduce from $R^{k_1} \neq R^{k_2}$ that $i_1 \neq i_2$. From the definition of b_{k_1} we deduce that $a_{\nu^{-1}(\{i_1\})} = a'_{\nu^{-1}(\{i_1\})}$, and likewise $a_{\nu^{-1}(\{i_2\})} = a'_{\nu^{-1}(\{i_2\})}$. Similarly to the proof of why I_{full} has no duplicate facts, this implies that $f(t) = g(t)$ for all $t \in T \backslash \{i_1\}$ and for all $t \in T \backslash \{i_2\}$. As $i_1 \neq i_2$, we conclude that $f = g$, so that $F_f = F_g$. Hence, we have $|I| = |I_{\mathsf{full}}| = \Omega(N^x)$.

Let us now show that I has tame overlaps for Σ_{UFD}. Consider two facts F, F' of I that overlap, and let $O := \mathsf{OVL}(F, F')$. We first claim that there exists $\emptyset \subsetneq K' \subseteq K$, such that, letting $X' := \{\lambda(R^k) \mid R^k \in K'\}$, we have $\mathsf{OVL}(F, F') = \{R^i \in \mathsf{Pos}(R) \mid X' \subseteq \mu(R^i)\}$. Indeed, letting F_f and F_g be the facts of I_{full} used to create F and F', we previously showed the existence of $\emptyset \subsetneq X \subseteq T$ such that $\mathsf{OVL}(F_f, F_g) = \{R^i \in \mathsf{Pos}(R_{\mathsf{full}}) \mid X \subseteq \nu(i)\}$. Our definition of F and F' from F_f and F_g makes it clear that we can satisfy the condition by taking $K' := \lambda^{-1}(X)$, so that $X' = X$.

Consider now $R^p \in \mathsf{Pos}(R) \backslash O$. We cannot have $X' \subseteq \mu(R^p)$, otherwise $R^p \in O$. Hence, there exists $R^k \in K'$ such that $\lambda(R^k) \notin \mu(R^p)$. This implies that $R^k \to R^p$ does not hold in Σ_{UFD}. However, as $R^k \in K'$, we have $\lambda(R^k) \in \mu(R^q)$ for all $R^q \in O$, so that $R^k \to O$ holds in Σ_{UFD}. This proves that $O = \mathsf{OVL}(F, F')$ is tame. Hence, I has the tame overlaps property, which concludes the proof of Theorem VIII.24.

The only thing left is to show that we can enforce the disjointness condition in the Dense Interpretations Theorem, namely:

COROLLARY VIII.26. *We can assume in the Dense Interpretations Theorem (Theorem VIII.11) the following **disjointness condition** on the resulting instance I: each element occurs at exactly one position of the relation R. Formally, for all $a \in \mathrm{dom}(I)$, there exists exactly one $R^p \in \mathsf{Pos}(R)$ such that $a \in \pi_{R^p}(I)$.*

PROOF. Let I be the instance constructed in the proof of the Dense Interpretations Theorem, and consider the instance I' whose domain is $\{(a, R^p) \mid a \in \mathrm{dom}(I), R^p \in \mathsf{Pos}(\sigma)\}$ and which contains for every fact $F = R(\boldsymbol{a})$ of I a fact $F' = R(\boldsymbol{b})$ such that $b_p = (a_p, R^p)$ for every $R^p \in \mathsf{Pos}(\sigma)$. Clearly this defines a bijection ϕ from the facts of I to the facts of I', and for any facts F, F' of I', $\mathsf{OVL}(F, F') = \mathsf{OVL}(\phi^{-1}(F), \phi^{-1}(F'))$. Thus any violation of the FDs Σ_{FD} in I' would witness one in I. Of course, $|\mathrm{dom}(I')| = |\sigma| \cdot |\mathrm{dom}(I)|$, so to achieve a constant factor of K between the domain size and instance size with the disjointness condition, we need to use the proof of the Dense Interpretations Theorem with a constant factor of $K' := |\sigma| \cdot K$. $\square$

IX. BLOWING UP CYCLES

We are now ready to prove the Universal Models Theorem, which concludes the proof of our Main Theorem (Theorem III.3):

THEOREM III.6 (UNIVERSAL MODELS). *The class of finitely closed UIDs and FDs has finite universal models for CQ: for every conjunction Σ of FDs Σ_{FD} and UIDs Σ_{UID} closed under finite implication, for any $k \in \mathbb{N}$, for every finite instance I_0 that satisfies Σ_{FD}, there exists a finite k-sound superinstance I of I_0 that satisfies Σ.*

To do this, we must ensure k-soundness for arbitrary Boolean CQs rather than just acyclic CQs.

Intuitively, the only cyclic CQs that hold in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ either have an acyclic self-homomorphic match (so they are implied by an acyclic CQ that also holds) or have all cycles matched to elements of I_0. Hence, in a k-sound instance for CQ, no other cyclic queries should be true. Our way to ensure this is by a cycle blowup process: starting with the superinstance constructed by Theorem VIII.1, which satisfies Σ and is k-sound for ACQ, we build its product with a group of high girth. The standard way to do so, inspired by [Otto 2002], is presented in Section IX.1.

The problem is that this blowup process may create FD violations. We work around this problem using some additional properties ensured by our construction. In Section IX.2, we accordingly show the Cautious Models Theorem, a variant of Theorem VIII.1 with additional properties. Section IX.2 is the only part of this section that depends on the details of the previous sections.

We then apply a slightly different blowup construction to that model, as described in Section IX.3, which ensures that no FD violations are created. This blowup no longer depends on the specifics of the construction, and does not depend on the specific UIDs and FDs that hold; in particular, the blowup constructions do not even require that the UIDs and FDs are finitely closed.

IX.1. Simple Product

We first define a simple notion of product, which we can use to extend k-soundness from ACQ to CQ, but which may introduce FD violations. Let us first introduce preliminary notions:

Definition IX.1. A **group** $G = (S, \cdot)$ over a finite set S consists of:

— an associative **product law** $\cdot : S \times S \to S$;
— a **neutral element** $e \in S$ such that $e \cdot x = x \cdot e = x$ for all $x \in S$;
— an **inverse law** $\cdot^{-1} : S \to S$ such that $x \cdot x^{-1} = x^{-1} \cdot x = e$ for all $x \in S$.

We say that G is **generated** by $X \subseteq S$ if all elements of S can be written as a product of elements of X and $X^{-1} := \{x^{-1} \mid x \in X\}$.

Given a group $G = (S, \cdot)$ generated by X, assuming $|S| > 2$, the **girth** of G under X is the length of the shortest non-empty word $\boldsymbol{w}$ of elements of X and X^{-1} such that $w_1 \cdots w_n = e$ and $w_i \neq w_{i+1}^{-1}$ for all $1 \leq i < n$.

The following result, originally from [Margulis 1982], is proven for $|X| > 1$ in, e.g., [Otto 2012] (Section 2.1), and is straightforward for $|X| = 1$ (take $\mathbb{Z}/n\mathbb{Z}$):

LEMMA. *For all $n \in \mathbb{N}$ and finite non-empty set X, there is a finite group $G = (S, \cdot)$ generated by X with girth $\geq n$ under X. We call G an **n-acyclic group generated by** X.*

In other words, in an n-acyclic group generated by X, there is no short product of elements of X and their inverses which evaluates to e, except those that include a factor $x \cdot x^{-1}$.

