HAL
open science

# Rewriting the Infinite Chase 

Michael Benedikt, Maxime Buron, Stefano Germano, Kevin Kappelmann, Boris Motik

## To cite this version:

Michael Benedikt, Maxime Buron, Stefano Germano, Kevin Kappelmann, Boris Motik. Rewriting the Infinite Chase. VLDB 2022-48th International Conference on Very Large Databases, Sep 2022, Sydney, Australia. pp.3045-3057, 10.14778/3551793.3551851 . lirmm-03714595

HAL Id: lirmm-03714595
https://hal-lirmm.ccsd.cnrs.fr/lirmm-03714595
Submitted on 5 Jul 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. publics ou privés.

Distributed under a Creative Commons Attribution - NonCommercial - NoDerivatives 4.0
International License

# Rewriting the Infinite Chase 

Michael Benedikt<br>Oxford University<br>Oxford, United Kingdom<br>michael.benedikt@cs.ox.ac.uk

Maxime Buron<br>LIRMM, Inria, Univ. of Montpellier<br>Montpellier, France<br>maxime.buron@inria.fr

Stefano Germano<br>Oxford University<br>Oxford, United Kingdom<br>stefano.germano@cs.ox.ac.uk

Kevin Kappelmann<br>Technical University of Munich<br>Munich, Germany<br>kevin.kappelmann@tum.de

Boris Motik<br>Oxford University<br>Oxford, United Kingdom<br>boris.motik@cs.ox.ac.uk


#### Abstract

Guarded tuple-generating dependencies (GTGDs) are a natural extension of description logics and referential constraints. It has long been known that queries over GTGDs can be answered by a variant of the chase-a quintessential technique for reasoning with dependencies. However, there has been little work on concrete algorithms and even less on implementation. To address this gap, we revisit Datalog rewriting approaches to query answering, where GTGDs are transformed to a Datalog program that entails the same base facts on each base instance. We show that the rewriting can be seen as containing "shortcut" rules that circumvent certain chase steps, we present several algorithms that compute the rewriting by simulating specific types of chase steps, and we discuss important implementation issues. Finally, we show empirically that our techniques can process complex GTGDs derived from synthetic and real benchmarks and are thus suitable for practical use.

\section*{PVLDB Reference Format:}

Michael Benedikt, Maxime Buron, Stefano Germano, Kevin Kappelmann, and Boris Motik. Rewriting the Infinite Chase. PVLDB, 15(11): XXX-XXX, 2022. doi:XX.XX/XXX.XX

\section*{PVLDB Artifact Availability:}

The source code, data, and/or other artifacts have been made available at https://github.com/KRR-Oxford/Guarded-saturation.


## 1 INTRODUCTION

Tuple-generating dependencies (TGDs) are a natural extension of description logics and referential constraints, and they are extensively used in databases. For example, they are used in data integration to capture semantic restrictions on data sources, mapping rules between data sources and the mediated schema, and constraints on the mediated schema. A fundamental computational problem in such applications is query answering under TGDs: given a query $Q$, a collection of facts $I$, and a set of TGDs $\Sigma$, find all the answers to $Q$ that logically follow from $I$ and $\Sigma$. This problem has long been seen as a key component of a declarative data integration systems

[^0][25,33], and it also arises in answering querying using views and accessing data sources with restrictions [20, 26, 37].

The chase is a quintessential technique for reasoning with TGDs. It essentially performs "forward reasoning" by extending a set of given facts $I$ to a set $I^{\prime}$ of all facts implied by $I$ and a set of TGDs $\Sigma$. To answer a query, one can compute $I^{\prime}$ using the chase and then evaluate the query in $I^{\prime}$. Unfortunately, the chase does not necessarily terminate, and in fact query answering for general TGDs is undecidable. Considerable effort was devoted to identifying classes of TGDs for which query answering is decidable. One line of work has focused on TGDs where the chase terminates; weakly-acyclic TGDs [21] are perhaps the best-known such class. Another line of work focused on guarded TGDs (GTGDs). GTGDs are interesting since they can capture common constraints used in data integration, and ontologies expressed in variants of description logic (DL) [6] can be translated directly into GTGDs. Example 1.1 illustrates the use of GTGDs used in a practical data integration scenario.

Example 1.1. The IEC Common Information Model (CIM) is an open model for describing power generation and distribution networks. It is frequently used as a semantic layer in applications that integrate data about power systems [22]. CIM is defined in UML, but its formal semantics has been provided by a translation into an OWL ontology. The domain of CIM is described using classes and properties, which correspond to unary and binary relations, respectively. Moreover, semantic relationships between classes and properties are represented as OWL axioms, many of which can be translated into GTGDs. A significant portion of CIM describes power distribution equipment using GTGDs such as (1)-(4).

$$
\begin{array}{r}
\text { ACEquipment }(x) \rightarrow \exists y \text { hasTerminal }(x, y) \wedge \operatorname{ACTerminal}(y) \\
\qquad \operatorname{ACTerminal}(x) \rightarrow \operatorname{Terminal}(x) \\
\text { hasTerminal }(x, z) \wedge \operatorname{Terminal}(z) \rightarrow \operatorname{Equipment}(x)
\end{array}
$$

Data integration is then achieved by populating the vocabulary using mappings, which can be seen queries over the data sources that produce a set of facts called a base instance. A key issue in data integration is dealing with incompleteness of data sources. For example, it is not uncommon that one data source mentions two switches $\mathrm{sw}_{1}$ and $\mathrm{sw}_{2}$, while another data source provides information about connected terminals only for switch $\mathrm{sw}_{1}$.

$$
\begin{array}{r}
\text { ACEquipment }\left(\mathrm{sw}_{1}\right) \quad \text { ACEquipment }\left(\mathrm{sw}_{2}\right) \\
\text { hasTerminal }\left(\mathrm{sw}_{1}, \operatorname{trm}_{1}\right) \quad \text { ACTerminal }\left(\operatorname{trm}_{1}\right) \tag{6}
\end{array}
$$

GTGDs can be used to complete the data. For example, if a user asks to list all pieces of equipment known to the system, both $\mathrm{sw}_{1}$ and $\mathrm{sw}_{2}$ will be returned, even though the base instance does not explicitly classify either switch as a piece of equipment.

Even though the chase for GTGDs does not necessarily terminate, query answering for GTGDs is decidable [34]. To prove decidability, one can argue that the result of a chase is tree-like-that is, the facts derived by the chase can be arranged into a particular kind of tree. Next, one can develop a finite representation of potentially infinite trees. One possibility is to describe the trees using a finite tree automaton, so query answering can be reduced to checking automaton emptiness. While theoretically elegant, this method is not amenable to practical use: building the automaton and the emptiness test are both complex and expensive, and such algorithms always exhibit worst-case complexity. Alternatively, one can use blocking to identify a tree prefix sufficient for query evaluation. Blocking is commonly used in description logic reasoning [6], and it was later lifted to guarded logic [27]. However, blocking was shown to be impractical for query answering: the required tree prefix can be much larger than the base instance $I$ so, as $I$ grows in size, the size of the tree prefix becomes unmanageable.

More promising query answering techniques for GTGDs are based on Datalog rewriting [35]. The idea was initially proposed by Marnette [35], and it was later extended to broader classes of TGDs [10, 24] and settings [11]. The main idea is to transform an input set of GTGDs $\Sigma$ into a set rew $(\Sigma)$ of Datalog rules such that $\Sigma$ and $\operatorname{rew}(\Sigma)$ entail the same base facts on each base instance. Thus, given a base instance $I$, instead of computing the chase of $I$ and $\Sigma$ (which may not terminate), we compute the chase $I^{\prime}$ of $I$ and $\operatorname{rew}(\Sigma)$. Since Datalog rules essentially correspond to existential-free TGDs, $I^{\prime}$ is always finite and it can be computed using optimized Datalog engines. Moreover, $\Sigma$ and $\operatorname{rew}(\Sigma)$ entail the same base facts on $I$, so we can answer any existential-free conjunctive query (i.e., queries where all variables are answer variables) by evaluating in $I^{\prime}$. The restriction to existential-free queries is technical: existentially quantified variables in a query can be matched to objects introduced by existential quantification, and these are not preserved in a Datalog rewriting. However, practical queries are typically existential-free since all query variables are usually answer variables.

Example 1.2. A Datalog program consisting of rules (2)-(3) and (7) is a rewriting of GTGDs (1)-(4).

$$
\begin{equation*}
\text { ACEquipment }(x) \rightarrow \text { Equipment }(x) \tag{7}
\end{equation*}
$$

Rule (7) is a logical consequence of GTGDs (1)-(3), and it provides a "shortcut" for the inferences of the other GTGDs.

The advantage of rewriting-based approaches is scalability in the size of the base instance $I$. Such techniques have been implemented and practically validated in the context of description logics [28, 29], but practical algorithms have not yet been proposed for GTGDs. This raises several theoretical and practical questions.
How to compute the Datalog rules needed for completeness? Existing Datalog rewriting algorithms often prove their correctness indirectly. For example, completeness of a rewriting algorithm for description logics [29] uses a proof-theoretic argument, which does not provide an intuition about why the algorithm actually works.

Our first contribution is to relate Datalog rewriting approaches to the chase. Towards this goal, we introduce the one-pass variant of the chase, which we use to develop a general completeness criterion for Datalog rewriting algorithms. This, in turn, provides us with a better understanding of how rewriting algorithms work, and it allows us to discover new algorithms in a systematic way.
What does the space of rewriting algorithms look like? Computing the rewriting $\operatorname{rew}(\Sigma)$ usually requires extending $\Sigma$ with certain logical consequences of $\Sigma$. We show that we can select the relevant consequences using different criteria. Some methods require deriving TGDs with existential quantifiers in the head, others generate Datalog rules directly, and yet other methods derive logical implications with function symbols. We relate all of these methods to the one-pass chase mentioned earlier, and we provide theoretical worst-case guarantees about their performance.
How do we ensure scalability of rewriting algorithms? Implementations of Datalog rewriting algorithms have thus far been mainly considered in the setting of description logics [29, 38]. To the best of our knowledge, we provide the first look at optimization and implementation of Datalog rewriting algorithms for GTGDs. We achieve scalability by developing and combining various indexing and redundancy elimination techniques.
How do we evaluate rewriting algorithms? We provide a benchmark for GTGD query answering algorithms, and we use it to evaluate our methods. To the best of our knowledge, this is the first attempt to evaluate query answering techniques for GTGDs.
Summary of contributions. We give an extensive account of Datalog rewriting for GTGDs. In particular, we develop a theoretical framework that allows us to understand, motivate, and show completeness of rewriting algorithms. Moreover, we present several concrete algorithms, establish worst-case complexity bounds, and discuss their relationships. We complement this theoretical analysis with a discussion of how to adapt techniques from first-order theorem proving to the setting of GTGDs. Finally, we empirically evaluate our techniques using an extensive benchmark. All proofs and the details of one algorithm are given in the appendix of this paper. Our implementation and a more detailed account of our experimental results can be found online [13].

## 2 RELATED WORK

Answering queries via rewriting has been extensively considered in description logics. For example, queries over ontologies in the DL-Lite family of languages can be rewritten into first-order queries [18], and fact entailment for $\mathcal{S H} I Q$ ontologies can be rewritten to disjunctive Datalog [29]. These techniques provide the foundation for the Ontop [17] and KAON2 [38] systems, respectively.

In the context of TGDs, first-order rewritings were considered in data integration systems with inclusion and key dependencies [16]. Datalog rewritings have been considered for GTGDs [35] and their extensions such as frontier-guarded TGDs [11], and nearly frontierguarded and nearly guarded TGDs [24]. The focus in these studies was to identify complexity bounds and characterize expressivity of TGD classes rather than provide practical algorithms. Existing implements of query answering for TGDs use first-order rewriting for linear TGDs [48], chase variants for TGDs with terminating chase
[14], chase with blocking for warded TGDs [12], chase with the magic sets transformation for shy TGDs [3], and Datalog rewriting for separable and weakly separable TGDs [49]. These TGD classes are all different from GTGDs, and we are unaware of any attempts to implement and evaluate GTGD rewriting algorithms.

Our algorithms are related to resolution-based decision procedures for variants of guarded logics [19, 23,50]. Moreover, our characterization of Datalog rewritings is related to a chase variant used to answer queries over data sources with access patterns [4]. Finally, a variant of the one-pass chase from Section 4 was generalized to the broader context of disjunctive GTGDs [31].

## 3 PRELIMINARIES

In this section, we recapitulate the well-known definitions and notation that we use to formalize our technical results.

TGDs. Let consts, vars, and nulls be pairwise disjoint, infinite sets of constants, variables, and labeled nulls, respectively. A term is a constant, a variable, or a labeled null; moreover, a term is ground if it does not contain a variable. For $\alpha$ a formula or a set thereof, consts $(\alpha)$, vars $(\alpha)$, nulls $(\alpha)$, and terms $(\alpha)$ are the sets of constants, free variables, labeled nulls, and terms, respectively, in $\alpha$.

A schema is a set of relations, each of which is associated with a nonnegative integer arity. A fact is an expression of the form $R(\vec{t})$, where $R$ is an $n$-ary relation and $\vec{t}$ is a vector of $n$ ground terms; moreover, $R(\vec{t})$ is a base fact if $\vec{t}$ contains only constants. An instance $I$ is a finite set of facts, and $I$ is a base instance if it contains only base facts. An atom is an expression of the form $R(\vec{t})$, where $R$ is an $n$-ary relation and $\vec{t}$ is a vector of $n$ terms not containing labeled nulls. Thus, each base fact is an atom. We often treat conjunctions as sets of conjuncts; for example, for $\gamma$ a conjunction of facts and $I$ an instance, $\gamma \subseteq I$ means that each conjunct of $\gamma$ is contained $I$.

A tuple generating dependency (TGD) is a first-order formula of the form $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$, where $\beta$ and $\eta$ are conjunctions of atoms, $\eta$ is not empty, the free variables of $\beta$ are $\vec{x}$, and the free variables of $\eta$ are contained in $\vec{x} \cup \vec{y}$. Conjunction $\beta$ is the body and formula $\exists \vec{y} \eta$ is the head of the TGD. We often omit $\forall \vec{x}$ when writing a TGD. A TGD is full if $\vec{y}$ is empty; otherwise, the TGD is non-full. A TGD is in head-normal form if it is full and its head contains exactly one atom, or it is non-full and each head atom contains at least one existentially quantified variable. Each TGD can be easily transformed to an equivalent set of TGDs in head-normal form. A full TGD in head-normal form is a Datalog rule, and a Datalog program is a finite set of Datalog rules. The head-width (hwidth) and the body-width (bwidth) of a TGD are the numbers of variables in the head and body, respectively; these are extended to sets of TGDs by taking the maxima over all TGDs. The notion of an instance satisfying a TGD is inherited from first-order logic. A base fact $F$ is entailed by an instance $I$ and a finite set of TGDs $\Sigma$, written $I, \Sigma \neq F$, if $F \in I^{\prime}$ holds for each instance $I^{\prime} \supseteq I$ that satisfies $\Sigma$.

A substitution $\sigma$ is a function that maps finitely many variables to terms. The domain and the range of $\sigma$ are dom $(\sigma)$ and $\mathrm{rng}(\sigma)$, respectively. For $\gamma$ a term, a vector of terms, or a formula, $\sigma(\gamma)$ is obtained by replacing each free occurrence of a variable $x$ in $\gamma$ such that $x \in \operatorname{dom}(\sigma)$ with $\sigma(x)$.
Fact Entailment for Guarded TGDs. Fact entailment for general TGDs is semidecidable, and many variants of the chase can be used
to define a (possibly infinite) set of facts that is homomorphically contained in each modef of a base instance and a set of TGDs.

Fact entailment is decidable for guarded TGDs (GTGDs): a TGD $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$ is guarded if $\beta$ contains an atom (called a guard) that contains all variables of $\vec{x}$. Note that a guard need not be unique in $\beta$. Let $\Sigma$ be a finite set of GTGDs. We say that a set of ground terms $G$ is $\Sigma$-guarded by a fact $R(\vec{t})$ if $G \subseteq \vec{t} \cup$ consts $(\Sigma)$. Moreover, $G$ is $\Sigma$-guarded by a set of facts $I$ if $G$ is $\Sigma$-guarded by some fact in $I$. Finally, a fact $S(\vec{u})$ is $\Sigma$-guarded by a fact $R(\vec{t})$ (respectively a set of facts $I$ ) if $\vec{u}$ is $\Sigma$-guarded by $R(\vec{t})$ (respectively $I$ ).
By adapting the reasoning techniques for guarded logics [5, 47] and referential database constraints [30], fact entailment for GTGDs can be decided by a chase variant that works on tree-like structures. A chase tree $T$ consists of a directed tree, one tree vertex that is said to be recently updated, and a function mapping each vertex $v$ in the tree to a finite set of facts $T(v)$. A chase tree $T$ can be transformed to another chase tree $T^{\prime}$ in the following two ways.

- One can apply a chase step with a GTGD $\tau=\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$ in head-normal form. The precondition is that there exist a vertex $v$ in $T$ and a substitution $\sigma$ with domain $\vec{x}$ such that $\sigma(\beta) \subseteq T(v)$. The result of the chase step is obtained as follows.
- If $\tau$ is full (and thus $\eta$ is a single atom), then chase tree $T^{\prime}$ is obtained from $T$ by making $v$ recently updated in $T^{\prime}$ and setting $T^{\prime}(v)=T(v) \cup\{\sigma(\eta)\}$.
- If $\tau$ is not full, then $\sigma$ is extended to a substitution $\sigma^{\prime}$ that maps each variable in $\vec{y}$ to a labeled null not occurring in $T$, and chase tree $T^{\prime}$ is obtained from $T$ by introducing a fresh child $v^{\prime}$ of $v$, making $v^{\prime}$ recently updated in $T^{\prime}$, and setting $T\left(v^{\prime}\right)=\sigma^{\prime}(\eta) \cup\left\{F \in T(v) \mid F\right.$ is $\Sigma$-guarded by $\left.\sigma^{\prime}(\eta)\right\}$.
- One can apply a propagation step from a vertex $v$ to a vertex $v^{\prime}$ in $T$. Chase tree $T^{\prime}$ is obtained from $T$ by making $v^{\prime}$ recently updated in $T^{\prime}$ and setting $T^{\prime}\left(v^{\prime}\right)=T\left(v^{\prime}\right) \cup S$ for some nonempty set $S$ satisfying $S \subseteq\left\{F \in T(v) \mid F\right.$ is $\Sigma$-guarded by $\left.T\left(v^{\prime}\right)\right\}$.
A tree-like chase sequence for a base instance $I$ and a finite set of GTGDs $\Sigma$ in head-normal form is a finite sequence of chase trees $T_{0}, \ldots, T_{n}$ such that $T_{0}$ contains exactly one root vertex $r$ that is recently updated in $T_{0}$ and $T_{0}(r)=I$, and each $T_{i}$ with $0<i \leq n$ is obtained from $T_{i-1}$ by a chase step with some $\tau \in \Sigma$ or a propagation step. For each vertex $v$ in $T_{n}$ and each fact $F \in T_{n}(v)$, this sequence is a tree-like chase proof of $F$ from $I$ and $\Sigma$. It is well known that $I, \Sigma \vDash F$ if and only if there exists a tree-like chase proof of $F$ from $I$ and $\Sigma$ (e.g., [34]). Example 4.3 in Section 4 illustrates these definitions. One can decide $I, \Sigma \neq F$ by imposing an upper bound on the size of chase trees that need to be considered [34].
Rewriting. A Datalog rewriting of a finite set of TGDs $\Sigma$ is a Datalog program $\operatorname{rew}(\Sigma)$ such that $I, \Sigma \mid=F$ if and only if $I, \operatorname{rew}(\Sigma) \vDash F$ for each base instance $I$ and each base fact $F$. If $\Sigma$ contains GTGDs only, then a Datalog rewriting $\operatorname{rew}(\Sigma)$ is guaranteed to exist (which is not the case for general TGDs). Thus, we can reduce fact entailment for GTGDs to Datalog reasoning, which can be solved using highly optimized Datalog techniques [1, 38]. For example, given a base instance $I$, we can compute the materialization of $\operatorname{rew}(\Sigma)$ on $I$ by applying the rules of $\operatorname{rew}(\Sigma)$ to $I$ up to a fixpoint. This will compute precisely all base facts entailed by $\operatorname{rew}(\Sigma)$ (and thus also by $\Sigma$ ) on $I$, and it can be done in time polynomial in the size of $I$.

Encoding Existentials by Function Symbols. It is sometimes convenient to represent existentially quantified values using functional terms. In such cases, we use a slightly modified notions of terms, atoms, and rules. It will be clear from the context which definitions we use in different parts of the paper.

We adjust the notion of a term as either a constant, a variable, or an expression of the form $f(\vec{t})$ where $f$ is an $n$-ary function symbol and $\vec{t}$ is a vector of $n$ terms. The notions of ground terms, (base) facts, and (base) instances, and atoms are the same as before, but they use the modified notion of terms. A rule is a first-order implication of the form $\forall \vec{x}[\beta \rightarrow H]$ where $\beta$ is a conjunction of atoms whose free variables are $\vec{x}$, and $H$ is an atom whose free variables are contained in $\vec{x}$; as for TGDs, we often omit $\forall \vec{x}$. A rule thus contains no existential quantifiers, but its head contains exactly one atom that can contain function symbols. Also, a Datalog rule, a function-free rule, and a full TGD in head-normal form are all synonyms. Finally, a base fact still contains only constants.

Skolemization allow us to replace existential quantifiers in TGDs by functional terms. Specifically, let $\tau=\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$, and let $\sigma$ be a substitution defined on each $y \in \vec{y}$ as $\sigma(y)=f_{\tau, y}(\vec{x})$ where $f_{\tau, y}$ is a fresh $|\vec{x}|$-ary Skolem symbol uniquely associated with $\tau$ and $y$. Then, the Skolemization of $\tau$ produces rules $\forall \vec{x}[\beta \rightarrow \sigma(H)]$ for each atom $H \in \eta$. Moreover, the Skolemization $\Sigma^{\prime}$ of a finite set of TGDs $\Sigma$ is the union of the rules obtained by Skolemizing each $\tau \in \Sigma$. It is well known that $I, \Sigma=F$ if and only if $I, \Sigma^{\prime} \vDash F$ for each base instance $I$ and each base fact $F$.

Unification. A unifier of atoms $A_{1}, \ldots, A_{n}$ and $B_{1}, \ldots, B_{n}$ is a substitution $\theta$ such that $\theta\left(A_{i}\right)=\theta\left(B_{i}\right)$ for $1 \leq i \leq n$. Such $\theta$ is a most general unifier (MGU) if, for each unifier $\sigma$ of $A_{1}, \ldots, A_{n}$ and $B_{1}, \ldots, B_{n}$, there exists a substitution $\rho$ such that $\sigma=\rho \circ \theta$ (where $\circ$ is function composition). An MGU is unique up to variable renaming if it exists, and it can be computed in time $O\left(\sum_{i=1}^{n}\left|A_{i}\right|+\left|B_{i}\right|\right)$ where $\left|A_{i}\right|$ and $\left|B_{i}\right|$ are the encoding sizes of $A_{i}$ and $B_{i}[41,42]$.

## 4 CHASE-BASED DATALOG REWRITING

Our objective is to develop rewriting algorithms that can handle complex GTGDs. Each algorithm will derive Datalog rules that provide "shortcuts" in tree-like chase proofs: instead of introducing a child vertex $v^{\prime}$ using a chase step with a non-full GTGD at vertex $v$, performing some inferences in $v^{\prime}$, and then propagating a derived fact $F$ back from $v^{\prime}$ to $v$, these "shortcuts" will derive $F$ in one step without having to introduce $v^{\prime}$. The main question is how to derive all "shortcuts" necessary for completeness while keeping the number of derivations low. In this section we lay the technical foundations that will allow us to study different strategies for deriving "shortcuts" in Section 5. We show that, instead of considering arbitrary chase proofs, we can restrict our attention to chase proofs that are one-pass according to Definition 4.1. Then, we identify the parts of such proofs that we need to be able to circumvent using "shortcuts". Finally, we present sufficient conditions that guarantee completeness of rewriting algorithms. We start by describing formally the structure of tree-like chase proofs.
Definition 4.1. A tree-like chase sequence $T_{0}, \ldots, T_{n}$ for a base instance $I$ and a finite set of GTGDs $\Sigma$ in head-normal form is onepass if, for each $0<i \leq n$, chase tree $T_{i}$ is obtained by applying one of the following two steps to the recently updated vertex $v$ of $T_{i-1}$ :


Figure 1: Tree-Like Chase Sequence for Example 4.3

- a propagation step copying exactly one fact from $v$ to its parent, or
- a chase step with a GTGD from $\Sigma$ provided that no propagation step from $v$ to the parent of $v$ is applicable.
Thus, each step in a tree-like chase sequence is applied to a "focused" vertex; steps with non-full TGDs move the "focus" from a parent to a child, and propagation steps move the "focus" in the opposite direction. Moreover, once a child-to-parent propagation takes place, the child cannot be revisited in further steps. Theorem 4.2 states a key property about chase proofs for GTGDs: whenever a proof exists, there exists a one-pass proof too. Example 4.3 illustrates important aspects of Definition 4.1 and Theorem 4.2.

