

**Dieses Dokument ist eine Zweitveröffentlichung (Postprint) /
This is a self-archiving document (accepted version):**

Muhammad Idris, Martín Ugarte, Stijn Vansummeren, Hannes Voigt, Wolfgang Lehner

General dynamic Yannakakis: conjunctive queries with theta joins under updates

Erstveröffentlichung in / First published in:





The VLDB Journal. 2020. 29, S. 619–653. Berlin: Springer. ISSN 0949-877X.

DOI: <https://doi.org/10.1007/s00778-019-00590-9>

Diese Version ist verfügbar / This version is available on:

<https://nbn-resolving.org/urn:nbn:de:bsz:14-qucosa2-864743>

General dynamic Yannakakis: conjunctive queries with theta joins under updates

Muhammad Idris^{1,2}  · Martín Ugarte³  · Stijn Vansummeren¹  · Hannes Voigt⁴  · Wolfgang Lehner² 

Received: 30 November 2018 / Revised: 30 August 2019 / Accepted: 23 October 2019 / Published online: 19 November 2019

Abstract

The ability to efficiently analyze changing data is a key requirement of many real-time analytics applications. In prior work, we have proposed general dynamic Yannakakis (GDYN), a general framework for dynamically processing acyclic conjunctive queries with θ -joins in the presence of data updates. Whereas traditional approaches face a trade-off between materialization of subresults (to avoid inefficient recomputation) and recomputation of subresults (to avoid the potentially large space overhead of materialization), GDYN is able to avoid this trade-off. It intelligently maintains a succinct data structure that supports efficient maintenance under updates and from which the full query result can quickly be enumerated. In this paper, we consolidate and extend the development of GDYN. First, we give full formal proof of GDYN's correctness and complexity. Second, we present a novel algorithm for computing GDYN query plans. Finally, we instantiate GDYN to the case where all θ -joins are inequalities and present extended experimental comparison against state-of-the-art engines. Our approach performs consistently better than the competitor systems with multiple orders of magnitude improvements in both time and memory consumption.

Keywords Incremental view maintenance · Dynamic query processing · Complex event processing · Theta joins · Inequalities · Acyclic joins

Communicated by Dr. Sihem Amer-Yahia.

M. Ugarte This work was done while the author was affiliated to ULB, Belgium.

H. Voigt This work was done while the author was affiliated to TU Dresden, Germany.

✉ Martín Ugarte
martinugarte@uc.cl

Muhammad Idris
midris@ulb.ac.be

Stijn Vansummeren
svsummer@ulb.ac.be

Hannes Voigt
hannes.voigt@neo4j.com

Wolfgang Lehner
wolfgang.lehner@tu-dresden.de
<https://www.db.inf.tu-dresden.de/Lehner>

¹ Université Libre de Bruxelles, Bruxelles, Belgium

² TU Dresden, Dresden, Germany

³ Pontifical Catholic University of Chile, Santiago, Chile

⁴ neo4j, München, Germany

1 Introduction

The ability to efficiently analyze changing data is a key requirement in stream processing [39], complex event recognition [16], business intelligence [35], and machine learning [44]. Generally, the analysis that needs to be kept up-to-date, or at least their basic elements, are specified in a query language. The main task is then to efficiently update the query results under data updates.

In this context, we tackle the problem of *dynamic query evaluation*, where a given query Q has to be evaluated against a database that is constantly changing. Concretely, when database db is updated to database $db + u$ under update u , the objective is to efficiently compute $Q(db + u)$, taking into consideration that $Q(db)$ was already evaluated and recomputations could be avoided. Dynamic query evaluation is of utmost importance if response time requirements for queries under concurrent data updates have to be met or if data volumes are so large that full re-evaluation is prohibitive.

In this paper, we focus on the problem of dynamic evaluation for conjunctive queries that feature multi-way θ -joins. The following example illustrates our setting. Assume that we wish to detect potential credit card fraud. Credit card

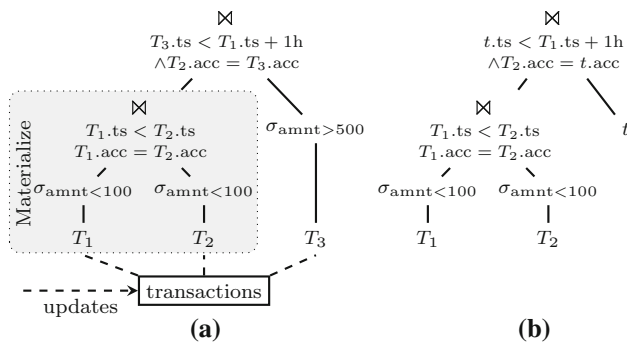


Fig. 1 **a** Example query for detecting fraudulent credit card activity. **b** Delta query of (a) to be executed upon insertion of new high-amount transaction tuple t

transactions specify their timestamp (ts), account number (acc), and amount ($amnt$). A typical fraud pattern is that, in a short period of time, a criminal tests a stolen credit card with a few small purchases to then make larger purchases (cf. [37]). Assuming that the short period of time is 1 h, this pattern could be detected by dynamically evaluating the query in Fig. 1a. Queries like this may exhibit arbitrary local predicates and multi-way joins with equality as well as inequality predicates.

Dynamic query evaluation has a rich history in data management and has been researched in the context of incremental view maintenance (IVM) [22,29,30], stream join processing [4,20,34,43], and complex event recognition (CER, also known as complex event processing) [11,15,31,46,49]. All of the existing techniques are based on recomputation of query (sub)results and/or on their materialization. We next illustrate the issues with recomputation and result materialization; a detailed literature review is given in Sect. 2.

The most extreme form of recomputation is of course full recomputation: simply evaluate Q on $db + u$ from scratch whenever an update u arrives on db . Clearly, this incurs the highest possible update processing cost, since it completely ignores the fact that $Q(db)$ was already evaluated and certain shared computations were performed. This can be solved by introducing the simplest form of materialization: simply store the current result $Q(db)$ and, whenever an update u arrives, evaluate the *delta query* ΔQ associated with Q [22]. ΔQ takes as input db and u and computes the update that needs to be applied to the materialized $Q(db)$ to obtain $Q(db + u)$. While this exploits certain shared computation, it unfortunately does so only to some extent.

Let us illustrate this by means of our example fraud query in Fig. 1a. Assume that u inserts a single new high-amount transaction t . Then, $\Delta Q(db, u)$ amounts to computing the join shown in Fig. 1b. While we can expect that this is more efficient than full recomputation of Q (since the join with T_3 is replaced by the join with a single tuple), observe that if we now get another update u' that inserts another high-amount

transaction t' then we are essentially *fully* recomputing the join of Fig. 1b, but now with t' instead of t . Hence, by fully recomputing ΔQ whenever a new update u' arrives, we are ignoring the fact that certain computation—most notably, the join between T_1 and T_2 common to t and t' —is shared and in principle need not be redone.

This can be solved by introducing more materialization: in addition, to storing $Q(db)$ also materialize the result of the join between T_1 and T_2 indicated in the shaded area in Fig. 1a and use additional delta queries to maintain this result as well as $Q(db)$. In the IVM literature, this approach is known as higher-order IVM (HIVM). This method is highly effective in practice and formally lowers the update processing complexity [29].

While more materialization can imply less recomputation, it has a serious drawback in terms of memory overhead: materialization of $Q(db)$ requires $\Omega(|Q(db)|)$ space, where $|db|$ denotes the size of db . Therefore, when $Q(db)$ is large, which is often the case in CER as well as in data preparation scenarios for training statistical models, materializing $Q(db)$ quickly becomes impractical, especially for main memory-based systems. Note that $|Q(db)|$ can be polynomial in $|db|$. HIVM is even more affected by this problem since it not only materializes the result of Q but also the results of partial joins, which can be larger than both db and $Q(db)$. For example, the shaded area of Fig. 1a builds the table of all pairs of *small* transactions that could be part of a credit card fraud. If we assume that there are N small transactions, all of the same account, this materialization will take $\Theta(N^2)$ space. This naturally becomes impractical when N grows.

In summary, in traditional techniques for dynamic query evaluation, there is a trade-off between recomputation and materialization: more materialization can mean less recomputation and hence faster update processing, but more memory consumption. In previous works [25,26], we have shown that this trade-off can be avoided by taking a different approach to dynamic query evaluation: instead of materializing $Q(db)$, we can build a succinct data structure that (1) supports updates efficiently and (2) *represents* $Q(db)$ in the sense that from it we can *generate* $Q(db)$ as efficiently as if it was materialized. In particular, the representation is equipped with index structures so that we can enumerate $Q(db)$ *with bounded delay* [38]: one tuple at a time, while spending only a small amount of work to produce each new tuple. This makes the enumeration competitive with enumeration from materialized query results.

In essence, we hence separate dynamic query processing into two stages: (1) an update stage where we only maintain under updates the (small) information that is necessary for result enumeration and (2) an enumeration stage where the query result is efficiently enumerated.

The main insight of [25,26] is that a practical family of algorithms for dynamic query evaluation based on this idea

can naturally be obtained by modifying Yannakakis' seminal algorithm for processing acyclic joins in the static setting [47]. In particular, instead of materializing $Q(db)$ and its subjoins (which can both be large), it suffices to materialize *semijoins* (which are of size linear in db) and build indexes on these semijoins as well as the base relations. The most general form of this modification, called *general dynamic Yannakakis* (GDYN), supports the class of acyclic generalized conjunctive queries (GCQs), which are acyclic conjunctive queries with θ -joins, that are evaluated under multiset semantics and support certain forms of aggregation. The representation of query results that underlies GDYN has several desirable properties:

- It allows to enumerate $Q(db)$ with bounded delay.
- It requires only $\mathcal{O}(|db|)$ space and is hence independent of the size of $Q(db)$.
- It features efficient maintenance under updates. When Q is a conjunctive query (with equijoins only), then we can update the representation of $Q(db)$ to a representation of $Q(db + u)$ in time $\mathcal{O}(|db| + |u|)$. In contrast, techniques based on (H)IVM may require $\Omega(|u| + |Q(db + u)|)$ time in the worst case. For the subclass of q -hierarchical queries [8], our update time is $\mathcal{O}(|u|)$. When Q consists of both equality and inequality joins ($<$, \leq) the update time increases. If Q has at most one inequality per pair of relations, the update time is $\mathcal{O}(M \log M)$, where $M = |db| + |u|$; otherwise, it is $\mathcal{O}(M^2)$. In contrast, existing techniques may require $\Omega(|db|^{k-1})$ time in the worst case, where k is the number of relations to be joined.

In this paper, we consolidate and expand our development of GDYN. Our contributions are as follows.

- (1) We give an intuitive, concise, and stand-alone description of GDYN and explain the main components behind its efficiency. In addition, we provide, for the first time, full formal proof of GDYN's correctness and complexity (Sect. 4).
- (2) Like most query evaluation algorithms, GDYN's operation is driven by the availability of a query plan. In our previous work, we have always assumed query plans to be explicitly given. In this paper, in contrast, we present, for the first time, an algorithm for efficiently computing GDYN query plans (Sect. 5).
- (3) Finally, we present our implementation of GDYN (Sect. 6) and extended experimental comparison of GDYN to state-of-the-art IVM and CER engines. We explore the full design space of queries with up to three joins. GDYN performs consistently better than the competitor systems with multiple orders of magnitude improvements in both time and memory consumption (Sect. 7).

We discuss related work in Sect. 2, introduce the required background in Sect. 3, and conclude in Sect. 8. Because of space restrictions, certain supporting material and proofs of auxiliary statements are given in the Appendix.

2 Related work

IVM The trade-off between materialization and recomputation explained in the Introduction is at the core of IVM [12,21,22,29,30]. IVM hence differs from GDYN as already explained in the Introduction.

CER There are two approaches to CER: relational and automaton based. Relational approaches (e.g., [31]) are similar to IVM. In contrast to the relational approaches, automaton approaches assume that event tuples are endowed with a timestamp and that the arrival order of event tuples corresponds to the timestamp order (i.e., there are no out-of-order events). They build an automaton to recognize the desired temporal patterns in the input stream. Broadly speaking, there are two automata-based recognition approaches. In the first approach, followed by [3,46], events are cached per state and once a final state is reached, a search through the cached events is done to recognize the complex events. While it is no longer necessary to check the temporal constraints (e.g., $T_1.ts < T_2.ts$) during the search, the additional constraints (in our fraud query example $T_1.acc = T_2.acc = T_3.acc$) must still be verified. At essence, this corresponds to fully recomputing a delta query since each event triggering a transition to a final state may cause re-evaluation of a subjoin on the cached data. In the second approach, followed by [11,14,15,49], partial runs are materialized according to the automaton's topology. For our example query, this means that, just like HIVM, the join in the shaded area of Fig. 1a is materialized and maintained, so it is available when a large amount transaction arrives. This approach hence shares with HIVM its high memory overhead.

Stream Joins The goal of stream join processing is to produce join results incrementally as new tuples are added to the input or old tuples retracted. To see how GDYN differs from stream joins, we discern two classes of stream join algorithms.

Algorithms in the first class are designed to produce output tuples as soon as possible, without blocking for more input to become available. They are based on the symmetric hash join [45] and its variants [41,43]. Algorithms in this class favor full recomputation of delta queries and hence ignore the opportunity to reduce redundant recomputation by additional materialization. Moreover, these algorithms are limited to processing equijoins only.

Algorithms in the second class focus on window-based stream joins [20,27,34,40]. Like the automata-based approaches in CER, they assume that each tuple is endowed with

a timestamp attribute and are restricted to the setting where tuples follow a FIFO paradigm: new tuples arrive in increasing timestamp order and tuples with the oldest timestamp are deleted first. This property is crucial for the algorithm's proposed optimizations to work. GDYN, in contrast, makes no FIFO assumption and can deal with arbitrary updates, including out-of-order updates. Furthermore, we note that window-based joins are a strict subclass of the class of all inequality joins, since in a window-based join, only a single temporal attribute will be compared across all relations. As such, queries like $R_1 \bowtie_{R_1.amnt < R_2.amnt} R_2 \bowtie_{R_2.ts < R_3.ts} R_3$ that inequality join across multiple unrelated attributes are not considered by [20,27,34,40]. GDYN, in contrast, processes such queries intelligently. Finally, we note that, in contrast to GDYN, support for multi-way stream joins is limited: [27,34,40] consider only binary joins, while [20] treats multi-way joins, but makes the simplifying assumption that all comparisons are on the same, single attribute.

Because of the increasing widespread use of distributed compute engines such as Flink, Spark, and Storm in contemporary data analysis scenarios, there has also been much research on how to support stream joins in such engines (e.g., [44]). To the best of our knowledge, this work builds upon the above-mentioned centralized stream join algorithms (while tackling additional challenges such as distribution and fault-tolerance) and hence similarly differ from GDYN as described above. We leave the extension of GDYN to parallel and distributed settings as future work.

Query evaluation with constant delay enumeration has gained increasing attention in the last decade [6–8,10,25,32,33,33,36,38]. This setting, however, deals with equijoins only.

Inequality joins also related, although restricted to the static setting, is the practical evaluation of binary [17,18,23] and multi-way [9,48] inequality joins. Our work, in contrast, considers dynamic processing of multi-way θ -joins, with a specialization to inequality joins. Khayyat et al. [28] proposed fast multi-way inequality join algorithms based on sorted arrays and space efficient bit arrays. They focus on the case where there are exactly two inequality conditions per pairwise join. While they also present an incremental algorithm for pairwise joins, their algorithm makes no effort to minimize the update cost of multi-way joins. As a result, they either materialize subresults (implying a space overhead that can be more than linear) or recompute subresults.