We now explain how to take the product of a superinstance I of I_0 with such a finite group G. This ensures that any cycles in the product instance are large, because they project to cycles in G. We use a specific generator:

Definition IX.3. The **fact labels** of a superinstance I of I_0 are $\Lambda(I) := \{1_i^F \mid F \in I \backslash I_0, 1 \leq i \leq |F|\}$, where $|F|$ is the arity of the relation for fact F.

Now, we define the product of a superinstance I of I_0 with a group generated by $\Lambda(I)$. We make sure not to blow up cycles in I_0, so the result remains a superinstance of I_0:

Definition IX.4. Let I be a finite superinstance of I_0 and G be a finite group generated by $\Lambda(I)$. The **product of I by G preserving** I_0, written $(I, I_0) \otimes G$, is the finite instance with domain $\mathrm{dom}(I) \times G$ consisting of the following facts, for all $g \in G$:

— For every fact $R(\boldsymbol{a})$ of I_0, the fact $R((a_1, g), \ldots, (a_{|R|}, g))$.
— For every fact $F = R(\boldsymbol{a})$ of $I \backslash I_0$, the fact $R((a_1, g \cdot 1_1^F), \ldots, (a_{|R|}, g \cdot 1_{|R|}^F))$.

We identify (a, e) to a for $a \in \mathrm{dom}(I_0)$, so $(I, I_0) \otimes G$ is still a superinstance of I_0.

It will be simpler to reason about initial instances I_0 where each element has been *individualized* by the addition of a fresh fact that is unique to that element. We give a name to this notion:

Definition IX.5. An **individualizing** instance I_0 is such that, for each $a \in \mathrm{dom}(I_0)$, I_0 contains a fact $P_a(a)$ where P_a is a fresh unary predicate which does not occur in queries, in UIDs or in FDs.

An **individualizing superinstance** of an instance I_0 is a superinstance I_1 of I_0 that adds precisely one unary fact $P_a(a)$, for a fresh unary relation P_a, to each $a \in \mathrm{dom}(I_0)$, so that I_1 is individualizing. In particular, we have $\mathrm{dom}(I_0) = \mathrm{dom}(I_1)$, and I_0 and I_1 match for all relations of σ that occur in the query q and the constraints Σ.

We can now state the following property, which we will prove in the rest of this subsection:

LEMMA IX.6 (SIMPLE PRODUCT). *Let Σ be finitely closed FDs and UIDs, let I be a finite superinstance of an individualizing I_0 and let G be a finite $(2k+1)$-acyclic group generated by $\Lambda(I)$. If I is $(k \cdot (|\sigma| + 1))$-sound for ACQ, I_0, and Σ, then $(I, I_0) \otimes G$ is k-sound for CQ, I_0, and Σ.*

The following example illustrates the idea of taking the simple product of an instance with a group of high girth:

Example IX.7. Consider $F_0 := R(a, b)$ and $I_0 := \{F_0\}$, illustrated in solid black in the left part of Figure 8. Consider Σ_{UID} consisting of $\tau : R^2 \subseteq S^1$, $\tau' : S^2 \subseteq R^1$, τ^{-1}, and $(\tau')^{-1}$. Let $F := S(b, a)$, and $I := I_0 \sqcup \{F\}$, where F is a red dashed edge in the drawing. I satisfies Σ_{UID} and is sound for ACQ, but not for CQ: take for instance $q : \exists xy \, R(x, y) \wedge S(y, x)$, which is cyclic and holds in I while $(I_0, \Sigma_{\mathrm{UID}}) \not\models_{\mathrm{unr}} q$.

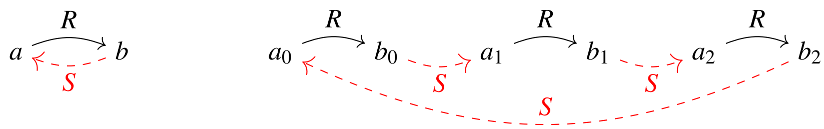

Fig. 8: Product with a group of large girth (see Example IX.7)

We have $\Lambda(I) = \{1^F_1, 1^F_2\}$. Identify 1^F_1 and 1^F_2 to 1 and 2 and consider the group $G := (\{0,1,2\},+)$ where $+$ is addition modulo 3. The group G has girth 2 under $\Lambda(I)$.

The product $I_\text{p} := (I, I_0) \otimes G$, writing pairs as subscripts for brevity, is $\{R(a_0,b_0), R(a_1,b_1), R(a_2,b_2), S(b_1,a_2), S(b_2,a_0), S(b_0,a_1)\}$. The right part of Figure 8 represents I_p. Here, I_p happens to be 5-sound for CQ.

We cannot directly use the simple product for our purposes, however, because $I_\text{p} := (I_\text{f}, I_0) \otimes G$ may violate Σ_UFD even though our instance I_f satisfies Σ_FD. Indeed, there may be a relation R, a UFD $\phi : R^p \to R^q$ in Σ_UFD, and two R-facts F and F' in $I_\text{f} \backslash I_0$ with $\pi_{R^p,R^q}(F) = \pi_{R^p,R^q}(F')$. In I_p there will be images of F and F' that overlap only on R^p, so they will violate ϕ.

Nevertheless, in the remainder of this subsection we prove the Simple Product Lemma, as it will be useful for our purposes later. Remember that a *match* of a CQ in an instance is witnessed by a homomorphism h, and that we also call the *match* the image of h. We start by proving an easy lemma:

LEMMA IX.8. *For any CQ q and instance I, if $I \models q$ with a witnessing homomorphism h that maps two different atoms of q to the same fact, then there is a CQ q' such that:*

— $|q'| < |q|$
— *q' **entails** q, meaning that for any instance I, if $I \models q'$ then $I \models q$*
— $I \models q'$

PROOF. Fix q, I, h, and let $A = R(\boldsymbol{x})$ and $A' = R(\boldsymbol{y})$ be the two atoms of q mapped to the same fact F by h. Necessarily A and A' are atoms for the same relation R of the fact F, and $h(A) = h(A')$ means that $h(x_i) = h(y_i)$ for all $R^i \in \mathsf{Pos}(R)$.

Let $\text{dom}(q)$ be the set of variables occurring in q. Consider the map f from $\text{dom}(q)$ to $\text{dom}(q)$ defined by $f(y_i) = x_i$ for all i, and $f(x) = x$ if x does not occur in A'. Observe that this ensures that $h(x) = h(f(x))$ for all $x \in \text{dom}(q)$. Let $q' = f(q)$ be the query obtained by replacing every variable x in q by $f(x)$, and, as $f(A') = f(A)$, removing one of those duplicate atoms so that $|q'| < |q|$. We claim that $h' := h_{|\text{dom}(q')}$ is a match of q' in I. Indeed, observe that any atom $f(A'')$ of q' is homomorphically mapped by h' to $h(A'')$ because $h'(f(x)) = h(x)$ for all x so $h'(f(A'')) = h(A'')$.