Theorem 4.2. For each base instance $I$, each finite set of GTGDs $\Sigma$ in head-normal form, and each base fact $F$ such that $I, \Sigma \vDash F$, there exists a one-pass tree-like chase proof of $F$ from $I$ and $\Sigma$.

Example 4.3. Let $I=\{A(a, b)\}$ and let $\Sigma$ contain GTGDs (8)-(13).

$$
\begin{align*}
A\left(x_{1}, x_{2}\right) & \rightarrow \exists y B\left(x_{1}, y\right) \wedge C\left(x_{1}, y\right)  \tag{8}\\
C\left(x_{1}, x_{2}\right) & \rightarrow D\left(x_{1}, x_{2}\right)  \tag{9}\\
B\left(x_{1}, x_{2}\right) \wedge D\left(x_{1}, x_{2}\right) & \rightarrow E\left(x_{1}\right)  \tag{10}\\
A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) & \rightarrow \exists y_{1}, y_{2} F\left(x_{1}, y_{1}\right) \wedge F\left(y_{1}, y_{2}\right)  \tag{11}\\
E\left(x_{1}\right) \wedge F\left(x_{1}, x_{2}\right) & \rightarrow G\left(x_{1}\right)  \tag{12}\\
B\left(x_{1}, x_{2}\right) \wedge G\left(x_{1}\right) & \rightarrow H\left(x_{1}\right) \tag{13}
\end{align*}
$$

A tree-like chase sequence for $I$ and $\Sigma$ is shown in Figure 1, and it provides a proof of the base fact $H(a)$ from $I$ and $\Sigma$. The recently updated vertex of each chase tree is shown in red. We denote the root vertex by $r$, and its left and right children by $v_{1}$ and $v_{2}$, respectively. The step producing $T_{7}$ from $T_{6}$ does not satisfy the requirements of one-pass chase: it propagates the fact $G(a)$ from $v_{2}$ to $v_{1}$, where the latter is a "sibling" of the former.

To obtain a one-pass chase sequence, we could try to "slow down" the propagation of $G(a)$ : we first propagate $G(a)$ from $v_{2}$ to $r$, and then from $r$ to $v_{1}$. The former step is allowed in one-pass chase, but the latter step is not: once we leave the subtree rooted at $v_{1}$, we are not allowed to revisit it later. Note, however, that $B\left(a, n_{1}\right)$ and $G(a)$ must occur jointly in a vertex of a chase tree in order to derive $H(a)$. Moreover, note that no reordering of chase steps will


Figure 2: One-Pass Chase Sequence Obtained from Figure 1
derive $H(a)$ : we must first produce $v_{1}$ to be able to derive $v_{2}$, and we must combine $G(a)$ from $v_{2}$ and $B\left(a, n_{1}\right)$ from $v_{1}$.

The solution, which is used in the proof of Theorem 4.2, is to replace propagation to the child by "regrowing" the entire subtree. In our example, we replace the steps producing $T_{7}$ and $T_{8}$ with the steps shown in Figure 2. Chase tree $T_{7}^{1}$ is obtained from $T_{6}$ by propagating $G(a)$ from $v_{2}$ to $r$. Then, instead of propagating $G(a)$ from $r$ to $v_{1}$, a new vertex $v_{3}$ is created in $T_{7}^{2}$ by reapplying (8) and fact $G(a)$ is pushed to $v_{3}$ as part of the chase step with a non-full GTGD. This allows $H(a)$ to be derived in vertex $v_{3}$ of $T_{8}^{1}$.

Fact $D\left(n_{3}\right)$ can be derived in vertex $v_{3}$, but this is not needed to prove $H(a)$. Moreover, our chase is oblivious [34]: a non-full TGD can be applied to the same facts several times, each time introducing a fresh vertex and fresh labeled nulls. The number of children of a vertex is thus not naturally bounded, and our objective is not to apply all chase steps exhaustively to obtain a universal model of $\Sigma$. Instead, we are interested only in chase proofs, which must only contain steps needed to demonstrate entailment of a specific fact. $\triangleleft$

One-pass chase proofs are interesting because they can be decomposed into loops as described in Definition 4.4.

Definition 4.4. For $T_{0}, \ldots, T_{n}$ a one-pass tree-like chase sequence for some I and $\Sigma$, a loop at vertex $v$ with output fact $F$ is a subsequence $T_{i}, \ldots, T_{j}$ with $0 \leq i<j \leq n$ such that

- $T_{i+1}$ is obtained by a chase step with a non-full GTGD,
- $T_{j}$ is obtained by a propagation step that copies $F$, and
- $v$ is the recently updated vertex of both $T_{i}$ and $T_{j}$.

The length of the loop is defined as $j-i$.
Example 4.5. Subsequence $T_{0}, T_{1}, T_{2}, T_{3}, T_{4}$ of the chase trees from Example 4.3 is a loop at the root vertex $r$ with output fact $E(a)$ : chase tree $T_{1}$ is obtained by applying a non-full GTGD to $r$, and chase tree $T_{4}$ is obtained by propagating $E(a)$ back to $r$. Analogously, $T_{4}, T_{5}, T_{6}, T_{7}^{1}$ is another loop at $r$ with output fact $G(a)$. Finally, $T_{7}^{1}, T_{7}^{2}, T_{8}^{1}, T_{9}$ is a loop at $r$ with output fact $H(a)$.

Thus, a loop is a subsequence of chase steps that move the "focus" from a parent to a child vertex, perform a series of inferences in the child and its descendants, and finally propagate one fact back to
the parent. If non-full TGDs are applied to the child, then the loop can be recursively decomposed into further loops at the child. The properties of the one-pass chase ensure that each loop is finished as soon as a fact is derived in the child that can be propagated to the parent, and that the vertices introduced in the loop are not revisited at any later point in the proof. In this way, each loop at vertex $v$ can be seen as taking the set $T_{i}(v)$ as input and producing the output fact $F$ that is added to $T_{j}(v)$. This leads us to the following idea: for each loop with the input set of facts $T_{i}(v)$, a rewriting should contain a "shortcut" Datalog rule that derives the loop's output.

Example 4.6. One can readily check that rules (14)-(16) provide "shortcuts" for the three loops identified in Example 4.5.

$$
\begin{align*}
A\left(x_{1}, x_{2}\right) & \rightarrow E\left(x_{1}\right)  \tag{14}\\
A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) & \rightarrow G\left(x_{1}\right)  \tag{15}\\
A\left(x_{1}, x_{2}\right) \wedge G\left(x_{1}\right) & \rightarrow H\left(x_{1}\right) \tag{16}
\end{align*}
$$

Moreover, these are all relevant "shortcuts": the union of rules (14)(16) and the Datalog rules from Example 4.3-that is, rules (9), (10), (12), and (13)-is a rewriting of the set $\Sigma$ from Example 4.1.

These ideas are formalized in Proposition 4.7, which will provide us with a correctness criterion for our algorithms.

Proposition 4.7. A Datalog program $\Sigma^{\prime}$ is a rewriting of a finite set of GTGDs $\Sigma$ in head-normal form if

- $\Sigma^{\prime}$ is a logical consequence of $\Sigma$,
- each Datalog rule of $\Sigma$ is a logical consequence of $\Sigma^{\prime}$, and
- for each base instance I, each one-pass tree-like chase sequence $T_{0}, \ldots, T_{n}$ for $I$ and $\Sigma$, and each loop $T_{i}, \ldots, T_{j}$ at the root vertex $r$ with output fact $F$, there exist a Datalog rule $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(r)$ and $\sigma(H)=F$.

Intuitively, the first condition ensures soundness: rewriting $\Sigma^{\prime}$ should not derive more facts than $\Sigma$. The second condition ensures that $\Sigma^{\prime}$ can mimic direct applications of Datalog rules from $\Sigma$ at the root vertex $r$. The third condition ensures that $\Sigma^{\prime}$ can reproduce the output of each loop at vertex $r$ using a "shortcut" Datalog rule.

## 5 REWRITING ALGORITHMS

We now consider ways to produce "shortcut" Datalog rules satisfying Proposition 4.7. In Subsection 5.1 we present the ExbDR algorithm that manipulates GTGDs directly, and in Subsections 5.2 and 5.3 we present the SkDR and HypDR algorithms, respectively, that manipulate rules obtained by Skolemizing the input GTGDs. All of these algorithms can produce intermediate GTGDs/rules that are not necessarily Datalog rules. In Appendix E we present the FullDR algorithm that manipulates GTGDs, but derives only Datalog rules. However, the performance of FullDR proved to not be competitive, so we do not discuss it any further here.

Each algorithm is defined by an inference rule Inf that can be applied to several TGDs/rules to derive additional TGDs/rules. For simplicity, we use the same name for the rule and the resulting algorithm. Given a set of GTGDs $\Sigma$, the algorithm applies Inf to (the Skolemization of) $\Sigma$ as long as possible and then returns all produced Datalog rules. This process, however, can derive a large
number of TGDs/rules, so it is vital to eliminate TGDs/rules whenever possible. We next define notions of redundancy that can be used to discard certain TGDs/rules produced by Inf.

Definition 5.1. $A$ TGD $\tau_{1}=\forall \vec{x}_{1}\left[\beta_{1} \rightarrow \exists \vec{y}_{1} \eta_{1}\right]$ is a syntactic tautology if it is in head-normal form and $\beta_{1} \cap \eta_{1} \neq \emptyset$. TGD $\tau_{1}$ subsumes a TGD $\tau_{2}=\forall \vec{x}_{2}\left[\beta_{2} \rightarrow \exists \vec{y}_{2} \eta_{2}\right]$ if there exists a substitution $\mu$ such that $\operatorname{dom}(\mu)=\vec{x}_{1} \cup \vec{y}_{1}, \mu\left(\vec{x}_{1}\right) \subseteq \vec{x}_{2}, \mu\left(\vec{y}_{1}\right) \subseteq \vec{y}_{1} \cup \vec{y}_{2}, \mu(y) \neq \mu\left(y^{\prime}\right)$ for distinct $y$ and $y^{\prime}$ in $\vec{y}_{1}, \mu\left(\beta_{1}\right) \subseteq \beta_{2}$, and $\mu\left(\eta_{1}\right) \supseteq \eta_{2}$.

A rule $\tau_{1}=\forall \vec{x}_{1}\left[\beta_{1} \rightarrow H_{1}\right]$ is a syntactic tautology if $H_{1} \in \beta_{1}$. Rule $\tau_{1}$ subsumes a rule $\tau_{2}=\forall \vec{x}_{2}\left[\beta_{2} \rightarrow H_{2}\right]$ if there exists a substitution $\mu$ such that $\mu\left(\beta_{1}\right) \subseteq \beta_{2}$ and $\mu\left(H_{1}\right)=H_{2}$.

A TGD/rule $\tau$ is contained in a set of TGDs/rules $S$ up to redundancy if $\tau$ is a syntactic tautology or some $\tau^{\prime} \in S$ subsumes $\tau$.

The following example illustrates Definition 5.1.
Example 5.2. Rule $A(x) \wedge B(x) \rightarrow A(x)$ is a syntactic tautology: applying a chase step with it cannot produce a new fact. A non-full TGD in head-normal form cannot be a syntactic tautology since each head atom of such a TGD contains an existentially quantified variable that does not occur in the TGD body.

Rule $\tau_{1}=A\left(f\left(x_{1}\right), f\left(x_{1}\right)\right) \wedge B\left(x_{1}\right) \rightarrow B\left(f\left(x_{1}\right)\right)$ is subsumed by rule $\tau_{2}=A\left(x_{2}, x_{3}\right) \rightarrow B\left(x_{2}\right)$ using substitution $\mu_{1}$ that maps both $x_{2}$ and $x_{3}$ to $f\left(x_{1}\right)$. If $\tau_{1}$ derives $B(f(t))$ in one step from a set of facts $I$ by a substitution $\sigma$ where $\sigma\left(x_{1}\right)=t$, then $\tau_{2}$ also derives $B(f(t))$ from $I$ in one step by substitution $\sigma \circ \mu_{1}$. Thus, rule $\tau_{1}$ is not needed when rule $\tau_{2}$ is present, so $\tau_{1}$ can be discarded.

While syntactic tautologies and rule subsumption are standard in first-order theorem proving [8], subsumption of TGDs is more involved. TGD $\tau_{3}=A\left(x_{1}, x_{1}\right) \wedge B\left(x_{1}\right) \rightarrow \exists y_{1} C\left(x_{1}, y_{1}\right)$ is subsumed by TGD $\tau_{4}=A\left(x_{2}, x_{3}\right) \rightarrow \exists y_{2}, y_{3} C\left(x_{2}, y_{2}\right) \wedge D\left(x_{3}, y_{3}\right)$ by substitution $\mu_{2}$ where $\mu_{2}\left(x_{2}\right)=\mu_{2}\left(x_{3}\right)=x_{1}, \mu_{2}\left(y_{2}\right)=y_{1}$, and $\mu_{2}\left(y_{3}\right)=y_{3}$. The conditions on substitution $\mu_{2}$ in Definition 5.1 ensure that $y_{2}$ and $y_{3}$ are not mapped to each other or to $x_{1}$. Thus, as in the previous paragraph, the result of each chase step with $\tau_{3}$ and substitutions $\sigma$ and $\sigma^{\prime}$ can always be obtained (up to isomorphism) by a chase step with $\tau_{4}$ and substitutions $\sigma \circ \mu_{2}$ and $\sigma^{\prime} \circ \mu_{2}$.

In Definition 5.3 we formalize the notion of applying Inf exhaustively up to redundancy. The definition, however, does not say how to actually do it: we discuss this and other issues in Section 6.
Definition 5.3. For $\operatorname{Inf}$ an inference rule and $\Sigma$ a finite set of GTGDs, $\operatorname{Inf}(\Sigma)$ is the subset of all Skolem-free Datalog rules of $\Sigma^{\prime}$, where $\Sigma^{\prime}$ is the smallest set that contains up to redundancy each TGD/rule obtained by

- transforming $\Sigma$ into head-normal form if Inf manipulates TGDs or Skolemizing $\Sigma$ if $\operatorname{Inf}$ manipulates rules, and
- selecting an adequate number of premises in $\Sigma^{\prime}$, renaming any variables shared by distinct premises, applying Inf to the renamed premises, and transforming the result into head-normal form.


### 5.1 The Existential-Based Rewriting

As we discussed in Section 4, each loop $T_{i}, \ldots, T_{j}$ at vertex $v$ in a one-pass chase sequence can be seen as taking $T_{i}(v)$ as input and producing one fact included in $T_{j}(v)$ as output. Let $v^{\prime}$ be child of $v$ introduced in $T_{i+1}$. The idea behind the ExbDR algorithm is to derive all GTGDs such that, for each $k$ with $i<k \leq j$, all facts of $T_{k}\left(v^{\prime}\right)$


Figure 3: Deriving "shortcuts" for the loop $T_{0}-T_{4}$ in ExbDR
can be derived from the input $T_{i}(v)$ in one step. The output of the loop can then also be derived from $T_{i}(v)$ in one step by full GTGD, so this GTGD provides us with the desired loop "shortcut". Before formalizing this idea, we slightly adapt the notion of unification.

Definition 5.4. For $X$ a set of variables, an $X$-unifier and an $X$ MGU $\theta$ of atoms $A_{1}, \ldots, A_{n}$ and $B_{1}, \ldots, B_{n}$ are defined as in Section 3, but with the additional requirement that $\theta(x)=x$ for each $x \in X$.

It is straightforward to see that an $X$-MGU is unique up to the renaming of variables not contained in $X$, and that it can be computed as usual while treating variables in $X$ as if they were constants. We are now ready to formalize the ExbDR algorithm.

Definition 5.5. The Existential-Based Datalog Rewriting inference rule ExbDR takes two guarded TGDs

$$
\begin{aligned}
\tau & =\forall \vec{x}\left[\beta \rightarrow \exists \vec{y} \eta \wedge A_{1} \wedge \cdots \wedge A_{n}\right] \quad \text { with } n \geq 1 \text { and } \\
\tau^{\prime} & =\forall \vec{z}\left[A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}\right]
\end{aligned}
$$

and, for $\theta$ a $\vec{y}-M G U$ of $A_{1}, \ldots, A_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$, , if $\theta(\vec{x}) \cap \vec{y}=\emptyset$ and $\operatorname{vars}\left(\theta\left(\beta^{\prime}\right)\right) \cap \vec{y}=\emptyset$, it derives

$$
\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \exists \vec{y} \theta(\eta) \wedge \theta\left(A_{1}\right) \wedge \cdots \wedge \theta\left(A_{n}\right) \wedge \theta\left(H^{\prime}\right)
$$

Example 5.6. Consider again the set $\Sigma$ from Example 4.3. The idea behind the ExbDR algorithm is illustrated in Figure 3, which summarizes the steps of the loop $T_{0}, T_{1}, T_{2}, T_{3}, T_{4}$ from Figure 1. We denote the vertices by $r$ and $v_{1}$ as in Example 4.3.

Fact $A(a, b)$ is the input to the loop, and the first step of the loop derives $B\left(a, n_{1}\right)$ and $C\left(a, n_{1}\right)$ using GTGD (8). Next, GTGD (9) evolves vertex $v_{1}$ by deriving $D\left(a, n_{1}\right)$. To capture this, the ExbDR inference rule combines (8), the GTGD that creates $v_{1}$, with (9), the GTGD that evolves $v_{1}$. This produces GTGD (17), which derives all facts of $v_{1}$ from the input fact in one step. Vertex $v_{1}$ is evolved further using GTGD (10) to derive $E(a)$. To reflect this, the ExbDR inference rule combines (17) and (10) to produce (18), which again derives all facts of $v_{1}$ from the loop's input in one step.

$$
\begin{align*}
& A\left(x_{1}, x_{2}\right) \rightarrow \exists y B\left(x_{1}, y\right) \wedge C\left(x_{1}, y\right) \wedge D\left(x_{1}, y\right)  \tag{17}\\
& A\left(x_{1}, x_{2}\right) \rightarrow \exists y B\left(x_{1}, y\right) \wedge C\left(x_{1}, y\right) \wedge D\left(x_{1}, y\right) \wedge E\left(x_{1}\right) \tag{18}
\end{align*}
$$

Fact $E(a)$ does not contain the labeled null $n_{1}$ that is introduced when creating $v_{1}$, so it can be propagated to the root vertex $r$ as the output of the loop. This is reflected in (18): atom $E\left(x_{1}\right)$ does not contain any existential variables. Definition 5.3 requires each derived GTGD to be brought into head-normal, so (18) is broken up into (17) and (14). The latter GTGD is full, and it provides us with the desired shortcut for the loop.

Next, (12) and atom $F\left(x_{1}, y_{1}\right)$ of (11) produce (19), and transformation into head-normal form produces (11) and (15). Moreover, (8) and (13) produce (20), and transformation (20) into head-normal form produces (16) and (21).

$$
\begin{align*}
& A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) \rightarrow \exists y_{1}, y_{2} F\left(x_{1}, y_{1}\right) \wedge F\left(y_{1}, y_{2}\right) \wedge G\left(x_{1}\right)  \tag{19}\\
& A\left(x_{1}, x_{2}\right) \wedge G\left(x_{1}\right) \rightarrow \exists y B\left(x_{1}, y\right) \wedge C\left(x_{1}, y\right) \wedge H\left(x_{1}\right)  \tag{20}\\
& A\left(x_{1}, x_{2}\right) \wedge G\left(x_{1}\right) \rightarrow \exists y B\left(x_{1}, y\right) \wedge C\left(x_{1}, y\right) \tag{21}
\end{align*}
$$

GTGD (21) is subsumed by (8) so it can be dropped. No further inferences are possible after this, so all derived full GTGDs are returned as the rewriting of $\Sigma$.

Before proceeding, we present an auxiliary result showing certain key properties of the ExbDR inference rule.

Proposition 5.7. Each application of the ExbDR inference rule to $\tau, \tau^{\prime}$, and $\theta$ as in Definition 5.5 satisfies the following properties.

1. Some atom $A_{i}^{\prime}$ with $1 \leq i \leq n$ is a guard in $\tau^{\prime}$.
2. For each $1 \leq i \leq n$ such that $A_{i}^{\prime}$ is a guard of $\tau^{\prime}$, and for $\sigma$ the $\vec{y}$ $M G U$ of $A_{i}^{\prime}$ and the corresponding atom $A_{i}$ such that $\sigma(\vec{x}) \cap \vec{y}=\emptyset$, it is the case that $\operatorname{vars}\left(\sigma\left(A_{j}^{\prime}\right)\right) \cap \vec{y} \neq \emptyset$ for each $1 \leq j \leq n$.
3. The result is a GTGD whose body and head width are at most bwidth $(\Sigma)$ and hwidth $(\Sigma)$, respectively.
In the second claim of Proposition 5.7, $\sigma$ unifies only $A_{i}$ and $A_{i}^{\prime}$, whereas $\theta$ unifies all $A_{1}, \ldots, A_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$; thus, $\sigma$ and $\theta$ are not necessarily the same. The third claim is needed to prove termination of ExbDR.

Proposition 5.7 can be used to guide the application of the ExbDR inference rule. Consider an attempt to apply the ExbDR inference rule to two candidate GTGDs $\tau=\beta \rightarrow \exists \vec{y} \eta$ and $\tau^{\prime}=\beta^{\prime} \rightarrow H^{\prime}$. The first claim of Proposition 5.7 tells us that a guard of $\tau^{\prime}$ will definitely participate in the inference. Thus, we can choose one such guard $G^{\prime} \in \beta^{\prime}$ of $\tau^{\prime}$ and try to find a $\vec{y}$-MGU $\sigma$ of $G^{\prime}$ and a counterpart atom $G \in \eta$ from the head of $\tau$. Next, we need to check whether $\sigma(\vec{x}) \cap \vec{y}=\emptyset$; if not, there is no way for $\theta(\vec{x}) \cap \vec{y}=\emptyset$ to hold so the inference is not possible. By the second claim of Proposition 5.7, all candidates for the atoms participating in the inference will contain a variable that is mapped by $\sigma$ to a member of $\vec{y}$; thus, $S^{\prime}=\left\{\sigma\left(A^{\prime}\right) \mid A^{\prime} \in \beta^{\prime} \wedge \operatorname{vars}\left(\sigma\left(A^{\prime}\right)\right) \cap \vec{y} \neq \emptyset\right\}$ is the set of all relevant side atoms. Note that we apply $\sigma$ to the atoms in $S^{\prime}$ to simplify further matching. The next step is to identify the corresponding head atoms of $\tau$. To achieve this, for each atom $A^{\prime} \in S^{\prime}$ of the form $R\left(t_{1}, \ldots, t_{n}\right)$, we identify the set $C\left[A^{\prime}\right]$ of candidate counterpart atoms as the set of atoms of the form $R\left(s_{1}, \ldots, s_{n}\right) \in \sigma(\eta)$ such that, for each argument position $i$ with $1 \leq i \leq n$, if either $t_{i} \in \vec{y}$ or $s_{i} \in \vec{y}$, then $t_{i}=s_{i}$. Finally, we consider each possible combination $S$ of such candidates, and we try to find an MGU $\theta$ of sets $S$ and $S^{\prime}$. If unification succeeds, we derive the corresponding GTGD.

Theorem 5.8. Program $\operatorname{ExbDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in
 the maximum relation arity in $\Sigma, w_{b}=\operatorname{bwidth}(\Sigma), w_{h}=\operatorname{hwidth}(\Sigma)$, $c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Program $\operatorname{ExbDR}(\Sigma)$ can thus be large in the worst case. In Section 7 we show empirically that rewritings are suitable for practical
use. From a theoretical point of view, checking fact entailment via $\operatorname{ExbDR}(\Sigma)$ is worst-case optimal. To see why, let $r, a$, and $c$ be as in Theorem 5.8, and consider a base instance $I$ with $c^{\prime}$ constants. The fixpoint of $\operatorname{ExbDR}(\Sigma)$ on $I$ contains at most $r\left(c+c^{\prime}\right)^{a}$ facts, and it can be computed in time $O\left(r\left(c+c^{\prime}\right)^{a} \cdot|\operatorname{ExbDR}(\Sigma)|\right)$ : each rule $\tau \in \operatorname{ExbDR}(\Sigma)$ is guarded so we can apply a chase step with $\tau$ by matching a guard and then checking the remaining body atoms. Hence, we can compute $\operatorname{ExbDR}(\Sigma)$ and find its fixpoint in 2ExpTime, in ExpTime if the relation arity is fixed, and in PTime if $\Sigma$ is fixed (i.e., if we consider data complexity). These results match the lower bounds for checking fact entailment for GTGDs [34].