3 Preliminaries

Query Language Throughout the paper, let x, y, z, \dots denote *variables* (also commonly called *column names* or *attributes*). A *hyperedge* is a finite set of variables. We use

\bar{x}, \bar{y}, \dots to denote hyperedges. A *generalized conjunctive query* (GCQ) is an expression of the form

$$Q = \pi_{\bar{y}}(r_1(\bar{x}_1) \bowtie \dots \bowtie r_n(\bar{x}_n) \mid \bigwedge_{i=1}^m \theta_i(\bar{z}_i)). \quad (1)$$

Here r_1, \dots, r_n are *relation symbols*; $\bar{x}_1, \dots, \bar{x}_n$ are hyperedges (of the same arity as r_1, \dots, r_n); $\theta_1, \dots, \theta_m$ are predicates over $\bar{z}_1, \dots, \bar{z}_m$, respectively; and both \bar{y} and $\bigcup_{i=1}^m \bar{z}_i$ are subsets of $\bigcup_{i=1}^n \bar{x}_i$. We treat predicates abstractly: for our purpose, a predicate over \bar{x} is a (not necessarily finite) decidable set θ of tuples over \bar{x} . For example, $\theta(x, y) = x < y$ is the set of all tuples (a, b) satisfying $a < b$. We indicate that θ is a predicate over \bar{x} by writing $\theta(\bar{x})$. Throughout the paper, we consider only non-nullary predicates, i.e., predicates with $\bar{x} \neq \emptyset$.

Example 1 The following query is hence a GCQ.

$$\pi_{x,y,z}(r(x, y) \bowtie s(z, u) \bowtie t(v, w) \mid x < u, z < w).$$

Since, as usual, the natural join between relations that have a disjoint schema is simply their Cartesian product, this query asks to take the Cartesian product of $r(x, y)$, $s(z, u)$ and $t(v, w)$; subsequently, select those tuples that satisfy $x < u$ and $z < w$ and finally project on x, y, z . Likewise, the GCQ

$$\pi_{y,z}(r(x, y) \bowtie s(y, z) \bowtie t(z, v) \mid y < v)$$

asks to take the natural join of $r(x, y)$, $s(y, z)$ and $t(z, v)$ (where $r(x, y)$ and $s(y, z)$ equijoin on y , and $s(y, z)$ and $t(z, v)$ equijoin on z), subsequently select those tuples that satisfy $y < v$ and project on y, z .

We call \bar{y} the *output variables* of Q and denote it by $out(Q)$. If $\bar{y} = \bar{x}_1 \cup \dots \cup \bar{x}_n$, then Q is called a *full query* and we may omit the symbol $\pi_{\bar{y}}$ altogether for brevity. We denote by $full(Q)$ the full GCQ obtained from Q by setting $out(Q)$ to $\bar{x}_1 \cup \dots \cup \bar{x}_n$. The elements $r_i(\bar{x}_i)$ are called *atoms*. $at(Q)$ denotes the set of all atoms in Q and $pred(Q)$ the set of all predicates in Q . A *conjunctive query* (or CQ) is a GCQ where $pred(Q) = \emptyset$.

Semantics We evaluate GCQs over generalized multiset relations (GMRs for short) [25,29,30]. Let $dom(x)$ denotes the domain of variable x . As usual, a tuple over \bar{x} is a function \mathbf{t} that assigns a value from $dom(x)$ to every $x \in \bar{x}$. $\mathbb{T}[\bar{x}]$ denotes the set of all tuples over \bar{x} . A GMR over \bar{x} is a function $R: \mathbb{T}[\bar{x}] \rightarrow \mathbb{Z}$ mapping tuples over \bar{x} to integers such that $R(\mathbf{t}) \neq 0$ for finitely many tuples \mathbf{t} . In contrast to classical multisets, the multiplicity of a tuple in a GMR can hence be negative, allowing to treat insertions and deletions uniformly. We write $var(R)$ for \bar{x} ; $supp(R)$ for the set of all tuples with nonzero multiplicity in R ; $\mathbf{t} \in R$ to indicate $\mathbf{t} \in supp(R)$; and $|R|$ for $|supp(R)|$.

R	S	T	$R \bowtie T$
$\begin{array}{ c c c c } \hline x & y & z & \mathbb{Z} \\ \hline 1 & 2 & 2 & 2 \\ 2 & 4 & 6 & 3 \\ 1 & 2 & 3 & 3 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline u & v & \mathbb{Z} \\ \hline 4 & 5 & 5 \\ 2 & 3 & 4 \\ 1 & 4 & 2 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline u & v & \mathbb{Z} \\ \hline 4 & 5 & -4 \\ 2 & 1 & 6 \\ \hline \end{array}$	$\begin{array}{ c c c c c c } \hline x & y & z & u & v & \mathbb{Z} \\ \hline 1 & 2 & 2 & 4 & 5 & -8 \\ 1 & 2 & 2 & 2 & 1 & 12 \\ 2 & 4 & 6 & 4 & 5 & -12 \\ 2 & 4 & 6 & 1 & 6 & 18 \\ 1 & 2 & 3 & 4 & 5 & -12 \\ 1 & 2 & 3 & 2 & 1 & 18 \\ \hline \end{array}$
$\pi_y(R)$	$S + T$	$S - T$	$\sigma_{y < u}(R \bowtie S)$
$\begin{array}{ c c } \hline y & \mathbb{Z} \\ \hline 2 & 5 \\ 4 & 3 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline u & v & \mathbb{Z} \\ \hline 4 & 5 & 1 \\ 2 & 3 & 4 \\ 1 & 4 & 2 \\ 2 & 1 & 6 \\ \hline \end{array}$	$\begin{array}{ c c c } \hline u & v & \mathbb{Z} \\ \hline 4 & 5 & 9 \\ 2 & 3 & 4 \\ 1 & 4 & 5 \\ 2 & 1 & -6 \\ \hline \end{array}$	$\begin{array}{ c c c c c c } \hline x & y & z & u & v & \mathbb{Z} \\ \hline 1 & 2 & 2 & 4 & 5 & 10 \\ 1 & 2 & 3 & 4 & 5 & 15 \\ \hline \end{array}$

Fig. 2 Operations on GMRs

Figure 2 illustrates the operations of GMR union ($R + S$), minus ($R - S$), projection ($\pi_{\bar{z}}R$), natural join ($R \bowtie T$), and selection ($\sigma_P(R)$) which are defined similarly as in relational algebra with multiset semantics. See [25,30] for formal semantics. We stress that, as usual, if R and T have disjoint schema, then $R \bowtie T$ is simply their Cartesian product.

A GMR R is *positive* if $R(\mathbf{t}) > 0$ for all $\mathbf{t} \in \text{supp}(R)$. A *database* over a set \mathcal{A} of atoms is a function db that maps every atom $r(\bar{x}) \in \mathcal{A}$ to a positive GMR $db_{r(\bar{x})}$ over \bar{x} . Given a database db over the atoms occurring in query Q , the evaluation of Q over db , denoted $Q(db)$, is the GMR over \bar{y} constructed in the expected way: take the natural join of the GMRs in the database, do a selection over the result w.r.t. each predicate and finally project on \bar{y} . It is instructive to note that after evaluation, each result tuple has an associated multiplicity that counts the number of derivations for the tuple. In other words, the query language has built-in support for COUNT aggregations. We note that, in their full generality, GMRs can carry multiplicities that are taken from an arbitrary algebraic semiring structure (cf., [29]), which can be useful to describe the computation of more advanced aggregations over the result of a GCQ [2]. To keep the notation and discussion simple, we fix the ring \mathbb{Z} of integers throughout the paper, but our results generalize to arbitrary semirings and their associated aggregations.

Semijoins If Θ is a set of predicates, then we write $\sigma_{\Theta}R$ for $\sigma_{\bigwedge_{\theta \in \Theta} \theta}R$ and $R \bowtie_{\Theta} S$ for $\sigma_{\Theta}(R \bowtie S)$. If $\bar{z} \subseteq \text{var}(R)$ or $\bar{z} \subseteq \text{var}(S)$, then $\pi_{\bar{z}}(R \bowtie_{\Theta} S)$ is called a *semijoin*. We write $R \bowtie_{\Theta} S$ for the subGMR of R consisting of all tuples that have a joining tuple in S :

$$R \bowtie_{\Theta} S \in \mathbb{T}[\text{var}(R)] \rightarrow \mathbb{Z}:$$

$$\mathbf{t} \mapsto \begin{cases} R(\mathbf{t}) & \text{if } \mathbf{t} \in \pi_{\text{var}(R)}(R \bowtie_{\Theta} S) \\ 0 & \text{otherwise} \end{cases}$$

Often, S will contain only a single tuple \mathbf{t} with multiplicity 1. In that case, we simply write $R \bowtie_{\Theta} \mathbf{t}$.

Updates An *update* to a GMR R is simply a GMR ΔR over the same variables as R . Applying update ΔR to R yields the GMR $R + \Delta R$. An *update to a database* db is a collection u of (not necessarily positive) GMRs, one GMR $u_{r(\bar{x})}$ for every atom $r(\bar{x})$ of db , such that $db_{r(\bar{x})} + u_{r(\bar{x})}$ is positive.¹ We write $db + u$ for the database obtained by applying u atom-wise to db .

Computational Model We focus on dynamic query evaluation in main memory. We assume a model of computation where the space used by tuple values and integers, the time of arithmetic operations on integers, the time of operations on tuples (such as projecting a tuple on a subset of its variables, or taking the union of two tuples), and the time of memory lookups are all $\mathcal{O}(1)$. We further assume that hash tables have $\mathcal{O}(1)$ access and update times while requiring linear space. While it is well known that real hash table access is $\mathcal{O}(1)$ *expected* time and updates are $\mathcal{O}(1)$ *amortized*, complexity results that we establish for this simpler model can be expected to translate to average (amortized) complexity in real-life implementations [13].

A direct consequence of these assumptions is that, using standard database implementation techniques, every GMR R can be represented in our model by a data structure that allows (1) enumeration of R with delay $\mathcal{O}(1)$ (as defined in Sect. 4.1); (2) multiplicity lookups $R(\mathbf{t})$ in $\mathcal{O}(1)$ time given \mathbf{t} ; (3) single-tuple insertions and deletions in $\mathcal{O}(1)$ time; while (4) having size that is proportional to $|R|$. In addition, we assume it possible to sort GMRs by a given order on its tuples in $\mathcal{O}(|R| \log |R|)$ time, after which it allows enumeration in the given order with $\mathcal{O}(1)$. Single-tuple insertions that keep the GMR sorted become $\mathcal{O}(\log |R|)$ in this case.

4 General dynamic Yannakakis

In this section, we formulate GDYN, a dynamic version of the Yannakakis algorithm [47] that focuses on the evaluation of GCQ s. GDYN takes a nonstandard approach to dynamic query evaluation: instead of materializing $Q(db)$ and subjoins of Q , GDYN builds a succinct, efficiently updatable data structure that represents $Q(db)$ in the sense that from it we can *enumerate* $Q(db)$. Formally, a data structure D supports *enumeration* of a set E if there is a routine ENUM such that ENUM(D) outputs each element of E exactly once. Such enumeration occurs with delay d if the time until the first element is output; the time between any two consecu-

¹ Note that, in this framework, value modifications inside a tuple are modeled by deleting the tuple with the old value, and then reinserting the tuple, but now with the new value.

tive elements and the time between the last element and the termination of $\text{ENUM}(D)$, are all bounded by d . D supports enumeration of a GMR R if it supports enumeration of the set $E_R = \{(\mathbf{t}, R(\mathbf{t})) \mid \mathbf{t} \in \text{supp}(R)\}$.

When evaluating a GCQ Q over a database db , we will be interested in representing the elements of $Q(db)$ by means of a data structure D_{db} , such that we can enumerate $Q(db)$ from D_{db} . If, for every db , the delay to enumerate $Q(db)$ from D_{db} is sublinear in $|db|$, then we say that the enumeration occurs with *sublinear delay*. Similarly, if the delay is independent of $|db|$, then we say that the enumeration occurs with *constant delay* [38]. Note that this is sublinear/constant in data complexity [42]: the delay may still depend on the size of the query Q . This is reasonable since Q specifies the arity of the query result, and a larger arity inherently implies a longer delay between elements. Note that enumeration with constant delay is what we typically obtain by materializing $Q(db)$. For example, if we store the elements of $Q(db)$ in an array, then enumerating $Q(db)$ amounts to scanning the array where each element access is $\mathcal{O}(1)$.

We start the development of GDYN by first giving some intuition into constant delay enumeration (henceforth: CDE) in Sect. 4.1. The algorithm itself is stated in Sect. 4.2, while proofs of its correctness and analysis of its complexity is given in Sect. 4.3. Finally, we specialize GDYN to the case where all θ -joins are inequality joins in Sect. 4.4.

4.1 Intuition

In this section, we discuss how we can obtain CDE of the result $Q(db)$ of a GCQ Q . Of course, the simplest way to obtain this is simply to materialize $Q(db)$. Unfortunately, this requires memory proportional to $|Q(db)|$ which, depending on Q , can be of size polynomial in $|db|$. We hence desire other data structures to represent $Q(db)$ using less space, while still allowing CDE. Let us build some intuition on how this can be done by subsequently considering three queries of increasing complexity:

$$Q_1 = r(x, y) \bowtie s(y, z),$$

$$Q_2 = s(y, z) \bowtie t(z, v) \mid v < y, \quad \text{and}$$

$$Q_3 = r(x, y) \bowtie s(y, z) \bowtie t(z, v) \mid v < y.$$

Throughout our discussion, assume that the GMRs assigned to $r(x, y)$, $s(y, z)$, and $t(z, v)$ by input database db are R , S , and T , respectively.

It is instructive to start with the simple binary equijoin query $Q_1 = r(x, y) \bowtie s(y, z)$ and analyze why traditional join processing algorithms do not yield CDE. Suppose that we evaluate Q_1 using a simple in-memory hash join with R as probe relation and S as build relation. Assume that the corresponding hash index of S on y has already been

computed. Concretely, this hash index allows us to retrieve, for every y -tuple \mathbf{t} , in $\mathcal{O}(1)$ time a pointer to the GMR $S \times \mathbf{t}$ of all S -tuples that join with \mathbf{t} . Now observe that, when iterating over R to probe the index, we may have to visit an unbounded number of R -tuples that do not join with any S -tuple. Consequently, the delay between consecutive outputs may be as large as $|R|$. A similar analysis shows that other join algorithms, such as the sort-merge join, do not yield CDE.

How then can we obtain CDE for $r(x, y) \bowtie s(y, z)$? Intuitively, if we can ensure to only iterate over those R -tuples that have matching S -tuples, we trivially obtain constant delay since then every probe will yield at least one output tuple. As such, the key is to first compute the semijoin $\rho_{x,y} = \pi_{x,y}(R \bowtie S)$. We can then iterate over the elements of $\rho_{x,y}$, probing S in each iteration to generate the output with constant delay. Note that, because $\rho_{x,y}$ is a semijoin, the space needed to store $\rho_{x,y}$ is linear in $|db|$.