To see why q' entails q, observe that f defines a homomorphism from q to q', so that, for any instance I', if q' has a match h'' in I', then $h'' \circ f$ is a match of q in I'. □

Let us now prove the Simple Product Lemma. Fix the constraints Σ and the superinstance I of the individualizing I_0 such that I is $((|\sigma|+1) \cdot k)$-sound for ACQ, I_0, and Σ. Fix the $(2k+1)$-acyclic group G generated by $\Lambda(I)$. Consider $I_\text{p} := (I, I_0) \otimes G$, which is a superinstance of I_0, up to our identification of (a, e) to a for $a \in \text{dom}(I_0)$, where e is the neutral element of G. We must show that I_p is k-sound for CQ, I_0, and Σ.

We call a match h of a CQ q in I_p **pure-instance-cyclic** if every atom containing two occurrences of the same variable is mapped by h to a fact of $I_0 \times G$, and every Berge cycle of q contains an atom mapped by h to a fact of $I_0 \times G$. In particular, if q is in ACQ then any match h of q in I_p is vacuously pure-instance-cyclic. Our proof consists of two claims:

(1) If a CQ q with $|q| \leq k$ has a pure-instance-cyclic match h in I_p, then $\mathsf{Chase}(I_0, \Sigma_\text{UID}) \models q$.
(2) If a CQ q with $|q| \leq k$ has a match h in I_p which is not pure-instance-cyclic, then there is a CQ q' with $|q'| < |q|$ such that q' entails q and q' has a match in I_p.

The fact that I_{p} is k-sound for CQ clearly follows from the two claims: if a CQ q with $|q| \leq k$ has a match in I_{p}, then apply the second claim repeatedly until you obtain a CQ q' with $|q'| < |q| \leq k$, q' entails q, and q' has a pure-instance-cyclic match in I_{p}: this must eventually occur because the empty query is in ACQ. Then use the first claim to deduce that $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \models q'$, where it follows that $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \models q$. So it suffices to prove these two claims.

We start by proving the first claim. Let q be a CQ with $|q| \leq k$ that has a pure-instance-cyclic match h in I_{p}.

We partition the atoms of q between the atoms $\mathcal{A}$ matched by h to $I_0 \times G$ and the atoms $\mathcal{A}'$ which are not: we can then write q as $\exists \boldsymbol{x}\ \mathcal{A}(x) \wedge \mathcal{A}'(x)$. Let $\mathcal{A}''$ consist of the atom $P_a(x)$ for each variable x occurring in $\mathcal{A}'$ which is mapped by h to an element $a \in \mathrm{dom}(I_0 \times G)$, and let q' be the query $\exists \boldsymbol{x}\ \mathcal{A}'(x) \wedge \mathcal{A}''(x)$. As I_0 is individualizing, it is immediate that h is a match of q' in I_{p}.

We first claim that q' is in ACQ. Indeed, no Berge cycle in q' can use the atoms of $\mathcal{A}''$ as they are unary, and for the same reason no atom in $\mathcal{A}''$ contains two occurrences of the same variable. Further, $\mathcal{A}'$ does not contain any Berge cycle or atom with two occurrences of the same variable, by definition of h being pure-instance-cyclic. Hence, q' is indeed in ACQ. Further, we have $|q'| \leq k \cdot (|\sigma| + 1)$, as $|\mathcal{A}''| \leq |\sigma| \cdot |\mathcal{A}'|$ and we have $|\mathcal{A}'| \leq |q| \leq k$, so that $|q'| \leq k \cdot (|\sigma| + 1)$. Now, we know that $I \models q'$, as evidenced by the homomorphism pr from I_{p} to I defined by $\mathsf{pr} : (a, g) \mapsto a$ for all $a \in \mathrm{dom}(I)$ and $g \in G$. As I is $(k \cdot (|\sigma| + 1))$-sound for ACQ, and q' is an ACQ query that holds in I with $|q'| \leq k \cdot (|\sigma| + 1)$, we know that $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}}) \models q'$.

Now, as $\mathcal{A}''$ covers all variables of q', by definition of I_0 being individualizing, the only possible match of q' in the chase is the one that maps each variable x to the $a \in \mathrm{dom}(I_0)$ such that the atom $P_a(x)$ is in $\mathcal{A}''$. Further, as h matched $\mathcal{A}$ to facts of I_0 such that $h(x) = a$ where $P_a(x)$ occurs in $\mathcal{A}''$, we can clearly extend the match of q' in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$ to a match of q in $\mathsf{Chase}(I_0, \Sigma_{\mathrm{UID}})$. This concludes the proof of the first claim.

We now prove the second claim. Let q be a CQ with $|q| \leq k$ that has a match h in I_{p} which is not pure-instance-cyclic. Consider a Berge cycle C of q, of the form $A_1, x_1, A_2, x_2, \ldots, A_n, x_n$, where the A_i are pairwise distinct atoms and the x_i pairwise distinct variables, where the A_i are mapped by h to facts not in $I_0 \times G$, and where for all $1 \leq i \leq n$, variable x_i occurs at position q_i of atom A_i and position p_{i+1} of A_{i+1}, with addition modulo $n := |C|$. We assume without loss of generality that $p_i \neq q_i$ for all i. However, we do not assume that $n \geq 2$: either $n \geq 2$ and C is really a Berge cycle according to our previous definition, or $n = 1$ and variable x_1 occurs in atom A_1 at positions $p_1 \neq q_1$, which corresponds to the case where there are multiple occurrences of the same variable in an atom.

For $1 \leq i \leq n$, we write $F_i = R_i(\boldsymbol{a^i})$ the image of A_i by h in I_{p}; by definition of I_{p}, as F_i is not a fact of $I_0 \times G$, there is a fact $F'_i = R_i(\boldsymbol{b^i})$ of I and $g_i \in G$ such that $a^i_j = (b^i_j, g_i \cdot 1^{F'_i}_j)$ for $R^j_i \in \mathsf{Pos}(R_i)$. Now, for all $1 \leq i \leq n$, as $h(x_i) = a^i_{q_i} = a^{i+1}_{p_i+1}$ for all $1 \leq i \leq n$, we deduce by projecting on the second component that $g_i \cdot 1^{F'_i}_{q_i} = g_{i+1} \cdot 1^{F'_{i+1}}_{p_{i+1}}$, so that, by collapsing the equations of the cycle together, $1^{F'_1}_{q_1} \cdot (1^{F'_2}_{p_2})^{-1} \cdot \cdots \cdot 1^{F'_{n-1}}_{q_{n-1}} \cdot (1^{F'_n}_{p_n})^{-1} \cdot 1^{F'_n}_{q_n} \cdot (1^{F'_1}_{p_1})^{-1} = e$.