### 5.2 Using Skolemization

The ExbDR algorithm exhibits two drawbacks. First, each application of the ExbDR inference rule potentially introduces a head atom, so the rule heads can get very long. Second, each inference requires matching a subset of body atoms of $\tau^{\prime}$ to a subset of the head atoms of $\tau$; despite the optimizations outlined after Proposition 5.7, this can be costly, particularly when rule heads are long.

We would ideally derive GTGDs with a single head atom and unify just one body atom of $\tau^{\prime}$ with the head atom of $\tau$, but this does not seem possible if we stick to manipulating GTGDs. For example, atoms $C(y)$ and $D(y)$ of GTGD (17) refer to the same labeled null (represented by variable $y$ ), and this information would be lost if we split (17) into two GTGDs. We thus need a way to refer to the same existentially quantified object in different logical formulas. This can be achieved by replacing existentially quantified variables by Skolem terms, which in turns gives rise to the SkDR algorithm from Definition 5.10. Before presenting the algorithm, in Definition 5.9 we generalize the notion of guardedness to rules.

Definition 5.9. Rule $\forall \vec{x}[\beta \rightarrow H]$ is guarded if each function symbol in the rule is a Skolem symbol, the body $\beta$ contains a Skolem-free atom $A \in \beta$ such that $\operatorname{vars}(A)=\vec{x}$, and each Skolem term in the rule is of the form $f(\vec{t})$ where $\operatorname{vars}(f(\vec{t}))=\vec{x}$ and $\vec{t}$ is function-free.

Definition 5.10. The Skolem Datalog Rewriting inference rule SkDR takes two guarded rules

$$
\tau=\beta \rightarrow H \quad \text { and } \quad \tau^{\prime}=A^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}
$$

such that

- $\beta$ is Skolem-free and H contains a Skolem symbol, and
- $A^{\prime}$ contains a Skolem symbol, or $\tau^{\prime}$ is Skolem-free and $A^{\prime}$ contains all variables of $\tau^{\prime}$, and, for $\theta$ an $M G U$ of $H$ and $A^{\prime}$, it derives

$$
\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)
$$

Example 5.11. Skolemizing GTGDs (8) and (11) produces rules (22)-(23), and (24)-(25), respectively.

$$
\begin{align*}
A\left(x_{1}, x_{2}\right) & \rightarrow B\left(x_{1}, f\left(x_{1}, x_{2}\right)\right)  \tag{22}\\
A\left(x_{1}, x_{2}\right) & \rightarrow C\left(x_{1}, f\left(x_{1}, x_{2}\right)\right)  \tag{23}\\
A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) & \rightarrow F\left(x_{1}, g\left(x_{1}, x_{2}\right)\right)  \tag{24}\\
A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) & \rightarrow F\left(g\left(x_{1}, x_{2}\right), h\left(x_{1}, x_{2}\right)\right) \tag{25}
\end{align*}
$$

Intuitively, rules (22) and (23) jointly represent the facts introduced by the non-full GTGD (8): functional term $f\left(x_{1}, x_{2}\right)$ allows both
rules to "talk" about the same labeled nulls. This allows the SkDR inference rule to simulate the ExbDR inference rule while unifying just pairs of atoms. In particular, SkDR combines (22) and (10) to obtain (26); it combines (23) and (9) to obtain (27); and it combines (26) and (27) to obtain the "shortcut" rule (14).

$$
\begin{align*}
A\left(x_{1}, x_{2}\right) \wedge D\left(x_{1}, f\left(x_{1}, x_{2}\right)\right) & \rightarrow E\left(x_{1}\right)  \tag{26}\\
A\left(x_{1}, x_{2}\right) & \rightarrow D\left(x_{1}, f\left(x_{1}, x_{2}\right)\right) \tag{27}
\end{align*}
$$

The rules with Skolem-free bodies derived in this way allow us to reconstruct derivations in one step analogously to Example 5.6, and the rules with Skolem symbols in body atoms capture the intermediate derivation steps. For example, rules (26) and (28) capture the result of matching the first and the second body atom, respectively, of rule (10) to facts produced by rules (22) and (27), respectively. To complete the rewriting, SkDR combines (24) with (12) to obtain (15), and it combines (22) with (13) to derive (16).

However, SkDR also combines (10) and (27) into (28), which with (22) derives (14) the second time. These inferences are superfluous: they just process the two body atoms of (10) in a different order. Also, SkDR combines (12) and (25) into rule (29), which is a "deadend" in that it does not further contribute to a Datalog rule.

$$
\begin{align*}
A\left(x_{1}, x_{2}\right) \wedge B\left(x_{1}, f\left(x_{1}, x_{2}\right)\right) & \rightarrow E\left(x_{1}\right)  \tag{28}\\
A\left(x_{1}, x_{2}\right) \wedge E\left(x_{1}\right) \wedge E\left(g\left(x_{1}, x_{2}\right)\right) & \rightarrow G\left(g\left(x_{1}, x_{2}\right)\right) \tag{29}
\end{align*}
$$

Our HypDR algorithm in Subsection 5.3 can avoid these overheads, but at the expense of using more than two rules at a time.

Proposition 5.12 and Theorem 5.13 capture the relevant properties of the SkDR algorithm.

Proposition 5.12. Each application of the SkDR inference rule to rules $\tau$ and $\tau^{\prime}$ as in Definition 5.10 produces a guarded rule.

Theorem 5.13. Program $\operatorname{SkDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in time $O\left(b^{r^{d} \cdot\left(e+w_{b}+c\right)^{d a}}\right)$ for $r$ the number of relations in $\Sigma$, a the maximum relation arity in $\Sigma$, e the number of existential quantifiers in $\Sigma, w_{b}=\operatorname{bwidth}(\Sigma), c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

It is natural to wonder whether SkDR is guaranteed to be more efficient than ExbDR. We next show that neither algorithm is generally better: there exist families of inputs on which SkDR performs exponentially more inferences than ExbDR, and vice versa.

Proposition 5.14. There exists a family $\left\{\Sigma_{n}\right\}_{n \in \mathbb{N}}$ of finite sets of GTGDs such that the number of GTGDs derived by ExbDR is $O\left(2^{n}\right)$ times larger than the number of rules derived by SkDR on each $\Sigma_{n}$.

Proof. For each $n \in \mathbb{N}$, let $\Sigma_{n}$ contain the following GTGDs.

$$
\begin{align*}
A(x) & \rightarrow \exists \vec{y} B_{1}\left(x, y_{1}\right) \wedge \cdots \wedge B_{n}\left(x, y_{n}\right)  \tag{30}\\
B_{i}\left(x_{1}, x_{2}\right) \wedge C_{i}\left(x_{1}\right) & \rightarrow D_{i}\left(x_{1}, x_{2}\right) \text { for } 1 \leq i \leq n \tag{31}
\end{align*}
$$

On such $\Sigma_{n}$, ExbDR derives a GTGD of the form (32) for each subset $\left\{k_{1}, \ldots, k_{m}\right\} \subseteq\{1, \ldots, n\}$, and there are $2^{n}$ such TGDs. In contrast, the Skolemization of (30) consists of $n$ rules shown in equation (33), so SkDR derives just $n$ rules shown in equation (34).

$$
\begin{equation*}
A(x) \wedge \bigwedge_{i=1}^{m} C_{k_{i}}(x) \rightarrow \exists \vec{y} \bigwedge_{i=1}^{n} B_{i}\left(x, y_{i}\right) \wedge \bigwedge_{i=1}^{m} D_{k_{i}}\left(x, y_{k_{i}}\right) \tag{32}
\end{equation*}
$$

$$
\begin{align*}
A(x) & \rightarrow B_{i}\left(x, f_{i}(x)\right) \text { for } 1 \leq i \leq n  \tag{33}\\
A(x) \wedge C_{i}(x) & \rightarrow D_{i}\left(x, f_{i}(x)\right) \text { for } 1 \leq i \leq n \tag{34}
\end{align*}
$$

Proposition 5.15. There exists a family $\left\{\Sigma_{n}\right\}_{n \in \mathbb{N}}$ of finite sets of GTGDs such that the number of rules derived by SkDR is $O\left(2^{n}\right)$ times larger than the number of TGDs derived by $\operatorname{ExbDR}$ on each $\Sigma_{n}$.

Proof. For each $n \in \mathbb{N}$, let $\Sigma_{n}$ contain the following GTGDs.

$$
\begin{array}{r}
A(x) \rightarrow \exists y B_{1}(x, y) \wedge \cdots \wedge B_{n}(x, y) \\
B_{1}\left(x_{1}, x_{2}\right) \wedge \cdots \wedge B_{n}\left(x_{1}, x_{2}\right) \rightarrow C\left(x_{1}\right) \tag{36}
\end{array}
$$

On such $\Sigma_{n}$, ExbDR derives just GTGD (37) in one step. In contrast, the Skolemization of (35) consists of $n$ rules of the form (38) for each $1 \leq i \leq n$. Thus, SkDR combines these with (36) to derive $2^{n}-1$ rules of the form (39), one for each subset $\left\{k_{1}, \ldots, k_{m}\right\} \subsetneq\{1, \ldots, n\}$.

$$
\begin{align*}
& A(x) \rightarrow C(x)  \tag{37}\\
& A(x) \rightarrow B_{i}(x, f(x))  \tag{38}\\
& A(x) \wedge B_{k_{1}}(x, f(x)) \wedge \cdots \wedge B_{k_{m}}(x, f(x)) \rightarrow C(x) \tag{39}
\end{align*}
$$

### 5.3 Combining Several SkDR Steps into One

The SkDR algorithm can produce many rules with Skolem symbols in the body, which is the main reason for Proposition 5.15. We next present the HypDR algorithm, which uses the hyperresolution inference rule as a kind of "macro" to combine several SkDR steps into one. We show that this can be beneficial for several reasons.
Definition 5.16. The Hyperresolution Rewriting inference rule HypDR takes guarded rules

$$
\begin{aligned}
\tau_{1} & =\beta_{1} \rightarrow H_{1} \quad \ldots \quad \tau_{n}=\beta_{n} \rightarrow H_{n} \text { and } \\
\tau^{\prime} & =A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}
\end{aligned}
$$

such that

- for each $i$ with $1 \leq i \leq n$, conjunction $\beta_{i}$ is Skolem-free and atom $H_{i}$ contains a Skolem symbol, and
- rule $\tau^{\prime}$ is Skolem-free,
and, for $\theta$ an $M G U$ of $H_{1}, \ldots, H_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$, if conjunction $\theta\left(\beta^{\prime}\right)$ is Skolem-free, it derives

$$
\theta\left(\beta_{1}\right) \wedge \cdots \wedge \theta\left(\beta_{n}\right) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)
$$

Example 5.17. The HypDR inference rule simulates chase steps in the child vertex of a loop analogously to ExbDR: all body atoms matching a fact introduced in the child vertex are resolved in one step. We can see two benefits of this on our running example.

First, HypDR derives (27) from (23) and (9), and it derives (14) from (10), (22), and (27). Rule (14) is derived just once, and without intermediate rules (26) and (28). In other words, the HypDR inference rule does not resolve the body atoms of a rule in every possible order. As Proposition 5.20 below shows, this can reduce the number of derived rules by an exponential factor.

Second, HypDR derives only rules with Skolem-free bodies, and thus does not derive the "dead-end" rule (29). In other words, all consequences of HypDR derive in one step one fact in the child vertex of a loop from the loop's input $T_{i}(v)$.

The downside of HypDR is that more than two rules can participate in an inference. This requires more complex unification and selection of candidates that can participate in an inference.

Proposition 5.18 and Theorem 5.19 capture the properties of HypDR, and Proposition 5.20 compares it to SkDR.

Proposition 5.18. Each application of the HypDR inference rule to rules $\tau_{1}, \ldots, \tau_{n}$ and $\tau^{\prime}$ as in Definition 5.16 produces a guarded rule.

Theorem 5.19. Program $\operatorname{HypDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in
 maximum relation arity in $\Sigma$, e the number of existential quantifiers in $\Sigma, w_{b}=\operatorname{bwidth}(\Sigma), c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Proposition 5.20. There exists a family $\left\{\Sigma_{n}\right\}_{n \in \mathbb{N}}$ of finite sets of GTGDs such that SkDR derives $O\left(2^{n}\right)$ more rules than HypDR on each $\Sigma_{n}$.

Proof. For each $n \in \mathbb{N}$, let $\Sigma_{n}$ contain the following GTGDs.

$$
\begin{array}{r}
A(x) \rightarrow \exists y B(x, y) \\
B\left(x_{1}, x_{2}\right) \wedge C_{i}\left(x_{1}\right) \rightarrow D_{i}\left(x_{1}, x_{2}\right) \text { for } 1 \leq i \leq n \\
D_{1}\left(x_{1}, x_{2}\right) \wedge \cdots \wedge D_{n}\left(x_{1}, x_{2}\right) \rightarrow E\left(x_{1}\right) \tag{42}
\end{array}
$$

Skolemizing (40) produces (43). Thus, SkDR combines (43) with each (41) to derive each (44), and it uses (44) and (42) to derive $2^{n}-1$ rules of the form (45) for each set of indexes $I$ satisfying $\emptyset \subseteq I \subseteq\{1, \ldots, n\}$; note that none of these rules are redundant.

$$
\begin{array}{r}
A(x) \rightarrow B(x, f(x)) \\
A(x) \wedge C_{i}(x) \rightarrow D_{i}(x, f(x)) \text { for } 1 \leq i \leq n \\
A(x) \wedge \bigwedge_{i \in I} C_{i}(x) \wedge \bigwedge_{j \in\{1, \ldots, n\} \backslash I} D_{j}(x, f(x)) \rightarrow E(x) \tag{45}
\end{array}
$$

In contrast, HypDR derives each (44) just like SkDR, and it combines in one step (42) and all (44) to derive (45) for $I=\{1, \ldots, n\}$.

## 6 IMPLEMENTATION AND OPTIMIZATIONS

In this section, we discuss numerous issues that have to be addressed to make the computation of a rewriting practical.

Computing $\operatorname{lnf}(\Sigma)$ in Practice. Definition 5.3 does not specify how to compute the set $\Sigma^{\prime}$, and redundancy elimination makes this question nontrivial. When Inf derives a TGD/rule $\tau$, we can apply subsumption in two ways. First, we can discard $\tau$ if $\tau$ is subsumed by a previously derived TGD/rule; this is known as forward subsumption. Second, if $\tau$ is not discarded, we can discard each previously derived TGD/rule that is subsumed by $\tau$; this is known as backward subsumption. The set of derived TGD/rules can thus grow and shrink, so the application of Inf has to be carefully structured to ensure that all inferences are performed eventually.

We address this problem by a variant of the Otter loop [36] used in first-order theorem provers. The pseudo-code is shown in Algorithm 1. The algorithm maintains two sets of TGDs/rules: the worked-off set $\mathcal{W}$ contain TGDs/rules that have been processed by Inf, and the unprocessed set $\mathcal{U}$ contains TGDs/rules that are still to be processed. Set $\mathcal{W}$ is initially empty (line 1 ), and set $\mathcal{U}$ is initialized to the head-normal form of $\Sigma$ if Inf manipulates TGDs, or to the Skolemization of $\Sigma$ if Inf manipulates rules. The algorithm then processes each $\tau \in \mathcal{U}$ until $\mathcal{U}$ becomes empty (lines 3-10). It is generally beneficial to process shorter TGDs/rules first as that improves chances of redundancy elimination. After moving $\tau$ to $\mathcal{W}$ (line 5),

```
Algorithm 1 Computing \(\operatorname{lnf}(\Sigma)\) for \(\Sigma\) a finite set of GTGDs
    \(\mathcal{W}=\emptyset\)
    \(\mathcal{U}=\) the head-normal form or the Skolemization of \(\Sigma\)
    while \(\mathcal{U} \neq \emptyset\) do
        Choose some \(\tau \in \mathcal{U}\) and remove it from \(\mathcal{U}\)
        \(\mathcal{W}=\mathcal{W} \cup\{\tau\}\)
        Let \(\mathcal{E}\) be the result of applying \(\operatorname{Inf}\) to \(\tau\) and a subset of \(\mathcal{W}\)
                and transforming the result into head-normal form
        for each \(\tau^{\prime} \in \mathcal{E}\) do
            if \(\tau^{\prime}\) is not contained in \(\mathcal{W} \cup \mathcal{U}\) up to redundancy then
                Remove from \(\mathcal{W}\) and \(\mathcal{U}\) each \(\tau^{\prime \prime}\) subsumed by \(\tau^{\prime}\)
                \(\mathcal{U}=\mathcal{U} \cup\left\{\tau^{\prime}\right\}\)
    return \(\{\tau \in \mathcal{W} \mid \tau\) is a Skolem-free Datalog rule \(\}\)
```

the algorithm applies Inf to $\tau$ and $\mathcal{W}$ and transforms the results into head-normal form (line 6). The algorithm discards each resulting $\tau^{\prime} \in \mathcal{E}$ that is a syntactic tautology or is forward-subsumed by an element of $\mathcal{W} \cup \mathcal{U}$ (line 8). If $\tau^{\prime}$ is not discarded, the algorithm applies backward subsumption to $\tau^{\prime}, \mathcal{W}$, and $\mathcal{U}$ (line 9) and adds $\tau^{\prime}$ to $\mathcal{U}$ (line 10). When all TGDs/rules are processed, the algorithm returns all Skolem-free Datalog rules from $\mathcal{W}$ (line 11). The result of applying Inf to TGDs/rules in $\mathcal{W}$ is thus contained in $\mathcal{W} \cup \mathcal{U}$ up to redundancy at all times so, upon algorithm's termination, set $\mathcal{W}$ satisfies the condition on $\Sigma^{\prime}$ from Definition 5.3.
Checking Subsumption. Checking whether TGD/rule $\tau_{1}$ subsumes $\tau_{2}$ is NP-complete [32], and the main difficulty is in matching the variables of $\tau_{1}$ to the variables of $\tau_{2}$. Thus, we use an approximate check in our implementation. First, we normalize each TGD to use fixed variables $x_{1}, x_{2}, \ldots$ and $y_{1}, y_{2}, \ldots$ : we sort the body and head atoms by their relations using an arbitrary, but fixed ordering and breaking ties arbitrarily, and then we rename all variables so that the $i^{t h}$ distinct occurrence of a universally (respectively existentially) quantified variable from left to right is $x_{i}$ (respectively $y_{i}$ ). To see whether $\tau_{1}=\beta_{1} \rightarrow \exists \vec{y} \eta_{1}$ subsumes $\tau_{2}=\beta_{2} \rightarrow \exists \vec{y} \eta_{2}$, we determine whether $\beta_{1} \subseteq \beta_{2}$ and $\eta_{1} \supseteq \eta_{2}$ holds, which requires only polynomial time. We use a similar approximation for rules. Variable normalization ensures termination, and using a modified subsumption check does not affect the correctness of the rewriting: set $\mathcal{W}$ may contain more TGDs/rules than strictly necessary, but these are all logical consequences of (the Skolemization of) $\Sigma$.
Subsumption Indexing. Sets $\mathcal{W}$ and $\mathcal{U}$ can be large, so we use a variant of feature vector indexing [44] to retrieve subsumption candidates in $\mathcal{W} \cup \mathcal{U}$. For simplicity, we consider only TGDs in the following discussion, but rules can be handled analogously. Note that a TGD $\tau_{1}$ can subsume TGD $\tau_{2}$ only if the set of relations occurring in the body of $\tau_{1}$ (respectively the head of $\tau_{2}$ ) is a subset of the set of relations occurring in the body of $\tau_{2}$ (respectively the head of $\tau_{1}$ ). Thus, we can reduce the problem of retrieving subsumption candidates to the problem of, given a domain set $D$, a set $N$ of subsets of $D$, a subset $S \subseteq D$, and $\bowtie \in\{\subseteq, \supseteq\}$, retrieving each $S^{\prime} \in N$ satisfying $S^{\prime} \bowtie S$. The set-trie data structure [43] can address this problem. The idea is to order $D$ in an arbitrary, yet fixed way, so that we can treat each subset of $N$ as a word over $D$. We then index $N$ by constructing a trie over the words representing the elements
of $N$. Finally, we retrieve all $S^{\prime} \in N$ satisfying $S^{\prime} \bowtie S$ by traversing the trie, where the ordering on $D$ allows us to considerably reduce the number of vertices we visit during the traversal.

A minor issue is that retrieving TGDs that subsume a given TGD requires both subset and superset testing for body and head relations, respectively, and vice versa for retrieval of subsumed TGDs. To address this, we introduce a distinct symbol $R^{b}$ and $R^{h}$ for each relation $R$ occurring in $\Sigma$, and we represent each $\operatorname{TGD} \tau$ as a feature vector $F_{\tau}$ of these symbols corresponding to the body and head of $\tau$. Moreover, we combine in the obvious way the subset and superset retrieval algorithms. For example, when searching for a TGD $\tau^{\prime} \in \mathcal{W} \cup \mathcal{U}$ that subsumes a given TGD $\tau$, we use the subset retrieval for the symbols $R^{b}$ and the superset retrieval for symbols $R^{h}$. Finally, we order these symbols by the decreasing frequency of the order of the symbols' occurrence in the set $\Sigma$ of input TGDs, and moreover we order each $R^{b}$ before all $R^{h}$.
Relation Clustering. We observed that the subsumption indexes can easily get very large, so index traversal can become a considerable source of overhead. To reduce the index size, we group the symbols $R^{b}$ and $R^{h}$ into clusters $C^{b}$ and $C^{h}$, respectively. Then, the feature vector $F_{\tau}$ associated with each TGD $\tau$ consists of all clusters $C^{b}$ and $C^{h}$ that contain a relation occurring in the body and head, respectively, of $\tau$. We adapt the trie traversal algorithms in the obvious way to take into account this change. The number of clusters is computed using the average numbers of symbols and atoms in the input TGDs, and clusters are computed with the aim of balancing the number of TGDs stored in each leaf vertex.
Unification Indexing. We construct indexes over $\mathcal{W}$ that allow us to quickly identify TGDs/rules that can participate in an inference with some $\tau$. For TGDs, we maintain a hash table that maps each relation $R$ to a set of TGDs containing $R$ in the body, and another hash table that does the same but for TGD heads. To index rules, we use a variant of a path indexing [45]: each atom in a rule is represented as a sequence of relations and function symbols occurring in the atom, and such sequences are entered into two tries (one for body and one for head atoms). Then, given rule $\tau$, we consider each body and head atom $A$ of $\tau$, we convert $A$ into the corresponding sequence, and we use the sequence to query the relevant trie for all candidates participating in an inference with $\tau$ on $A$.
Cheap Lookahead Optimization. Consider an application of the ExbDR inference rule to GTGDs $\tau$ and $\tau^{\prime}$ as in Definition 5.5, producing a GTGD $\tau^{\prime \prime}$ where $\operatorname{vars}\left(\theta\left(H^{\prime}\right)\right) \cap \vec{y} \neq \emptyset$ and the relation of $H^{\prime}$ does not occur in the body of a GTGD in $\Sigma$. In each one-pass chase sequence for some base instance and $\Sigma$, no GTGD of $\Sigma$ can be applied to a fact obtained by instantiating $\theta\left(H^{\prime}\right)$, so deriving this fact is redundant. Consequently, we can drop such $\tau^{\prime \prime}$ as soon as we derive it in line 6. Analogously, when the SkDR inference rule is applied to rules $\tau$ and $\tau^{\prime}$ as in Definition 5.10, we can drop the resulting rule if $\theta\left(H^{\prime}\right)$ is not full and it contains a relation not occurring in the body of a GTGD in $\Sigma$.

## 7 EXPERIMENTAL EVALUATION

We implemented a system that can produce a Datalog rewriting of a set of GTGDs using our algorithms, and we conducted an empirical evaluation using a comprehensive collection of 428 synthetic and

Table 1: Input GTGDs at a Glance

| Inputs | \# Full TGDs |  |  |  | \# Non-Full TGDs |  |  |  |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
|  | Min | Max | Avg | Med | Min | Max | Avg | Med |
| 428 | 1 | 171,905 | 11,030 | 789 | 2 | 156,743 | 5,255 | 283 |

realistic inputs. Our objectives were to show that our algorithms can indeed rewrite complex GTGDs, and that the rewriting can be successfully processed by modern Datalog systems. In Subsection 7.1 we describe the test setting. Then, in Subsection 7.2 we discuss the rewriting experiments with GTGDs obtained from ontologies, and in Subsection 7.3 we validate the usefulness of the rewriting approach end-to-end. Finally, in Subsection 7.4 we discuss rewriting GTGDs of higher arity. Due to the very large number of inputs, we can only summarize our results in this paper; however, our complete evaluation results are available online [13].