CDE for queries that feature θ -joins can be obtained similarly. Consider $Q_2 = (s(y, z) \bowtie t(z, v) \mid y < v)$ which is a combination of an equijoin on z and inequality join on $y < v$. To obtain CDE for Q_2 , first compute the semijoin $\rho_{y,z} = \pi_{y,z}(S(y, z) \bowtie_{y < v} T(z, v))$ which consists of all tuples in S that have a matching tuple in T . Assume for a moment that we have a more powerful index structure I that allows, for any $\{y, z\}$ -tuple \mathbf{s} , to enumerate $T \bowtie_{y < v} \mathbf{s}$ with constant delay. We can then easily enumerate $Q_2(db)$ with constant delay by iterating over $\mathbf{s} \in \rho_{y,z}$, and for each such \mathbf{s} , probe I to produce the tuples $\mathbf{t} \in T \bowtie_{y < v} \mathbf{s}$, outputting $(\mathbf{s} \cup \mathbf{t}, S(\mathbf{s}) \times T(\mathbf{t}))$ for each such \mathbf{s} and \mathbf{t} . Since $T \bowtie_{y < v} \mathbf{s}$ allows CDE and multiplicity lookups are $\mathcal{O}(1)$, the entire procedure is CDE. The key question then is how we can build this more powerful index structure I . The solution is to group T on z ; subsequently, sort each group in descending order on v ; and create a normal (hash-based) index J that allows to find the group for each z -value. This now supports CDE of $T \bowtie_{y < v} \mathbf{s}$: first, use J to get a pointer to the group with z -value $\mathbf{s}(z)$ in $\mathcal{O}(1)$ time, and then enumerate this group with constant delay and in decreasing order on v . Yield the current tuple \mathbf{t} that is being enumerated in this fashion, provided that $\mathbf{s}(y) < \mathbf{t}(v)$. As soon as $\mathbf{s}(y) \geq \mathbf{t}(v)$, we know that all subsequent \mathbf{t} will fail the inequality, and we can hence terminate. The lower left of Fig. 3 illustrates S , T , $\rho_{y,z}$, and J .

CDE for queries that join more than two relations can be obtained similarly, but now by computing nested semijoins. Figure 3 illustrates how to obtain CDE for Q_3 . Concretely, we first ensure CDE of the subquery Q_2 of Q_3 as already explained above by computing the semijoin $\rho_{y,z} = \pi_{y,z}(S(y, z) \bowtie_{y < v} T(z, v))$ and suitably indexing T . Then CDE of $Q_3(db)$ is obtained by observing that $Q_3(db) \equiv R \bowtie Q_2(db)$, where $R \bowtie Q_2(db)$ is a binary equijoin, which can hence treated completely analogous as Q_1 . Concretely, compute the nested semijoin $\rho_{x,y} = \pi_{x,y}(R \bowtie \rho_{y,z})$, and

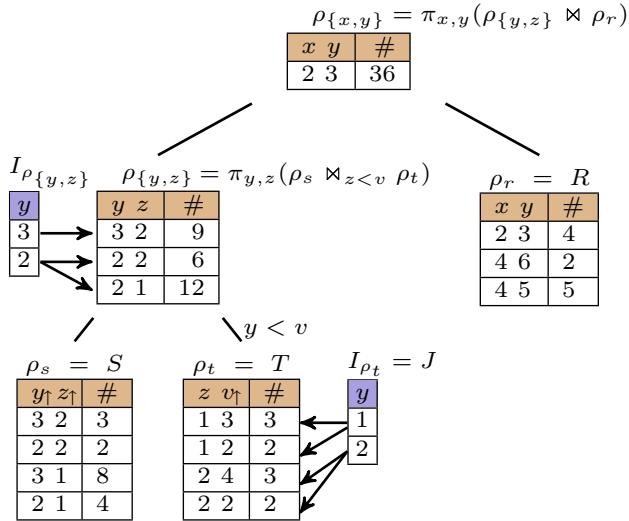


Fig. 3 Illustration of query Q_3

build a hash index on $\rho_{y,z}$ on y . (This index is depicted as $I_{\rho_{y,z}}$ in Fig. 3). Enumeration of $Q_3(db)$ is done by iterating over the tuples $\mathbf{r} \in \rho_{x,y}$, and for each such tuple \mathbf{r} , use $I_{\rho_{y,z}}$ to get a pointer to $\rho_{y,z} \bowtie \mathbf{r}$, which consists of all $\mathbf{s} \in \rho_{y,z}$ that equijoin with \mathbf{r} . Iterate over these \mathbf{s} with constant delay, and finally use the more advanced index on T to enumerate all tuples $\mathbf{t} \in T \bowtie_{y < v} \mathbf{s}$. For each such \mathbf{r}, \mathbf{s} , and \mathbf{t} , we output $(\mathbf{r} \cup \mathbf{s} \cup \mathbf{t}, R(\mathbf{r}) \times S(\mathbf{s}) \times T(\mathbf{t}))$. By construction of $\rho_{x,y}$, we are ensured that matching \mathbf{s} will exist for every \mathbf{r} . Similarly, matching \mathbf{t} exist for every \mathbf{s} by construction of $\rho_{y,z}$. Therefore, each tuple that we iterate over will produce a new output, and the entire enumeration of $Q(db)$ is CDE.

In conclusion As the examples above illustrate, we can obtain CDE for GCQs by computing (nested) semijoins, and suitably indexing both base relations and semijoin results for enumeration. Because the only additional relations that we compute are obtained by semijoining existing relations, the size of all additional GMRs that are stored is *linear* in the input *db*. Contrast this to techniques that materialize subjoin results, whose size may become polynomial in the database.

Updates We finish this section by remarking that, in the presence of updates, this approach is only valid if we materialize and maintain all required semijoin results. To speed up the maintenance of semijoin results under updates, it is sometimes beneficial to create additional indexes that help in incremental computation of the semijoins, as we illustrate next. Reconsider Q_3 as illustrated in Fig. 3. If we receive an update ΔT to T , then we need to correspondingly update $\rho_{y,z}$ from $\pi_{y,z}(S \bowtie_{y < v} T)$ to $\pi_{y,z}(S \bowtie_{y < v} (T + \Delta T))$. To that end, it suffices to compute $\Delta \rho_{y,z} = \pi_{y,z}(S \bowtie_{y < v} \Delta T)$, and add this to $\rho_{y,z}$. Computing $\Delta \rho_{y,z}$ by means of a nested loop join has $\Omega(|S| \times |\Delta T|)$ complexity. We can do better if we index S by sorting S lexicographically on (z, y) , in decreasing

order (this is actually how S is depicted in Fig. 3). $\Delta \rho'_{y,z}$ can then be computed by means of a hybrid form of sort-merge and index nested loop join. First, group ΔT on z and, per group, sort tuples in decreasing order on variable v . Create a hash index on ΔT to be able to quickly find each group by a given z -value. Second, iterate over the tuples in S in the given lexicographic order. For each z -group in S , find the corresponding group in ΔT by passing the z -value to the hash table. Let \mathbf{s} be the first tuple in the S -group. Then iterate over the tuples of the ΔT group in decreasing order on v , and sum up their multiplicities until $\mathbf{s}(y)$ becomes larger than v . Add \mathbf{s} to $\Delta \rho_{y,z}$, with its original multiplicity in S multiplied by the found sum (provided that it is nonzero). Then, consider the next tuple in the S -group, and continue summing from the current tuple in the ΔT group until $\mathbf{s}(y)$ becomes again larger than v , and add the result tuple with the correct multiplicity. Continue repeating this process for each tuple in the S -group, and for each group in S . Assuming that the index on S already existed, then the total cost of computing $\rho_{y,z}$ in this way is $\mathcal{O}(|S| + |\Delta T| + |\Delta T| \log |\Delta T|)$ since we scan S and ΔT only once, need to sort ΔT , and create a hash table for each group. This is much better than the $\mathcal{O}(|S| \times |\Delta T|)$ complexity of a nested loop.

4.2 The algorithm

We now turn to the general formulation of the dynamic Yannakakis algorithm. As exemplified in Sect. 4.1, GDYN obtains CDE by computing (nested) semijoins and indexing both these semijoins and the base relations. The order in which semijoins are computed and how they are indexed is recorded in a *dynamic query plan*, which is introduced next.

Dynamic Query Plans To simplify notation, we denote the set of all variables (resp. atoms, resp. predicates) that occur in an object X (such as a query) by $\text{var}(X)$ (resp. $\text{at}(X)$, resp. $\text{pred}(X)$). In particular, if X is itself a set of variables (a hyperedge), then $\text{var}(X) = X$. We extend this notion uniformly to labeled trees, e.g., if n is a node in tree T , then $\text{var}_T(n)$ denotes the set of variables occurring in the label of n and similarly for edges and trees themselves. If T is clear from the context, we omit subscripts from our notation.

Definition 1 A *dynamic query plan* (or simply: *plan*) is a tuple (T, N) where T is a binary generalized join tree, and N is a sibling-closed connex subset of T . A *generalized join tree* (GJT) is a node-labeled and edge-labeled directed tree $T = (V, E)$ such that:

- Every leaf is labeled by an atom.
- Every interior node n is labeled by a hyperedge and has at least one child c such that $\text{var}(n) \subseteq \text{var}(c)$. Such a child is called a *guard* of n .




Fig. 4 Two example plans. The connex sets are indicated by the shaded areas

- Whenever the same variable x occurs in the label of two nodes m and n of T , then x occurs in the label of each node on the unique path linking m and n . This condition is called the *connectedness condition*.
- Every edge $p \rightarrow c$ from parent p to child c in T is labeled by a set $\text{pred}(p \rightarrow c)$ of predicates. It is required that for every predicate $\theta(\bar{z}) \in \text{pred}(p \rightarrow c)$, we have $\text{var}(\theta) = \bar{z} \subseteq \text{var}(p) \cup \text{var}(c)$.

T is *binary* if every node in T has at most two children. A *connex subset* of T is a set $N \subseteq V$ that includes the root of T such that the subgraph of T induced by N is a tree. N is *sibling closed* if for every node $n \in N$ with a sibling m in T , m is also in N . The *frontier* of a connex set N is the subset $F \subseteq N$ consisting of those nodes in N that are leaves in the subtree of T induced by N .

Figure 4 shows two plans (T_1, N_1) and (T_2, N_2) . The set N_1 contains all nodes of T_1 and is therefore a sibling-closed connex subset of T_1 . Its frontier is $\{r(x, y), s(y, z), t(z, v)\}$. If we remove $\{y, z\}$ from N_1 the set is no longer connex. If instead we remove $r(x, y)$, the set is still connex, but no longer sibling closed. The set N_2 is a sibling-closed connex subset of T_2 , and its frontier is $\{\{y, z, w\}, \{u\}\}$. Removing either $\{y, z, w\}$ or $\{u\}$ makes N_2 no longer sibling closed.

Definition 2 Let T be a GJT and let N be a connex subset of T . Assume that $\{r_1(\bar{x}_1), \dots, r_n(\bar{x}_n)\}$ is the multiset of atoms occurring as labels in the leaves of T . Then, the query associated with T is the full join

$$\mathcal{Q}[T] = r_1(\bar{x}_1) \bowtie \dots \bowtie r_n(\bar{x}_n) \mid \bigwedge_{\theta(\bar{z}) \in \text{pred}(T)} \theta(\bar{z}),$$

and the query associated with (T, N) is the GCQ defined as $\mathcal{Q}[T, N] = \pi_{\text{var}(N)}(\mathcal{Q}[T])$.

To illustrate, referring to Fig. 4, we have that $\mathcal{Q}[T_1, N_1] = \pi_{x,y,z,v}(r(x, y) \bowtie s(y, z) \bowtie t(z, v) \mid y < v)$, which is \mathcal{Q}_3 from Sect. 4.1.

The data structure Following the intuition of Sect. 4.1, the GJT T of a plan (T, N) specifies the semijoin results that

need be maintained and indexed, while the connex set N drives the enumeration of query results. We formalize this next. Because in this section, we are introducing GDYN for arbitrary GCQs (with arbitrary join predicates θ), we first need to introduce general notions of an index.

Definition 3 (*Enumeration index*) Let R be a GMR, θ a predicate, and \bar{y} be a hyperedge. A data structure I that is of size linear in R and that allows, for any given \bar{y} -tuple \mathbf{t} , enumeration of $(R \bowtie_{\theta} \mathbf{t})$ with delay $\mathcal{O}(f(|R|))$ is called an *enumeration index of R by (θ, \bar{y}) with delay $f: \mathbb{N} \rightarrow \mathbb{N}$* .

For example, in Sect. 4.1, we have discussed how, by means of grouping and sorting, we can obtain an enumeration index of $T(z, v)$ on $(y < v, \{y, z\})$ with constant delay.

Definition 4 (*Join Index*) Let R be a GMR, θ be a predicate, and \bar{y}, \bar{z} be hyperedges such that $\bar{z} \subseteq \text{var}(R)$ or $\bar{z} \subseteq \bar{y}$. A data structure I that is of size linear in R and that allows, for any GMR S over \bar{y} , computation of $\pi_{\bar{z}}(R \bowtie_{\theta} S)$ in time $\mathcal{O}(g(|R|, |S|))$ is called a *join index of R by $(\theta, \bar{y}, \bar{z})$ with access time $g: \mathbb{N}^2 \rightarrow \mathbb{N}$* .

For example, in Sect. 4.1, we have discussed how, by means of grouping and sorting, we can obtain a join index of $R(y, z)$ by $(y < v, \{z, v\}, \{y, z\})$ whose access time is $\mathcal{O}(|R| + |S| \log |S|)$.

For both enumeration and join indexes, the *update time* is the time required to update the index to a corresponding new index on $R + \Delta R$, given update ΔR .

Definition 5 Let (T, N) be a plan and let db be a database over $at(T)$. The *T -reduct* (or *semijoin reduction*) of db is a collection ρ of GMRs, one GMR ρ_n for each node $n \in T$, defined inductively as follows. Let $\text{pred}(n)$ denote the set of all predicates on the edges from n to its children in T .

- If $n = r(x)$ is an atom, then $\rho_n = db_{r(x)}$.
- If n has a single child c , then $\rho_n = \pi_{\text{var}(n)} \sigma_{\text{pred}(n)} \rho_c$.
- Otherwise, n has two children c_1 and c_2 . In this case, we have $\rho_n = \pi_{\text{var}(n)} \sigma_{\text{pred}(n)} (\rho_{c_1} \bowtie \rho_{c_2})$. Note that, because n has a guard child, this is actually a semijoin.

A T -reduct needs to be augmented by suitable index structures to be used for both enumeration and maintenance under updates. Concretely for each node n with parent p in T , the following indexes are created:

- If n belongs to N , then we store an enumeration index P_n on ρ_n by $(\text{pred}(p \rightarrow n), \text{var}(p))$.
- If n is a node with a sibling m , then we store a join index S_n on ρ_n by $(\text{pred}(p), \text{var}(m), \text{var}(p))$.

The T -reduct ρ together with these indexes is called a (T, N) -representation for db , or (T, N) -rep for short.

Algorithm 1 GDYN: General Dynamic Yannakakis

```

1: function ENUMT,N( $\rho$ )
2:   for each  $\mathbf{t} \in \rho_{\text{root}(T)}$  do ENUMT,N( $\text{root}(T), \mathbf{t}, \rho$ )
3: function ENUMT,N( $n, \mathbf{t}, \rho$ )
4:   if  $n$  is in the frontier of  $N$  then yield ( $\mathbf{t}, \rho_n(\mathbf{t})$ )
5:   else if  $n$  has one child  $c$  then
6:     for each  $\mathbf{s} \in \rho_c \bowtie_{\text{pred}(n \rightarrow c)} \mathbf{t}$  do ENUMT,N( $c, \mathbf{s}, \rho$ )
7:   else if  $n$  has two children  $c_1$  and  $c_2$ 
8:     for each  $\mathbf{t}_1 \in \rho_{c_1} \bowtie_{\text{pred}(n \rightarrow c_1)} \mathbf{t}$  do
9:       for each  $\mathbf{t}_2 \in \rho_{c_2} \bowtie_{\text{pred}(n \rightarrow c_2)} \mathbf{t}$  do
10:        for each  $(\mathbf{s}_1, \mu) \in \text{ENUM}_{T,N}(c_1, \mathbf{t}_1, \rho)$  do
11:          for each  $(\mathbf{s}_2, \nu) \in \text{ENUM}_{T,N}(c_2, \mathbf{t}_2, \rho)$  do
12:            yield ( $\mathbf{s}_1 \cup \mathbf{s}_2, \mu \times \nu$ )
13: procedure UPDATET,N( $\rho, u$ )
14:   for each  $n \in \text{leaves}(T)$  labeled by  $r(\bar{x})$  do
15:      $\Delta_n \leftarrow u_{r(\bar{x})}$ 
16:   for each  $n \in \text{nodes}(T) \setminus \text{leaves}(T)$  do
17:      $\Delta_n \leftarrow \text{empty GMR over } \text{var}(n)$ 
18:   for each  $n \in \text{nodes}(T)$ , traversed bottom-up do
19:      $\rho_n \leftarrow \Delta_n$ 
20:     if  $n$  has a parent  $p$  and a sibling  $m$  then
21:        $\Delta_p \leftarrow \pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}(p)} \Delta_n)$ 
22:     else if  $n$  has parent  $p$  then
23:        $\Delta_p \leftarrow \pi_{\text{var}(p)} \sigma_{\text{pred}(p)} \Delta_n$ 

```

Reconsider the plan (T_1, N_1) from Fig. 4. Figure 3 depicts an example (T_1, N_1) -representation ρ for the database db composed of the GMRs shown at the leaves of the tree. $I_{\rho_{y,z}}$ and I_{ρ_t} illustrate the enumeration indexes; the join indexes are not illustrated.