As the girth of G under $\Lambda(I)$ is $\geq 2k+1$, and this product contains $2n \leq 2k$ elements, we must have either $1^{F'_i}_{q_i} = 1^{F'_{i+1}}_{p_{i+1}}$ for some i, or $1^{F'_i}_{p_i} = 1^{F'_i}_{q_i}$ for some i. The second case is impossible because we assumed that $p_i \neq q_i$ for all $1 \leq i \leq n$. Hence, necessarily $1^{F'_i}_{q_i} = 1^{F'_{i+1}}_{p_{i+1}}$, so in particular we must have $n > 1$ and $F'_i = F'_{i+1}$. Hence the atoms $A_i \neq A_{i+1}$ of q are mapped by h to the same fact $F'_i = F'_{i+1}$. We conclude by Lemma IX.8 that there is a strictly smaller q' which entails q and has a match in I_{p}, which is what we wanted to show. This concludes the proof of the second claim, and of the Simple Product Lemma.

IX.2. Cautiousness

As the simple product may cause FD violations, we will define a more refined notion of product, which intuitively does not attempt to blow up cycles within fact overlaps. In order to clarify this, however, we will first need to study in more detail the instance I_{f} to which we will apply the process, namely, the one that we constructed to prove Theorem VIII.1. We will consider a *quotient* of I_{f}:

Definition IX.9. The **quotient** $I/\sim$ of an instance I by an equivalence relation $\sim$ on $\mathrm{dom}(I)$ is defined as follows:

— $\mathrm{dom}(I/\sim)$ is the equivalence classes of $\sim$ on $\mathrm{dom}(I)$,
— $I/\sim$ contains one fact $R(\boldsymbol{A})$ for every fact $R(\boldsymbol{a})$ of I, where A_i is the $\sim$-class of a_i for all $R^i \in \mathsf{Pos}(R)$.

The **quotient homomorphism** $\chi_\sim$ is the homomorphism from I to $I/\sim$ defined by mapping each element of $\mathrm{dom}(I)$ to its $\sim$-class.

We quotient I_{f} by the equivalence relation $\simeq_k$ (recall Definition VI.2). The result may no longer satisfy Σ. However, it is still k-sound for ACQ, for the following reason:

LEMMA IX.10. *Any k-bounded simulation from an instance I to an instance I' defines a k-bounded simulation from $I/\simeq_k$ to I'.*

PROOF. Fix the instance I and the k-bounded simulation sim to an instance I', and consider $I'' := I/\simeq_k$. We show that there is a k-bounded simulation sim' from I'' to I, because $\mathsf{sim} \circ \mathsf{sim}'$ would then be a k-bounded simulation from I'' to I', the desired claim. We define $\mathsf{sim}'(A)$ for all $A \in I''$ to be a for any member $a \in A$ of the equivalence class A in I, and show that sim' thus defined is indeed a k-bounded simulation.

We will show the stronger result that $(I'',A) \leq_k (I,a)$ for all $A \in \mathrm{dom}(I'')$ and for any $a \in A$. We do it by proving, by induction on $0 \leq k' \leq k$, that $(I'',A) \leq_{k'} (I,a)$ for all $A \in \mathrm{dom}(I'')$ and $a \in A$. The case $k' = 0$ is trivial. Hence, fix $0 < k' \leq k$, assume that $(I'',A) \leq_{k'-1} (I,a)$ for all $A \in \mathrm{dom}(I'')$ and $a \in A$, and show that this is also true for k'. Choose $A \in \mathrm{dom}(I'')$, $a \in A$, we must show that $(I'',A) \leq_{k'} (I,a)$. To do so, consider any fact $F = R(\boldsymbol{A})$ of I'' such that $A_p = A$ for some $R^p \in \mathsf{Pos}(R)$. Let $F' = R(\boldsymbol{a}')$ be a fact of I that is a preimage of F by $\chi_{\simeq_k}$, so that $a'_q \in A_q$ for all $R^q \in \mathsf{Pos}(R)$. We have $a'_p \in A$ and $a \in A$, so that $a'_p \simeq_k a$ holds in I. Hence, in particular we have $(I,a'_p) \leq_{k'} (I,a)$ because $k' \leq k$, so there exists a fact $F'' = R(\boldsymbol{a}'')$ of I such that $a''_p = a$ and $(I,a'_q) \leq_{k'-1} (I,a''_q)$ for all $R^q \in \mathsf{Pos}(R)$. We show that F'' is a witness fact for F. Indeed, we have $a''_p = a$. Let us now choose $R^q \in \mathsf{Pos}(R)$ and show that $(I'',A_q) \leq_{k'-1} (I,a''_q)$. By induction hypothesis, as $a'_q \in A_q$, we have $(I'',A_q) \leq_{k'-1} (I,a'_q)$, and as $(I,a'_q) \leq_{k'-1} (I,a''_q)$, by transitivity we have indeed $(I'',A_q) \leq_{k'-1} (I,a''_q)$. Hence, we have shown that $(I'',A) \leq_{k'} (I,a)$.

By induction, we conclude that $(I'',A) \leq_k (I,a)$ for all $A \in \mathrm{dom}(I'')$ and $a \in A$, so that there is indeed a k-bounded simulation from I'' to I, which, as we have explained, implies the desired claim. □

Let us thus consider $I'_{\mathrm{f}} := I_{\mathrm{f}}/\simeq_k$ which is still k-sound for ACQ by the previous lemma, and consider the homomorphism $\chi_{\simeq_k}$ from I_{f} to I'_{f}. Our idea is to blow up cycles in I_{f} by a *mixed product* that only distinguishes facts that have a different image in I'_{f} by $\chi_{\simeq_k}$. This is sufficient to lift k-soundness from ACQ to CQ, and it will not create FD violations on facts that have the same image by $\chi_{\simeq k}$. Crucially, however, we can show from our construction that all overlapping facts of I_{f} have the same image by $\chi_{\simeq k}$. Let us formalize this condition:

Definition IX.11. Let I be an instance, let $I_1 \subseteq I$, and let f be any mapping with domain I. We say I is **cautious** for f and I_1 if for any two **overlapping** facts, namely, two facts $F = R(\boldsymbol{a})$ and $F' = R(\boldsymbol{b})$ of the same relation with $a_p = b_p$ for some $R^p \in \mathsf{Pos}(R)$, one of the following holds: $F, F' \in I_1$, or $f(a_p) = f(b_p)$ for all $R^p \in \mathsf{Pos}(R)$.