### 7.1 Input GTGDs, Competitors, \& Test Setting

Before discussing our results, we next describe our test setting.
Input GTGDs. We are unaware of any publicly available sets of GTGDs that we could readily use in our evaluation, so we derived the input GTGDs for our evaluation from the ontologies in the Oxford Ontology Library [40]. At the time of writing, this library contained 787 ontologies, each assigned a unique five-digit identifier. After removing closely-related ontology variants, we were left with 428 core ontologies. We loaded each ontology using the parser from the Graal system [9], discarded axioms that cannot be translated into GTGDs, and converted the remaining axioms into GTGDs. We used the standard translation of description logics into first-order logic [6], where each class corresponds to a unary relation, and each property corresponds to a binary relation. We thus obtained 428 sets of input GTGDs with properties shown in Table 1.

To evaluate our algorithms on TGDs containing relations of arity higher than two, we devised a way to "blow up" relation arity. Given a set of GTGDs and a blowup factor $b$, our method proceeds as follows. First, in each atom of each GTGD, it replaces each variable argument with $b$ fresh variables uniquely associated with the variable; for example, for $b=2$, atom $A(x, y)$ is transformed into atom $A\left(x_{1}, x_{2}, y_{1}, y_{2}\right)$. Next, the method randomly introduces fresh head and body atoms over the newly introduced variables; in doing so, it ensures that the new atoms do not introduce patterns that would prevent application of the ExbDR inference rule.
Competitors. We compared the ExbDR, SkDR, and HypDR algorithms, implemented as described in Section 6. As noted in Section 2, no existing system we are aware of implements a Datalog rewriting algorithm for GTGDs. However, the KAON2 system [29, 38, 39] can rewrite GTGDs obtained from OWL ontologies, so we used KAON2 as a baseline in our experiments with OWL-based GTGDs. We made sure that all inputs to KAON2 and our algorithms include only GTGDs that all methods can process.
Test Setting. We conducted all experiments on a laptop with an Intel Core i5-6500 CPU @ 3.20 GHz and 16 GB of RAM, running Ubuntu 20.04.4 LTS and Java 11.0.15. In each test run, we loaded a set of TGDs, measured the wall-clock time required to compute

Time (s)


|  |  | ExbDR | SkDR | HypDR | KAON2 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| \# of Processed Inputs |  | 367 | 377 | 382 | 362 |
| Max. Processed Input Size |  | 185,515 | 324,092 | 324,092 | N/A |
| Max. Output Size |  | 196,594 | 124,846 | 124,846 | 61,964 |
| Max. Size Blowup |  | 8.95 | 8.85 | 8.85 | N/A |
| Max. Body Atoms in Output \# Blowup $\geq 1.5$ |  | 7 | 6 | 6 | 4 |
|  |  | 26 | 14 | 16 | N/A |
| $\begin{aligned} & \text { © } \\ & \text { E } \end{aligned}$ | Min. | 0.05 | 0.05 | 0.04 | 0.21 |
|  | Max. | 582.18 | 584.79 | 404.34 | 547.53 |
|  | Avg. | 23.23 | 14.34 | 6.38 | 18.66 |
|  | Med. | 0.82 | 0.52 | 0.55 | 0.49 |


|  | time $(Y) /$ time $(X) \geq 10$ |  |  |  | $X$ and $Y$ both fail |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | ExbDR SkDR HypDR KAON2 |  |  |  | ExbDR SkDR HypDR KAON2 |  |  |  |
| ExbDR |  | 19 | 0 | 19 | 61 |  |  |  |
| SkDR | 37 |  | 0 | 26 | 33 | 51 |  |  |
| HypDR | 37 | 12 |  | 31 | 35 | 43 | 46 |  |
| KAON2 | 35 | 15 | 0 |  | 37 | 47 | 46 | 66 |

Figure 4: Results for TGDs Derived from Ontologies
the rewriting of a set of GTGDs, and saved the produced Datalog rewriting. We used a timeout of ten minutes for each test run.

### 7.2 Experiments with GTGDs from Ontologies

We computed the Datalog rewriting of GTGDs obtained from OWL ontologies using our three algorithms and KAON2. Figure 4 shows the number of inputs that each algorithm processed in a given time, provides information about the inputs and outputs of each system, and compares the performance among systems. The input size for ExbDR is the number of GTGDs after transforming the input into head-normal form, and for SkDR and HypDR it is the number of rules after Skolemization. Input size is not available for KAON2 since this system reads an OWL ontology and transforms it into GTGDs internally. The output size is the number of Datalog rules in the rewriting. Finally, the blowup is the ratio of the output and the input sizes. Each input GTGD contained at most seven body atoms. Out of 428 inputs, 349 were processed within the ten minute limit by our three systems, and 334 inputs were processed by all four systems. Moreover, 32 inputs, each containing between 20,270 and 221,648 GTGD, were not processed by any system.
Discussion. As one can see in Figure 4, all algorithms were able to compute the rewriting of large inputs containing $100 \mathrm{k}+$ GTGDs. Moreover, for the vast majority of inputs that were successfully
processed, the size of the rewriting and the number of body atoms in the rewriting are typically of the same order of magnitude as the input. Hence, the worst-case exponential blowup from Theorems 5.8, 5.13, and 5.19 does not appear in practice: the size of the rewriting seems to be determined primarily by the input size.
Relative Performance. No system can be identified as the best in general, but HypDR seems to offer the best performance on average. The algorithm was able to process most inputs; it was at least $35 \%$ faster than the other systems on the slowest input; it was never slower by an order of magnitude; there were only 14 inputs that could be processed by some other algorithm but not HypDR; and the output of HypDR does not differ significantly from the output of SkDR. This is in line with our motivation for HypDR outlined in Example 5.17. Specifically, HypDR derives rules with just one head atom, but it does not derive intermediate rules with functional body atoms. The main source of overhead in HypDR seems to be more complex selection of rules participating in an inference.

Impact of Subsumption. All algorithms spend a considerable portion of their running time checking TGD/rule subsumption, so it is natural to wonder whether this overhead is justified. To answer this question, we ran our three approaches using a modification of Algorithm 1: we replaced the check for containment up to redundancy in line 8 with just checking $\tau^{\prime} \notin \mathcal{W} \cup \mathcal{U}$, and we removed line 9 . Note that our normalization of variables described in Section 6 still guarantees termination. This change significantly increased the number of derivations: the numbers of derived TGDs/rules increased on average by a factor of 104, 185, and 103 on ExbDR, SkDR, and HypDR, respectively. Interestingly, this increase did not affect the performance uniformly. While SkDR was able to process 12 inputs an order of magnitude faster, ExbDR and HypDR timed out on 72 and 17 additional inputs, respectively. This, we believe, is due to how different inference rules select inference candidates. The SkDR rule is applied to just pairs of rules, and candidate pairs can be efficiently retrieved using unification indexes. In contrast, ExbDR requires matching several head atoms with as many body atoms, which makes developing a precise index for candidate pair retrieval difficult; thus, as the number of derived TGDs increases, the number of false candidates retrieved from the index increases as well. Finally, HypDR can be applied to an arbitrary number of rules, so selecting inference candidates clearly becomes more difficult as the number of derived rules increases.
Impact of Structural Transformation. KAON2 uses structural transformation [7] to simplify ontology axioms before translating them into GTGDs. For example, axiom $A \sqsubseteq \exists B . \exists C . D$ is transformed into $A \sqsubseteq \exists B . X$ and $X \sqsubseteq \exists C . D$ for $X$ a fresh class. The resulting axioms have simpler structure, which is often beneficial to performance. To see how this transformation affects our algorithms, we reran our experiments while transforming the input axioms in the same way as in KAON2. This indeed improved the performance of SkDR by one order of magnitude on 22 ontologies, and it did not hurt the performance of HypDR. The main challenge is to generalize this transformation to arbitrary GTGDs: whereas description logic axioms exhibit syntactic nesting that lends itself naturally to this transformation, it is less clear how to systematically apply this

Table 2: Computing the Fixpoint of the Rewriting

| Ont. ID | \# Rules | \# Input Facts | \# Output Facts | Time (s) |
| :--- | ---: | ---: | ---: | ---: |
| 00387 | 63,422 | $4,403,105$ | $51,439,424$ | 53 |
| 00448 | 67,986 | $5,510,444$ | $107,235,697$ | 110 |
| 00470 | 75,146 | $10,532,943$ | $141,396,446$ | 242 |
| 00471 | 78,977 | $11,077,423$ | $128,954,126$ | 253 |
| 00472 | 75,146 | $10,533,008$ | $141,396,576$ | 279 |
| 00473 | 78,977 | $11,077,459$ | $128,954,198$ | 291 |
| 00573 | 113,959 | $9,197,254$ | $155,118,592$ | 206 |
| 00682 | 68,461 | $5,183,460$ | $105,431,952$ | 101 |
| 00684 | 81,553 | $6,057,017$ | $66,981,628$ | 109 |
| 00686 | 124,846 | $10,402,324$ | $166,366,039$ | 238 |

transformation to TGDs, where heads and bodies consist of "flat" conjunctions. We leave this question for future work.

### 7.3 End-to-End Experiments

To validate our approach end-to-end, we selected ten inputs where ExbDR produced the largest rewritings. For each of these, we generated a large base instance using WatDiv [2], and we computed the fixpoint of the rewriting and the instance using the RDFox [46] Datalog system v5.4. Table 2 summarizes our results.

All programs used in this experiment are at least several orders of magnitude larger than what is usually encountered in practical applications of Datalog, but RDFox nevertheless computed the fixpoint of all rewritings in a few minutes. Moreover, although the fixpoints seem to be an order of magnitude larger than the base instance, this is not a problem for highly optimized systems such as RDFox. Hence, checking fact entailment via rewritings produced by our algorithms is feasible in practice.

### 7.4 GTGDs With Relations of Higher Arity

Finally, we computed the rewriting of GTGDs obtained by blowing up relation arity as described in Subsection 7.1 using a blowup factor of five. We did not use KAON2 since this system supports relations of arity at most two. Figure 5 summarizes our results. Out of 428 inputs, 187 were processed within the ten minute limit by our three systems, and 128 inputs were not processed by any system.

While HypDR performed best on GTGDs derived from ontologies, Figure 5 shows it to be worst-performing on higher-arity GTGDs: it successfully processed only 199 inputs within the ten minute timeout, whereas SkDR and ExbDR processed 238 and 274 inputs, respectively. This is mainly due to additional body atoms introduced by our "blowup" method: these increase the number of rules participating in an application of the HypDR inference rule, which makes selecting the participating rules harder.

This experiment proved to be more challenging, as most problems discussed in Section 6 became harder. For example, in ExbDR, higher arity of atoms increases the likelihood that an atom retrieved through a unification index does not unify with a given atom, and that the atoms of the selected GTGDs cannot be successfully matched. Subsumption indexing is also more difficult for similar reasons. However, the inputs used in this experiment consist of a large numbers of GTGDs with relations of arity ten, so they can be seen as a kind of a "stress test". Our algorithms were able to

Time (s)


|  | ExbDR | SkDR | HypDR |
| :--- | ---: | ---: | ---: |
| \# of Processed Inputs | 274 | 238 | 199 |
| Max. Processed Input Size | 69,046 | 182,569 | 38,362 |
| Max. Output Size | 58,749 | 171,832 | 38,335 |
| Max. Size Blowup | 9.00 | 5.84 | 5.84 |
| \# Blowup $\geq 1.5$ | 26 | 5 | 3 |
| Min. | 0.06 | 0.05 | 0.04 |
|  | Avg. | 591.82 | 504.49 |
| Med. | 26.70 | 38.39 | 17.05 |



Figure 5: Results for TGDs with Higher-Arity Relations
process more than half of such inputs, which leads us to believe that they can also handle more well-behaved GTGDs used in practice.

## 8 CONCLUSION

We presented several algorithms for rewriting a finite set of guarded TGDs into a Datalog program that entails the same base facts on each base instance. Our algorithms are based on a new framework that establishes a close connection between Datalog rewritings and a particular style of the chase. In future, we plan to generalize our framework to wider classes of TGDs, such as frontier-guarded TGDs, as well as provide rewritings for conjunctive queries under certain answer semantics. Moreover, we shall investigate whether the extension of our framework to disjunctive guarded TGDs [31] can be used to obtain practical algorithms for rewriting disjunctive guarded TGDs into disjunctive Datalog programs.

## ACKNOWLEDGMENTS

This work was funded by the EPSRC grants OASIS (EP/S032347/1), AnaLOG (EP/P025943/1), Concur (EP/V050869/1), and QUINTON (EP/T022124/1). For the purpose of Open Access, the author has applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

## REFERENCES

[1] Shqiponja Ahmetaj, Magdalena Ortiz, and Mantas Simkus. 2018. Rewrit ing Guarded Existential Rules into Small Datalog Programs. In ICDT. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 4:1-4:24.
[2] Günes Aluç, Olaf Hartig, M. Tamer Özsu, and Khuzaima Daudjee. 2014. Diversified Stress Testing of RDF Data Management Systems. In ISWC. Springer, 197-212.
[3] Mario Alviano, Nicola Leone, Marco Manna, Giorgio Terracina, and Pierfrancesco Veltri. 2012. Magic-Sets for Datalog with Existential Quantifiers. In Datalog 2.0. Springer, 31-43.
[4] Antoine Amarilli and Michael Benedikt. 2022. When Can We Answer Queries Using Result-Bounded Data Interfaces? Log. Methods Comput. Sci. 18, 2 (2022), 14:1--14:81.
[5] Hajnal Andréka, Johan van Benthem, and István Németi. 1998. Modal languages and bounded fragments of predicate logic. Fournal of Philosophical Logic 27 (1998), 217-274.
[6] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, and P. F. Patel-Schneider (Eds.). 2007. The Description Logic Handbook: Theory, Implementation and Applications (2nd ed.). Cambridge University Press, Cambridge, UK.
[7] M. Baaz, U. Egly, and A. Leitsch. 2001. Normal Form Transformations. In Handbook of Automated Reasoning, A. Robinson and A. Voronkov (Eds.). Vol. I. Elsevier Science, Amsterdam, The Netherlands, Chapter 5, 273-333.
[8] Leo Bachmair and Harald Ganzinger. 2001. Resolution Theorem Proving. Handbook of automated reasoning 1 (2001), 19-99.
[9] J.-F. Baget, M. Leclère, M.-L. Mugnier, S. Rocher, and C. Sipieter. 2015. Graal: A Toolkit for Query Answering with Existential Rules. In RuleML. Springer, 328-344.
[10] Jean-François Baget, Marie-Laure Mugnier, Sebastian Rudolph, and Michaël Thomazo. 2011. Walking the complexity lines for generalized guarded existential rules. In IFCAI. AAAI Press, 712-717.
[11] Vince Bárány, Michael Benedikt, and Balder Ten Cate. 2013. Rewriting Guarded Negation Queries. In MFCS. Springer, 98-110.
[12] Luigi Bellomarini, Emanuel Sallinger, and Georg Gottlob. 2018. The Vadalog System: Datalog-based Reasoning for Knowledge Graphs. Proc. VLDB Endow. 11, 9 (2018), 975-987.
[13] Michael Benedikt, Maxime Buron, Stefano Germano, Kevin Kappelmann, and Boris Motik. 2021. Guarded Saturation. GitHub. Retrieved July 4, 2022 from https://krr-oxford.github.io/Guarded-saturation/
[14] Michael Benedikt, George Konstantinidis, Giansalvatore Mecca, Boris Motik, Paolo Papotti, Donatello Santoro, and Efthymia Tsamoura. 2017. Benchmarking the Chase. In PODS. ACM, 37-52.
[15] Andrea Calì, Georg Gottlob, and Michael Kifer. 2013. Taming the Infinite Chase Query Answering under Expressive Relational Constraints. Fournal of Artificial Intelligence Research 48 (2013), 115-174.
[16] Andrea Calì, Domenico Lembo, and Riccardo Rosati. 2003. Query rewriting and answering under constraints in data integration systems. In IFCAI. Morgan Kaufmann, 16-21.
[17] Diego Calvanese, Benjamin Cogrel, Sarah Komla-Ebri, Roman Kontchakov, Davide Lanti, Martin Rezk, Mariano Rodriguez-Muro, and Guohui Xiao. 2017. Ontop Answering SPARQL queries over relational databases. Semantic Web 8, 3 (2017), 471-487.
[18] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. 2007. Tractable Reasoning and Efficient Query Answering in Description Logics: The DL-Lite Family. 7. Autom. Reason. 39, 3 (2007), 385-429.
[19] Hans de Nivelle. 1998. A Resolution Decision Procedure for the Guarded Fragment. In CADE. Springer, 191-204.
[20] A. Deutsch, L. Popa, and V. Tannen. 2006. Query reformulation with constraints. SIGMOD Record 35, 1 (2006), 65-73.
[21] R. Fagin, P. G. Kolaitis, R. J. Miller, and L. Popa. 2005. Data exchange: semantics and query answering. Theoretical Computer Science 336, 1 (2005), 89-124.
[22] Mohamed Gaha, Arnaud Zinflou, Christian Langheit, Alexandre Bouffard, Mathieu Viau, and Luc Vouligny. 2013. An Ontology-Based Reasoning Approach for Electric Power Utilities. In RR. Springer, 95-108.
[23] H. Ganzinger and H. de Nivelle. 1999. A Superposition Decision Procedure for the Guarded Fragment with Equality. In Proc. of the 14th IEEE Symposium on

Logic in Computer Science (LICS '99). IEEE Computer Society, 295-305.
[24] Georg Gottlob, Sebastian Rudolph, and Mantas Simkus. 2014. Expressiveness of Guarded Existential Rule Languages. In PODS. ACM, 27-38.
[25] Alon Halevy, Anand Rajaraman, and Joann Ordille. 2006. Data Integration: The Teenage Years. In VLDB. ACM, 9-16.
[26] Alon Y. Halevy. 2001. Answering queries using views: A survey. VLDB fournal 10, 4 (2001), 270-294.
[27] Colin Hirsch. 2002. Guarded Logics: Algorithms and Bisimulation. Ph.D. Dissertation. RWTH Aachen, Aachen, Germany. Retrieved July 4, 2022 from http://www.umbrialogic.com/hirsch-thesis.pdf
28] Ullrich Hustadt, Boris Motik, and Ulrike Sattler. 2004. Reducing SHIQ Description Logic to Disjunctive Datalog Programs. In KR. AAAI Press, 152-162.
[29] Ullrich Hustadt, Boris Motik, and Ulrike Sattler. 2007. Reasoning in description logics by a reduction to disjunctive datalog. Journal of Automated Reasoning 39, 3 (2007), 351-384.
[30] David S. Johnson and Anthony C. Klug. 1984. Testing Containment of Conjunctive Queries under Functional and Inclusion Dependencies. $\mathcal{F C S S}$ 28, 1 (1984), 167-189.
[31] Kevin Kappelmann. 2019. Decision Procedures for Guarded Logics. CoRR abs/1911.03679 (2019), 92. Retrieved July 4, 2022 from http://arxiv.org/abs/ 1911.03679
[32] Deepak Kapur and Paliath Narendran. 1986. NP-Completeness of the Set Unification and Matching Problems. In CADE. Springer, 489-495.
[33] Alon Y. Levy. 2000. Logic-Based Techniques in Data Integration. Kluwer Academic Publishers, Norwell, MA, USA, 575-595.
[34] Thomas Lukasiewicz, Andrea Cali, and Georg Gottlob. 2012. A General DatalogBased Framework for Tractable Query Answering over Ontologies. Fournal of Web Semantics 14, 0 (2012), 57-83.
[35] Bruno Marnette. 2012. Resolution and Datalog Rewriting Under Value Invention and Equality Constraints. CoRR abs/1212.0254 (2012), 12.
[36] William McCune and Larry Wos. 1997. Otter-The CADE-13 Competition Incarnations. Journal of Automated Reasoning 18, 2 (1997), 211-220.
[37] M. Meier. 2014. The backchase revisited. VLDB f. 23, 3 (2014), 495-516.
[38] Boris Motik. 2006. Reasoning in description logics using resolution and deductive databases. Ph.D. Dissertation. Karlsruhe Institute of Technology, Karlsruhe, Germany. Retrieved July 4, 2022 from http://digbib.ubka.uni-karlsruhe.de/ volltexte/1000003797
[39] Boris Motik. 2022. The KAON2 System. Karslruhe Institute of Technology. Retrieved July 4, 2022 from http://kaon2.semanticweb.org/
[40] Oxford KR group. 2021. Oxford Ontology Library. Oxford University. Retrieved July 4, 2022 from http://krr-nas.cs.ox.ac.uk/ontologies/
[41] Mike Paterson and Mark N. Wegman. 1978. Linear Unification. F. Comput. Syst. Sci. 16, 2 (1978), 158-167.
[42] John Alan Robinson. 1965. A machine-oriented logic based on the resolution principle. $7 A C M$ 12, 1 (1965), 23-41.
[43] Iztok Savnik. 2013. Index Data Structure for Fast Subset and Superset Queries. In CD-ARES. Springer, 134-148.
[44] Stephan Schulz. 2013. Simple and Efficient Clause Subsumption with Feature Vector Indexing. In Automated Reasoning and Mathematics. Springer, 45-67.
[45] Mark E. Stickel. 1989. The Path-Indexing Method for Indexing Terms. Technical Report. SRI.
[46] Oxford Semantic Technologies. 2022. The RDFox System. Oxford Semantic Technologies. Retrieved July 4, 2022 from https://www.oxfordsemantic.tech/
[47] Moshe Y. Vardi. 1997. Why Is Modal Logic so Robustly Decidable?. In DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 31. American Mathematical Society, 149-184.
[48] Roberto De Virgilio, Giorgio Orsi, Letizia Tanca, and Riccardo Torlone. 2012. NYAYA: A System Supporting the Uniform Management of Large Sets of Semantic Data. In ICDE. IEEE Computer Society, 1309-1312.
[49] Zhe Wang, Peng Xiao, Kewen Wang, Zhiqiang Zhuang, and Hai Wan. 2021. Query Answering for Existential Rules via Efficient Datalog Rewriting. In IFCAI. ijcai.org, 1933-1939.
[50] Sen Zheng and Renate A. Schmidt. 2020. Deciding the Loosely Guarded Fragment and Querying Its Horn Fragment Using Resolution. In AAAI. AAAI Press, 30803087.

## A PROOFS FOR SECTION 4: ONE-PASS CHASE PROOFS

In Section 4 we introduced the notion of a one-pass chase proof, which allows us to establish a completeness criterion for saturations that is tied to the chase. We provide details of the proofs in this appendix.

## A. 1 Proof of Theorem 4.2: Existence of One-Pass Chase Proofs

Theorem 4.2. For each base instance $I$, each finite set of GTGDs $\Sigma$ in head-normal form, and each base fact $F$ such that $I, \Sigma \vDash F$, there exists a one-pass tree-like chase proof of $F$ from I and $\Sigma$.

Throughout this section, we fix an arbitrary base instance $I$ and a finite set of GTGDs $\Sigma$. It is known that $I, \Sigma \vDash F$ if and only if there exists a tree-like chase proof of $F$ from $I$ and $\Sigma$. We next prove Theorem 4.2 by showing that each such proof can be transformed to a one-pass chase proof of $F$ from $I$ and $\Sigma$. This argument was developed jointly with Antoine Amarilli, and it is related to proofs by Amarilli and Benedikt [4] and Kappelmann [31]; however, note that Definition 4.1 imposes slightly stronger conditions on one-pass chase sequences than related definitions in those works.

Towards our goal, we first state two basic properties of tree-like chase sequences. The first claim is a variation of the well-known fact that any chase tree produced for GTGDs represents a tree decomposition [15]. The second claim captures the idea that, as the chase progresses, facts may be added within a vertex, but this will not produced new guarded sets of terms.

Lemma A.1. Let $T_{0}, \ldots, T_{n}$ be an arbitrary tree-like chase sequence for $I$ and $\Sigma$.

1. For each $0 \leq i \leq n$, all vertices $v_{1}$ and $v_{2}$ in $T_{i}$, each set $G$ of ground terms that is $\Sigma$-guarded by $T_{i}\left(v_{1}\right)$ and by $T_{i}\left(v_{2}\right)$, and each vertex $v_{3}$ on the unique path in $T_{i}$ between $v_{1}$ and $v_{2}$, set $G$ is $\Sigma$-guarded by $T_{i}\left(v_{3}\right)$.
2. For each $0 \leq i \leq n$, each vertex $v$ in $T_{i}$, each set $G$ of ground terms that is $\Sigma$-guarded by $T_{i}(v)$, and each $0 \leq j \leq i$ such that $T_{j}$ contains $v$, set $G$ is $\Sigma$-guarded by $T_{j}(v)$.