It is important to observe that, since a T -reduct constructs only semijoins of database GMRs, and projections thereof, each $|\rho_n|$ is linear in the size of db . Consequently, the indexes are also of size linear in db and hence the entire (T, N) -rep is linear in db .

Proposition 1 $|\rho_n| \leq \max_{r(\bar{x}) \in \text{at}(T)} |db_{r(\bar{x})}|$, for every $n \in T$.

Given these definitions, the enumeration and maintenance algorithms that form GDYN are shown in Algorithm 1. They operate as follows.

Enumeration To enumerate from a (T, N) -rep, we iterate over the reductions ρ_n with $n \in N$ in a nested fashion, starting at the root and proceeding top down. When n is the root, we iterate over all tuples in ρ_n . For every such tuple \mathbf{t} , we iterate only over the tuples in the children c of n that are compatible with \mathbf{t} (i.e., tuples in ρ_c that join with \mathbf{t} and satisfy $\text{pred}(n \rightarrow c)$). Note that such tuples can be enumerated directly using the enumeration index P_c . This procedure continues until we reach nodes in the frontier of N at which time the output tuple can be constructed. The pseudocode is given by the routine ENUM in Algorithm 1.

Update processing To maintain a (T, N) -rep under update u , we traverse the nodes of T in a bottom-up fashion. At

each node n , we have to compute the update Δ_n to apply to ρ_n and its associated indexes. For leaf nodes, this update is given by the update u itself. For interior nodes, Δ_n can be computed from the update and the original reduct of its children. UPDATE in Algorithm 1 gives the pseudocode. Here, line 21 is implemented by using the join index S_m on ρ_m by $(\text{pred}(p), \text{var}(n), \text{var}(p))$. Line 23 can be implemented by a straightforward hash-based aggregation. We assume here that, as a side effect of the presented modifications, the associated indexes are also updated.

4.3 Correctness and complexity

We next prove correctness of the enumeration and update procedures, and bound their complexity. We start with enumeration. Throughout this section, let (T, N) be a plan, let db be a database over $\text{at}(T)$, and assume that we have (T, N) -rep of db with T -reduct ρ . Given a node $n \in T$, we denote the subtree of T rooted at n by T_n , and the subset of all nodes of N that are in T_n by N_n . The following lemma relates the GMRs at each node n of a T -reduct to the query induced by the subtree of T at n . Here, we write $\mathcal{Q}[T_n, n]$ as a shorthand for $\mathcal{Q}[T_n, \{n\}]$. Recall from Definition 2 that this is the query that joins all atoms in T_n (w.r.t. all predicates in T_n), and subsequently projects on $\text{var}(n)$.

Lemma 1 $\rho_n = \mathcal{Q}[T_n, n](db)$, for every node $n \in T$.

The proof by induction is detailed in “Appendix A”.

To show correctness of enumeration, we need the following additional lemma regarding the subroutine of Algorithm 1 (line 3). The proof is again by induction and detailed in “Appendix A”.

Lemma 2 For every node $n \in N$ and every tuple \mathbf{t} in ρ_n , ENUM_{T,N}(n, \mathbf{t}, ρ) enumerates $\mathcal{Q}[T_n, N_n](db) \bowtie \mathbf{t}$.

Finally, we require the following insights, also proved in “Appendix A”.

Lemma 3 1. $\mathcal{Q}(db)$ is a positive GMR, for any GCQ Q and any database db .

2. If R is a positive GMR over \bar{x} and $\bar{y} \subseteq \bar{x}$, then $\mathbf{t}[\bar{y}] \in \pi_{\bar{y}} R$ for every tuple $\mathbf{t} \in R$.

We note that item (2) is not true for when R has both positive and negative multiplicities, since multiplicities of opposite sign could cancel each other out when projecting, thereby removing $\mathbf{t}[\bar{y}]$ from $\pi_{\bar{y}}(R)$.

Proposition 2 If ρ is a T -reduct of db , then ENUM_{T,N}(ρ) enumerates $\mathcal{Q}[T, N](db)$.

Proof Let r be the root of T . By Lemma 1, we have $\rho_r = \mathcal{Q}[T_r, r](db) = \mathcal{Q}[T, r](db) = \pi_{\text{var}(r)} \mathcal{Q}[T](db)$. Furthermore, $\pi_{\text{var}(r)} \mathcal{Q}[T](db) = \pi_{\text{var}(r)} \pi_{\text{var}(N)} \mathcal{Q}[T](db)$

since $\text{var}(r) \subseteq \text{var}(N)$ as $r \in N$. Therefore, $\rho_r = \pi_{\text{var}(r)}\pi_{\text{var}(N)}\mathcal{Q}[T](db) = \pi_{\text{var}(r)}\mathcal{Q}[T, N](db)$. We conclude that ρ_r is a projection of $\mathcal{Q}[T, N](db)$, and hence by Lemma 3 that every tuple in $\mathcal{Q}[T, N](db)$ has a compatible tuple in ρ_r . As such, $\mathcal{Q}[T, N](db)$ equals the disjoint union $\bigcup_{t \in \rho_r} \mathcal{Q}[T, N](db) \times t$. By Lemma 2, this is exactly what $\text{ENUM}_{T,N}(\rho)$ enumerates. \square

We now analyze the complexity of $\text{ENUM}_{T,N}$. First, observe that by definition of T -reducts, it is the case that for every parent node p with a child n and every $t \in \rho_p$ there exists a tuple in $\rho_n \times_{\text{pred}(p \rightarrow n)} t$. Hence, every tuple that we iterate over will eventually produce a new output tuple. This ensures that we do not risk wasting time in iterating over tuples that in the end yield no output. As such, the time needed for $\text{ENUM}_{T,N}(\rho)$ to produce a single new output is dominated by the time taken to iterate over the tuples in $\rho_n \times_{\text{pred}(p \rightarrow n)} t$, where p is the parent of n . Since we can use the enumeration index P_n to do so efficiently, the efficiency of the entire enumeration will depend on the delay incurred by accessing the enumeration indexes. The following proposition formalizes this insight.

Proposition 3 Assume that every enumeration index has enumeration delay f , where f is a monotone function. Then, using these indexes, $\text{ENUM}_{T,N}(\rho)$ enumerates $\mathcal{Q}[T, N](db)$ with delay $\mathcal{O}(|N|f(M))$ where $M = \max_{r(\bar{x}) \in \text{at}(T)} |db_{r(\bar{x})}|$. Thus, the total time required to execute $\text{ENUM}_{T,N}(\rho)$ is $\mathcal{O}(|\mathcal{Q}[T, N](db)| |N|f(M))$.

Proof $\text{ENUM}_{T,N}(\rho)$ correctly enumerates $\mathcal{Q}[T, N](db)$ by Proposition 2. As such, it suffices to show that the delay satisfies the given bounds. To that end, we show that for every node $n \in N$ and $t \in \rho_n$, the call $\text{ENUM}_{T,N}(n, t, \rho)$ will produce outputs with delay $\mathcal{O}(|N_n|f(M))$. We proceed by induction on $|N_n|$. If $|N_n| = 1$, then n is in the frontier of N and the delay is clearly constant as the algorithm will only yield $(t, \rho_n(t))$ (line 4). Now assume that $|N_n| > 1$. Then, n is not in the frontier of N . If n has a single child c , then line 6 is executed, and the enumeration index P_c allows us to iterate over $\rho_c \times_{\text{pred}(n)} t$ with delay $\mathcal{O}(f(|\rho_c|))$, which is $\mathcal{O}(f(M))$ by Proposition 1. For each element s of this enumeration, the algorithm calls $\text{ENUM}_{T,N}(c, s, \rho)$, which by induction hypothesis produces output elements with delay $\mathcal{O}(|N_c|f(M))$. Hence, the maximum delay between two outputs is $\mathcal{O}(f(|M|) + |N_c|f(M)) = \mathcal{O}((|N_c| + 1)f(M)) = \mathcal{O}(|N_n|f(M))$. For the case in which n has two children c_1 and c_2 , lines 7–12 are executed. By similar reasoning, it is easy to show that the maximum delay between outputs is

$$\begin{aligned} &\mathcal{O}(f(|M|)) + \mathcal{O}(|N_{c_1}|f(M)) \\ &+ \mathcal{O}(f(|M|)) + \mathcal{O}(|N_{c_2}|f(M)) \end{aligned}$$

which is bounded by $\mathcal{O}((|N_{c_1}| + |N_{c_2}| + 2)f(M)) = \mathcal{O}(2|N_n|f(M)) = \mathcal{O}(|N_n|f(M))$. \square

In particular, if all enumeration indexes are with constant delay (i.e., $f(M) = \mathcal{O}(1)$), then GDYN enumerates $\mathcal{Q}[T, N](db)$ with delay $\mathcal{O}(|N|)$, which is also constant in data complexity.

Update processing We next turn our attention to the update procedure UPDATE of Algorithm 1. Since this is straightforward to prove correct, we focus on its complexity. Since UPDATE uses the join indexes available in the (T, N) -rep during its execution we will hence bound the running time of UPDATE in terms of the join index access and update times. We first require the following insight, which bounds the running time under the condition that the GMRs computed by UPDATE have a certain bounded size. The proof is in “Appendix A”.

Proposition 4 Assume that all join indexes in the (T, N) -rep have access time g , and that all indexes (join and enumeration) have update time h , where g and h are monotone functions. Further assume that, during the entire execution of UPDATE, K and U bound the size of ρ_n and Δ_n , respectively, for all n . Then, $\text{UPDATE}_{T,N}(\rho, u)$ runs in time $\mathcal{O}(|T| \cdot (U + h(K, U) + g(K, U)))$.

We next bound the size of ρ_n and Δ_n throughout the execution of UPDATE.

Proposition 5 During the entire execution of UPDATE, we have $|\rho_n| \leq M$ and $|\Delta_n| \leq 4M$ for every $n \in T$, where $M = \max_{r(\bar{x}) \in \text{at}(T)} |db_{r(\bar{x})}| + |u_{r(\bar{x})}|$.

Proof We first establish the bound on $|\rho_n|$ during execution. Before execution, ρ is a T -reduct of db . Hence, by Proposition 1, $|\rho_n| \leq \max_{r(\bar{x})} |db_{r(\bar{x})}| \leq M$ before execution starts. Now note that the only line that updates ρ_n is line 19, executed while visiting node n in the bottom-up traversal of T . This line is only applied once for every node n . Hence, since at the end of execution the collection of modified GMRs ρ_m for $m \in N$ form a T -reduct of $db + u$, we know that after executing line 19, ρ_n contains exactly the content for the T -reduct of $db + u$. Hence, by Proposition 1, $|\rho_n| \leq \max_{r(\bar{x})} |(db + u)_{r(\bar{x})}| \leq \max_{r(\bar{x})} |db_{r(\bar{x})}| + |u_{r(\bar{x})}| = M$.

We are now ready to establish the bounds on $|\Delta_n|$. Clearly, $|\Delta_n| \leq M$ during the initialization of Δ_n done in lines 14–17. Now consider that we are executing the bottom-up traversal of T in lines 18–23 and that n is the currently visited node. We have already established that both before and after applying the update Δ_n to ρ_n , we have $|\rho_n| \leq M$ (line 19). This implies that $|\Delta_n| \leq 2M$: in the worst case, Δ_n deletes all existing tuples in ρ_n and adds M new ones. To see that $|\Delta_p| \leq 4M$ after executing line 21, we consider two cases. If n is visited before m in the bottom-up traversal of T , then Δ_p is necessarily empty before executing line 21 and hence $|\Delta_p| = |\pi_{\text{var}(p)}(\rho_m \times_{\text{pred}(p)} \Delta_n)|$. Because, by definition of GJTs, p has either m or n as a guard, it follows that every

tuple in $\pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}(p)} \Delta_n)$ is either a projected version of some tuple in ρ_m , or a projected version of some tuple in Δ_n . As such, $|\Delta_p| \leq \max(|\rho_m|, |\Delta_n|) = 2M$. If, on the other hand, m is visited before n in the bottom-up traversal of T , then Δ_p necessarily contains the result computed during executing line 21 while visiting m . By the reasoning of the previous case, $|\Delta_p| \leq 2M$ before executing line 21 while visiting n . Furthermore, by the same reasoning, we know that $|\pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}(p)} \Delta_n)| \leq 2M$. Hence, at most $2M$ new tuples can be added to Δ_p . As such, $|\Delta_p| \leq 4M$ after executing line 21 when visiting n . To see that also $|\Delta_p| \leq 4M$ after executing line 23 when visiting n , it suffices to observe that Δ_p starts out empty, and $|\Delta_n| \leq 2M$ (as already established). \square

Combining Propositions 4 and 5, we obtain

Theorem 1 Assume that all join indexes in the (T, N) -rep have access time g and that all indexes (join and enumeration) have update time h , where g and h are monotone functions. $\text{UPDATE}_{T,N}(\rho, u)$ runs in time

$$\mathcal{O}(|T| \cdot (4M + h(M, 4M) + g(M, 4M))).$$

where $M = \max_{r(\bar{x}) \in \text{at}(T)} |db_{r(\bar{x})}| + |u_{r(\bar{x})}|$.

4.4 IEDyn

GDYN provides a general framework for dynamic query processing in the presence of arbitrary θ -joins. In this section, we instantiate GDYN to the specific setting where queries mention only inequality predicates ($<$, \leq , $>$, \geq).² We refer to this instantiation as IEDYN. Concretely, IEDYN uses the following data structures for its enumeration and join indexes. Let R be a GMR over \bar{x} , θ be a conjunction of inequalities, and \bar{y}, \bar{z} be hyperedges such that $\bar{z} \subseteq \bar{x}$ or $\bar{z} \subseteq \bar{y}$. The data structure underlying the enumeration and join index of R by (θ, \bar{y}) resp. $(\theta, \bar{y}, \bar{z})$ depends on the number of inequalities in θ .

(1) *No inequality* In this case, θ is hence equivalent to `true`, and the enumeration and join index only have to deal with equijoins. Concretely, the enumeration index of R on (θ, \bar{y}) is obtained by creating a traditional (hash-based) index of R on the variables that \bar{x} and \bar{y} have in common. Then, for any \bar{y} -tuple \mathbf{t} , $R \bowtie \mathbf{t}$ can be enumerated with constant delay by first using the hash index to find the corresponding $(\bar{x} \cap \bar{y})$ -group of R and enumerating the elements of that group.