We conclude the subsection by presenting a strengthening of Theorem VIII.1. This is the only point in this section where we rely on the details of the process of the previous sections:

THEOREM IX.12 (CAUTIOUS MODELS). *For any finitely closed Σ formed of UIDs Σ_{UID} and FDs Σ_{FD}, instance I_0, and $k \in \mathbb{N}$, we can build a finite superinstance I_{f} of an instance I_1 such that:*

— *I_{f} satisfies Σ;*
— *I_{f} is k-sound for Σ, ACQ, and I_1;*
— *I_1 is an individualizing superinstance of a disjoint union of copies of I_0;*
— *I_{f} is cautious for $\chi_{\simeq_k}$ and I_1.*

We will use the Cautious Models Theorem in the next subsection. For now, let us show how to prove it. Fix Σ, I_0, and $k \in \mathbb{N}$. Let $I_{0,\mathrm{i}}$ be an individualizing superinstance of I_0, and apply k UID chase rounds with the UIDs of Σ_{UID} to $I_{0,\mathrm{i}}$ to obtain $I'_{0,\mathrm{i}}$. Apply the Sufficiently Envelope-Saturated Proposition to $I'_{0,\mathrm{i}}$ to obtain an aligned superinstance J of a disjoint union $I''_{0,\mathrm{i}}$ of copies of $I'_{0,\mathrm{i}}$. Now, modify J to J' and $I''_{0,\mathrm{i}}$ to I'_1 by replacing the copies of the facts of $I_{0,\mathrm{i}} \backslash I_0$ by new individualizing facts (i.e., make the individualizing facts unique across copies of $I'_{0,\mathrm{i}}$). This ensures by definition that I'_1 is the result of applying k UID chase rounds to an individualizing superinstance of a disjoint union of copies of I_0. Further, the modification to J' can be done so as to ensure that J' is an aligned superinstance of I'_1; the k chase rounds applied when defining $I'_{0,\mathrm{i}}$ ensure that the sim mapping can still be defined notwithstanding the change in the individualizing facts. Further, we have $|J'| = |J|$, so J' is still sufficiently envelope-saturated.

We now apply the Envelope-Thrifty Completion Proposition to the aligned superinstance J' of I'_1 to obtain a superinstance J_{f} of I'_1 which is k-sound for Σ, ACQ, and I'_1, and that satisfies Σ. Now, define I_1 from I'_1 by removing the facts created in the k UID chase rounds, so it is by definition an individualizing superinstance of a disjoint union of copies of I_0. As I'_1 is the result of applying chase rounds to I_1, I_{f} is also k-sound for Σ, ACQ and I_1. Hence, I_{f} satisfies the first three conditions that we have to show in the Cautious Models Theorem, and the third is satisfied by definition. The only thing left is to show the last one, namely:

LEMMA IX.13 (CAUTIOUSNESS). *I_{f} is cautious for $\chi_{\simeq_k}$ and I_1.*

We show the Cautiousness Lemma in the rest of the subsection, which concludes the proof of the Cautious Models Theorem.

We first show that overlapping facts in $J_{\mathrm{f}} = (I_{\mathrm{f}}, \mathsf{sim}_{\mathrm{f}})$ are cautious for the sim mapping that we construct, in terms of $\simeq_k$-classes. Formally, let $I_{\mathrm{c}} := \mathsf{Chase}(I_1, \Sigma_{\mathrm{UID}})$, and let $\chi'_{\simeq_k}$ be the homomorphism from I_{c} to $I_{\mathrm{c}}/\simeq_k$. We claim:

LEMMA IX.14. *I_{f} is cautious for $\chi'_{\simeq_k} \circ \mathsf{sim}$ and I_1.*

In other words, whenever two facts $F = R(\boldsymbol{a})$ and $F' = R(\boldsymbol{b})$ have non-empty overlap in I_{f} and are not both in I_1, then, for any position $R^p \in \mathsf{Pos}(R)$, we have $\mathsf{sim}(a_p) \simeq_k \mathsf{sim}(b_p)$ in I_{c}.

PROOF. We first check that this claim holds on the result J' of the Sufficiently Envelope-Saturated Proposition (with our modifications to the individualizing facts). J' is a disjoint union of instances J_D for each fact class $D \in \mathsf{AFactCl}$. If D is safe, no facts overlap in J_D except possibly fact pairs in the copy of I_0, hence, in I_1. For unsafe D, in Lemma VIII.8, the only facts with non-empty overlap in J_D are fact pairs in some copy of I_0, hence in I_1, or they are the facts $f'(F_i)$, which all map to $\simeq_k$-equivalent sim-images by construction. So the claim holds on J'.

Second, it suffices to show that the claim is preserved by envelope-thrifty chase steps. By their definition, whenever we create a new fact F_{n} for a fact class D, the only elements of F_{n} that can be part of an overlap between F_{n} and an existing fact are envelope elements, appearing at the one position at which they appear in $\mathcal{E}(D)$. Then, by condition 4 of the definition of envelopes (Definition VIII.3), we deduce that the two overlapping facts achieve the same fact class. □

Returning to the proof of the Cautiousness Lemma, we now show that two elements in J_{f} having $\simeq_k$-equivalent sim images in I_{c} must themselves be $\simeq_k$-equivalent in J_{f}. We do it by showing that, in fact, for any $a \in \mathrm{dom}(I_\mathrm{f})$, not only do we have $(I_\mathrm{f}, a) \leq_k (I_\mathrm{c}, \mathsf{sim}(a))$, as required by the k-bounded simulation sim, but we also have the reverse: $(I_\mathrm{c}, \mathsf{sim}(a)) \leq_k (I_\mathrm{f}, a)$; in fact, we even have a *homomorphism* from I_c to I_f that maps $\mathsf{sim}(a)$ to a. The existence of this homomorphism is thanks to our specific definition of sim, and on the directionality condition of aligned superinstances; further, it only holds for the final result I_f, which satisfies Σ_UID; it is not respected at intermediate steps of the process.

To prove this, and conclude the proof of the Cautiousness Lemma, remember the forest structure on the UID chase (Definition VI.19). We define the **ancestry** $\mathcal{A}_F$ of a fact F in I_c as I_1 plus the facts of the path in the chase forest that leads to F; if $F \in I_1$ then $\mathcal{A}_F$ is just I_1. The **ancestry** $\mathcal{A}_a$ of $a \in \mathrm{dom}(I_\mathrm{c})$ is that of the fact where a was introduced.

We now claim the following lemma about J_f, which relies on the directionality condition:

LEMMA IX.15. *For any $a \in \mathrm{dom}(I_\mathrm{f})$, there is a homomorphism h_a from $\mathcal{A}_{\mathsf{sim}(a)}$ to I_f such that $h_a(\mathsf{sim}(a)) = a$.*

PROOF. We prove that this property holds on I_f, by first showing that it is true of J' constructed by our modification of the Sufficiently Envelope-Saturated Solutions Proposition. This is clearly the case because the instances created by Lemma VIII.8 are just truncations of the chase where some elements are identified at the last level.

Second, we show that the property is maintained by envelope-thrifty steps; in fact, by any thrifty chase steps (Definition VI.12) Consider a thrifty chase step where, in a state $J_1 = (I_1, \mathsf{sim}_1)$ of the construction of our aligned superinstance, we apply a UID $\tau : R^p \subseteq S^q$ to a fact $F_\mathrm{a} = R(\boldsymbol{a})$ to create a fact $F_\mathrm{n} = S(\boldsymbol{b})$ and obtain the aligned superinstance $J_2 = (I_2, \mathsf{sim}_2)$. Consider the chase witness $F_\mathrm{w} = S(\boldsymbol{b}')$. By Lemma VI.13, b'_q is the exported element between F_w and its parent in $\mathsf{Chase}(I_0, \Sigma_\mathrm{UID})$. So we know that for any $S^r \neq S^q$, we have $\mathcal{A}_{b'_r} = \mathcal{A}_{b'_q} \sqcup \{F_\mathrm{w}\}$.