Proof of Claim 1. The proof is by induction on $i$ with $0 \leq i \leq n$. For $i=0$, chase tree $T_{0}$ contains just one vertex so the claim holds trivially. Now assume that the property holds for some $0 \leq i<n$ and consider ways in which $T_{i+1}$ can be derived from $T_{i}$. First, $T_{i+1}$ can be obtained by applying a chase step to $T_{i}$ at vertex $v$ with some GTGD $\tau \in \Sigma$. Let $v_{1}$ be the recently updated vertex of $T_{i+1}$. Thus, $v_{1}$ is either $v$ or a fresh child of $v$. Moreover, consider each fact $R(\vec{t})$ derived by the step, each set of ground terms $G \subseteq \vec{t}$, each vertex $v_{2}$ such that $G$ is $\Sigma$-guarded by $T_{i+1}\left(v_{2}\right)$, and each vertex $v_{3}$ on the unique path in $T_{i+1}$ between $v_{1}$ and $v_{2}$. If $G$ contains a labeled null that is freshly introduced in $T_{i+1}$, the claim holds trivially because $v_{2}$ and $v_{3}$ are necessarily the same as $v_{1}$. Otherwise, $\tau$ is guarded, so $T_{i}(v)$ contains a fact $S(\vec{u})$ such that $G \subseteq \vec{u}$. But then, $G$ is $\Sigma$-guarded by $T_{i}\left(v_{3}\right)$ by the induction assumption. Moreover, $T_{i}\left(v_{3}\right) \subseteq T_{i+1}\left(v_{3}\right)$ ensures that $G$ is $\Sigma$-guarded by $T_{i+1}\left(v_{3}\right)$, as required. Second, $T_{i+1}$ can be obtained by applying a propagation step to $T_{i}$, but then the property clearly holds.

Proof of Claim 2. The proof is by induction on $i$ with $0 \leq i \leq n$. The base case for $i=0$ is trivial. For the induction step, assume that the property holds for some $i$. If $T_{i+1}$ is obtained from $T_{i}$ by a chase step with a non-full GTGD, then the claim clearly holds for $T_{i+1}$ because the step introduces a fresh vertex that does not occur in any $T_{j}$ with $0 \leq j \leq i$. Otherwise, $T_{i+1}$ is obtained by extending some $T_{i}(v)$, so consider an arbitrary fact $F \in T_{i+1}(v) \backslash T_{i}(v)$. Clearly, $F$ is $\Sigma$-guarded by $T_{i}(v)$ : if the step involves a full GTGD, then a body atom of the GTGD is matched to a fact $F^{\prime} \in T_{i}(v)$ such that $F$ is $\Sigma$-guarded by $F^{\prime}$; moreover, if the step involves propagation, then by definition there exists a fact $F^{\prime} \in T_{i}(v)$ such that $F$ is $\Sigma$-guarded by $F^{\prime}$. Thus, each set of ground terms $G$ that is $\Sigma$-guarded by $T_{i+1}(v)$ is also $\Sigma$-guarded by $F^{\prime} \in T_{i}(v)$, so the claim holds.

In the rest of the proof, we show how to convert an arbitrary tree-like chase proof into a one-pass one through a series of transformations. Before proceeding, we next describe formally the types of chase sequence that we consider in our transformations.

## Definition A.2.

- A chase sequence is local if each propagation step in the sequence copies just one fact to either the parent or a child vertex.
- A chase sequence is rootward if each propagation step in the sequence copies just one fact from a child to its parent.
- A chase sequence is almost one-pass if it is rootward and each chase or propagation step is applied to the recently updated vertex or an ancestor thereof, and a chase step is applied only if a propagation step is not applicable to the recently updated vertex or an ancestor thereof.
Note that facts can still be copied from a parent to a child in a rootward chase sequence, but this can be done only in chase steps with non-full GTGDs that introduce a child. Furthermore, the use of "almost" in the "almost one-pass" reflects the caveat that, in an almost one-pass chase sequence, a step can be applied to an ancestor of the recently updated vertex, thus "jumping rootward" in the tree, whereas such steps are forbidden in a one-pass chase sequence.

We capture formally the relationship between the chase sequences produced by our transformations using the notion introduced in Definition A. 3 .

Definition A.3. A chase tree $T$ is a subset of a chase tree $T^{\prime}$, written $T \subseteq T^{\prime}$, if the tree of $T$ is a subtree of $T^{\prime}$ (i.e., the root of $T$ is the root of $T^{\prime}$, and whenever vertex $v$ is a parent of vertex $v^{\prime}$ in $T$, then $v$ is a parent of $v^{\prime}$ in $\left.T^{\prime}\right)$, and $T(v) \subseteq T^{\prime}(v)$ holds for each vertex $v$ of $T$.

We are now ready to present our transformations, which we capture in a series of lemmas. We next summarize the main intuitions.

- In Lemma A.4, we show that an arbitrary chase sequence can be transformed into a local chase sequence by "slowing down" propagation steps so that facts are copied only between vertices that are adjacent in a chase tree.
- In Lemma A.5, we show that each local chase sequence can be transformed into a rootward chase sequence. Intuitively, instead of propagating a fact from a parent to a child, we "regrow" a clone of the relevant child and the entire subtree underneath. The relevant fact is then copied as part of the chase step with the non-full GTGD that "regrows" the child's clone.
- In Lemma A.6, we show that each rootward chase sequence can be transformed to an almost one-pass chase sequence. The main difficulty arises due to the fact that steps in a rootward chase sequence can be applied to arbitrary vertices. We address this problem by shuffling and regrowing parts of the chase trees.
- Finally, in Lemma A.7, we show that each almost one-pass chase proof can be transformed to a one-pass chase proof by pruning irrelevant parts of the chase sequence.
Lemma A.4. For each tree-like chase sequence $T_{0}, \ldots, T_{n}$ for I and $\Sigma$, there exists a local tree-like chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{m}$ for I and $\Sigma$ such that $T_{n} \subseteq \bar{T}_{m}$.

Proof. Each propagation step in $T_{0}, \ldots, T_{n}$ that copies more than one fact can clearly be "expanded" into several steps, each copying just one fact. Moreover, due to Claim 1 of Lemma A.1, each propagation step that copies a fact $F$ between vertices $v$ and $v^{\prime}$ that are further apart can be "expanded" into several steps that propagate $F$ to all vertices on the unique path between $v$ and $v^{\prime}$.

Lemma A.5. For each local tree-like chase sequence $T_{0}, \ldots, T_{n}$ for I and $\Sigma$, there exists a rootward tree-like chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{m}$ for $I$ and $\Sigma$ such that
(S1) $T_{n} \subseteq \bar{T}_{m}$, and
(S2) for each vertex $v$ in $T_{n}$ that is introduced by a chase step with a non-full $G T G D \tau \in \Sigma$ and substitutions $\sigma$ and $\sigma^{\prime}$, vertex $v$ is introduced into some $\bar{T}_{k}$ with $0 \leq k \leq m$ by a chase step with the same $\tau, \sigma$, and $\sigma^{\prime}$.

Proof. Let $T_{0}, \ldots, T_{n}$ be an arbitrary local tree-like chase sequence for $I$ and $\Sigma$. We prove the claim by induction on $0 \leq i \leq n$. The induction base $i=0$ holds trivially. For the induction step, we assume that the claim holds for some $i$ with $0 \leq i<n$. By the inductive assumption, there exists a rootward chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ for $I$ such that $T_{i} \subseteq \bar{T}_{j}$ and property (S2) holds. Let $v$ be the vertex of $T_{i}$ to which a chase or propagation step is applied to derive $T_{i+1}$. By Definition A.3, chase tree $\bar{T}_{j}$ contains vertex $v$ and $T_{i}(v) \subseteq \bar{T}_{j}(v)$ holds. We now consider ways in which $T_{i+1}$ can be derived from $T_{i}$.

Assume that $T_{i+1}$ is obtained from $T_{i}$ by a chase step with non-full TGD $\tau \in \Sigma$, and let $v^{\prime}$ be the child of $v$ introduced by the step. Without loss of generality, we can choose $v^{\prime}$ and the fresh labeled nulls such that they do not occur in $\bar{T}_{j}$. Now let $\bar{T}_{j+1}$ be obtained from $\bar{T}_{j}$ by adding $v^{\prime}$ as a child of $v$ and setting $\bar{T}_{j+1}\left(v^{\prime}\right)=T_{i+1}\left(v^{\prime}\right)$. Clearly, $\bar{T}_{0}, \ldots, \bar{T}_{j}, \bar{T}_{j+1}$ is a rootward chase sequence such that $T_{i+1} \subseteq \bar{T}_{j+1}$ and property (S2) hold, as required.

Assume that $T_{i+1}$ is obtained from $T_{i}$ by a chase step with a full TGD $\tau \in \Sigma$ deriving a fact $F$, or by a rootward propagation step that copies a fact $F$ from $T_{i}(v)$ to the parent of $v$. Let $v^{\prime}$ be the recently updated vertex of $T_{i+1}$. Chase tree $\bar{T}_{j}$ clearly contains $v^{\prime}$. If $F \in \bar{T}_{j}\left(v^{\prime}\right)$, then sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ satisfies the inductive property. If $T_{i+1}$ is obtained from $T_{i}$ by a propagation step, then $F$ is $\Sigma$-guarded by $T_{i}\left(v^{\prime}\right)$. But then, $T_{i}\left(v^{\prime}\right) \subseteq \bar{T}_{j}\left(v^{\prime}\right)$ ensures that $F$ is also $\Sigma$-guarded by $\bar{T}_{j}\left(v^{\prime}\right)$ and thus the propagation step is applicable to vertices $v$ and $v^{\prime}$ in $\bar{T}_{j}$. Now let $\bar{T}_{j+1}$ to be the same as $\bar{T}_{j}$ but with $\bar{T}_{j+1}\left(v^{\prime}\right)=\bar{T}_{j}\left(v^{\prime}\right) \cup\{F\}$ and with $v^{\prime}$ being the recently updated vertex. Clearly, $\bar{T}_{0}, \ldots, \bar{T}_{j}, \bar{T}_{j+1}$ is a rootward chase sequence satisfying $T_{i+1} \subseteq \bar{T}_{j+1}$, as required. Moreover, property (S2) holds by the induction hypothesis.

The only remaining case is when $T_{i+1}$ is obtained from $T_{i}$ by applying a propagation step that copies one fact $F$ to a child $v^{\prime}$ of $v$. By Definition A.3, chase tree $\bar{T}_{j}$ contains vertex $v^{\prime}$ and $T_{i}\left(v^{\prime}\right) \subseteq \bar{T}_{j}\left(v^{\prime}\right)$ holds. Sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ satisfies the inductive property if $F \in \bar{T}_{j}\left(v^{\prime}\right)$ holds, so we next assume $F \notin \bar{T}_{j}\left(v^{\prime}\right)$. We next show that we can simulate propagation by "replaying" the chase steps that generate $v^{\prime}$ and all of its descendants. Towards this goal, let $T_{k}$ be the chase tree in the original sequence where $v^{\prime}$ is first introduced by applying a chase step with the non-full GTGD $\tau=\beta \rightarrow \exists \vec{y} \eta \in \Sigma$, and let $\sigma$ and $\sigma^{\prime}$ be substitutions used in the step. By the inductive property (S2), there exists $\ell_{0}$ with $0<\ell_{0} \leq j$ such that $v^{\prime}$ is introduced in $\bar{T}_{\ell_{0}}$ as the result of applying a chase step with the same non-full TGD $\tau$ and substitutions $\sigma$ and $\sigma^{\prime}$. Finally, let $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$ be the subsequence of $\bar{T}_{0}, \ldots, \bar{T}_{j}$ consisting of precisely those chase trees that were obtained by applying a chase or a propagation step to $v^{\prime}$ or a descendant of $v^{\prime}$. In other words, the chase steps producing $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$ are exactly the steps that we need to "replay" to simulate the propagation of $F$ from $v$ to $v^{\prime}$.

Our objective is to "replay" the steps producing $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$ so that they introduce exactly the same vertices and labeled nulls, which is needed because property ( S 1 ) talks about exact containment of the final chase trees of the two sequences (rather than containment up to isomorphism). A technical issue is that these vertices and labeled nulls already occur in the sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$; thus, if we extended this sequence directly, we could not "reapply" the chase steps with non-full GTGDs, which by definition introduce fresh vertices and labeled nulls. To get around this, we first perform the following renaming step. Let $N$ be the set of labeled nulls introduced by the chase steps with non-full TGDs in subsequence $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$, and let $W$ be the set of introduced vertices (thus, $W$ contains $v^{\prime}$ and all of its descendants). Moreover, let $\bar{U}_{0}, \ldots, \bar{U}_{j}$ be the chase sequence obtained by uniformly replacing in $\bar{T}_{0}, \ldots, \bar{T}_{j}$ each labeled null in $N$ with a distinct, fresh labeled null, and by uniformly replacing each vertex $w \in W$ by a fresh vertex.

We next describe the chase trees that will be produced by "replaying" the steps producing the subsequence $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$. Intuitively, we must "graft" the results of these steps onto $\bar{U}_{j}$ : for $v^{\prime}$ or a descendant of $v^{\prime}$ we take the results of the chase steps in the subsequence, and for each other vertex we copy the content from $\bar{U}_{j}$. Formally, let $\bar{V}_{0}, \ldots, \bar{V}_{m}$ be the sequence obtained from the subsequence $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$ using the following steps.
(R1) For each $0 \leq p \leq m$ and each vertex $w$ in $\bar{U}_{j}$ such that $w$ is neither $v^{\prime}$ nor a descendant of $v^{\prime}$ in $\bar{U}_{j}$, we set $\bar{V}_{p}(w)=\bar{U}_{j}(w)$.
(R2) For each $0 \leq p \leq m$ and each vertex $w$ that occurs in $\bar{T}_{\ell_{p}}$ such that $w$ is $v^{\prime}$ or a descendant of $v^{\prime}$ in $\bar{T}_{\ell_{p}}$, we set $\bar{V}_{p}(w)=\bar{T}_{\ell_{p}}(w)$.
(R3) We add to $\bar{V}_{0}\left(v^{\prime}\right)$ each fact $G \in \bar{U}_{j}(v)$ that is $\Sigma$-guarded by $\sigma^{\prime}(\eta)$.
(R4) We analogously extend each $\bar{V}_{p}$ with $1 \leq p \leq m$ to ensure that each chase step with a non-full GTGD correctly propagates all relevant facts to a child.

We now argue that $\bar{U}_{0}, \ldots, \bar{U}_{j}, \bar{V}_{0}, \ldots, \bar{V}_{m}$ is a rootward chase sequence that satisfies properties (S1) and (S2). Towards this goal, we make the following observations.

- Sequence $\bar{U}_{0}, \ldots, \bar{U}_{j}$ is a rootward chase sequence produced by the same steps as $\bar{T}_{0}, \ldots, \bar{T}_{j}$, but with the vertices in $W$ and labeled nulls in $N$ uniformly renamed. Also, due to step (R4), $\bar{V}_{0}, \ldots, \bar{V}_{m}$ is a rootward chase sequence produced by the same steps as $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$.
- Chase tree $\bar{V}_{0}$ coincides with $\bar{U}_{j}$ on each vertex that is not $v^{\prime}$ or a descendant of $v^{\prime}$. Moreover, $\bar{U}_{j}$ does not contain a labeled null in $N$, and it does not contain $v^{\prime}$ or a descendant of $v^{\prime}$; thus, $\bar{V}_{0}$ can be seen as the result of applying to $\bar{U}_{j}$ a chase step with the non-full GTGD $\tau$ and substitutions $\sigma$ and $\sigma^{\prime}$ that introduces vertex $v^{\prime}$ as a child of $v$.
- We now show that property (S2) is satisfied-that is, that $T_{i+1} \subseteq \bar{V}_{m}$ holds. Towards this goal, consider an arbitrary vertex $w$ occurring in $T_{i+1}$; by the induction assumption, we have $T_{i}(w) \subseteq \bar{T}_{j}(w)$. If $w$ is neither $v^{\prime}$ nor a descendant thereof, then neither $w$ nor a labeled null occurring in $T_{i}(w)$ was renamed in $\bar{U}_{j}$, so we have $\bar{T}_{j}(w)=\bar{U}_{j}(w)=\bar{V}_{m}(w)$, where the last equality is ensured by step (R1); thus, $T_{i+1}(w)=T_{i}(w) \subseteq \bar{V}_{m}(w)$ holds, as required. Now assume that $w$ is $v^{\prime}$ or a descendant thereof. Then, $\bar{T}_{j}(w)=\bar{T}_{\ell_{m}}(w)$ holds by the fact that $\bar{T}_{\ell_{m}}$ is the last place in $\bar{T}_{0}, \ldots, \bar{T}_{j}$ where $v^{\prime}$ or a descendant of $v^{\prime}$ was modified, and $\bar{T}_{\ell_{m}}(w)=\bar{V}_{m}(w)$ holds by step (R2); putting it all together, we have $T_{i}(w) \subseteq \bar{V}_{m}(w)$. Now if $w$ is not $v^{\prime}$ (i.e., $w$ is a descendant of $v^{\prime}$ ), then $T_{i+i}(w)=T_{i}(w) \subseteq \bar{V}_{m}(w)$ holds, as required. We finally consider the case when $w$ is $v^{\prime}$, so $T_{i+1}\left(v^{\prime}\right)=T_{i}\left(v^{\prime}\right) \cup\{F\}$. Since the propagation step is applicable to $T_{i}$, fact $F$ is $\Sigma$-guarded by $T_{i}\left(v^{\prime}\right)$. By Claim 2 of Lemma A.1, fact $F$ is also $\Sigma$-guarded by $T_{k}\left(v^{\prime}\right)$. Finally, by the definition of a chase step with a non-full TGD, fact $F$ is $\Sigma$-guarded by $\sigma^{\prime}(\eta)$. But then, step (R3) ensures $F \in \bar{V}_{0}\left(v^{\prime}\right) \subseteq \bar{V}_{m}\left(v^{\prime}\right)$. Consequently, $T_{i+i}\left(v^{\prime}\right) \subseteq \bar{V}_{m}\left(v^{\prime}\right)$ holds, as required.
- We now show that property ( S 2 ) is satisfied. To this end, consider an arbitrary vertex $w$ in $T_{i+1}$ introduced by a chase step with a non-full GTGD $\tau$ and substitutions $\sigma$ and $\sigma^{\prime}$. If $w$ is not $v^{\prime}$ or a descendant thereof, then the labeled nulls introduced by the chase step are not renamed in $\bar{U}_{0}, \ldots, \bar{U}_{j}$, so the claim holds by the induction assumption. Otherwise, the chase steps producing $\bar{V}_{0}, \ldots, \bar{V}_{m}$ are exactly the same as the chase steps producing $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{m}}$, so the claim holds by the induction assumption too.

Lemma A.6. For each rootward tree-like chase sequence $T_{0}, \ldots, T_{n}$ for $I$ and $\Sigma$, there exists an almost one-pass chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{m}$ for $I$ and $\Sigma$ such that $T_{n} \subseteq \bar{T}_{m}$.

Proof. Let $T_{0}, \ldots, T_{n}$ be an arbitrary rootward tree-like chase sequence for $I$ and $\Sigma$. The induction base $i=0$ holds trivially. For the induction step, we assume that the claim holds for some $i$ with $0 \leq i<n$. By the inductive assumption, there exists an almost one-pass chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ for $I$ and $\Sigma$ such that $T_{i} \subseteq \bar{T}_{j}$ holds. Now assume that $T_{i+1}$ is obtained by applying a chase or a propagation step to some vertex $v$ of $T_{i}$, and let $k$ be the maximal number such that $0 \leq k \leq j$ and $v$ is recently updated in $\bar{T}_{k}$. Such $k$ clearly exists since $v$ occurs in $\bar{T}_{j}$, and $T_{i}(v) \subseteq \bar{T}_{k}(v)$ holds because $k$ is maximal. We now consider ways in which $T_{i+1}$ can be derived from $T_{i}$.

Assume that $T_{i+1}$ is obtained from $T_{i}$ by a chase step with non-full GTGD $\tau \in \Sigma$ and substitutions $\sigma$ and $\sigma^{\prime}$, and let $v^{\prime}$ be the child of $v$ introduced by the step. Without loss of generality, we can choose $v^{\prime}$ and the fresh labeled nulls such that they do not occur in $\bar{T}_{j}$. We shall now "move" this chase step so that it is performed immediately after $\bar{T}_{k}$. Towards this goal, we describe the chase trees that are obtained by this move. For each $p$ with $k \leq p \leq j$, let $\bar{U}_{p}$ be the chase tree obtained from $\bar{T}_{p}$ by adding vertex $v^{\prime}$ and letting $\bar{U}_{p}\left(v^{\prime}\right)=T_{i+1}\left(v^{\prime}\right)$. We now argue that $\bar{T}_{0}, \ldots, \bar{T}_{k}, \bar{U}_{k}, \ldots, \bar{U}_{j}$ is an almost one-pass chase sequence satisfying the conditions of the lemma.

- Chase tree $\bar{U}_{k}$ can be seen as obtained from $\bar{T}_{k}$ by a chase step with $\tau$ and substitutions $\sigma$ and $\sigma^{\prime}$. Moreover, for each $p$ with $k \leq p<j$, chase tree $\bar{U}_{p+1}$ is obtained from $\bar{U}_{p}$ in the same way as $\bar{T}_{p+1}$ is obtained from $\bar{T}_{p}$. Thus, all preconditions of all chase steps are satisfied.
- Chase tree $\bar{U}_{k}$ is obtained from $\bar{T}_{k}$ by applying the chase step to the recently updated vertex $v$ of $\bar{T}_{k}$. Moreover, if $k<j$, then $\bar{T}_{k+1}$ is obtained from $\bar{T}_{k}$ by applying a step to $v$ or an ancestor of $v$, and so $\bar{U}_{k+1}$ is obtained from $\bar{U}_{k}$ by applying a step to an ancestor of the recently updated vertex of $\bar{U}_{k}$. Thus, the sequence is almost one-pass.
- The construction clearly satisfies $T_{i+1} \subseteq \bar{U}_{j}$.

In the rest of this proof we consider the case when $T_{i+1}$ is obtained from $T_{i}$ by a chase step with a full GTGD $\tau \in \Sigma$ deriving a fact $F$, or by a propagation step that copies a fact $F$ from $T_{i}(v)$ to the parent of $v$. Let $v^{\prime}$ be the recently updated vertex of $T_{i+1}$. Chase tree $\bar{T}_{j}$ clearly
contains $v^{\prime}$. If $F \in \bar{T}_{k}\left(v^{\prime}\right)$, then sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ satisfies the inductive property, so we next assume that $F \notin \bar{T}_{k}\left(v^{\prime}\right)$ holds. We shall now transform $\bar{T}_{0}, \ldots, \bar{T}_{j}$ so that this step is applied immediately after $\bar{T}_{k}$, and fact $F$ is propagated towards the root as far as possible. Since this will move the recently updated vertex towards the root, we will then "reapply" all relevant steps from $\bar{T}_{0}, \ldots, \bar{T}_{j}$ to "regrow" the relevant part of the sequence. In each case, we specify the structure of the chase trees and discuss the steps that produce these trees.

Let $\bar{U}_{0}$ be obtained from $\bar{T}_{k}$ by adding $F$ to $\bar{T}_{k}(v)$. We argue that $\bar{U}_{0}$ can be seen as being obtained from $\bar{T}_{k}$ by the same step that produces $T_{i+1}$ from $T_{i}$.

- If $T_{i+1}$ is obtained from $T_{i}$ by a chase step with a full GTGD, then $F \notin \bar{T}_{k}\left(v^{\prime}\right)$ holds ensures that the same step is applicable to $\bar{T}_{k}\left(v^{\prime}\right)$ (where $v^{\prime}=v$ ).
- If $T_{i+1}$ is obtained from $T_{i}$ by a propagation step, then $F$ is $\Sigma$-guarded by $T_{i}\left(v^{\prime}\right)$. But then, $T_{i}\left(v^{\prime}\right) \subseteq \bar{T}_{j}\left(v^{\prime}\right)$ ensures that $F$ is also $\Sigma$-guarded by $\bar{T}_{j}\left(v^{\prime}\right)$, and Claim 2 of Lemma A. 1 ensures that $F$ is $\Sigma$-guarded by $\bar{T}_{k}\left(v^{\prime}\right)$. Thus, the propagation step is applicable to vertices $v$ and $v^{\prime}$ in $\bar{T}_{k}$.
Moreover, let $\bar{U}_{1}, \ldots, \bar{U}_{s}$ be the chase trees obtained by propagating $F$ starting from $\bar{U}_{0}$ towards the root using local steps as long as possible. Clearly, $\bar{T}_{0}, \ldots, \bar{T}_{k}, \bar{U}_{0}, \bar{U}_{1}, \ldots, \bar{U}_{s}$ is a correctly formed almost one-pass chase sequence. Let $v^{\prime \prime}$ be the recently updated vertex of $\bar{U}_{s}$,

We cannot simply append the step producing $T_{k+1}$ after $\bar{U}_{s}$ because this step might not be applicable to $v^{\prime \prime}$ or an ancestor thereof. Thus, to obtain the chase sequence satisfying the claim of the lemma, we shall find a place in sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$ where vertex $v^{\prime \prime}$ is introduced, and we shall "replay" all steps from that point onwards. In doing so, we shall use chase steps that introduce the same vertices and labeled nulls, so we will first need to rename these in the sequence $\bar{T}_{0}, \ldots, \bar{T}_{k}, \bar{U}_{0}, \bar{U}_{1}, \ldots, \bar{U}_{s}$.