To obtain the join index by $(\theta, \bar{y}, \bar{z})$, we discern two cases. If $\bar{z} \subseteq \bar{y}$, then the same index as for enumeration is reused except that in addition, for each $(\bar{x} \cap \bar{y})$ -group of R , we

cache the sum of all multiplicities in that group. This allows to evaluate $\pi_{\bar{z}}(R \bowtie S)$ in time $\mathcal{O}(|S|)$ independently of $|R|$, as follows. Initialize an empty GMR to hold the join result. Group S on the variables in $\bar{x} \cap \bar{y}$ in $\mathcal{O}(|S|)$ time (using hashing). For each group in S , use the index on R to locate the corresponding group of R in $\mathcal{O}(1)$ time and retrieve the cached sum-of-multiplicities μ of that group. Then, iterate over the tuples \mathbf{s} of the S -group one by one and add $S(\mathbf{s}) \times \mu$ to the multiplicity of $\mathbf{s}[\bar{z}]$ in the result GMR. (If $\mathbf{s}[\bar{z}]$ has not been added to the result before, this multiplicity is zero.) Repeat this process for all tuples in the S -group, and for every group in S . Since we only scan $|S|$ once, the total time is $\mathcal{O}(|S|)$.

If $\bar{z} \subseteq \bar{x}$, no special data structure is required: we can compute $\pi_{\bar{z}}(R \bowtie S)$ in $\mathcal{O}(|R| + |S|)$ time by first computing $\pi_{\bar{x} \cap \bar{y}} S$ in $\mathcal{O}(|S|)$ time and then repeating the above process with the roles of R and S reversed.

We conclude that in this case the enumeration delay is $\mathcal{O}(1)$, the access time is $\mathcal{O}(|S|)$ if $\bar{z} \subseteq \bar{x}$ and $\mathcal{O}(|R| + |S|)$ otherwise, and the update time is $\mathcal{O}(|\Delta R|)$.

(2) *Single inequality* Assume that $\theta = x > y$ with $x \in \bar{x}$ and $y \in \bar{y}$ (the other cases $x < y, x \leq y, x \geq y$ are similar). Then, we build a hash-based index I of R on $\bar{x} \cap \bar{y}$, sorting each group in descending order on x . In Sect. 4.1, we have illustrated, by means of example, that this realizes an enumeration index by (θ, \bar{y}) with constant delay. At the end of the same section, we have also illustrated that, when $\bar{z} \subseteq \bar{x}$, this also realizes a join index of R by $(\theta, \bar{y}, \bar{z})$ with access time $\mathcal{O}(|R| + |S| + |S| \log |S|) = \mathcal{O}(|R| + |S| \log |S|)$.³ When $\bar{z} \subseteq \bar{y}$, the same procedure, but using S for the outer loop, and R for the inner loop, can be used to realize a join index of R by $(\theta, \bar{y}, \bar{z})$ with access time $\mathcal{O}(|R| + |S| \log |S|)$.

Note that, because we need to keep data sorted the update time of these indexes is $\mathcal{O}(|\Delta R| \log(|R| + |\Delta R|))$.

We conclude that in this case the enumeration delay is $\mathcal{O}(1)$, the access time is $\mathcal{O}(|R| + |S| \log |S|)$, and the update time $\mathcal{O}(|\Delta R| \log(|R| + |\Delta R|))$.

(3) *Multiple inequalities* Assume that $\theta = x_1 > y_1 \wedge x_2 > y_2 \wedge \dots \wedge x_k > y_k$ with $x_1, \dots, x_k \in \bar{x}$ and $y_1, \dots, y_k \in \bar{y}$. (The reasoning where some of the $>$ are replaced by $<$ is completely analogous.) Then, as with the case with single inequalities, we build a hash-based index I of R on $\bar{x} \cap \bar{y}$ but now sort each group lexicographically on (x_1, \dots, x_k) (each x_i in descending order). In addition, for each group and each i ($1 \leq i \leq k$), we record the smallest x_i -value present in the group.

The fact that we have multiple inequalities complicates matters, in the sense that enumeration delay becomes logarithmic instead of constant. We can see this as follows. To

² Note that such queries may also contain equijoins by sharing variables between atoms.

³ Strictly speaking, we described in section that R needs to be sorted lexicographically, first on $\bar{x} \cap \bar{y}$ and then on x . The grouping + sorting of the enumeration index obtains the same result.

enumerate $R \bowtie_{\theta} \mathbf{t}$ given \bar{y} -tuple \mathbf{t} , we first use I to obtain a pointer to $R \bowtie \mathbf{t}[\bar{x} \cap \bar{y}]$ in $\mathcal{O}(1)$ time. Initialize (m_1, \dots, m_k) such that m_i is the smallest x_i -value in the group. Then, start enumerating $R \bowtie \mathbf{t}[\bar{x} \cap \bar{y}]$ with constant delay and in decreasing lexicographic order. Yield the current pair (\mathbf{s}, μ) that is being enumerated, provided that $\mathbf{s}[x_i] > \mathbf{t}[y_i]$ for all i . In contrast to the case where there is a single inequality, however, we cannot deduce that all subsequent \mathbf{s} will fail θ . The problem is then to find the next tuple in the group that occurs after \mathbf{s} in sorted order, but is larger than \mathbf{t} . Concretely, let i be the smallest index such that $\mathbf{s}(x_i) \leq \mathbf{t}(y_i)$ yet $\mathbf{s}(x_j) > \mathbf{t}(x_j)$ for all $j < i$. Let \mathbf{s}' be the tuple obtained from \mathbf{s} by setting $\mathbf{s}(x_j) := m_j$ for all $j \geq i$. Then, using binary search, find the next tuple that lexicographically larger or equal than \mathbf{s}' and recontinue enumeration from there. This binary search takes $\mathcal{O}(\log |R|)$ time, which causes the logarithmic delay.

Having multiple inequalities also complicates the update processing, since the sorted order no longer can be exploited to speed up join computation. In this case, therefore, we simply do a nested loop join per group, which yields a total access time of $\mathcal{O}(|R| \times |S|)$. Designing an join index with better access time is an interesting avenue for future work.

We conclude that in this case the enumeration delay is $\mathcal{O}(\log |R|)$, the access time is $\mathcal{O}(|R| \times |S|)$, and the update time is $\mathcal{O}(|\Delta R| \log(|R| + |\Delta R|))$.

Complexity of IEDYN By plugging in the above-mentioned delay into Proposition 3 and the above-mentioned access time and update time into Corollary 1, we obtain the following complexity of IEDYN.

Theorem 2 Assume that (T, N) is a plan in which all predicates are inequalities and let all enumeration and join indexes be as described above. Then, ENUM enumerates with delay $\mathcal{O}(|N| \log(\max_{r(\bar{x})} |db_{r(\bar{x})}|))$ and UPDATE processes updates in time $\mathcal{O}(|T| M^2)$ where $M = \max_{r(\bar{x}) \in \text{at}(T)} |db_{r(\bar{x})}| + |u_{r(\bar{x})}|$.⁴ If T is such that each edge is labeled by at most one predicate, then the enumeration delay is $\mathcal{O}(|N|)$ and update time is $\mathcal{O}(|T| M \log M)$. If T has no inequalities, the update time is $\mathcal{O}(|T| M)$.

A *simple* GJT is a GJT without predicates where $\text{var}(p) \subseteq \text{var}(n)$ for every node n with parent p . In a simple GJT, every child is hence a guard of its parent. For simple GJTs, the update processing time is optimal, in the following sense.

Theorem 3 If T is simple, then UPDATE processes updates in time $\mathcal{O}(|T| \max_{r(\bar{x})} |u_{r(\bar{x})}|)$, which is independent of $|db|$.

Proof Using the fact that every node is a guard of its parent, it is straightforward to prove by induction on the height of a

⁴ In the conference version of this paper [26], there was an incorrect claim: we stated that updates could be processed in time $\mathcal{O}(M \cdot \log(M))$ in the general case of multiple inequalities. We then found a bug in our proof and we currently do not know if this bound can be achieved.

node n in T (defined as the length of the shortest path from n to a leaf in T) that $|\Delta \rho_n| \leq \max_{r(\bar{x})} |u_{r(\bar{x})}|$, for each n . Since T does not contain any predicates, all join indexes that are created are as described in the paragraph “(1) No inequality” above. In particular, since every node is guard of its parent, we have that for every join index S_m by $(\text{pred}(p), \text{pred}(n), \text{var}(p))$ that is created we have $\text{var}(p) \subseteq \text{pred}(m)$. For this particular case, the access time to execute the semijoin $\pi_{\text{var}(p)}(\rho_m \bowtie \Delta_n)$ is $\mathcal{O}(|\Delta|_n) = \mathcal{O}(\max_{r(\bar{x})} |u_{r(\bar{x})}|)$. By now plugging in this access time and the linear index update time into Proposition 4, the result follows. \square

Berkholz et al. [8] show that, unless the Online Matrix-Vector Multiplication conjecture [24] is false, the class of conjunctive queries that allow both (1) constant delay enumeration of query results (in data complexity) and (2) update processing time that is linear in $|u|$ (again in data complexity) for every update u , is exactly the class of so-called *q-hierarchical queries*. While we forego a formal definition of this class, we show in [25] that a CQ Q is *q-hierarchical* if, and only if, there exists a plan (T, N) for Q such that T is simple. Since, by the results above, GDYN has both constant delay enumeration and update time $\mathcal{O}(|u|)$ (in data complexity) for exactly these queries, GDYN hence meets the theoretical lower bound.

5 Computing query plans

We say that (T, N) is a *plan* for GCQ Q , or that Q has plan (T, N) , if Q and $Q[T, N]$ are the same query, up to reordering of atoms and predicates, i.e., if $\#at(Q) = \#at(T)$, $\text{pred}(Q) = \text{pred}(T)$, and $\text{out}(Q) = \text{var}(N)$. Here, $\#at(X)$ denotes the multiset of atoms occurring in object X .

By the results of Sect. 4, it follows that we can dynamically process a given GCQ Q by first computing a plan for Q and subsequently applying GDYN on that plan. In this section we show *how* to compute a plan for Q by describing two algorithms.

1. The first algorithm computes a *GJT pair* for Q . Here, a *GJT pair* is a pair (T', N') defined exactly like a query plan, except that T' need not be binary and N' need not be sibling closed. A query plan is hence a particular kind of GJT pair. We call (T', N') a *GJT pair* for Q if $\#at(Q) = \#at(T')$, $\text{pred}(Q) = \text{pred}(T')$, and $\text{out}(Q) = \text{var}(N')$.
2. The second algorithm transforms this GJT pair into an equivalent query plan. Here, two GJT pairs (T, N) and (T', N') are *equivalent* if $\#at(T) = \#at(T')$, $\text{pred}(T) = \text{pred}(T')$, and $\text{var}(N) = \text{var}(N')$.

Clearly, the plan resulting from the composition of the two algorithms must be a plan for Q .

Before describing these two algorithms, it is important to emphasize that there are GCQs for which no GJT pair exists (and, consequently, for which no query plans exists). This is illustrated in Example 2 below. In particular, for full conjunctive queries (i.e., GCQs without θ -joins and projections), the results of [25] imply that a GJT pair exists for a full CQ Q if, and only if, Q is *acyclic*, a well-studied class of queries [1,47]. Similarly, the results imply that for conjunctive queries (with projections, but still without θ -joins) a GJT pair exists if and only if the query is *free-connex acyclic*, another well-studied class [6]. The existing definitions of acyclicity and free-connex acyclicity are given for CQs only. Given the previous discussion, we extend these notions to GCQs as follows.

Definition 6 A GCQ Q is *free-connex acyclic* if it has a GJT pair. It is *acyclic* if $\text{full}(Q)$ has a GJT pair. A GCQ that is not acyclic is *cyclic*.

In particular, every full GCQ that is acyclic is also free-connex acyclic. Also note that, since $\text{out}(\text{full}(Q)) = \text{var}(Q)$, a GJT pair will exist for $\text{full}(Q)$ if and only if there exists a GJT T for Q , i.e., a GJT with $\#at(Q) = \#at(T)$ and $\text{pred}(Q) = \text{pred}(T)$. Indeed, if T is a GCQ for Q , then (T, N) with N the set of all nodes in T , is a GJT pair for Q : clearly N is connex and $\text{var}(N) = \text{var}(T) = \text{out}(Q)$. For this reason, free-connex acyclicity is a stronger requirement than acyclicity: acyclicity only requires that a GJT for Q exists while free-connex acyclicity requires, in addition, that there exists a connex subset with $\text{out}(Q) = \text{var}(N)$.

Example 2 The trees T_1 and T_2 depicted in Fig. 4 are GTJs for the full GCQs

$$Q_1 = (r(x, y) \bowtie s(y, z) \bowtie t(z, v) \mid y < v), \quad \text{and} \\ Q_2 = (r(x, y) \bowtie s(y, z, w) \bowtie t(u, v) \mid x < z \wedge w < u),$$

respectively. These queries are hence acyclic. In contrast, $r(x, y) \bowtie s(y, z) \bowtie t(x, z)$ (also known as the triangle query) is the prototypical cyclic join query.

Let $Q'_2 = \pi_{y,z,w,u}(Q_2)$. Q'_2 is free-connex acyclic since the pair (T_2, N_2) of Fig. 4 is a GJT pair for Q_2 . In contrast, there is no GJT pair for $Q'_1 = \pi_{y,z}(Q_1)$ that contains tree T_1 . Indeed, observe that any connex set of T_1 must include the root, which includes $x \notin \text{out}(Q'_1)$. Finally, it can be verified that there is no GJT pair for $\pi_{x,v}(Q_1)$; this query is hence not free-connex acyclic.

Given that plans only exist for free-connex acyclic queries, it is desirable to be able to check free-connex acyclicity. In this respect, we develop an algorithm in Sect. 5.1 that checks whether Q is free-connex acyclic and if so computes a GJT pair for Q . This algorithm hence realizes step (1) above. Subsequently, step (2) above is realized in Sect. 5.2 where

we discuss how to transform GJT pairs into equivalent query plans.

5.1 Computing GJT pairs

The canonical algorithm for checking acyclicity of normal conjunctive queries is the GYO algorithm [1]. The algorithm described in this section is a generalization of the GYO algorithm that checks free-connex acyclicity in addition to normal acyclicity and deals with GCQs featuring θ -join predicates instead of CQs that have equality joins only. We first recall the classical GYO algorithm and then formulate its extension.

5.1.1 Classical GYO

The GYO algorithm operates on *hypergraphs*. A *hypergraph* H is a set of nonempty hyperedges. Recall from Sect. 3 that a hyperedge is just a finite set of variables. Every GCQ is associated with a hypergraph as follows.

Definition 7 Let Q be a GCQ. The hypergraph of Q , denoted $\text{hyp}(Q)$, is the hypergraph

$$\text{hyp}(Q) = \{\bar{x} \mid r(\bar{x}) \text{ is an atom of } Q \text{ with } \bar{x} \neq \emptyset\}.$$

The GYO algorithm checks acyclicity of a normal conjunctive query Q by constructing $\text{hyp}(Q)$ and repeatedly removing *ears* from this hypergraph. If ears can be removed until only the empty hypergraph remains, then the query is acyclic; otherwise, it is cyclic.

An *ear* in a hypergraph H is a hyperedge e for which we can divide its variables into two groups: (1) those that appear exclusively in e and (2) those that are contained in another hyperedge ℓ of H . A variable that appears exclusively in a single hyperedge is also called an *isolated variable*. Thus, ear removal corresponds to executing the following two reduction operations.

- Remove isolated variables: select a hyperedge e in H and remove isolated variables from it; if e becomes empty, remove e altogether from H .
- Subset elimination: remove hyperedge e from H if there exists another hyperedge ℓ for which $e \subseteq \ell$.

The *GYO reduction* of a hypergraph is the hypergraph that is obtained by executing these operations until no further operation is applicable. The following result is standard; see, e.g., [1] for a proof.

Proposition 6 A CQ Q is acyclic if and only if the GYO reduction of $\text{hyp}(Q)$ is the empty hypergraph.