We must build the desired homomorphism h_a for all $a \in \mathrm{dom}(I_2) \setminus \mathrm{dom}(I_1)$. Indeed, for $a \in \mathrm{dom}(I_1)$, by hypothesis on I_1, there is a homomorphism h_a from $\mathcal{A}_{\mathsf{sim}_1(a)}$ to I_1 with $h_a(\mathsf{sim}_1(a)) = a$, and as $\mathsf{sim}_2(a) = \mathsf{sim}_1(a)$, we can use h_a as the desired homomorphism from $\mathcal{A}_{\mathsf{sim}_2(a)}$ to I_2. So let us pick $b \in \mathrm{dom}(I_2) \setminus \mathrm{dom}(I_1)$ and construct h_b. By construction of I_2, b must occur in the new fact F_n; further, by definition of thrifty chase steps, we have defined $\mathsf{sim}_2(b) := b'_r$ for some S^r where $b_r = b$. Now, as $a_p = b_q$ is in $\mathrm{dom}(I_1)$, we know that there is a homomorphism h_{b_q} from $\mathcal{A}_{\mathsf{sim}(b_q)} = \mathcal{A}_{b'_q}$ to I_1 such that we have $h_{b_q}(b'_q) = b_q$. We extend h_{b_q} to the homomorphism h_b from $\mathcal{A}_{b'_r} = \mathcal{A}_{b'_q} \sqcup \{F_\mathrm{w}\}$ to I_2 such that $h_b(b'_r) = b$, by setting $h_b(F_\mathrm{w}) := F_\mathrm{n}$ and $h_b(F) := h(F)$ for any other F of $\mathcal{A}_{b'_r}$; we can do this because, by definition of the UID chase, F_w shares no element with the other facts of $\mathcal{A}_{b'_r}$ (that is, with $\mathcal{A}_{b'_q}$), except b'_q for which our definition coincides with the existing image of b'_q by h_{b_q}. This proves the claim. □

This allows us to deduce the following, which is specific to J_f, and relates to the universality of the chase I_c:

COROLLARY IX.16. *For any $a \in \mathrm{dom}(I_\mathrm{f})$, there is a homomorphism h_a from I_c to I_f such that $h_a(\mathsf{sim}(a)) = a$.*

PROOF. Choose $a \in \mathrm{dom}(I_\mathrm{f})$ and let us construct h_a. Let h'_a be the homomorphism from $\mathcal{A}_{\mathsf{sim}(a)}$ to I_f with $h'_a(\mathsf{sim}(a)) = a$ whose existence was proved in Lemma IX.15. Now start by setting $h_a := h'_a$, and extend h'_a to be the desired homomorphism, fact by fact, using the property that $I_\mathrm{f} \models \Sigma_\mathrm{UID}$: for any $b \in \mathrm{dom}(I_\mathrm{c})$ not in the domain of h'_a but which was introduced in a fact F whose exported element c is in the current domain of h'_a, let us extend h'_a to the elements of F in the following way: consider the parent fact F' of F in I_c and its match by h'_a in I_f, let τ be the UID used to create F' from F, and $c' \in \mathrm{dom}(I_\mathrm{c})$ be the exported element between F and F' (so $h'_a(c')$ is defined). We know that $c := h'_a(c')$ occurs in I_f at all positions where c' occurs in I_c. Hence, because $I_\mathrm{f} \models \tau$, there must be a suitable fact F'' in I_f to extend h'_a to all elements of F by setting $h'_a(F) := F''$, which is consistent with the image of c previously defined in h'_a. The (generally infinite) result of this process is the desired homomorphism h_a. □

We are now ready to show our desired claim:

LEMMA IX.17. *For any $a,b \in \mathrm{dom}(I_\mathrm{f})$, if $\mathsf{sim}(a) \simeq_k \mathsf{sim}(b)$ in I_c, then $a \simeq_k b$ in I_f.*

PROOF. Fix $a,b \in \mathrm{dom}(I_\mathrm{f})$. We have $(I_\mathrm{f}, a) \leq_k (I_\mathrm{c}, \mathsf{sim}(a))$ because sim is a k-bounded simulation; we have $(I_\mathrm{c}, \mathsf{sim}(a)) \leq_k (I_\mathrm{c}, \mathsf{sim}(b))$ because $\mathsf{sim}(a) \simeq_k \mathsf{sim}(b)$; and we have $(I_\mathrm{c}, \mathsf{sim}(b)) \leq_k (I_\mathrm{f}, b)$ by Corollary IX.16 as witnessed by h_b. By transitivity, we have $(I_\mathrm{f}, a) \leq_k (I_\mathrm{f}, b)$. The other direction is symmetric, so the desired claim follows. □

The Cautiousness Lemma (Lemma IX.13) follows immediately from Lemma IX.14 and Lemma IX.17. This concludes the proof of the Cautious Models Theorem.

IX.3. Mixed Product

Using the Cautious Models Theorem, we now define the notion of mixed product, which uses the same fact label for facts with the same image by $h := \chi_{\simeq_k}$:

Definition IX.18. Let I be a finite superinstance of I_1 with a homomorphism h to another finite superinstance I' of I_1 such that $h_{|I_1}$ is the identity and $h_{|(I \backslash I_1)}$ maps to $I' \backslash I_1$. Let G be a finite group generated by $\Lambda(I')$.

The **mixed product** of I by G via h preserving I_1, written $(I, I_1) \otimes^h G$, is the finite superinstance of I_1 with domain $\mathrm{dom}(I) \times G$ consisting of the following facts, for every $g \in G$:

— For every fact $R(\boldsymbol{a})$ of I_1, the fact $R((a_1, g), \ldots, (a_{|R|}, g))$.
— For every fact $F = R(\boldsymbol{a})$ of $I \backslash I_1$, the fact $R((a_1, g \cdot 1_1^{h(F)}), \ldots, (a_{|R|}, g \cdot 1_{|R|}^{h(F)}))$.

We now show that the mixed product preserves UIDs and FDs when cautiousness is assumed.

LEMMA IX.19 (MIXED PRODUCT PRESERVATION). *For any UID or FD τ, if $I \models \tau$ and I is cautious for h, then $(I, I_1) \otimes^h G \models \tau$.*

PROOF. Write $I_\mathrm{m} := (I, I_1) \otimes^h G$ and write I' for the range of h as before.

If τ is a UID, the claim is immediate even without the cautiousness hypothesis. (In fact, the analogous claim could even be proven for the simple product.) Indeed, for any $a \in \mathrm{dom}(I)$ and $R^p \in \mathsf{Pos}(\sigma)$, if $a \in \pi_{R^p}(I)$ then $(a, g) \in \pi_{R^p}(I_\mathrm{m})$ for all $g \in G$; conversely, if $a \notin \pi_{R^p}(I)$ then $(a, g) \notin \pi_{R^p}(I_\mathrm{m})$ for all $g \in G$. Hence, letting $\tau : R^p \subseteq S^q$ be a UID of Σ_UID, if there is $(a, g) \in \mathrm{dom}(I_\mathrm{m})$ such that $(a, g) \in \pi_{R^p}(I_\mathrm{m})$ but $(a, g) \notin \pi_{S^q}(I_\mathrm{m})$ then $a \in \pi_{R^p}(I)$ but $a \notin \pi_{S^q}(I)$. Hence any violation of τ in I_m implies the existence of a violation of τ in I, so we conclude because $I \models \tau$.