Let $\ell$ be the smallest integer such that $\bar{T}_{\ell}$ contains $v^{\prime \prime}$. Clearly, $\ell \leq k$ holds. Now let $N$ be the set of labeled nulls introduced by applying a chase step to $v^{\prime \prime}$ or a descendant thereof, and let $W$ the the set of descendants of $v^{\prime \prime}$ in the sequence $\bar{T}_{0}, \ldots, \bar{T}_{j}$. Moreover, let $\bar{T}_{0}^{\prime}, \ldots, \bar{T}_{k}^{\prime}, \bar{U}_{0}^{\prime}, \bar{U}_{1}^{\prime}, \ldots, \bar{U}_{s}^{\prime}$ be the chase sequence obtained by uniformly replacing in $\bar{T}_{0}, \ldots, \bar{T}_{k}, \bar{U}_{0}, \bar{U}_{1}, \ldots, \bar{U}_{s}$ each labeled null in $N$ with a distinct, fresh labeled null, and by uniformly replacing each vertex $w \in W$ by a fresh vertex.

We now transform chase trees $\bar{T}_{\ell+1}, \ldots, \bar{T}_{j}$ into chase tress $\bar{V}_{\ell+1}, \ldots, \bar{V}_{j}$ that reflect the result of "replaying" after $\bar{U}_{s}^{\prime}$ the steps producing the former sequence. Intuitively, each $\bar{V}_{p}$ is a "union" of $\bar{U}_{s}$ and $\bar{T}_{p}$. Formally, for each $p$ with $\ell<p \leq j$, we define $\bar{V}_{p}$ as follows.
(S1) The chase tree $\bar{V}_{p}$ contains the union of the vertices of $\bar{U}_{s}^{\prime}$ and $\bar{T}_{p}$.
(S2) For each vertex $w$ occurring only in $\bar{U}_{s}^{\prime}\left(\right.$ resp. $\left.\bar{T}_{p}\right)$, we define $\bar{V}_{p}(w)=\bar{U}_{s}^{\prime}(w)\left(\right.$ resp. $\left.\bar{V}_{p}(w)=\bar{T}_{p}(w)\right)$.
(S3) For each vertex $w$ occurring in both $\bar{U}_{s}^{\prime}$ and $\bar{T}_{p}$, we define $\bar{V}_{p}(w)=\bar{U}_{s}^{\prime}(w) \cup \bar{T}_{p}(w)$.
(S4) If $\bar{T}_{p}$ is obtained by applying to a vertex $w$ of $\bar{T}_{p-1}$ a chase step with a non-full GTGD $\tau=\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$ and substitutions $\sigma$ and $\sigma^{\prime}$, then, for $w^{\prime}$ the child of $w$ introduced by the step, we extend $\bar{V}_{p}\left(w^{\prime}\right)$ with each fact $G \in \bar{V}_{p-1}(w)$ that is $\Sigma$-guarded by $\sigma^{\prime}(\eta)$.

We now argue that $\bar{T}_{0}^{\prime}, \ldots, \bar{T}_{k}^{\prime}, \bar{U}_{0}^{\prime}, \bar{U}_{1}^{\prime}, \ldots, \bar{U}_{s}^{\prime}, \bar{V}_{\ell+1}, \ldots, \bar{V}_{j}$ contains an almost one-pass chase sequence for $\Sigma$ and $I$ that satisfies the conditions of this lemma.

- Sequence $\bar{T}_{0}^{\prime}, \ldots, \bar{T}_{k}^{\prime}, \bar{U}_{0}^{\prime}, \bar{U}_{1}^{\prime}, \ldots, \bar{U}_{s}^{\prime}$ is clearly a valid almost one-pass chase sequence.
- For $\ell \leq p<j$, either $\bar{V}_{p+1}$ is obtained from $\bar{V}_{p}$ (or $\bar{U}_{s}^{\prime}$ in case $p=\ell$ ) by the same step that produces $\bar{T}_{p+1}$ from $\bar{T}_{p}$, or the step is not applicable. In the latter case, we can simply drop such $\bar{V}_{p+1}$ from the sequence. By dropping all such $\bar{V}_{p+1}$, we clearly obtain a valid almost one-pass chase sequence.
- We have $T_{i} \subseteq \bar{T}_{j}$ by the induction assumption, and steps (S1)-(S3) clearly ensure $T_{i} \subseteq \bar{V}_{j}$. Moreover, $T_{i+1}$ differs from $T_{i}$ only in vertex $v^{\prime}$, where $T_{i+1}\left(v^{\prime}\right)=T\left(v^{\prime}\right) \cup\{F\}$ holds. Our construction, however, clearly ensures $F \in \bar{U}_{s}^{\prime}\left(v^{\prime \prime}\right)$, and step (S4) ensures that $F$ is propagated in each chase step with a non-full GTGD introducing a vertex on the unique path from $v^{\prime \prime}$ to $v^{\prime}$. Thus, $T_{i+1} \subseteq \bar{V}_{j}$ holds.

Lemma A.7. For each base fact $F$ and each almost one-pass tree-like chase proof of $F$ from I and $\Sigma$, there exists a one-pass tree-like chase proof of $F$ from $I$ and $\Sigma$.

Proof. Consider an arbitrary base fact $F$ and an arbitrary almost one-pass tree-like chase proof $T_{0}, \ldots, T_{n}$ of $F$ from $I$ and $\Sigma$. Since $F$ is a base fact, without loss of generality we can assume that $F$ occurs in the facts of the root vertex. Now let $T_{i}$ be the first chase tree that contains $F$ in the root, and let $W$ be the set containing each non-root vertex $v$ occurring in any of the chase trees such that no propagation step is applied to $v$. We transform this proof to a one-pass proof as follows. First, we delete each $T_{j}$ with $i<j \leq n$. Next, we delete in each remaining $T_{j}$ each vertex $v \in W$ and each descendant of $v$. Finally, we delete each remaining $T_{j}$ that is equal to $T_{j+1}$. After this transformation, every vertex has a propagation step applied to it. It is straightforward to see that the result is a one-pass tree-like chase sequence. Moreover, since $F$ occurs in the root, the sequence is a tree-like chase proof of $F$ from $I$ and $\Sigma$.

## A. 2 Proof of Proposition 4.7: Rewriting Criterion Using One-pass Chase Proofs

Proposition 4.7. A Datalog program $\Sigma^{\prime}$ is a rewriting of a finite set of GTGDs $\Sigma$ in head-normal form if

- $\Sigma^{\prime}$ is a logical consequence of $\Sigma$,
- each Datalog rule of $\Sigma$ is a logical consequence of $\Sigma^{\prime}$, and
- for each base instance $I$, each one-pass tree-like chase sequence $T_{0}, \ldots, T_{n}$ for $I$ and $\Sigma$, and each loop $T_{i}, \ldots, T_{j}$ at the root vertex $r$ with output fact $F$, there exist a Datalog rule $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(r)$ and $\sigma(H)=F$.

Proof. Let $\Sigma$ and $\Sigma^{\prime}$ be as specified in the proposition, let $I$ be an arbitrary base instance, and let $F$ be an arbitrary base fact. Since $\Sigma^{\prime}$ is a logical consequence of $\Sigma$, it is clear that $I, \Sigma^{\prime} \vDash F$ implies $I, \Sigma \vDash F$. Thus, we assume that $I, \Sigma \vDash F$ holds, and we prove that $I, \Sigma^{\prime} \vDash F$ holds as well. By Theorem 4.2, there exists a one-pass tree-like chase proof $T_{0}, \ldots, T_{n}$ of $F$ from $I$ and $\Sigma$. Without loss of generality, we can assume that $F$ is produced in the last step of the proof, and so the recently updated vertex of $T_{n}$ is root vertex $r$. Let $i_{0}<\cdots<i_{m}$ be exactly the indexes between 0 and $n$ such that the recently updated vertex of $T_{i_{j}}$ is $r$. We next construct a tree-like chase sequence $\bar{T}_{0}, \ldots, \bar{T}_{k}$ for $I$ and $\Sigma^{\prime}$ such that $T_{n}(r) \subseteq \bar{T}_{k}(r)$. To formalize our inductive construction of this chase sequence, we shall also construct a sequence of indexes $\ell_{0}, \ldots, \ell_{m}$ such that $\ell_{m}=k$ and, for each $j$ with $0 \leq j \leq m$, we have $T_{i_{j}}(r) \subseteq \bar{T}_{\ell_{j}}(r)$; in other words, each index $\ell_{j}$ helps us establish the inductive property by relating $T_{i_{j}}$ and $\bar{T}_{\ell_{j}}$. For the base case, $i_{0}=0$ holds by the definition of a tree-like chase proof; thus, we set $\bar{T}_{0}=T_{0}$ and $\ell_{0}=0$, and the required property clearly holds. For the inductive step, we consider arbitrary $0<j \leq m$ such that the claim holds for $j-1$, and assume that the sequence constructed thus far is $\bar{T}_{\ell_{0}}, \ldots, \bar{T}_{\ell_{j-1}}$. We have the following two cases.

- The recently updated vertex of $T_{i_{j}-1}$ is $r$. Thus, $i_{j}-1=i_{j-1}$, and $T_{i_{j}}$ is obtained from $T_{i_{j-1}}$ by a chase step with a full GTGD $\beta \rightarrow H \in \Sigma$ producing a fact $G \in T_{i_{j}}(r)$. The second condition of the proposition ensures that $\beta \rightarrow H$ is a logical consequence of $\Sigma^{\prime}$, so $G$ can be derived from $\bar{T}_{\ell_{1}}(r)$ and the Datalog rules of $\Sigma^{\prime}$ using $\wp$ steps. We then define $\ell_{j}=\ell_{j-1}+\wp$, and we append the corresponding steps to obtain the sequence $\bar{T}_{0}, \ldots, \bar{T}_{\ell_{j-1}}, \ldots, \bar{T}_{\ell_{j}}$.
- Otherwise, $T_{i_{j-1}}, \ldots, T_{i_{j}}$ is a loop at the root vertex $r$ with some output fact $G \in T_{i_{j}}(r)$. The third condition of the proposition ensures that there exists a Datalog rule $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(r)$ and $\sigma(H)=G$. We define $\ell_{j}=\ell_{j-1}+1$, and we define $\bar{T}_{\ell_{j}}$ as the the chase tree containing just the root vertex $r$ such that $\bar{T}_{\ell_{j}}(r)=\bar{T}_{\ell_{j-1}}(r) \cup\{G\}$; thus, $\bar{T}_{\ell_{j}}$ is obtained from $\bar{T}_{\ell_{j-1}}$ by applying the Datalog rule $\beta \rightarrow H \in \Sigma^{\prime}$ to the root vertex $r$. Moreover, $T_{i_{j}}(r) \subseteq \bar{T}_{\ell_{j}}(r)$ clearly holds, as required.


## A. 3 Proof of Proposition A.8: Properties of One-Pass Chase Proofs

We finally prove a property that will be needed in the proofs in Appendices B-E. This property intuitively ensures that, as soon as a fact $F$ is derived in the child vertex of a loop such that $F$ does not contain any null values introduced by the child, the loop is completed and $F$ is propagated to the parent vertex.

Proposition A.8. For each loop $T_{i}, \ldots, T_{j}$ at a vertex $v$ in a one-pass tree-like chase proof for some $I$ and $\Sigma$, for $i<k \leq j$, and for $v^{\prime}$ the vertex introduced in $T_{i+1}$, set terms $\left(T_{k}\left(v^{\prime}\right)\right) \backslash \operatorname{nulls}\left(T_{i+1}\left(v^{\prime}\right)\right)$ is $\Sigma$-guarded by $T_{i}(v)$.

Proof. Consider an arbitrary loop $T_{i}, \ldots, T_{j}$ at a vertex $v$ in a one-pass tree-like chase proof for some $I$ and $\Sigma$, and let $v^{\prime}$ be the child of $v$ introduced by the chase step producing $T_{i+1}$. We prove the claim by induction on $k$ with $i<k \leq j$. For the induction base $k=i+1$, the definition of a chase step with a non-full GTGD clearly ensures this claim for $T_{i+1}\left(v^{\prime}\right)$. For the induction step, consider an arbitrary $k$ such that the claim holds. Our claim holds trivially if $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right)$, so we assume that $T_{k+1}\left(v^{\prime}\right) \backslash T_{k}\left(v^{\prime}\right)$ contains exactly one fact $F$, which can be derived in one of the following two ways.

- Assume $F$ is obtained by a propagation step to vertex $v^{\prime}$. Then, $F$ is $\Sigma$-guarded by $T_{k}\left(v^{\prime}\right)$, so terms $(F) \subseteq \operatorname{terms}\left(T_{k}\left(v^{\prime}\right)\right) \cup \operatorname{consts}(\Sigma)$ holds.
- Assume $F$ is obtained by applying a full GTGD $\tau \in \Sigma$ to $T_{k}\left(v^{\prime}\right)$ using a substitution $\sigma$. Then, $\tau$ contains a guard atom $A$ in the body such that $\sigma(A) \subseteq T_{k}\left(v^{\prime}\right)$; moreover, the head $\tau$ contains all variables of $A$, and so we have terms $(F) \subseteq \operatorname{terms}\left(T_{k}\left(v^{\prime}\right)\right) \cup \operatorname{consts}(\Sigma)$.
Either way, we have terms $\left(T_{k+1}\left(v^{\prime}\right)\right) \subseteq \operatorname{terms}\left(T_{k}\left(v^{\prime}\right)\right) \cup \operatorname{consts}(\Sigma)$. By the induction assumption, set terms $\left(T_{k}\left(v^{\prime}\right)\right) \backslash \operatorname{nulls}\left(T_{i+1}\left(v^{\prime}\right)\right)$ is $\Sigma$-guarded by $T_{i}(v)$, and so set terms $\left(T_{k_{1}}\left(v^{\prime}\right)\right) \backslash \operatorname{nulls}\left(T_{i+1}\left(v^{\prime}\right)\right)$ is also $\Sigma$-guarded by $T_{i}(v)$, as required.


## B PROOFS FOR ExbDR

## B. 1 Proof of Proposition 5.7: Properties of ExbDR

Proposition 5.7. Each application of the ExbDR inference rule to $\tau, \tau^{\prime}$, and $\theta$ as in Definition 5.5 satisfies the following properties.

1. Some atom $A_{i}^{\prime}$ with $1 \leq i \leq n$ is a guard in $\tau^{\prime}$.
2. For each $1 \leq i \leq n$ such that $A_{i}^{\prime}$ is a guard of $\tau^{\prime}$, and for $\sigma$ the $\vec{y}-M G U$ of $A_{i}^{\prime}$ and the corresponding atom $A_{i}$ such that $\sigma(\vec{x}) \cap \vec{y}=\emptyset$, it is the case that $\operatorname{vars}\left(\sigma\left(A_{j}^{\prime}\right)\right) \cap \vec{y} \neq \emptyset$ for each $1 \leq j \leq n$.
3. The result is a GTGD whose body and head width are at most bwidth $(\Sigma)$ and hwidth $(\Sigma)$, respectively.

Proof of Claim 1. Let $G$ be a guard for $\tau^{\prime}$. For the sake of a contradiction, assume that $G$ is not one of the atoms $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$-that is, $G \in \beta^{\prime}$. Since $n \geq 1$, atom $A_{1}^{\prime}$ in the body of $\tau^{\prime}$ is matched to $A_{1}$ in the head of $\tau$. Since $\tau$ is in head-normal form, $A_{1}$ contains at least one variable $y \in \vec{y}$. Moreover, the conditions of the ExbDR inference rule ensure $\theta(y)=y$. Since $y$ does not occur in $A_{1}^{\prime}$ and $\theta$ unifies $A_{1}^{\prime}$ and
$A_{1}$, atom $A_{1}^{\prime}$ contains at some position a variable $z$ such that $\theta(z)=y$. Since $G$ is a guard for $\tau^{\prime}$, variable $z$ occurs in $G$. Therefore, we have $\operatorname{vars}(\theta(G)) \cap \vec{y} \neq \emptyset$, which contradicts the requirement $\operatorname{vars}\left(\theta\left(\beta^{\prime}\right)\right) \cap \vec{y}=\emptyset$ of the ExbDR inference rule.

Proof of Claim 2. Consider arbitrary $i$ such that $1 \leq i \leq n$ and $A_{i}^{\prime}$ is a guard of $\tau^{\prime}$, and let $\sigma$ be an MGU of $A_{i}^{\prime}$ and the corresponding atom $A_{i}$ of $\tau$. Since $\theta$ is a unifier of $A_{i}^{\prime}$ and $A_{i}$ as well as of other pairs of atoms, there clearly exists a substitution $\rho$ such that $\theta=\rho \circ \sigma$. Now consider an arbitrary $A_{j}^{\prime}$ with $1 \leq j \leq n$ in $\tau^{\prime}$. Substitution $\theta$ matches $A_{j}^{\prime}$ to the corresponding atom $A_{j}$ in the head of $\tau$. Since TGD $\tau$ is in head-normal form, atom $A_{j}$ contains at least one variable $y \in \vec{y}$. Since $\theta(y)=y$, we necessarily have $y \in \operatorname{vars}\left(\theta\left(A_{j}^{\prime}\right)\right)$. Consequently, atom $A_{j}^{\prime}$ contains some variable $z$ such that $\theta(z)=y$. Since $A_{i}^{\prime}$ is a guard for $\tau^{\prime}$, variable $z$ occurs in $A_{i}^{\prime}$. Now assume for the sake of a contradiction that $\sigma(z) \neq y$. Then $\sigma(z)=\sigma(x)$ for some $x \in \operatorname{vars}\left(A_{i}\right)$ and $y=\theta(z)=\rho(\sigma(z))=\rho(\sigma(x))=\theta(x)$. However, this contradicts the requirement $\theta(\vec{x}) \cap \vec{y}=\emptyset$ of the ExbDR inference rule.

Proof of Claim 3. By Claim 1, there exists $i$ with $1 \leq i \leq n$ such that atom $A_{i}^{\prime}$ is a guard for $\tau^{\prime}$. Thus, $\operatorname{vars}\left(\beta^{\prime}\right) \cup \operatorname{vars}\left(H^{\prime}\right) \subseteq \operatorname{vars}\left(A_{i}^{\prime}\right)$. The ExbDR inference rule ensures vars $\left(\theta\left(\beta^{\prime}\right)\right) \cap \vec{y}=\emptyset$, which in turn ensures $\operatorname{vars}\left(\theta\left(\beta^{\prime}\right)\right) \subseteq \operatorname{vars}\left(\theta\left(A_{i}^{\prime}\right)\right) \backslash \vec{y}$. Now let $G$ be a guard for $\tau$. We clearly have $\operatorname{vars}(\theta(\beta)) \subseteq \operatorname{vars}(\theta(G))$. Moreover, $\theta\left(y_{i}\right)=y_{i}$ and $\theta(x) \cap \vec{y}=\emptyset$ ensure $\operatorname{vars}(\theta(\eta)) \cup \operatorname{vars}\left(\theta\left(A_{i}\right)\right) \subseteq \operatorname{vars}(\theta(G)) \cup \vec{y}$. Thus, $\theta(G)$ is a guard for the TGD produced by the ExbDR inference rule. Finally, since $G$ contains all variables of $\tau$, the widths of the resulting TGD and $\tau$ are equal.

## B. 2 Proof of Theorem 5.8: Correctness and Complexity of ExbDR

Theorem 5.8. Program $\operatorname{ExbDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in time
 $c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Proof of Correctness. Let $\Sigma^{\prime}$ be the set $\Sigma$ closed under the ExbDR inferences rule as specified in Definition 5.3. It is straightforward to see that $\Sigma^{\prime}$ is a logical consequence of $\Sigma$, so $\operatorname{ExbDR}(\Sigma)$ is also a logical consequence of $\Sigma$. Moreover, $\operatorname{ExbDR}(\Sigma)$ contains each full GTGD of $\Sigma$ up to redundancy, so each full GTGD of $\Sigma$ is logically entailed by $\operatorname{ExbDR}(\Sigma)$. We next consider an arbitrary base instance $I$ and a one-pass tree-like chase sequence for $I$ and $\Sigma$, and we show the following property:
$(\star)$ for each loop $T_{i}, \ldots, T_{j}$ in the sequence at some vertex $v$ with output fact $F$, there exist a full GTGD $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(v)$ and $F=\sigma(H)$.
Since $\operatorname{ExbDR}(\Sigma)$ contains all full TGDs of $\Sigma^{\prime}$ and this property holds for the root vertex $r$, Proposition 4.7 ensures that $\operatorname{ExbDR}(\Sigma)$ is a rewriting of $\Sigma$.

Our proof is by induction on the length of the loop. The base case and the inductive step have the same structure, so we consider them jointly. Thus, consider an arbitrary loop $T_{i}, T_{i+1}, \ldots, T_{j-1}, T_{j}$ at vertex $v$ in the sequence, and assume that the claim holds for all shorter loops in the sequence. By the definition of the loop, chase tree $T_{i+1}$ is obtained from $T_{i}$ by applying a chase step to some non-full TGD $\forall \vec{x}\left[\beta_{0} \rightarrow \exists \vec{y} \eta_{0}\right] \in \Sigma$. Let $\sigma_{0}$ and $\sigma_{0}^{\prime}$ be the substitutions used in this chase step, let $N=\operatorname{nulls}\left(\operatorname{rng}\left(\sigma_{0}^{\prime}\right)\right) \backslash \operatorname{nulls}\left(\operatorname{rng}\left(\sigma_{0}\right)\right)$, let $v^{\prime}$ be the child of $v$ introduced in $T_{i+1}$, and let $S \subseteq T_{i}(v)$ be the facts that are copied to $T_{i+1}\left(v^{\prime}\right)$ because they are $\Sigma$-guarded by $\sigma_{0}^{\prime}\left(\eta_{0}\right)$. Thus, we have $T_{i+1}\left(v^{\prime}\right)=S \cup \sigma_{0}^{\prime}\left(\eta_{0}\right)$. By Proposition A. 8 and the fact that a chase step is applied only if propagation to the parent is not applicable, the output fact of the loop is added to $T_{j-1}\left(v^{\prime}\right)$ in step $j-1$, and in $T_{j}$ this fact is propagated back to $T_{j}(v)$. In other words, for each $k$ with $i<k<j-1$, each fact in $T_{k}\left(v^{\prime}\right) \backslash S$ contains at least one labeled null from $N$, or the fact would be $\Sigma$-guarded by $T_{i}(v)$ and thus propagated back to vertex $v$. We show that, in the loop $T_{i}, T_{i+1}, \ldots, T_{j-1}, T_{j}$ fixed above, the following property holds for each $k$ with $i<k<j-1$ :
$(\diamond)$ there exist a GTGD $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta] \in \Sigma^{\prime}$, a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(v)$, and a substitution $\sigma^{\prime}$ that extends $\sigma$ by mapping $\vec{y}$ to fresh labeled nulls such that $T_{k}\left(v^{\prime}\right) \subseteq S \cup \sigma^{\prime}(\eta)$.
We prove $(\diamond)$ by induction on $k$. We have already proved the base case $k=i+1$ above. For the inductive step, assume that $(\diamond)$ holds for some $k$, so there exists a GTGD $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta] \in \Sigma^{\prime}$ and substitutions $\sigma$ and $\sigma^{\prime}$ satisfying $(\diamond)$ for $k$. Now consider $T_{k+1}$. Property ( $\diamond$ ) holds by the inductive hypothesis if $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right)$-that is, if the step involves a descendant of $v^{\prime}$. Otherwise, $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right) \cup\{G\}$ where fact $G$ is obtained in one of the following two ways.