5.1.2 GYO reduction for GCQs

In order to extend the GYO reduction to check free-connex acyclicity (not simply acyclicity) of GCQs (not simply standard CQs), we will: (1) redefine the notion of being an ear to take into account the predicates and (2) transform the GYO reduction into a two-stage procedure. The first stage allows to check that a connex set with exactly $out(Q)$ can exist while the first stage and second stage combined check that the query is acyclic.

Our algorithm operates on *hypergraph triplets* instead of hypergraphs, which are defined as follows.

Definition 8 A *hypergraph triplet* is a triple

$$\mathcal{H} = (hyp(\mathcal{H}), out(\mathcal{H}), pred(\mathcal{H}))$$

with $hyp(\mathcal{H})$ a hypergraph, $out(\mathcal{H})$ a hyperedge, and $pred(\mathcal{H})$ a set of predicates.

Intuitively, the variables in $out(\mathcal{H})$ will correspond to the output variables of a query and the set $pred(\mathcal{H})$ will contain predicates that need to be taken into account when removing ears. Every GCQ is, therefore, naturally associated with a hypergraph triplet as follows.

Definition 9 The hypergraph triplet of a GCQ Q , denoted $\mathcal{H}(Q)$, is the triplet $(hyp(Q), out(Q), pred(Q))$.

In order to extend the notion of an ear, we require the following definitions. Let \mathcal{H} be a hypergraph triplet. Variables that occur in $out(\mathcal{H})$ or in at least two hyperedges in $hyp(\mathcal{H})$ are called *equijoin variables* of \mathcal{H} . We denote the set of all equijoin variables of \mathcal{H} by $jev(\mathcal{H})$ and abbreviate $jev_{\mathcal{H}}(e) = e \cap jev(\mathcal{H})$. A variable x is *isolated* in \mathcal{H} if it is not an equijoin variable and is not mentioned in any predicate, i.e., if $x \notin jev(\mathcal{H})$ and $x \notin var(pred(\mathcal{H}))$. We denote the set of isolated variables of \mathcal{H} by $isol(\mathcal{H})$ and abbreviate $isol_{\mathcal{H}}(e) = e \cap isol(\mathcal{H})$. The *extended variables* of hyperedge e in \mathcal{H} , denoted $ext_{\mathcal{H}}(e)$ is the set of all variables of predicates that mention some variable in e , except the variables in e themselves:

$$ext_{\mathcal{H}}(e) = \bigcup \{var(\theta) \mid \theta \in pred(\mathcal{H}), var(\theta) \cap e \neq \emptyset\} \setminus e.$$

Finally, a hyperedge e is a *conditional subset* of hyperedge ℓ w.r.t. \mathcal{H} , denoted $e \sqsubseteq_{\mathcal{H}} \ell$, if $jev_{\mathcal{H}}(e) \subseteq \ell$ and $ext_{\mathcal{H}}(e \setminus \ell) \subseteq \ell$. We omit subscripts from our notation if the triplet is clear from the context.

Example 3 In Fig. 5, we depict several hypergraph triplets. There, hyperedges in \mathcal{H} are depicted by colored regions and variables in $out(\mathcal{H})$ are underlined. We use dashed lines to connect variables that appear together in a predicate. So, in \mathcal{H}_1 , we have predicates θ_1, θ_2 with $var(\theta_1) = \{t, v\}$ and

$var(\theta_2) = \{x, y\}$. Now consider triplet \mathcal{H}_1 in particular. It is the hypergraph triplet $\mathcal{H}(Q)$ for the following GCQ Q :

$$Q = \pi_{t,u,z,w}(r_1(s, t, u) \bowtie r_2(t, u) \bowtie r_3(u, w, x) \bowtie r_4(s, v) \bowtie r_5(w, z, y) \mid t < v \wedge x < y).$$

Moreover, $jev(\mathcal{H}_1) = \{s, t, u, w, z\}$ and $isol(\mathcal{H}_1) = \emptyset$. Furthermore, $ext_{\mathcal{H}_1}(\{v\}) = \{t\}$ since $\theta_1 = t < v$ shares variables with $\{v\}$. Finally $jev_{\mathcal{H}_1}(\{s, v\}) = \{s\} \subseteq \{s, t, u\}$ and $ext_{\mathcal{H}_1}(\{s, v\} \setminus \{s, t, u\}) = ext_{\mathcal{H}_1}(\{v\}) = \{t\} \subseteq \{s, t, u\}$. Therefore, $\{s, v\} \sqsubseteq_{\mathcal{H}_1} \{s, t, u\}$. Similarly, $\{t, u\} \sqsubseteq_{\mathcal{H}_1} \{s, t, u\}$.

We define ears in our context as follows.

Definition 10 A hyperedge e is an ear in a hypergraph triplet \mathcal{H} if $e \in hyp(\mathcal{H})$ and either

1. We can divide its variables into two: (a) those that are isolated and (b) those that form a conditional subset of another hyperedge $\ell \in hyp(\mathcal{H}) \setminus \{e\}$; or
2. e consists only of nonjoin variables, i.e., $jev(e) = \emptyset$ and $ext(e) = \emptyset$.

Note that case (2) allows for $\theta \in pred(\mathcal{H})$ with $var(\theta) \subseteq e$. We call predicates that are covered by a hyperedge in this sense *filters* because they correspond to filtering a single GMR instead of θ -joining two GMRs. If, in case (2), there is no filter θ with $var(\theta) \subseteq e$, then $e = isol_{\mathcal{H}}(e)$. Similar to the classical GYO reduction, we can view ear removal as a rewriting process on triplets, where we consider the following reduction operations.

- (ISO) Remove isolated variables: select a hyperedge $e \in hyp(\mathcal{H})$ and remove a nonempty set $X \subseteq isol_{\mathcal{H}}(e)$ from it. If e becomes empty, remove it from $hyp(\mathcal{H})$.
- (CSE) Conditional subset elimination: remove hyperedge e from $hyp(\mathcal{H})$ if it is a conditional subset of another hyperedge f in $hyp(\mathcal{H})$. Also update $pred(\mathcal{H})$ by removing all predicates θ with $var(\theta) \cap (e \setminus f) \neq \emptyset$.
- (FLT) Filter elimination: select $e \in hyp(\mathcal{H})$ and a nonempty subset of predicates $\Theta \subseteq pred(\mathcal{H})$ with $var(\Theta) \subseteq e$. Remove all predicates in Θ from $pred(\mathcal{H})$.

We write $\mathcal{H} \rightsquigarrow \mathcal{I}$ to denote that triplet \mathcal{I} is obtained from triplet \mathcal{H} by applying a single such operation, and $\mathcal{H} \rightsquigarrow^* \mathcal{I}$ to denote that \mathcal{I} is obtained by a sequence of zero or more of such operations.

Example 4 For the hypergraph triplets illustrated in Fig. 5, we have $\mathcal{H}_1 \rightsquigarrow \mathcal{H}_2 \rightsquigarrow \mathcal{H}_3 \rightsquigarrow \mathcal{H}_4$ and $\mathcal{H}_5 \rightsquigarrow \mathcal{H}_6 \rightsquigarrow \mathcal{H}_7 \rightsquigarrow \mathcal{H}_8 \rightsquigarrow \mathcal{H}_9 \rightsquigarrow \mathcal{H}_{10} \rightsquigarrow \mathcal{H}_{11}$. For each reduction, it is illustrated in the figure which set of isolated variables is removed, or which conditional subset is removed.




Fig. 5 Illustration of GYO reduction for GCQs. Colored regions depict hyperedges. Variables in *out* are underlined. Variables occurring in the same predicate are connected by dashed lines (color figure online)

We write $\mathcal{H} \downarrow$ to denote \mathcal{H} is in *normal form*, i.e., that no operation is applicable on triplet \mathcal{H} . Note that, because each operation removes at least one variable, hyperedge, or predicate, we will always reach a normal form after a finite number of operations. Furthermore, while multiple different reduction steps may be applicable on a given triplet \mathcal{H} , the order in which we apply them does not matter:

Proposition 7 (Confluence) *Whenever $\mathcal{H} \rightsquigarrow^* \mathcal{I}_1$ and $\mathcal{H} \rightsquigarrow^* \mathcal{I}_2$, there exists \mathcal{J} such that $\mathcal{I}_1 \rightsquigarrow^* \mathcal{J}$ and $\mathcal{I}_2 \rightsquigarrow^* \mathcal{J}$.*

Because the proof is technical but not overly enlightening, we defer it to “Appendix B.1”. A direct consequence is that normal forms are unique: if $\mathcal{H} \rightsquigarrow^* \mathcal{I}_1 \downarrow$ and $\mathcal{H} \rightsquigarrow^* \mathcal{I}_2 \downarrow$, then $\mathcal{I}_1 = \mathcal{I}_2$.

Let \mathcal{H} be a triplet. The residual of \mathcal{H} , denoted $\tilde{\mathcal{H}}$, is the triplet $(\text{hyp}(\mathcal{H}), \emptyset, \text{pred}(\mathcal{H}))$, i.e., the triplet where $\text{out}(\mathcal{H})$ is set to \emptyset . A triplet is *empty* if it equals $(\emptyset, \emptyset, \emptyset)$.

Our main result in this section states that to check whether a GCQ Q is free-connex acyclic, it suffices to start from $\mathcal{H}(Q)$ and do a two-stage reduction: the first from $\mathcal{H}(Q)$ until a normal form $\mathcal{I} \downarrow$ is reached, and the second from the residual of $\mathcal{I} \downarrow$, until another normal form \mathcal{J} is reached.⁵

Theorem 4 *Let Q be a GCQ. Assume $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$. Then, the following hold.*

1. Q is acyclic if, and only if, \mathcal{J} is the empty triplet.

⁵ Note that because we set $\text{out}(\mathcal{I}) = \emptyset$ on the residual, new variables may become isolated and therefore more reductions steps may be possible on the normal form of \mathcal{I} .

2. Q is free-connex acyclic if, and only if, \mathcal{J} is the empty triplet and $\text{var}(\text{hyp}(\mathcal{I})) = \text{out}(Q)$.
3. For every GJT T of Q and every connex subset N of T , it holds that $\text{var}(\text{hyp}(\mathcal{I})) \subseteq \text{var}(N)$.

We devote Sect. 5.1.3 to the proof of this theorem.

Example 5 Figure 5 illustrates the two-stage sequence of reductions starting from $\mathcal{H}(Q)$ with Q the GCQ of Example 3. Note that $\mathcal{H}(Q) = \mathcal{H}_1$ and \mathcal{H}_5 is the residual of \mathcal{H}_4 . Because we end with the empty triplet Q is acyclic, but it is not free-connex since $\text{out}(Q) \subsetneq \text{var}(\mathcal{H}_4)$.

Theorem 4 gives us a decision procedure for checking free-connex acyclicity of GCQ Q . From its proof in Sect. 5.1.3, we can actually derive an algorithm for constructing a GJT pair for Q . At its essence, this algorithm starts with the set of atoms appearing in Q and subsequently uses the sequence of reduction steps from Theorem 4 to construct a GJT from it, at the same time checking free-connex acyclicity. Every reduction step causes new nodes to be added to the partial GJT constructed so far. We will refer to such partial GJTs as *generalized join forests* (GJF).

Definition 11 [GJF] A *generalized join forest* is a set F of pairwise disjoint GJTs s.t. for distinct trees $T_1, T_2 \in F$, we have $\text{var}(T_1) \cap \text{var}(T_2) = \text{var}(n_1) \cap \text{var}(n_2)$ where n_1 and n_2 are the roots of T_1 and T_2 .

Every GJF encodes a hypergraph as follows.

Definition 12 The hypergraph $\text{hyp}(F)$ associated with GJF F is the hypergraph that has one hyperedge for every nonempty root node in F ,

$hyp(F) = \{var(n) \mid n \text{ root node in } F, var(n) \neq \emptyset\}$.

The GJT construction algorithm does not manipulate hypergraph triplets directly. Instead, it manipulates *GJF triplets*. A GJF triplet is defined like a hypergraph triplet, except that it has a GJF instead of a hypergraph.

Definition 13 A *GJF triplet* is a triple $\mathbb{F} = (forest(\mathbb{F}), out(\mathbb{F}), \Theta_{\mathbb{F}})$ with $forest(\mathbb{F})$ a GJF, $out(\mathbb{F})$ a hyperedge, and $\Theta_{\mathbb{F}}$ a set of predicates. Every GJF triplet \mathbb{F} induces a hypergraph triplet $\mathcal{H}(\mathbb{F}) = (hyp(forest(\mathbb{F})), out(\mathbb{F}), \Theta_{\mathbb{F}})$.

The algorithm for constructing a GJT pair for a given GCQ Q is now shown in Algorithm 2. It starts in line 2 by initializing the GJF triplet \mathbb{F} to $\mathbb{F} = (forest(Q), out(Q), pred(Q))$. Here, $forest(Q)$ is the GJF obtained by creating, for every atom $r(\bar{x})$ that occurs $k > 0$ times in Q , k corresponding leaf nodes labeled by $r(\bar{x})$. In lines 3–4, Algorithm 2 then performs the first phase of reduction steps of Theorem 4. To this end, it checks whether a reduction operation is applicable to $\mathcal{H}(\mathbb{F})$ and, if so, *enacts* this operation by modifying \mathbb{F} as follows.

- (ISO) If the reduction operation on the hypergraph triplet $\mathcal{H}(\mathbb{F})$ is to remove a nonempty subset X of isolated variables from hyperedge e , then \mathbb{F} is modified as follows. Let n_1, \dots, n_k be all the root nodes in $forest(\mathbb{F})$ that are labeled by e . Merge the corresponding trees into one tree by creating a new node n with $var(n) = e$ and attaching n_1, \dots, n_k as children to it with $pred(n \rightarrow n_i) = \emptyset$ for $1 \leq i \leq k$. Then, enact the removal of X by creating a new node p with $var(p) = e \setminus X$ and attaching n as child to it with $pred(p \rightarrow n) = \emptyset$.
- (CSE) If the reduction operation on $\mathcal{H}(\mathbb{F})$ is to remove a hyperedge e because it is a conditional subset of another hyperedge ℓ , then \mathbb{F} is modified as follows. Let n_1, \dots, n_k (resp. m_1, \dots, m_l) be all the root nodes in $forest(\mathbb{F})$ that are labeled by e (resp. ℓ) and let T_1, \dots, T_k (resp. U_1, \dots, U_l) be their corresponding trees. Similar to the previous case, merge the T_i (resp. U_j) into a single tree with new root n labeled by e (resp. m labeled by ℓ). Then, enact the removal of e by creating a new node p with $var(p) = \ell$ and attaching n and m as children with $pred(p \rightarrow n) = \{\theta \in pred(\mathbb{F}) \mid var(\theta) \cap (e \setminus \ell) \neq \emptyset\}$ and $pred(p \rightarrow m) = \emptyset$.
- (FLT) If the reduction operation on $\mathcal{H}(\mathbb{F})$ is to remove nonempty set of predicates Θ because there exists a hyperedge e with $var(\Theta) \subseteq e$, then \mathbb{F} is modified as follows. Let n_1, \dots, n_k be all the root nodes in $forest(\mathbb{F})$ that are labeled by e . Merge the corresponding trees into one tree by creating a new root n labeled by e and attaching n_1, \dots, n_k as children with $pred(n \rightarrow n_i) = \Theta$. Enact the removal of Θ by removing all $\theta \in \Theta$ from $\Theta(\mathbb{F})$.