Assume now that τ is a FD $\phi : R^L \to R^r$. Assume by contradiction that there are two facts $F_1 = R(\boldsymbol{a})$ and $F_2 = R(\boldsymbol{b})$ in I_m that violate ϕ, i.e., we have $a_l = b_l$ for all $l \in L$, but $a_r \neq b_r$. Write $a_i = (v_i, f_i)$ and $b_i = (w_i, g_i)$ for all $R^i \in \mathsf{Pos}(R)$. Consider $F'_1 := R(\boldsymbol{v})$ and $F'_2 := R(\boldsymbol{w})$ the facts of I that are the images of F_1 and F_2 by the homomorphism from I_m to I that projects on the first component. As $I \models \tau$, F'_1 and F'_2 cannot violate ϕ, so as $v_l = w_l$ for all $l \in L$, we must have $v_r = w_r$. Now, as I is cautious for h and F'_1 and F'_2 overlap (take any $R^{l_0} \in R^L$), either $F'_1, F'_2 \in I_1$ or $h(F'_1) = h(F'_2)$.

In the first case, by definition of the mixed product, there are $f, g \in G$ such that $f_i = f$ and $g_i = g$ for all $R^i \in \mathsf{Pos}(R)$. Thus, taking any $l_0 \in L$, as we have $a_{l_0} = b_{l_0}$, we have $f_{l_0} = g_{l_0}$, so $f = g$, which implies that $f_r = g_r$. Hence, as $v_r = w_r$, we have $(v_r, f_r) = (w_r, g_r)$, contradicting the fact that $a_r \neq b_r$.

In the second case, as h is the identity on I_1 and maps $I \backslash I_1$ to $I' \backslash I_1$, $h(F'_1) = h(F'_2)$ implies that either F'_1 and F'_2 are both facts of I_1 or they are both facts of $I \backslash I_1$; but we have already excluded the former possibility in the first case, so we assume the latter. By definition of the mixed product, there are $f, g \in G$ such that $f_i = f \cdot 1_i^{h(F'_1)}$ and $g_i = g \cdot 1_i^{h(F'_2)}$ for all $R^i \in \mathsf{Pos}(R)$. Picking any $l_0 \in L$, from $a_{l_0} = b_{l_0}$, we deduce that $f \cdot 1_{l_0}^{h(F_1')} = g \cdot 1_{l_0}^{h(F_2')}$; as $h(F_1') = h(F_2')$, this simplifies to $f = g$. Hence, $f_r = g_r$ and we conclude like in the first case. □

Second, we show that $h : I \to I'$ lifts to a homomorphism from the mixed product to the simple product, so we can rely on the result of the Simple Product Lemma.

LEMMA IX.20 (MIXED PRODUCT HOMOMORPHISM). *There is a homomorphism from* $(I, I_1) \otimes^h G$ *to* $(I, I_1) \otimes G$.

PROOF. We use the homomorphism $h : I \to I_1$ to define the homomorphism h' from $I_\mathrm{m} := (I, I_1) \otimes^h G$ to $I_\mathrm{p} := (I, I_1) \otimes G$ by $h'((a, g)) := (h(a), g)$ for every $(a, g) \in \mathrm{dom}(I) \times G$.

Consider a fact $F = R(\boldsymbol{a})$ of I_m, with $a_i = (v_i, g_i)$ for all $R^i \in \mathsf{Pos}(R)$. Consider its image $F' = R(\boldsymbol{v})$ by the homomorphism from I_m to I obtained by projecting to the first component, and the image $h(F')$ of F' by the homomorphism h. As $h_{|I_1}$ is the identity and $h_{|(I \setminus I_1)}$ maps to $I_1 \setminus I_1$, $h(F')$ is a fact of I_1 iff F' is. Now by definition of the simple product it is clear that I_p contains the fact $h'(F)$: it was created in I_p from $h(F')$ for the same choice of $g \in G$. This shows that h' is indeed a homomorphism, which concludes the proof. □

We can now conclude our proof of the Universal Models Theorem (Theorem III.6). Let I_1 be the individualizing union of disjoint copies of I_0 and I_f be the superinstance of I_1 given by the Cautious Models Theorem applied to $k' := k \cdot (|\sigma| + 1)$. As I_1 is individualizing, we know that each element of I_1 is alone in its $\simeq_{k'}$-class in I_f, so the restriction of $I_\mathrm{f}/\simeq_{k'}$ to $\chi_{\simeq_{k'}}(I_1)$ is actually I_1 up to isomorphism; so we define I_f' to be $I_\mathrm{f}/\simeq_{k'}$ modified by identifying $\chi_{\simeq_{k'}}(I_1)$ to I_1; it is a finite superinstance of I_1. Let h be the homomorphism from I_f to I_f' obtained by modifying $\chi_{\simeq_{k'}}$ accordingly, which ensures that $h_{|I_1}$ is the identity and $h_{|I_\mathrm{f} \setminus I_1}$ maps to $I_\mathrm{f}' \setminus I_1$.

Let G be a $(2k+1)$-acyclic group generated by $\Lambda(I_\mathrm{f}')$, and consider $I_\mathrm{p} := (I_\mathrm{f}', I_1) \otimes G$. As I_f was k'-sound for ACQ, I_1 and Σ, so is I_f' by Lemma IX.10, so, as I_1 is individualizing, I_p is k-sound for CQ, I_1 and Σ by the Simple Product Lemma. However, as we explained, it may be the case that $I_\mathrm{p} \not\models \Sigma$. We therefore construct $I_\mathrm{m} := (I_\mathrm{f}, I_1) \otimes^h G$. By the Mixed Product Homomorphism Lemma, I_m has a homomorphism to I_p, so it is also k-sound for CQ, I_1 and Σ. Now, as I_1 is an individualizing superinstance of a disjoint union of copies of I_0, and as the fresh relations in the individualizing superinstance I_1 do not occur in queries or in constraints, it is clear that I_m is also k-sound for CQ, I_0 and Σ. Further, by the conditions ensured by the Cautious Models Theorem, I_f is cautious for h and I_1. So, by the Mixed Product Preservation Lemma, we have $I_\mathrm{m} \models \Sigma$ because $I_\mathrm{f} \models \Sigma$.

Hence, the mixed product I_m is a finite k-universal instance for CQ, I_0 and Σ. This concludes the proof of the Universal Models Theorem, and hence of our main theorem (Theorem III.3).

X. CONCLUSION

In this work we have developed the first techniques to build finite models on arbitrary arity schemas that satisfy both referential constraints and number restrictions, while controlling which CQs are satisfied. We have used these techniques to prove that finite open-world query answering for CQs, UIDs and FDs is finitely controllable up to finite closure of the dependencies. This allowed us to isolate the complexity of FQA for UIDs and FDs.

As presented the constructions are quite specific to dependencies, but in future work we will look to extend them to constraint languages containing disjunction, with the goal of generalizing to higher arity the rich arity-2 constraint languages of, e.g., [Ibáñez-García et al. 2014; Pratt-Hartmann 2009], while maintaining the decidability of FQA.