- A full GTGD in $\Sigma$ derives $G$ from $T_{k}\left(v^{\prime}\right)$. Set $\Sigma^{\prime}$ contains this GTGD up to redundancy, so by Definition 5.1 there exist a full GTGD $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\rho$ such that $\rho\left(\beta^{\prime \prime}\right) \subseteq T_{k}\left(v^{\prime}\right)$ and $\rho\left(H^{\prime}\right)=G$.
- Fact $G$ is the output of a loop at vertex $v^{\prime}$. But then, this loop is shorter than $T_{i}, \ldots, T_{j}$ so, by property $(\downarrow)$, there exists a full GTGD $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\rho$ such that $\rho\left(\beta^{\prime \prime}\right) \subseteq T_{k}\left(v^{\prime}\right)$ and $\rho\left(H^{\prime}\right)=G$.
Since $\eta$ is in head-normal form, each atom in $\sigma^{\prime}(\eta)$ contains at least one labeled null of $N$. Now let $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$ be the atoms of $\beta^{\prime \prime}$ that are matched to the atoms in $\sigma^{\prime}(\eta)$. Atom $\rho\left(H^{\prime}\right)$ contains at least one labeled null of $N$, so $n \geq 1$. Thus, we can assume that $\beta^{\prime \prime} \rightarrow H^{\prime}$ is of the form $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ where $\left\{\rho\left(A_{1}^{\prime}\right), \ldots, \rho\left(A_{n}^{\prime}\right)\right\} \subseteq \sigma^{\prime}(\eta)$ and $\rho\left(\beta^{\prime}\right) \subseteq S$. Also, since $\beta^{\prime \prime} \rightarrow H^{\prime}$ is guarded, at least one of $A_{i}^{\prime}$ is a guard for $\beta^{\prime \prime} \rightarrow H^{\prime}$. Let $A_{1}, \ldots, A_{n}$ be the atoms of $\eta$ such that $\sigma^{\prime}\left(A_{i}\right)=\rho\left(A_{i}^{\prime}\right)$ for $1 \leq i \leq n$. Since $\sigma^{\prime}$ maps each $y \in \vec{y}$ to a distinct labeled null that does not occur in $T_{i}$, we have $\sigma^{\prime}(\vec{x}) \cap \sigma^{\prime}(\vec{y})=\emptyset$. Thus, there exists a $\vec{y}$-MGU $\theta$ of $A_{1}, \ldots, A_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$ satisfying $\theta(\vec{x}) \cap \vec{y}=\emptyset$. Conjunction $\rho\left(\beta^{\prime}\right)$ does not contain a labeled null of $N$, so vars $\left(\theta\left(\beta^{\prime}\right)\right) \cap \vec{y}=\emptyset$ holds. Thus, the preconditions of the ExbDR inference rule are satisfied for $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta]$ and $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$, so the ExbDR rule derives $\tau=\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \exists \vec{y} \theta(\eta) \wedge \theta\left(H^{\prime}\right)$.

Moreover, some $A_{i}^{\prime}$ is a guard so all variables of $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ participate in unification, and thus we can extend $\sigma$ and $\sigma^{\prime}$ to substitutions $\zeta$ and $\zeta^{\prime}$, respectively, covering these variables such that $\zeta(\theta(\beta)) \cup \zeta\left(\theta\left(\beta^{\prime}\right)\right) \subseteq T_{i}(v)$ and $T_{k+1}\left(v^{\prime}\right) \subseteq S \cup \zeta^{\prime}(\theta(\eta)) \cup \zeta^{\prime}\left(\theta\left(H^{\prime}\right)\right)$. Set $\Sigma^{\prime}$ contains $\tau$ up to redundancy. Since $G \notin T_{k}\left(v^{\prime}\right)$, GTGD $\tau$ is not a syntactic tautology, so there exists a GTGD $\forall \vec{x}_{1}\left[\beta_{1} \rightarrow \exists \vec{y}_{1} \eta_{1}\right] \in \Sigma^{\prime}$ and substitution $\mu$ such that $\operatorname{dom}(\mu)=\vec{x}_{1} \cup \vec{y}_{1}, \mu\left(\vec{x}_{1}\right) \subseteq \vec{x}_{2}, \mu\left(\vec{y}_{1}\right) \subseteq \vec{y}_{1} \cup \vec{y}_{2}$ and $\mu(y) \neq \mu\left(y^{\prime}\right)$ for distinct $y$ and $y^{\prime}$ in $\vec{y}_{1}$, and $\mu\left(\beta_{1}\right) \subseteq \theta(\beta) \wedge \theta\left(\beta^{\prime}\right)$ and $\mu\left(\eta_{1}\right) \supseteq \theta(\eta) \wedge \theta\left(H^{\prime}\right)$. Now let $\sigma_{1}$ be the substitution defined as $\sigma_{1}(x)=\zeta(\mu(x))$ on each $x \in \vec{x}$, and let $\sigma_{1}^{\prime}$ be the extension of $\sigma_{1}$ to $\vec{y}_{1}$ such that $\sigma_{1}^{\prime}(y)=\zeta^{\prime}(\mu(y))$ for each $y \in \vec{y}$. Clearly, $\sigma_{1}\left(\beta_{1}\right) \subseteq T_{k}\left(v^{\prime}\right)$ and $T_{k+1}\left(v^{\prime}\right) \subseteq S \cup \sigma_{1}^{\prime}\left(\eta_{1}\right)$ hold, so property $(\diamond)$ is satisfied.

To complete the proof, consider now the derivation of $T_{j-1}$. By property ( $\diamond$ ), there exists a GTGD $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta] \in \Sigma^{\prime}$ and substitutions $\sigma$ and $\sigma^{\prime}$ such that $\sigma(\beta) \subseteq T_{i}(v)$ and $T_{j-2}\left(v^{\prime}\right)=S \cup \sigma^{\prime}(\eta)$. Then, as above, $\Sigma^{\prime}$ contains a full TGD of the form $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ that satisfies $\rho\left(A_{1}^{\prime}\right) \cup \cdots \cup \rho\left(A_{n}^{\prime}\right) \subseteq \sigma^{\prime}(\eta)$ and $\rho\left(\beta^{\prime}\right) \subseteq S$ for some substitution $\rho$. A minor difference is that $\rho\left(H^{\prime}\right)$ does not contain a labeled null introduced by $\sigma^{\prime}\left(\eta_{0}\right)$, so $n=0$ is possible; however, in such a case, this TGD immediately satisfies property $(\boldsymbol{*})$. Moreover, if $n>0$, then $\beta \rightarrow \exists \vec{y} \eta$ can again be resolved with $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ to produce

$$
\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \exists \vec{y} \theta(\eta) \wedge \theta\left(H^{\prime}\right) \in \Sigma^{\prime}
$$

satisfying $\operatorname{vars}\left(\theta\left(H^{\prime}\right)\right) \cap \vec{y}=\emptyset$. This TGD is transformed into head-normal form by Definition 5.3, so $\forall \vec{x}\left[\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)\right]$ is contained in $\Sigma^{\prime}$ up to redundancy. But then, $\Sigma^{\prime}$ contains a full GTGD that satisfies property $(\boldsymbol{*})$ by the same argument as above.

Proof of Complexity. Fix $\Sigma, r, w_{b}, w_{h}, c$, and $a$ as stated in the theorem. The number of different body atoms of arity $a$ constructed using $r$ relations, $w_{b}$ variables, and $c$ constants is clearly bounded by $\ell_{b}=r \cdot\left(w_{b}+c\right)^{a}$. Moreover, by the third claim of Proposition 5.7, the number of variables in the head of each TGD is bounded by $w_{h}$, so the number of head atoms is bounded by $\ell_{h}=r \cdot\left(w_{h}+c\right)^{a}$. The body (resp. head) of each GTGD corresponds to a subset of these atoms, so number of different GTGDs up to variable renaming is bounded by $\wp=2^{\ell_{b}} \cdot 2^{\ell_{h}}$. Thus, the ExbDR inference rule needs to be applied to at most $\wp^{2}=2^{2\left(\ell_{b}+\ell_{h}\right)}$ pairs of GTGDs. For each such pair, one might need to consider each possible way to match the $\ell_{b}$ body atoms of $\tau^{\prime}$ to $\ell_{h}$ head atoms of $\tau$, and there are at most $\left(\ell_{h}\right)^{\ell_{b}} \leq 2^{\ell_{b} \cdot \ell_{h}}$ of these. Consequently, unifier $\theta$ may need to be computed at most $2^{2\left(\ell_{b}+\ell_{h}\right)} \cdot 2^{\ell_{b} \cdot \ell_{h}} \leq 2^{5 \cdot \ell_{b} \cdot \ell_{h}}=32^{\ell_{b} \cdot \ell_{h}}$ times. To check whether TGD $\tau_{1}=\forall \vec{x}_{1}\left[\beta_{1} \rightarrow \vec{y}_{1} \eta_{1}\right]$ is subsumed by $\tau_{2}=\forall \vec{x}_{2}\left[\beta_{2} \rightarrow \vec{y}_{2} \eta_{2}\right]$, we can proceed as follows. First, we consider all possible ways to match an atom of $\beta_{1}$ to an atom of $\beta_{2}$; since both conjunctions contain at most $\ell_{b}$ atoms, there are at most $\ell_{b} \ell_{b} \leq 2^{\ell_{b}{ }^{2}}$ such matchings. Second, we analogously consider each of at most $2^{\left(\ell_{h}\right)^{2}}$ ways to match an atom of $\eta_{2}$ to an atom of $\eta_{1}$. Once all atoms have been matched, we try to find a substitution $\mu$ satisfying Definition 5.1 in linear time. Thus, a subsumption check for pairs of TGDs takes at most $2^{\ell_{b}{ }^{2}} \cdot 2^{\ell_{h}{ }^{2}}$ steps. Finally, unification of atoms requires time that is linear in $a$, and all other steps require linear time too.

## C PROOFS FOR SkDR

## C. 1 Proof of Proposition 5.12: Properties of SkDR

We reuse results by de Nivelle [19] about unification of atoms in guarded rules. The variable depth [19, Definition 3] of an atom is defined as -1 if the atom is ground, or as the maximum number of nested function symbols that contain a variable of the atom. Moreover, an atom is weakly covering [19, Definition 6] if each nonground functional subterm of the atom contains all variables of the atom. Finally, de Nivelle [19, Theorem 1] says that, for $\theta$ an MGU of weakly covering atoms $A$ and $B$, atom $C=\theta(A)=\theta(B)$ is also weakly covering, the variable depth of $C$ is bounded by the variable depth of $A$ and $B$, and the number of variables of $C$ is bounded by the number of variables of $A$ and $B$ too.

Proposition 5.12. Each application of the SkDR inference rule to rules $\tau$ and $\tau^{\prime}$ as in Definition 5.10 produces a guarded rule.
Proof. Consider arbitrary rules $\tau=\beta \rightarrow H$ and $\tau^{\prime}=A^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ and an MGU $\theta$ of $H$ and $A^{\prime}$ satisfying the preconditions of the SkDR inference rule. Atom $H$ thus contains a Skolem symbol, and rule $\tau$ is guarded; consequently, atom $H$ is weakly covering, it contains a term of the form $f(\vec{t})$ where $\vec{t}$ consists of constants and all variables of the rule, and the variable depth of $H$ is at most one. The corresponding atom $A^{\prime}$ can be of the following two forms.

- Atom $A^{\prime}$ is Skolem-free. But then, $A^{\prime}$ contains all variables of $\tau^{\prime}$, and it is clearly weakly covering. By de Nivelle [19, Theorem 1], atom $\theta\left(A^{\prime}\right)$ is weakly covering and has variable depth at most one; consequently, each atom in rule $\theta\left(A^{\prime}\right) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ is weakly covering and has variable depth at most one. Moreover, the variable depth of $\theta(H)$ is also at most one, which can be only if $\theta$ maps each variable in $H$ to another variable or a constant. Thus, each atom in rule $\theta(\beta) \rightarrow \theta(H)$ is weakly covering and has variable depth at most one; moreover, $\theta(\beta)$ contains an atom that contains all variables of the rule. But then, rule $\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ is guarded, as required.
- Atom $A^{\prime}$ contains a Skolem symbol. But then, $A^{\prime}$ is weakly covering by the definition of guarded rules, and its variable depth is at most one. By de Nivelle [19, Theorem 1], atom $\theta(H)=\theta\left(A^{\prime}\right)$ is weakly covering and has variable depth at most one, which can be the case only if $\theta$ maps all variables to other variables or constants. Consequently, rules $\theta(\beta) \rightarrow \theta(H)$ and $\theta\left(A^{\prime}\right) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ are both guarded. But then, rule $\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ is guarded, as required.


## C. 2 Proof of Theorem 5.13: Correctness and Complexity of SkDR

Theorem 5.13. Program $\operatorname{SkDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in time
 $w_{b}=\operatorname{bwidth}(\Sigma), c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Proof of Correctness. Let $\Sigma$ be an arbitrary finite set of GTGDs, and let $\Sigma^{\prime}$ be the set of rules obtained from $\Sigma$ as specified in Definition 5.3. It is straightforward to see that $\Sigma^{\prime}$ is a logical consequence of the Skolemization of $\Sigma$, so $\operatorname{SkDR}(\Sigma)$ is also a logical consequence of $\Sigma$. Moreover, $\operatorname{SkDR}(\Sigma)$ contains each full TGD of $\Sigma$ up to redundancy, so each full TGD of $\Sigma$ is logically entailed by $\operatorname{SkDR}(\Sigma)$. We next consider an arbitrary base instance $I$ and a one-pass tree-like chase sequence for $I$ and $\Sigma$, and we show the following property:
$(\star)$ for each loop $T_{i}, \ldots, T_{j}$ in the sequence at some vertex $v$ with output fact $F$, there exist a Skolem-free rule $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i}(v)$ and $F=\sigma(H)$.
Since $\operatorname{SkDR}(\Sigma)$ contains all Skolem-free rules of $\Sigma^{\prime}$ and this property holds for the root vertex $r$, Proposition 4.7 ensures that $\operatorname{SkDR}(\Sigma)$ is a rewriting of $\Sigma$.

Our proof is by induction in the length of the loop. The base case and the inductive step have the same structure, so we consider them jointly. Thus, consider an arbitrary loop $T_{i}, T_{i+1}, \ldots, T_{j-1}, T_{j}$ at vertex $v$ in the sequence, and assume that the claim holds for all shorter loops in the sequence. By the definition of a loop, chase tree $T_{i+1}$ is obtained from $T_{i}$ by applying a chase step to some non-full GTGD $\tau \in \Sigma$ and substitution $\gamma$. Let $v^{\prime}$ be the child of $v$ introduced in $T_{i+1}$, let $S \subseteq T_{i}(v)$ be the facts that are copied to $T_{i+1}\left(v^{\prime}\right)$ because they are $\Sigma$-guarded by the instantiated head of $\tau$, let $N=\left\{n_{1}, \ldots, n_{m}\right\}$ be the set of labeled nulls introduced in the chase step for the existentially quantified variables $y_{1}, \ldots, y_{m}$ of $\tau$, let $v$ be a function that maps each labeled null $n_{i}$ to the ground term $f_{i}(\gamma(\vec{x}))$ where $f_{i}$ is the symbol used in the Skolemization of $y_{i}$. For $U$ a set of facts, let $v(U)$ be the result of replacing each occurrence of a labeled null $n \in \operatorname{dom}(v)$ in $U$ with $v(n)$ and eliminating any duplicate facts in the result. Clearly, the inverse function $v^{-}$is well-defined, and we define $v^{-}(U)$ for $U$ a set of facts in the obvious way. By Proposition A. 8 and the fact that propagation is applied eagerly, the output fact of the loop is added to $T_{j-1}\left(v^{\prime}\right)$ in step $j-1$, and in $T_{j}$ this fact is propagated back to $T_{j}(v)$. In other words, for each $k$ with $i<k<j-1$, each fact in $T_{k}\left(v^{\prime}\right) \backslash S$ contains at least one labeled null from $N$, or the fact would be $\Sigma$-guarded by $T_{i}(v)$ and would thus be propagated back to vertex $v$. We now show that the following property holds for each $k$ with $i<k \leq j-1$ :
$(\diamond)$ for each fact $G \in T_{k}\left(v^{\prime}\right) \backslash S$, there exist a rule $\beta \rightarrow H \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\beta$ is Skolem-free, $\sigma(\beta) \subseteq v\left(T_{i}(v)\right)$, and $\sigma(H)=v(G)$.
Property $(\diamond)$ implies $(\diamond)$ : fact $F$ does not contain a labeled null from $N$, so the rule $\beta \rightarrow H \in \Sigma^{\prime}$ whose existence is implied by $(\diamond)$ for $k=j-1$ is actually a Skolem-free rule that satisfies ( $\downarrow$ ).

We next prove property $(\diamond)$ by a nested induction on $k$. For the base case $k=i+1$, property $(\diamond)$ holds due to the fact that $\Sigma^{\prime}$ contains the rules obtained by Skolemizing GTGD $\tau$. For the inductive step, assume that $(\diamond)$ holds for some $k$ and consider the possible ways to obtain $T_{k+1}$ from $T_{k}$. Property $(\diamond)$ holds by the inductive hypothesis if $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right)$-that is, if the step involves a descendant of $v^{\prime}$. Otherwise, $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right) \cup\{G\}$ where fact $G$ is obtained in one of the following two ways.

- A full TGD in $\Sigma$ derives $G$ from $T_{k}\left(v^{\prime}\right)$. Set $\Sigma^{\prime}$ contains this TGD up to redundancy, so by Definition 5.1 there exist a Skolem-free rule $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\sigma^{\prime}$ such that $\sigma^{\prime}\left(\beta^{\prime \prime}\right) \subseteq v\left(T_{k}\left(v^{\prime}\right)\right)$ and $\sigma^{\prime}\left(H^{\prime}\right)=v(G)$.
- Fact $G$ is the output of a loop at vertex $v^{\prime}$. But then, this loop is shorter than $T_{i}, \ldots, T_{j}$ so, by property $(\boldsymbol{*})$, there exist a Skolem-free rule $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\sigma^{\prime}$ such that $\sigma^{\prime}\left(\beta^{\prime \prime}\right) \subseteq v\left(T_{k}\left(v^{\prime}\right)\right)$ and $\sigma^{\prime}\left(H^{\prime}\right)=v(G)$.
Now let $W=\left\{B^{\prime} \in \beta^{\prime \prime} \mid \sigma^{\prime}\left(B^{\prime}\right) \notin S\right\}$. We next show that set $\Sigma^{\prime}$ contains up to redundancy the result of "resolving away" each atom $B^{\prime} \in W$. A slight complication arises due to the fact that the SkDR inference rule considers only two rules at a time, and that the result of each inference is contained in $\Sigma^{\prime}$ up to redundancy. Thus, we will achieve our goal by showing that the SkDR inference rule can be applied up to $n=|W|$ times. Our proof is by induction on $1 \leq \ell \leq n$. Towards this goal, we shall define $n$ rules $\beta_{\ell}^{\prime \prime} \rightarrow H_{\ell}^{\prime}$, substitutions $\sigma_{\ell}^{\prime}$, and sets of atoms $W=W_{0} \supsetneq \cdots \supsetneq W_{n}$ for $\ell$ with $0 \leq \ell \leq n$ satisfying the following invariant:
(*) $\sigma_{\ell}^{\prime}\left(\beta_{\ell}^{\prime \prime}\right) \subseteq S \cup\left\{\sigma^{\prime}\left(B^{\prime}\right) \mid B^{\prime} \in W_{\ell}\right\}$ and $\sigma_{\ell}^{\prime}\left(H_{\ell}^{\prime}\right)=\sigma^{\prime}\left(H^{\prime}\right)$.
For $\ell=n$, we have $W_{n}=\emptyset$, and so property $(*)$ implies property $(\diamond)$, as required. Our construction proceeds as follows.
For the base case $\ell=0$, property $(*)$ clearly holds for $\beta_{0}^{\prime \prime}=\beta^{\prime \prime}, \sigma_{0}^{\prime}=\sigma^{\prime}$, and let $W_{0}=W$. For the induction step, assume that ( $*$ ) holds for some $0 \leq \ell<n$, so $\beta_{\ell}^{\prime \prime} \rightarrow H_{\ell}^{\prime}, \sigma_{\ell}^{\prime}$, and $W_{\ell}$ satisfying $(*)$ have been defined. First, assume that there exists $B^{\prime} \in W_{\ell}$ such that $\sigma^{\prime}\left(B^{\prime}\right) \notin \sigma_{\ell}^{\prime}\left(\beta_{\ell}^{\prime}\right)$. Then, property (*) clearly holds for $\beta_{\ell+1}^{\prime \prime}=\beta_{\ell}^{\prime \prime}, H_{\ell+1}^{\prime}=H_{\ell}^{\prime}, \sigma_{\ell+1}^{\prime}=\sigma_{\ell}^{\prime}$, and $W_{\ell+1}=W_{\ell} \backslash\left\{B^{\prime}\right\}$. Otherwise, we consider the following possibilities.
- If rule $\beta_{\ell}^{\prime \prime} \rightarrow H_{\ell}^{\prime}$ is Skolem-free, the rule is of the form $A_{\ell}^{\prime} \wedge \beta_{\ell}^{\prime} \rightarrow H_{\ell}^{\prime}$ where $A_{\ell}^{\prime}$ contains all variables of the rule.
- Otherwise, rule $\beta_{\ell}^{\prime \prime} \rightarrow H_{\ell}^{\prime}$ is of the form $A_{\ell}^{\prime} \wedge \beta_{\ell}^{\prime} \rightarrow H_{\ell}^{\prime}$ where atom $A_{\ell}^{\prime}$ contains a Skolem symbol, in which case this atom contains all variables of the rule.
Either way, there exists $B^{\prime} \in W_{\ell}$ such that $\sigma_{\ell}^{\prime}\left(A_{\ell}^{\prime}\right)=\sigma^{\prime}\left(B^{\prime}\right)$ where $\sigma^{\prime}\left(B^{\prime}\right) \in T_{k}\left(v^{\prime}\right) \backslash S$. Thus, by property $(\diamond)$, these exist a rule $\beta_{\ell} \rightarrow H_{\ell} \in \Sigma^{\prime}$ and a substitution $\sigma_{\ell}$ such that $\beta_{\ell}$ is Skolem-free, $\sigma_{\ell}\left(\beta_{\ell}\right) \subseteq v\left(T_{i}(v)\right)$, and $\sigma_{\ell}\left(H_{\ell}\right)=\sigma^{\prime}\left(B^{\prime}\right)$; the last observation ensures that $H_{\ell}$ contains a Skolem symbol. Moreover, there exists an MGU $\theta_{\ell}$ of $H_{\ell}$ and $A_{\ell}^{\prime}$, so the SkDR inference rule is applicable to $\beta_{\ell} \rightarrow H_{\ell}$ and $A_{\ell}^{\prime} \wedge \beta_{\ell}^{\prime} \rightarrow H_{\ell}^{\prime}$, and $\Sigma^{\prime}$ contains rule $\theta_{\ell}\left(\beta_{\ell}\right) \wedge \theta_{\ell}\left(\beta_{\ell}^{\prime}\right) \rightarrow \theta_{\ell}\left(H_{\ell}\right)$ up to redundancy. Now let $\zeta_{\ell}=\left(\sigma_{\ell} \cup \sigma_{\ell}^{\prime}\right) \circ \theta$ be the composition of $\sigma_{\ell} \cup \sigma_{\ell}^{\prime}$ and $\theta$; note that substitution $\sigma_{\ell} \cup \sigma_{\ell}^{\prime}$ is correctly defined because rules $\beta_{\ell} \rightarrow H_{\ell}$ and $A_{\ell}^{\prime} \wedge \beta_{\ell}^{\prime} \rightarrow H_{\ell}^{\prime}$ do not share variables. Moreover, let $W_{\ell+1}=W_{\ell} \backslash\left\{B^{\prime}\right\}$. We clearly have $\zeta_{\ell}\left(\theta_{\ell}\left(\beta_{\ell}\right)\right) \subseteq S, \zeta_{\ell}\left(\theta_{\ell}\left(\beta_{\ell}^{\prime}\right)\right) \subseteq S \cup\left\{\sigma^{\prime}\left(C^{\prime}\right) \mid C^{\prime} \in W_{\ell+1}\right\}$, and $\zeta_{\ell}\left(\theta_{\ell}(H)\right)=\sigma^{\prime}\left(H^{\prime}\right)$. Since $G \notin T_{k}\left(v^{\prime}\right)$, rule $\theta_{\ell}\left(\beta_{\ell}\right) \wedge \theta_{\ell}\left(\beta_{\ell}^{\prime}\right) \rightarrow \theta_{\ell}\left(H_{\ell}\right)$ is not a syntactic tautology. Thus, by Definition 5.1, there exist a rule $\beta_{\ell+1}^{\prime \prime} \rightarrow H_{\ell+1}^{\prime} \in \Sigma^{\prime}$ and substitution $\mu_{\ell+1}$ such that $\mu_{\ell+1}\left(\beta_{\ell+1}^{\prime \prime}\right) \subseteq \theta_{\ell}\left(\beta_{\ell}\right) \cup \theta_{\ell}\left(\beta_{\ell}^{\prime}\right)$ and $\mu_{\ell+1}\left(H_{\ell+1}^{\prime}\right)=\theta_{\ell}\left(H_{\ell}^{\prime}\right)$. Now let $\sigma_{\ell+1}^{\prime}$ be the substitution defined on each variable $x$ in $\beta_{\ell+1}^{\prime \prime} \rightarrow H_{\ell+1}^{\prime}$ such that $\sigma_{\ell+1}^{\prime}(x)=\zeta_{\ell}\left(\mu_{\ell+1}(x)\right)$. Then, property (*) clearly holds for $\beta_{\ell+1}^{\prime \prime} \rightarrow H_{\ell+1}^{\prime}, \sigma_{\ell+1}^{\prime}$, and $W_{\ell+1}$, as required.