Algorithm 2 Compute a GJT pair

```

1: Input: A GCQ  $Q$ .
2:  $\mathbb{F} \leftarrow (forest(Q), out(Q), pred(Q))$ 
3: while a reduction step is applicable to  $\mathcal{H}(\mathbb{F})$  do
4:   enact the reduction on  $\mathbb{F}$ 
5:  $X \leftarrow$  set of all root nodes in  $\mathbb{F}$ 
6: set  $pred(\mathbb{F}) := \emptyset$ 
7: while a reduction step is applicable to  $\mathcal{H}(\mathbb{F})$  do
8:   enact the reduction on  $\mathbb{F}$ 
9: if  $\mathcal{H}(\mathbb{F})$  is not the empty triplet then
10:   error " $Q$  is not acyclic"
11: else
12:    $T \leftarrow$  tree obtained by connecting all root nodes of  $\mathbb{F}$ 's forest to
       a new root, labeled by  $\emptyset$ 
13:    $N \leftarrow$  all nodes in  $X$  and their ancestors in  $T$ 
14:   return  $(T, N)$ 

```

It is straightforward to check that these modifications of the forest triplet \mathbb{F} faithfully enact the corresponding operations on $\mathcal{H}(\mathbb{F})$, in the following sense.

Lemma 4 Let \mathbb{F} be a forest triplet and assume $\mathcal{H}(\mathbb{F}) \rightsquigarrow \mathcal{I}$. Let \mathbb{G} be the result of enacting this reduction operation on \mathbb{F} . Then \mathbb{G} is a valid forest triplet and $\mathcal{H}(\mathbb{G}) = \mathcal{I}$.

We continue the explanation of Algorithm 2. In line 5, Algorithm 2 records the set of root nodes obtained after the first stage of reductions. It then sets $out(\mathbb{F}) = \emptyset$ in line 6 and continues with the second stage of reductions in lines 7, 8. It then employs Theorem 4 to check acyclicity of Q . If Q is not acyclic, it reports this in lines 9, 10. If Q is acyclic, then we know by Theorem 4 that $\mathcal{H}(\mathbb{F})$ has become the empty triplet. Note that $\mathcal{H}(\mathbb{F})$ can be empty only if all the roots of \mathbb{F} 's join forest are labeled by the empty set of variables. As such, we can transform this forest into a join tree T by linking all of these roots to a new unique root, also labeled \emptyset . This is done in line 12. In line 13, the set of nodes N is computed and consists of all nodes identified at the end of the first stage (line 5) plus all of their parents in T .

We will prove in Sect. 5.1.3 that Algorithm 2 is correct, in the following sense.

Theorem 5 Given a GCQ Q , Algorithm 2 reports an error if Q is cyclic. Otherwise, it returns a GJT pair (T, N) with T a GJT for Q . If Q is free-connex acyclic, then (T, N) is GJT pair for Q . Otherwise, $out(Q) \subsetneq var(N)$, but $var(N)$ is minimal in the sense that for every other GJT pair (T', N') with T' a GJT for Q , we have $var(N) \subseteq var(N')$.

It is straightforward to check that this algorithm runs in polynomial time in the size of Q .

Example 6 In Fig. 6, we show a GJT T and use this GJT to illustrate a number of GJFs F_1, \dots, F_{10} in the following way: let level 1 be the leaf nodes, level 2 the parents of the leaves, and so on. Then, we take GJF F_i to be the set of all trees




Fig. 6 GJT Construction by GYO reduction

rooted at nodes at level i , for $1 \leq i \leq 10$, and with each level i , we mention the set of remaining predicates θ_i for $1 \leq i \leq k$ where k is the number of predicates in Q . Nodes (resp. predicates with each F_i) labeled by “•” in Fig. 6 indicates that the node (and hence tree, resp. predicates) was already present in F_{i-1} and did not change. These should hence not be interpreted as new nodes (resp. predicates changed). With this coding of forests, it is easy to see that for all $1 \leq i \leq 9$, $F_i = \text{hyp}(\mathcal{H}_i)$ with \mathcal{H}_i illustrated in Fig. 5 (note here that the hypergraph of residual of \mathcal{H}_4 , i.e., \mathcal{H}_5 is the same as \mathcal{H}_4 , hence we do not show the corresponding F_5). Furthermore, $\text{pred}(F_i) = \text{pred}(Q) \setminus \text{pred}(\mathcal{H}_i)$ with Q the GCQ from Example 3. As such, the tree illustrates the sequence of GJF triplets that is obtained by enacting the hypergraph reductions illustrated in Fig. 5. For example, let $\mathbb{F}_1 = (F_1, \text{out}(Q), \text{pred}(Q))$. After enacting the removal of hyperedge $\{t, u\}$ from \mathcal{H}_1 to obtain \mathcal{H}_2 , we obtain $\mathbb{F}_2 = (F_2, \text{out}(Q), \text{pred}(Q))$. Here, F_2 is obtained by merging the single-node trees (i.e., labeled by the atoms in Q) $\{s, t, u\}$ and $\{t, u\}$ in to a single tree with root $\{s, t, u\}$. The shaded area illustrates the nodes in the connex subset N computed by Algorithm 2.

We stress that Algorithm 2 is nondeterministic in the sense that the pair (T, N) returned depends on the order in which the reduction operations are performed.

5.1.3 Correctness

To prove Theorems 4 and 5 we show some propositions.

Proposition 8 *Let Q be a GCQ. Assume $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$. If \mathcal{J} is the empty triplet, then, when run on Q ,*

Algorithm 2 returns a pair (T, N) s.t. T is a GJT for Q and $\text{var}(N) = \text{var}(\text{hyp}(\mathcal{I}))$.

Proof Assume that \mathcal{J} is the empty triplet. Algorithm 2 starts in line 3 by initializing $\mathbb{F} = (\text{forest}(Q), \text{out}(Q), \text{pred}(Q))$. Clearly, $\mathcal{H}(\mathbb{F}) = \mathcal{H}(Q)$ at this point. Algorithm 2 subsequently modifies \mathbb{F} throughout its execution. Let \mathbb{H} denote the initial version of \mathbb{F} ; let \mathbb{I} denote the version of \mathbb{F} when executing line 5; let $\tilde{\mathbb{I}}$ denote the version of \mathbb{F} after executing line 6 and let \mathbb{J} denote the version of \mathbb{F} when executing line 9. By repeated application of Lemma 4, we know that $\mathcal{H}(Q) = \mathcal{H}(\mathbb{H}) \rightsquigarrow^* \mathcal{H}(\mathbb{I})$. Furthermore, $\mathcal{H}(\mathbb{I})$ is in normal form. Since also $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and normal forms are unique, $\mathcal{H}(\mathbb{I}) = \mathcal{I}$. Therefore, $\mathcal{H}(\tilde{\mathbb{I}}) = \tilde{\mathcal{I}}$. Again by repeated application of Lemma 4, we know that $\tilde{\mathcal{I}} = \mathcal{H}(\tilde{\mathbb{I}}) \rightsquigarrow^* \mathcal{H}(\mathbb{J})$. Moreover, $\mathcal{H}(\mathbb{J})$ is in normal form. Since also $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$ and normal forms are unique, $\mathcal{H}(\mathbb{J}) = \mathcal{J}$. As \mathcal{J} is empty, we will execute lines 12–14. Since \mathcal{J} is the empty hypergraph triplet, every root of every tree in $\text{forest}(\mathbb{J})$ must be labeled by \emptyset . By definition of join forests, no two distinct trees in $\text{forest}(\mathbb{J})$ hence share variables. As such, the tree T obtained in line 12 by linking all of these roots to a new unique root, also labeled \emptyset , is a valid GJT.

We claim that T is a GJT for Q . Indeed, observe that $\text{at}(T) = \text{at}(Q)$ and the number of times that an atom occurs in Q equals the number of times that it occurs as a label in T . This is because initially $\text{forest}(\mathbb{H}) = \text{forest}(Q)$ and by enacting reduction steps, we never remove nor add nodes labeled by atoms. Furthermore $\text{pred}(T) = \text{pred}(Q)$. This is because initially $\text{pred}(\mathbb{H}) = \text{pred}(Q)$ yet $\Theta_{\mathbb{J}}$ is empty. This means that, for every $\theta \in \text{pred}(Q)$, there was some reduction step that removed θ from the set of predicates of the current GJF triplet \mathbb{F} . However, when enacting reduction steps we only remove predicates after we have added them to $\text{forest}(\mathbb{F})$. Therefore, every predicate in $\text{pred}(Q)$ must occur in T . Conversely, during enactment of reduction steps, we never add predicates to $\text{forest}(\mathbb{F})$ that are not in $\Theta_{\mathbb{F}}$, so all predicates in T are also in $\text{pred}(Q)$. Thus, T is a GJT for Q .

It remains to show that N is a connex subset of T and $\text{var}(N) = \text{var}(\text{hyp}(\mathcal{I}))$. To this end, let X be the set of all root nodes of $\text{forest}(\mathbb{I})$, as computed in Line 5. Since \mathbb{J} is obtained from $\tilde{\mathbb{I}}$ by a sequence of reduction enactments, and since such enactments only add new nodes and never delete them, X is a subset of nodes of $\text{forest}(\mathbb{J})$ and therefore also of T . As computed in line 5, N consists of X and all ancestors of nodes of X in T . Then, N is a connex subset of T by definition. Furthermore, since $\mathcal{H}(\mathbb{I}) = \mathcal{I}$, $\text{hyp}(\text{forest}(\mathbb{I})) = \text{hyp}(\mathcal{I})$. Thus, $\text{var}(X) = \text{var}(\text{hyp}(\mathbb{I})) = \text{var}(\text{hyp}(\mathcal{I}))$. Hence, to establish that $\text{var}(N) = \text{var}(\text{hyp}(\mathcal{I}))$, it suffices to show that $\text{var}(X) = \text{var}(N)$. Since $X \subseteq N$ the inclusion $\text{var}(X) \subseteq \text{var}(N)$ is immediate. To also establish $\text{var}(N) \subseteq \text{var}(X)$, let n be an arbitrary but fixed node in N . If $n \in X$, then clearly $\text{var}(n) \subseteq \text{var}(X)$. If $n \notin X$, then n was created during the

sequence of reduction enactments in that transform $\tilde{\mathbb{I}}$ into \mathbb{J} . Now note that, whenever a new node m is created during a reduction enactment on a GJF \mathbb{G} , there exists a root node of $\text{forest}(\mathbb{G})$ that contains all variables of m . From this observation and the fact that n was created during a sequence of reduction enactments that start from $\tilde{\mathbb{I}}$, it follows that there is some root node r in $\tilde{\mathbb{I}}$ with $\text{var}(n) \subseteq \text{var}(r)$. Then, because X contains all root nodes of $\tilde{\mathbb{I}}$, also $\text{var}(n) \subseteq \text{var}(X)$. Therefore, $\text{var}(N) = \text{var}(X) = \text{var}(\text{hyp}(\mathcal{I}))$. \square

Corollary 1 (Soundness) *Let Q be a GCQ and assume that $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$. Then:*

1. *If \mathcal{J} is the empty triplet then Q is acyclic.*
2. *If \mathcal{J} is the empty triplet and $\text{var}(\text{hyp}(\mathcal{I})) = \text{out}(Q)$, then Q is free-connex acyclic.*

To also show completeness, we will interpret a GJT T for a GCQ Q as a “parse tree” that specifies the two-stage sequence of reduction steps that can be done on $\mathcal{H}(Q)$ to reach the empty triplet. Not all GJTs will allow us to do so easily, however, and we will, therefore, restrict our attention to those GJTs that are *canonical*.

Definition 14 (Canonical) A GJT T is *canonical* if:

1. its root is labeled by \emptyset ;
2. every leaf node n is the child of an internal node m with $\text{var}(n) = \text{var}(m)$;
3. for all internal nodes n and m with $n \neq m$, we have $\text{var}(n) \neq \text{var}(m)$; and
4. for every edge $m \rightarrow n$ and all $\theta \in \text{pred}(m \rightarrow n)$, we have $\text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) \neq \emptyset$.

A connex subset N of T is *canonical* if every node in it is an interior node of T . A GJT pair (T, N) is canonical if both T and N are canonical.

The following proposition, proven in “Appendix B”, shows that we may restrict our attention to canonical GJT pairs without loss of generality.

Proposition 9 *For every GJT pair, there exists an equivalent canonical pair.*

We also require the following auxiliary notions and insights. First, if (T, N) is a GJT pair, then define the *hypergraph associated with* (T, N) , denoted $\text{hyp}(T, N)$, to be the hypergraph formed by node labels in N ,

$$\text{hyp}(T, N) = \{\text{var}_T(n) \mid n \in N, \text{var}_T(n) \neq \emptyset\}.$$

Further, define $\text{pred}(T, N)$ to be the set of all predicates occurring on edges between nodes in N . For a hyper-edge z , define the *hypergraph triplet* of (T, N) w.r.t. z ,

denoted $\mathcal{H}(T, N, \bar{z})$ to be the hypergraph triplet $(\text{hyp}(T, N), \bar{z}, \text{pred}(T, N))$.

The following technical Lemma shows that we can use canonical pairs as “parse” trees to derive a sequence of reduction steps. Its proof can be found in “Appendix B”.

Lemma 5 *Let (T, N_1) and (T, N_2) be canonical GJT pairs with $N_2 \subseteq N_1$. Then, $\mathcal{H}(T, N_1, \bar{z}) \rightsquigarrow^* \mathcal{H}(T, N_2, \bar{z})$ for every $\bar{z} \subseteq \text{var}(N_2)$.*

We require the following additional lemma, proven in “Appendix B”:

Lemma 6 *Let H_1 and H_2 be two hypergraphs such that for all $e \in H_2$ there exists $\ell \in H_1$ such that $e \subseteq \ell$. Then, $(H_1 \cup H_2, \bar{z}, \Theta) \rightsquigarrow^* (H_1, \bar{z}, \Theta)$, for every hyperedge \bar{z} and set of predicates Θ .*

We these tools in hand we can prove completeness.

Proposition 10 *Let Q be a GCQ, let T be a GJT for Q and let N be a connex subset of T with $\text{out}(Q) \subseteq \text{var}(N)$. Assume that $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$. Then, \mathcal{J} is the empty triplet and $\text{var}(\text{hyp}(\mathcal{I})) \subseteq \text{var}(N)$.*

Proof By Proposition 9, we may assume without loss of generality that (T, N) is a canonical GJT pair. Let A be the set of all of T ’s interior nodes. Clearly, A is a connex subset of T and $\text{var}(A) \subseteq \text{var}(Q)$. Furthermore, because for every atom $r(\bar{x})$ in Q there is a leaf node l in T labeled by $r(\bar{x})$ (as T is a GJT for Q), which has a parent interior node n_l labeled \bar{x} (because T is canonical), also $\text{var}(Q) \subseteq \text{var}(A)$. Therefore, $\text{var}(A) = \text{var}(Q)$. By the same reasoning, $\text{hyp}(Q) \subseteq \text{hyp}(T, A)$. Therefore, $\text{hyp}(T, A) = \text{hyp}(T, A) \cup \text{hyp}(Q)$. Furthermore, because every interior node in a GJT has a guard descendant, and the leaves of T are all labeled by atoms in Q , we know that for every node $n \in A$ there exists some hyperedge $f \in \text{hyp}(Q)$ such that $\text{var}(n) \subseteq \text{var}(f)$. In addition, we claim that $\text{pred}(T, A) = \text{pred}(Q)$. Indeed, $\text{pred}(T, A) \subseteq \text{pred}(Q)$ since T is a GJT for Q . The converse inclusion follows from canonicity properties (2) and (4): because leaf nodes in a canonical GJT have a parent labeled by the same hyperedge, there can be no predicates on edges to leaf nodes in T . Thus, all predicates in T are on edges between interior nodes, i.e., in $\text{pred}(T, A)$. Then, because every predicate in Q appears somewhere in T (since T is a GJT for Q), we have $\text{pred}(Q) \subseteq \text{pred}(T, A)$. From all of the observations made so far and Lemma 6, we obtain:

$$\begin{aligned} \mathcal{H}(T, A, \text{out}(Q)) &= (\text{hyp}(T, A), \text{out}(Q), \text{pred}(T, A)) \\ &= (\text{hyp}(T, A) \cup \text{hyp}(Q), \text{out}(Q), \text{pred}(T, A)) \\ &\rightsquigarrow^* (\text{hyp}(Q), \text{out}(Q), \text{pred}(T, A)) \\ &= (\text{hyp}(Q), \text{out}(Q), \text{pred}(Q)) = \mathcal{H}(Q) \end{aligned}$$

Thus $\mathcal{H}(T, A, \text{out}(Q)) \rightsquigarrow^* \mathcal{H}(Q) \rightsquigarrow^* \mathcal{I}$. Furthermore, because (T, N) is also canonical with $N \subseteq A$ and $\text{out}(Q) \subseteq \text{var}(N)$, we have $\mathcal{H}(T, A, \text{out}(Q)) \rightsquigarrow^* \mathcal{H}(T, N, \text{out}(Q))$ by Lemma 5. Then, because reduction is confluent (Proposition 7), we obtain that $\mathcal{H}(T, N, \text{out}(Q))$ and \mathcal{I} can be reduced to the same triplet. Because \mathcal{I} is in normal form, necessarily $\mathcal{H}(T, N, \text{out}(Q)) \rightsquigarrow^* \mathcal{I}$. Since reduction steps can only remove nodes and hyperedges (and never add them), $\text{var}(\text{hyp}(\mathcal{I})) \subseteq \text{var}(N)$.