Acknowledgements. This work was supported in part by the Engineering and Physical Sciences Research Council, UK (EP/G004021/1) and the French ANR NormAtis project. We are very grateful to Balder ten Cate, Thomas Gogacz, Andreas Pieris, and Pierre Senellart for comments on earlier drafts, and to the anonymous reviewers of LICS for their valuable feedback.

REFERENCES

Serge Abiteboul, Richard Hull, and Victor Vianu. 1995. *Foundations of Databases*. Addison-Wesley.

William Ward Armstrong. 1974. Dependency structure of data base relationships. In *IFIP Congress*.

Vince Bárány, Georg Gottlob, and Martin Otto. 2010. Querying the guarded fragment. In *LICS*.

Andrea Calì, Georg Gottlob, and Andreas Pieris. 2012. Towards more expressive ontology languages: The query answering problem. *Artif. Intel.* 193 (2012).

Andrea Calì, Domenico Lembo, and Riccardo Rosati. 2003a. On the decidability and complexity of query answering over inconsistent and incomplete databases. In *PODS*.

Andrea Calì, Domenico Lembo, and Riccardo Rosati. 2003b. Query rewriting and answering under constraints in data integration systems. In *IJCAI*.

Marco A Casanova, Ronald Fagin, and Christos H Papadimitriou. 1984. Inclusion dependencies and their interaction with functional dependencies. *JCSS* 28, 1 (1984).

Stavros S. Cosmadakis, Paris C. Kanellakis, and Moshe Y. Vardi. 1990. Polynomial-time implication problems for unary inclusion dependencies. *JACM* 37, 1 (1990).

Alin Deutsch, Alan Nash, and Jeff Remmel. 2008. The chase revisited. In *SIGMOD*.

Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. 2003. Data Exchange: Semantics and Query Answering. In *ICDT*.

Yazmín Ibáñez-García, Carsten Lutz, and Thomas Schneider. 2014. Finite Model Reasoning in Horn Description Logics. In *KR*.

David S. Johnson and Anthony C. Klug. 1984. Testing Containment of Conjunctive Queries under Functional and Inclusion Dependencies. *JCSS* 28, 1 (1984).

Grigori A Margulis. 1982. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica* 2 (1982), 71–78.

Adrian Onet. 2013. The Chase Procedure and its Applications in Data Exchange. In *Data Exchange, Integration, and Streams*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 1–37.

Martin Otto. 2002. Modal and guarded characterisation theorems over finite transition systems. In *LICS*.

Martin Otto. 2012. Highly acyclic groups, hypergraph covers, and the guarded fragment. *JACM* 59, 1 (2012), 5.

Ian Pratt-Hartmann. 2009. Data-complexity of the two-variable fragment with counting quantifiers. *Inf. Comput.* 207, 8 (2009).

Riccardo Rosati. 2006. On the decidability and finite controllability of query processing in databases with incomplete information. In *PODS*.

Riccardo Rosati. 2008. Finite model reasoning in DL-Lite. In *ESWC*.

Riccardo Rosati. 2011. On the finite controllability of conjunctive query answering in databases under open-world assumption. *JCSS* 77, 3 (2011).

A. DETAILS ABOUT THE UID CHASE AND UNIQUE WITNESS PROPERTY

Recall the *Unique Witness Property*:

> For any element $a \in \mathrm{dom}(\mathsf{Chase}(I, \Sigma_{\mathrm{UID}}))$ and position R^p of σ, if two facts of $\mathsf{Chase}(I, \Sigma_{\mathrm{UID}})$ contain a at position R^p, then they are both facts of I.

We first give an example showing why this may not be guaranteed by the first round of the UID chase. Consider the instance $I = \{R(a), S(a)\}$ and the UIDs $\tau_1 : R^1 \subseteq T^1$ and $\tau_2 : S^1 \subseteq T^1$, where T is a binary relation. Applying a round of the UID chase creates the instance $\{R(a), S(a), T(a, b_1), T(a, b_2)\}$, with $T(a, b_1)$ being created by applying τ_1 to the active fact $R(a)$, and $T(a, b_2)$ being created by applying τ_2 to the active fact $S(a)$.

By contrast, the core chase would create only one of these two facts, because it would consider that two new facts are *equivalent*: they have the same exported element occurring at the same position. In general, the core chase keeps only one fact within each class of facts that are equivalent in this sense.

However, after one chase round by the core chase, there is no longer any distinction between the UID chase and the core chase, because the following property holds on the result I' of a chase round (be it by the core chase or by the UID chase) on any instance I'': (*) for any $\tau \in \Sigma_{\mathrm{UID}}$ and element $a \in \mathsf{Wants}(I', \tau)$, a occurs in only one fact of I'. This is true because Σ_{UID} is transitively closed, so we know that no UID of Σ_{UID} is applicable to an element of $\mathrm{dom}(I'')$ in I'; hence the only elements that witness violations occur in the one fact where they were introduced in I'.

We now claim that (*) implies that the Unique Witness Property holds when we chase by the core chase for the first round and the UID chase for subsequent rounds. Indeed, assume to the contrary that $a \in \mathrm{dom}(\mathsf{Chase}(I, \Sigma_{\mathrm{UID}}))$ violates the Property.

If $a \in \mathrm{dom}(I)$, because Σ_{UID} is transitively closed, after the first chase round on I, we no longer create any fact that involves a. Hence, each one of F_1 and F_2 is either a fact of I or a fact created in the first round of the chase (which is a chase round by the core chase). However, if one of F_1 and F_2 is in I, then it witnesses that we could not have $a \in \mathsf{Wants}(I, R^p)$, so it is not possible that the other fact was created in the first chase round. It cannot be the case either that F_1 and F_2 were both created in the first chase round, by definition of the core chase. Hence, F_1 and F_2 are necessarily both facts of I.

If $a \in \mathrm{dom}(\mathsf{Chase}(I, \Sigma_{\mathrm{UID}})) \setminus \mathrm{dom}(I)$, assume that a occurs at position R^p in two facts F_1, F_2. As $a \notin \mathrm{dom}(I)$, none of them is a fact of I. We then show a contradiction. It is not possible that one of those facts was created in a chase round before the other, as otherwise the second created fact could not have been created because of the first created fact. Hence, both facts must have been created in the same chase round. So there was a chase round from I'' to I' where we had $a \in \mathsf{Wants}(I'', R^p)$ and both F_1 and F_2 were created respectively from active facts F'_1 and F'_2 of I'' by UIDs $\tau_1 : S^q \subseteq R^p$ and $\tau_2 : T^r \subseteq R^p$. But then, by property (*), a occurs in only one fact, so as it occurs in F'_1 and F'_2 we have $F'_1 = F'_2$. Further, as $a \notin \mathrm{dom}(I)$, F'_1 and F'_2 are not facts of I either, so by definition of the UID chase and of the core chase, it is easy to see a occurs at only one position in $F'_1 = F'_2$. This implies that $\tau_1 = \tau_2$. Hence, we must have $F_1 = F_2$, a contradiction. This concludes the proof of the Unique Witness Property.