Proof of Complexity. Fix $\Sigma, r, w_{b}, e, c$, and $a$ as stated in the theorem. Skolemizing a GTGD $\forall \vec{x}[\beta \rightarrow \exists \vec{y} \eta$ ] produces guarded rules in which each atom is of the form $R\left(t_{1}, \ldots, t_{n}\right)$ such that each $t_{i}$ is a constant, a variable from $\vec{x}$, or a term of the form $f(\vec{x})$ where $f$ is a Skolem symbol. Moreover, each atom obtained from $R\left(t_{1}, \ldots, t_{n}\right)$ by the SkDR inference rule is obtained by replacing a variable in $\vec{x}$ with another variable or a constant. Thus, atom $R\left(t_{1}, \ldots, t_{n}\right)$ cannot contain more than $|\vec{x}|$ variables. Since the number of different symbols obtained by Skolemization is clearly bounded by $e$, the number of different atoms of such form is bounded by $\ell=r \cdot\left(w_{b}+e+c\right)^{a}$. The body of each guarded rule corresponds to a subset of these atoms, so the number of different rules up to variable remaining is bounded by $2^{\ell} \cdot \ell \leq 2^{\ell} \cdot 2^{\ell}=2^{2 \ell}=\wp$. By Definition 5.3, the result of applying the SkDR inference rule is retained in set $\Sigma^{\prime}$ only if the set does not contain a variable renaming of the result. Thus, the SkDR inference rule needs to be applied to at most $\wp^{2}=2^{4 \ell}$ pairs of rules. For each pair, one might need to unify at most $\ell$ body atoms of one rule with the head atom of the other rule, so the unifier $\theta$ may need to be computed at most $\wp^{2} \cdot \ell \leq \wp^{2} \cdot 2^{\ell}=2^{5 \ell}=32^{\ell}$ times. We can check subsumption between a pair of rules analogously to Theorem 5.8: for each of at most $2^{\ell^{2}}$ ways to match the body atoms of one rule to the body atoms of another rule, we try to find a substitution $\mu$ satisfying Definition 5.1. Finally, unification of atoms requires time that is linear in $a$, and all other steps require linear time too.

## D PROOFS FOR HypDR

## D. 1 Proof of Proposition 5.18: Properties of HypDR

Proposition 5.18. Each application of the HypDR inference rule to rules $\tau_{1}, \ldots, \tau_{n}$ and $\tau^{\prime}$ as in Definition 5.16 produces a guarded rule.
Proof. Consider arbitrary rules $\tau_{i}=\beta_{i} \rightarrow H_{i}$ with $1 \leq i \leq n$ such that $\beta_{i}$ is Skolem-free and $H_{i}$ contains a Skolem symbol, a Skolem-free rule $\tau^{\prime}=A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$, and an MGU $\theta$ of $H_{1}, \ldots, H_{n}$ and $A_{1}^{\prime}, \cdots A_{n}^{\prime}$ satisfying the preconditions of the HypDR inference rule. Rule $\tau_{1}$ contains a term with a Skolem symbol in the head, and this term is unified with a variable, say $x$, occurring in a Skolem-free body atom $A_{1}^{\prime}$ of rule $\tau^{\prime}$. Moreover, rule $\tau^{\prime}$ is guarded, so the body of the rule contains a Skolem-free atom $G$ that contains all variables of the rule; thus, $G$ also contains $x$. Since $\theta(x)$ contains a Skolem symbol, $\theta(G)$ contains a Skolem symbol too. However, $\theta\left(\beta^{\prime}\right)$ is Skolem-free, so $G$ must be one of the atoms $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$ from the body of rule $\tau^{\prime}$ that are participating in the HypDR inference rule. But then, we can show that the result of the inference is guarded analogously to the proof of Proposition 5.12.

## D. 2 Proof of Theorem 5.19: Correctness and Complexity of HypDR

Theorem 5.19. Program $\operatorname{HypDR}(\Sigma)$ is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in time
 $w_{b}=\operatorname{bwidth}(\Sigma), c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Proof of Correctness. The correctness proof for HypDR is almost identical to the correctness proof in Theorem 5.13, so we outline just the difference. In particular, we wish to prove properties $(\diamond)$ and $(\diamond)$ exactly as stated in Theorem 5.13 using the same proof structure. In the proof of property $(\diamond)$, we establish existence of a Skolem-free rule $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\sigma^{\prime}$ such that $\sigma^{\prime}\left(\beta^{\prime \prime}\right) \subseteq v\left(T_{k}\left(v^{\prime}\right)\right)$ and $\sigma^{\prime}\left(H^{\prime}\right)=v(G)$ in exactly the same way. The difference to the proof of Theorem 5.13 is that we "resolve away" all relevant body atoms of $\beta^{\prime \prime}$ in one step. To this end, let $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$ be precisely the atoms of $\beta^{\prime \prime}$ such that $\sigma^{\prime}\left(A_{i}^{\prime}\right) \notin S$ for each $1 \leq i \leq n$. Thus, we can assume that the rule is of the form $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$, and $\sigma^{\prime}\left(\beta^{\prime}\right) \subseteq S$ clearly holds. By property $(\diamond)$, for each $1 \leq \ell \leq n$, there exist a rule $\beta_{\ell} \rightarrow H_{\ell} \in \Sigma^{\prime}$ and substitution $\sigma_{\ell}$ such that $\beta_{\ell}$ is Skolem-free and $\sigma_{\ell}\left(H_{\ell}\right)=\sigma^{\prime}\left(A_{\ell}^{\prime}\right)$; the last observation ensures that $H_{\ell}$ contains a Skolem symbol. Finally, there exists an MGU $\theta$ of $H_{1}, \ldots, H_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$. Since $\sigma^{\prime}\left(\beta^{\prime}\right) \subseteq S$, conjunction $\theta\left(\beta^{\prime}\right)$ is Skolem-free. Thus, the HypDR inference rule is applicable to $\beta_{1} \rightarrow H_{1}, \ldots, \beta_{n} \rightarrow H_{n}$ and $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$, so set $\Sigma^{\prime}$ contains rule $\theta\left(\beta_{1}\right) \wedge \cdots \wedge \theta\left(\beta_{n}\right) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ up to redundancy. Since no premises share variables, substitution $\sigma_{1} \cup \cdots \cup \sigma_{n} \cup \sigma^{\prime}$ is correctly defined, so let $\zeta$ be the composition of $\sigma_{1} \cup \cdots \cup \sigma_{n} \cup \sigma^{\prime}$ and $\theta$. Clearly, we have $\zeta\left(\theta\left(\beta_{1}\right)\right) \cup \cdots \cup \zeta\left(\theta\left(\beta_{n}\right)\right) \cup \zeta\left(\theta\left(\beta^{\prime}\right)\right) \subseteq S$ and $\zeta\left(\theta\left(H^{\prime}\right)\right)=\sigma^{\prime}\left(H^{\prime}\right)=v(G)$. Since $G \notin T_{k}\left(v^{\prime}\right)$, rule $\theta\left(\beta_{1}\right) \wedge \cdots \wedge \theta\left(\beta_{n}\right) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right)$ is not a syntactic tautology so, by Definition 5.1, there exist a rule $\beta \rightarrow H \in \Sigma^{\prime}$ and substitution $\mu$ such that $\mu(\beta) \subseteq \theta\left(\beta_{1}\right) \cup \cdots \cup \theta\left(\beta_{n}\right) \cup \theta\left(\beta^{\prime}\right)$ and $\mu(H)=\theta\left(H^{\prime}\right)$. Let $\sigma$ be the substitution defined on each variable $x$ in $\beta \rightarrow H$ such that $\sigma(x)=\zeta(\mu(x))$. Then, $\sigma(\beta) \subseteq S$ and $\sigma(H)=\sigma^{\prime}\left(H^{\prime}\right)=v(G)$, as required for property $(\diamond)$.

Proof of Complexity. Fix $\Sigma, r, w_{b}, e, c$, and $a$ as stated in the theorem. In the same way as in the complexity proof of Theorem 5.13, the number of different atoms can be bounded by $\ell=r \cdot\left(w_{b}+e+c\right)^{a}$, and the number of different rules can be bounded by $\wp=2^{2 \ell}$. Now we can apply the HypDR inference rule as follows: we choose one of the $\wp$ rules that plays the role of $\tau^{\prime}$ and then, for each of the at most $\ell$ body atoms in $\tau^{\prime}$, we select one of the $\wp$ rules that play the role of rules $\tau_{i}$. Hence, there are at most $\wp \cdot \wp^{\ell}=\wp^{\ell+1}$ different applications of the HypDR inference rule. Thus, we may need to compute the unifier $\theta$ at most $\left(2^{2 \ell}\right)^{\ell+1}=2^{2 \ell^{2}+2 \ell} \leq 2^{3 \ell^{2}}$ times. Finally, the times needed for subsumption checking, unification, and all other steps can be bounded analogously as in the complexity proof of Theorem 5.13.

## E THE FullDR ALGORITHM: CREATING DATALOG RULES DIRECTLY

The algorithms presented in the body of the paper all create the Datalog rules needed for the final rewriting as well as intermediate non-full TGDs or rules that are discarded after all inferences are performed. We now present an algorithm that produces only Datalog rules. Similar algorithms have appeared in the prior literature [4]. After presenting such an algorithm, we explain the shortcomings of this approach.

Definition E.1. The Full Datalog Rewriting inference rule FullDR can be applied in two ways, depending on the types of TGDs it takes.

- The (COMPOSE) variant of the FullDR inference rule takes full TGDs

$$
\tau=\forall \vec{x}[\beta \rightarrow A] \quad \text { and } \quad \tau^{\prime}=\forall \vec{z}\left[A^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}\right]
$$

and a substitution $\theta$ such that

- $\theta(A)=\theta\left(A^{\prime}\right)$,
- $\operatorname{dom}(\theta)=\vec{x} \cup \vec{z}$, and
- $\operatorname{rng}(\theta) \subseteq \vec{w} \cup$ consts $(\tau) \cup \operatorname{consts}\left(\tau^{\prime}\right)$ where $\vec{w}$ is a vector of hwidth $(\Sigma)+\mid$ consts $(\Sigma) \mid$ variables different from $\vec{x} \cup \vec{y} \cup \vec{z}$,
and it derives

$$
\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right) .
$$

- The (PROPAGATE) variant of the FullDR inference rule takes TGDs

$$
\tau=\forall \vec{x}\left[\beta \rightarrow \exists \vec{y} \eta \wedge A_{1} \wedge \cdots \wedge A_{n}\right] \quad \text { and } \quad \tau^{\prime}=\forall \vec{z}\left[A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}\right]
$$

and a substitution $\theta$ such that

- $\theta\left(A_{i}\right)=\theta\left(A_{i}^{\prime}\right)$ for each $i$ with $1 \leq i \leq n$,
$-\operatorname{dom}(\theta)=\vec{x} \cup \vec{z}$,
- $\operatorname{rng}(\theta) \subseteq \vec{w} \cup \vec{y} \cup \operatorname{consts}(\tau) \cup$ consts $\left(\tau^{\prime}\right)$ where $\vec{w}$ is a vector of hwidth $(\Sigma)+\mid$ consts $(\Sigma) \mid$ variables different from $\vec{x} \cup \vec{y} \cup \vec{z}$,
- $\theta(\vec{x}) \cap \vec{y}=\emptyset$, and
$-\operatorname{vars}\left(\theta\left(\beta^{\prime}\right)\right) \cap \vec{y}=\emptyset$ and $\operatorname{vars}\left(\theta\left(H^{\prime}\right)\right) \cap \vec{y}=\emptyset$,
and it derives

$$
\theta(\beta) \wedge \theta\left(\beta^{\prime}\right) \rightarrow \theta\left(H^{\prime}\right) .
$$

Theorem E.2. Program Fulldr( $\Sigma$ ) is a Datalog rewriting of a finite set of GTGDs $\Sigma$. Moreover, the rewriting can be computed in time $O\left(b^{r^{d} \cdot(w+c)^{d a}}\right)$ for $r$ the number of relations in $\Sigma$, a the maximum relation arity in $\Sigma, w=\operatorname{width}(\Sigma), c=|\operatorname{consts}(\Sigma)|$, and some $b$ and $d$.

Proof of Correctness. The proof follows the same structure as the correctness proof of Theorem 5.8: we show that property ( $\uparrow$ ) holds for each loop on a one-pass tree-like chase sequence for $I$ and $\Sigma$; a minor difference is that the TGD whose existence is implied by $(*)$ is not necessarily guarded, but has width bounded by hwidth $(\sigma)$. To this end, we consider an arbitrary loop $T_{i}, T_{i+1}, \ldots, T_{j-1}, T_{j}$ at vertex $v$ in the sequence, and assume that the claim holds for all shorter loops in the sequence. By the definition of the loop, chase tree $T_{i+1}$ is obtained from $T_{i}$ by applying a chase step to some non-full TGD $\forall \vec{x}\left[\beta_{0} \rightarrow \exists \vec{y} \eta_{0}\right] \in \Sigma$. Let $\sigma_{0}$ and $\sigma_{0}^{\prime}$ be the substitutions used in this chase step, and let $v^{\prime}$ be the child of $v$ introduced in $T_{i+1}$. Note that $T_{i+1}\left(v^{\prime}\right)$ contains at most hwidth $(\Sigma)+|\operatorname{consts}(\Sigma)|$ distinct terms. We show by another induction on $k$ that the following property holds for each $k$ with $i<k \leq j-1$ :
$(\diamond)$ for each fact $G \in T_{k}\left(v^{\prime}\right) \backslash T_{i+1}\left(v^{\prime}\right)$, there exist a full TGD $\forall \vec{x}[\beta \rightarrow H] \in \Sigma^{\prime}$ of width at most width $(\Sigma)$ and a substitution $\sigma$ such that $\sigma(\beta) \subseteq T_{i+1}\left(v^{\prime}\right)$ and $\sigma(H)=G$.
For the base case $k=i+1$, property $(\diamond)$ holds vacuously because $T_{k}\left(v^{\prime}\right) \backslash T_{i+1}\left(v^{\prime}\right)=\emptyset$. For the inductive step, assume that ( $\diamond$ ) holds for some $k$ and consider the possible ways to obtain $T_{k+1}$ from $T_{k}$. Property $(\diamond)$ holds by the inductive hypothesis if $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right)$-that is, if the step involves a descendant of $v^{\prime}$. Otherwise, $T_{k+1}\left(v^{\prime}\right)=T_{k}\left(v^{\prime}\right) \cup\{G\}$ where fact $G$ is obtained in one of the following two ways.

- A full TGD in $\Sigma$ derives $G$ from $T_{k}\left(v^{\prime}\right)$. Set $\Sigma^{\prime}$ contains this TGD up to redundancy, so by Definition 5.1 there exist a full TGD $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma\left(\beta^{\prime \prime}\right) \subseteq T_{k}\left(v^{\prime}\right)$ and $\sigma\left(H^{\prime}\right)=G$.
- Fact $G$ is the output of a loop at vertex $v^{\prime}$. But then, this loop is shorter than $T_{i}, \ldots, T_{j}$ so, by property $(\boldsymbol{*})$, there exists a full TGD $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and a substitution $\sigma$ such that $\sigma\left(\beta^{\prime \prime}\right) \subseteq T_{k}\left(v^{\prime}\right)$ and $\sigma\left(H^{\prime}\right)=G$.
Either way, the width of rule $\beta^{\prime \prime} \rightarrow H^{\prime}$ is bounded by width $(\Sigma)$, and we can assume that $\beta^{\prime \prime} \rightarrow H^{\prime}$ is of the form $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ where $\sigma\left(A_{\ell}^{\prime}\right) \in T_{k}\left(v^{\prime}\right) \backslash T_{i+1}\left(v^{\prime}\right)$ for each $1 \leq \ell \leq n$, and $\sigma\left(\beta^{\prime}\right) \subseteq T_{i+1}\left(v^{\prime}\right)$. By property ( $\diamond$ ), for each $1 \leq \ell \leq n$ there exist a full TGD $\beta_{\ell} \rightarrow H_{\ell} \in \Sigma^{\prime}$ and a substitution $\sigma_{\ell}$ such that $\sigma_{\ell}\left(\beta_{\ell}\right) \subseteq T_{i+1}\left(v^{\prime}\right)$ and $\sigma_{\ell}\left(H_{\ell}\right)=\gamma\left(A_{\ell}^{\prime}\right)$. Moreover, set rng $\left(\sigma_{\ell}\right)$ clearly contains at most hwidth $(\Sigma)+|\operatorname{consts}(\Sigma)|$ distinct terms. But then, there exist substitutions $\theta_{1}, \ldots, \theta_{n}$ that allow us to iteratively compose each $\beta_{\ell} \rightarrow H_{\ell}$ with $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ to obtain a full TGD subsumed by some $\tau \in \Sigma^{\prime}$ and substitution $\sigma_{1}$ such that $\tau$ and $\sigma_{1}$ satisfy property ( () .

To complete the proof, consider an arbitrary fact $F \in T_{j-1}\left(v^{\prime}\right) \backslash T_{i+1}\left(v^{\prime}\right)$ that is propagated from $v^{\prime}$ to $v$ in $T_{j}$, and let $\beta^{\prime \prime} \rightarrow H^{\prime} \in \Sigma^{\prime}$ and $\sigma$ be the TGD and substitution whose existence is guaranteed by property $(\diamond)$. Now if $\sigma\left(\beta^{\prime \prime}\right) \subseteq T_{i+1}\left(v^{\prime}\right) \backslash \sigma_{0}^{\prime}\left(\eta_{0}\right)$, then TGD $\beta^{\prime \prime} \rightarrow H^{\prime}$ satisfies property $(\uparrow)$. Otherwise, we can assume that the rule is of the form $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ where $\sigma\left(A_{i}^{\prime}\right) \in \sigma_{0}^{\prime}\left(\eta_{0}\right)$ for each $1 \leq i \leq n$, and $\sigma\left(\beta^{\prime}\right) \subseteq T_{i+1}\left(v^{\prime}\right) \backslash \sigma_{0}^{\prime}\left(\eta_{0}\right)$. Moreover, $\operatorname{rng}(\sigma)$ clearly contains at most hwidth $(\Sigma)+|\operatorname{consts}(\Sigma)|$ distinct terms. But then, there exists a substitution $\theta$ that allows us to apply the (PROPAGATE) variant of the FullDR inference rule to $\beta_{0} \rightarrow \exists \vec{y} \eta_{0}$ and $A_{1}^{\prime} \wedge \cdots \wedge A_{n}^{\prime} \wedge \beta^{\prime} \rightarrow H^{\prime}$ to obtain a full TGD subsumed by some TGD $\tau \in \Sigma^{\prime}$ and substitution $\sigma_{1}$ such that $\tau$ and $\sigma_{1}$ satisfy property ( $\downarrow$ ).

Proof of Complexity. The proof is analogous to the proof of complexity of Theorem 5.8. In particular, the (PROPAGATE) variant of the FulIDR inference rule is analogous to the ExbDR inference rule, so we can bound in the same way the number of candidate rule pairs and possible ways to match body atoms of $\tau^{\prime}$ to head atoms of $\tau$ by $32^{\ell^{2}}$, where $\ell=r \cdot(w+c)^{a}$. Once a candidate pair of $\tau$ and $\tau^{\prime}$ has been selected, we need to consider all possible substitutions $\theta$. Each such $\theta$ is defined on at most $2 w$ variables $\vec{x} \cup \vec{z}$. Moreover, each variable is mapped to one of the $w+c$ variables or to one of the $c$ constants in consts $(\Sigma)$. Hence, there are at most $(w+2 c)^{2 w} \leq 4^{(w+2 c) \cdot w} \leq 4^{(w+c)^{2}}$ different substitutions $\theta$. Consequently, the (PROPAGATE) variant of the FulIDR inference rule can be applied at most $32^{\ell^{2}} \cdot 4^{(w+c)^{2}}<32^{n \cdot(w+c)^{2 a+1}}$ times. Applications of the (COMPOSE) variant can be bounded analogously. Finally, the times needed for subsumption checking, unification, and all other steps can be bounded analogously as in the complexity proof of Theorem 5.8, with a minor difference that only body atoms need to be matched in the subsumption checks.

The FullDR algorithm has several obvious weak points. First, it considers all possible ways to compose Datalog rules as long as this produces a rule with at most hwidth $(\Sigma)+|\operatorname{consts}(\Sigma)|$ variables. This may seem unnecessary, but the (COMPOSE) variant of the FullDR inference rule cannot be simply dropped while retaining completeness. To understand why, consider an arbitrary loop $T_{i}, \ldots, T_{j}$ at vertex $v$ with output fact $F$ in a one-pass chase proof: the (PROPAGATE) variant of the FullDR inference reflects only the chase step that derives the loop's output $F$, so the (COMPOSE) variant is needed to reflects the chase steps that produce facts derived in the child of $v$ that are not propagated back to $v$. Second, it is not clear how one efficiently obtains the atoms $A_{1}, \ldots, A_{n}$ and $A_{1}^{\prime}, \ldots, A_{n}^{\prime}$ participating in the (PROPAGATE) variant. Third, the number of substitutions $\theta$ in the (COMPOSE) and (PROPAGATE) variants of the FullDR inference rule can be very large. Example E. 3 illustrates this problem for the (COMPOSE) variant, but one can show analogously that the (PROPAGATE) variant suffers from the same issues.

Example E.3. Consider the steps of the FullDR algorithm on GTGDs (46)-(48).

$$
\begin{align*}
R\left(x_{1}, x_{2}\right) & \rightarrow \exists y_{1}, y_{2} S\left(x_{1}, x_{2}, y_{1}, y_{2}\right) \wedge T\left(x_{1}, x_{2}, y_{2}\right)  \tag{46}\\
S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) & \rightarrow U\left(x_{4}\right)  \tag{47}\\
T\left(z_{1}, z_{2}, z_{3}\right) \wedge U\left(z_{3}\right) & \rightarrow P\left(z_{1}\right) \tag{48}
\end{align*}
$$

The (COMPOSE) variant of the FullDR inference rule should be applied to GTGDs (47) and (48), but it is not clear which unifier $\theta$, identifying variables $z_{i}$ in the latter with variables $x_{i}$ in the former, one should use. The standard resolution inference rule from first-order theorem proving would consider only the MGU $\theta$ that maps $z_{3}$ to $x_{4}$; however, this would produce the resolvent $S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) \wedge T\left(z_{1}, z_{2}, x_{4}\right) \rightarrow P\left(z_{1}\right)$ containing more than hwidth $(\Sigma)=4$ variables, so this rule is not derived by the (COMPOSE) variant. Eliminating the upper bound on the number of variables is not a solution: doing so would allow the derivation of full TGDs with an unbounded number of variables, which would prevent termination. Instead, the (COMPOSE) variant requires us to consider every possible substitution $\theta$ that maps variables $x_{1}, \ldots, x_{4}, z_{1}, \ldots, z_{3}$ to at most hwidth $(\Sigma)$ variables. Consequently, $7^{4}=2401$ substitutions deriving rules such as

$$
\begin{align*}
& S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) \wedge T\left(x_{1}, x_{2}, x_{4}\right) \rightarrow P\left(x_{1}\right),  \tag{49}\\
& S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) \wedge T\left(x_{2}, x_{1}, x_{4}\right) \rightarrow P\left(x_{2}\right),  \tag{50}\\
& S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) \wedge T\left(x_{1}, x_{3}, x_{4}\right) \rightarrow P\left(x_{1}\right),  \tag{51}\\
& S\left(x_{1}, x_{2}, x_{3}, x_{4}\right) \wedge T\left(x_{3}, x_{1}, x_{4}\right) \rightarrow P\left(x_{3}\right), \tag{52}
\end{align*}
$$

need to be considered, which is clearly infeasible in practice.
Nevertheless, we implemented FullDR using the subsumption and indexing techniques described in Section 6. Unsurprisingly, we did not find FullDR competitive in our experiments. In fact, FullDR timed out on 173 ontologies, and there are only three ontologies where another algorithm reached the timeout but FullDR did not. For this reason, we do not discuss the results with FullDR in Section 7.


[^0]:    This work is licensed under the Creative Commons BY-NC-ND 4.0 International
    License. Visit https://creativecommons.org/licenses/by-nc-nd/4.0/ to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
    Proceedings of the VLDB Endowment, Vol. 15, No. 11 ISSN 2150-8097.
    doi:XX.XX/XXX.XX