It remains to show that \mathcal{J} is the empty triplet. Hereto, first verify the following. For any hypergraph triplets \mathcal{U} and \mathcal{V} , if $\mathcal{U} \rightsquigarrow^* \mathcal{V}$ then also $\tilde{\mathcal{U}} \rightsquigarrow^* \tilde{\mathcal{V}}$. From this, $\mathcal{H}(T, A, \text{out}(Q)) \rightsquigarrow^* \mathcal{I}$, and the fact that $\mathcal{H}(T, A, \emptyset)$ is the residual of $\mathcal{H}(T, A, \text{out}(Q))$, we conclude $\mathcal{H}(T, A, \emptyset) \rightsquigarrow^* \tilde{\mathcal{I}}$. Then, because $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J}$, it follows that $\mathcal{H}(T, A, \emptyset) \rightsquigarrow^* \mathcal{J}$. Let r be T 's root node, which is labeled by \emptyset since T is canonical. Then, $\{r\}$ is a connex subset of T . By Lemma 5, $\mathcal{H}(T, A, \emptyset) \rightsquigarrow^* \mathcal{H}(T, \{r\}, \emptyset)$. Now observe that the hypergraph of $\mathcal{H}(T, \{r\}, \emptyset)$ is empty and its predicate set is also empty. Therefore, $\mathcal{H}(T, \{r\}, \emptyset)$ is the empty hypergraph triplet. In particular, it is in normal form. But, since \mathcal{J} is also in normal form and normal forms are unique, \mathcal{J} must also be the empty triplet. \square

Corollary 2 (Completeness) *Let Q be a GCQ. Assume that $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I} \downarrow$ and $\tilde{\mathcal{I}} \rightsquigarrow^* \mathcal{J} \downarrow$.*

1. *If Q is acyclic, then \mathcal{J} is the empty triplet.*
2. *If Q is free-connex acyclic, then \mathcal{J} is the empty triplet and $\text{var}(\text{hyp}(\mathcal{I})) = \text{out}(Q)$.*
3. *For every GJT T of Q and every connex subset N of T it holds that $\text{var}(\text{hyp}(\mathcal{I})) \subseteq \text{var}(N)$.*

Proof (1) Since Q is acyclic, there exists a GJT T for Q . Let N be the set of all of T 's nodes. Then, N is a connex subset of T and $\text{out}(Q) \subseteq \text{var}(N) = \text{var}(Q)$. The result then follows from Proposition 10.

(2) Since Q is free-connex acyclic, there exists a GJT pair (T, N) compatible with Q . In particular, $\text{var}(N) = \text{out}(Q)$. By Proposition 10, \mathcal{J} is the empty triplet, and $\text{var}(\text{hyp}(\mathcal{I})) \subseteq \text{var}(N) = \text{out}(Q)$. It remains to show $\text{out}(Q) \subseteq \text{var}(\text{hyp}(\mathcal{I}))$. First, verify the following: A reduction step on a hypergraph triplet \mathcal{H} never removes any variable in $\text{out}(\mathcal{H})$ from $\text{hyp}(\mathcal{H})$, nor does it modify $\text{out}(\mathcal{H})$. Then, since $\text{out}(\mathcal{H}(Q)) = \text{out}(Q) \subseteq \text{var}(Q) \subseteq \text{var}(\text{hyp}(\mathcal{H}(Q)))$, and $\mathcal{H}(Q) \rightsquigarrow^* \mathcal{I}$, we obtain $\text{out}(Q) \subseteq \text{var}(\text{hyp}(\mathcal{I}))$.

(3) Follows directly from Proposition 10. \square

Theorem 4 follows directly from Corollaries 1 and 2. Theorem 5 follows from Theorem 4 and Proposition 8.

5.2 Transforming GJT pairs to query plans

Let us call a GJT pair (T, N) *binary* if T is binary and *sibling closed* if N is sibling closed. A query plan is hence a binary and sibling-closed GJT pair. In this section, we prove the following result.

Proposition 11 *Every GJT pair can be transformed in polynomial time into an equivalent plan.*

We prove Proposition 11 in two steps. First, we show that any pair (T, N) can be transformed in polynomial time into an equivalent sibling-closed pair. Next, we show that any sibling-closed GJT pair (T, N) can be converted in polynomial time into an equivalent plan. Proposition 11 hence follows by composing these two transformations. Throughout this section, let $\text{ch}_T(n)$ denote the set of children of n in T .

Sibling-closed transformation We say that $n \in T$ is a *violation node* in a GJT pair (T, N) if $n \in N$ and some, but not all children of n are in N . A violator is of *type 1* if some node in $\text{ch}_T(n) \cap N$ is a guard of n . It is of *type 2* otherwise. We now define two operations on (T, N) that remove violators of type 1 and type 2, respectively. The sibling-closed transformation is then obtained by repeatedly applying these operators until all violators are removed.

The first operator is applicable when n is a type 1 violator. It returns the pair (T', N') obtained as follows:

- Since n is a type 1 violator, some $g \in \text{ch}_T(n) \cap N$ is a child guard of n (i.e., $\text{var}(n) \subseteq \text{var}(g)$).
- Because every node has a guard, there is some leaf node l that is a descendant guard of g (i.e., $\text{var}(g) \subseteq \text{var}(l)$). Possibly, l is g itself.
- Now create a new node p between node l and its parent with label $\text{var}(p) = \text{var}(l)$. Since l is a descendant guard of n and g , p becomes a descendant guard of n and g as well. Detach all nodes in $\text{ch}_T(n) \setminus N$ from n and attach them as children to p , preserving their edge labels. This effectively moves all subtrees rooted at nodes in $\text{ch}_T(n) \setminus N$ from n to p . Denote by T' the final result.
- If l was not in N , then $N' = N$. Otherwise, $N' = N \setminus \{l\} \cup \{p\}$.

We write $(T, N) \xrightarrow{1,n} (T', N')$ to indicate that (T', N') can be obtained by applying the above-described operation on node n .

Example 7 Consider the GJT pair (T, N) from Fig. 7 where N is indicated by the nodes in the shaded area. Let us denote the root node by n and its guard child with label $\{y, z, w\}$ by g . The node $l = h(y, z, w, t)$ is a descendant guard of g . Since $s(y, z, m)$ is not in N , n is violator of type 1. After




Fig. 7 Illustration of the sibling-closed transform: removal of type-1 violator. The connex sets N and N' are indicated by the shaded areas

applying the operation 1 for the choice of guard node g and descendant guard node l , (T', N') shows the resulting valid sibling-closed GJT.

Lemma 7 Let n be a violator of type 1 in (T, N) and assume $(T, N) \xrightarrow{1,n} (T', N')$. Then, (T', N') is a GJT pair and it is equivalent to (T, N) . Moreover, the number of violators in (T', N') is strictly smaller than the number of violators in (T, N) .

We prove this lemma in “Appendix C”. The second operator is applicable when n is a type 2 violator. When applied to n in (T, N) , it returns the pair (T', N') obtained as follows:

- Since n is a type 2 violator, no node in $\text{ch}_T(n) \cap N$ is a guard of n . Since every node has a guard, there is some $g \in \text{ch}_T(n) \setminus N$ which is a guard of n .
- Create a new child p of n with label $\text{var}(p) = \text{var}(n)$; detach all nodes in $\text{ch}_T(n) \setminus N$ (including g) from N , and add them as children of p , preserving their edge labels. This moves all subtrees rooted at nodes in $\text{ch}_T(n) \setminus N$ from n to p . Denote by T' the final result.
- Set $N' = N \cup \{p\}$.

We write $(T, N) \xrightarrow{2,n} (T', N')$ to indicate that (T', N') was obtained by applying this operation on n .

Example 8 Consider the GJT pair (T, N) in Fig. 8. Let us denote the root node by n . Since its guard child $h(y, z, w, t)$ is not in N , n is violator of type 2. After applying operation 2 on n , (T', N') shows the resulting valid sibling-closed GJT.

Lemma 8 Let n be a violator of type 2 in (T, N) and assume $(T, N) \xrightarrow{2,n} (T', N')$. Then, (T', N') is a GJT pair and it is equivalent to (T, N) . Moreover, the number of violators in (T', N') is strictly smaller than the number of violators in (T, N) .

The proof can be found in “Appendix C”.

Proposition 12 Every GJT pair can be transformed in polynomial time into an equivalent sibling-closed pair.




Fig. 8 Illustration of the sibling-closed transform: removal of type-2 violator. The connex sets N and N' are indicated by the shaded areas




Fig. 9 Binarizing a k -ary node n

Proof The two operations introduced above remove violators, one at a time. By repeatedly applying these operations until no violator remains we obtain an equivalent pair without violators, which must hence be sibling closed. Since each operator can clearly be executed in polynomial time and the number of times that we must apply an operator is bounded by the number of nodes in the GJT pair, the removal takes polynomial time. \square

Binary transformation Next, we show how to transform a sibling-closed pair (T, N) into an equivalent binary and sibling-closed pair (T', N') . The idea here is to “binarize” each node n with $k > 2$ children as shown in Fig. 9. There, we assume without loss of generality that c_1 is a guard child of n . The binarization introduces $k - 2$ new intermediate nodes m_1, \dots, m_{k-2} , all with $\text{var}(m_i) = \text{var}(n)$. Note that, since c_1 is a guard of n and $\text{var}(m_i) = \text{var}(n)$, it is straightforward to see that c_1 will be a guard of m_1 , which will be a guard of m_2 , which will be a guard of m_3 , and so on. Finally, m_{k-2} will be a guard of n . The connex set N is updated as follows. If none of n ’s children are in N , i.e., n is a frontier node, set $N' = N$. Otherwise, since N is sibling closed, all children of n are in N , and we set $N' = N \cup \{m_1, \dots, m_{k-2}\}$. Clearly, N' remains a sibling-closed connex subset of T' and $\text{var}(N') = \text{var}(N)$. We may hence conclude:

Lemma 9 By binarizing a single node in a sibling-closed GJT pair (T, N) as shown in Fig. 9, we obtain an equivalent GJT pair (T', N') that has strictly fewer nonbinary nodes than (T, N) .

Binarizing a single node is a polynomial-time operation. Then, by iteratively binarizing nonbinary nodes until all nodes have become binary we hence obtain:

Proposition 13 *Every sibling-closed GJT pair can be transformed in polynomial time into an equivalent plan.*

6 Implementation

We have implemented IEDYN, the instantiation of GDYN to setting where all θ -joins are inequality joins described in Sect. 4.4, as a query compiler that generates executable code in the Scala programming language. The generated code instantiates a (T, N) -rep and defines *trigger functions* that maintain the (T, N) -rep under updates. Off-the-shelf Scala collection libraries are used to implement the required indexes, and we take care to share the data structures between the join and enumeration indexes whenever possible. An important optimization used by our implementation lies in observing that, for the nodes in the connex set N that are not in the frontier of N , we never use the multiplicities stored in ρ_n during enumeration. As such, we also do not need to compute these multiplicities during update processing. While our theoretical framework supports batch updates, our implementation is currently limited to single-tuple updates.

Our implementation supports two modes of operation: *push based* and *pull based*. In both modes, the system maintains the T -rep under updates. In the *push-based mode* the system generates, on its output stream, the delta result $\Delta Q(db, u)$ after each single-tuple update u . To do so, it uses a modified version of enumeration (Algorithm 1) that we call *delta enumeration*. Similarly to how Algorithm 1 enumerates $Q(db)$, delta enumeration enumerates $\Delta Q(db, u)$ with constant delay (if Q has at most one inequality per pair of atoms) resp. logarithmic delay (otherwise). To do so, it uses both (1) the T -reduct GMRs ρ_n and (2) the delta GMRs $\Delta\rho_n$ that are computed by Algorithm 1 when processing u . In this case, however, one also needs to index the $\Delta\rho_n$ similarly to ρ_n . In the *pull-based mode*, in contrast, the system only maintains the (T, N) -rep under updates but does not generate any output stream. Nevertheless, at any time a user can ENUM (Algorithm 1) to obtain the current output.

We have described in Sect. 4 how IEDYN can process free-connex acyclic GCQs under updates. It should be noted that our implementation also supports the processing of general acyclic GCQs that are not necessarily free-connex. This is done using the following simple strategy. Let Q be acyclic but not free-connex. First, compute a free-connex acyclic approximation Q_F of Q . Q_F can always be obtained from Q by extending the set of output variables of Q . In the worst case, we need to add all variables, and Q_F becomes the full join underlying Q . Then, use IEDYN to maintain a (T, N) -rep

for Q_F . When operating in push-based mode, for each update u , we use the (T, N) -rep to delta-enumerate $\Delta Q_F(db, u)$ and project each resulting tuple to materialize $\Delta Q(db, u)$ in an array. Subsequently, we copy this array to the output. Note that the materialization of $\Delta Q(db, u)$ here is necessary since the delta enumeration can produce duplicate tuples after projection. When operating in pull-based mode, we materialize $Q(db)$ in an array and use delta enumeration of Q_F to maintain the array under updates. Of course, under this strategy, we require $\Omega(|Q(db)|)$ space in the worst case, just like (H)IVM would, but we avoid the (partial) materialization of delta queries. Note the distinction between the two modes: in push-based mode $\Delta Q(db, u)$ is materialized (and discarded once the output is generated), while in pull-based mode $Q(db)$ is materialized upon requests.

Our query compiler computes query plans using the algorithm of Sect. 5. Whenever we have the choice between enacting multiple reduction steps, we first enact using (ISO), then using (FLT), and finally using (CSE). This corresponds to the usual heuristics of pushing down projections and selections. If multiple applications of (CSE) are possible, we prefer those where the hyperedge to be eliminated has no extended variables. This corresponds to pushing down equi-semijoins so that inequality semijoins are hopefully executed over GMRs of reduced cardinality.

7 Experimental evaluation

In this section, we experimentally compare IEDYN against competing state-of-the-art HIVM and CER systems.

Queries Because the effectiveness of DYN for equijoin queries has already been documented [25], we focus on inequality join queries during our experimental analysis. Since there is no industry-strength established benchmark suite for such queries, we perform a systematic and in-depth exploration of the design space of inequality joins of up to three relations. Concretely, we evaluate IEDYN on the queries listed in Table 1. Here, Q_1 – Q_7 are full join queries (i.e., queries without projections). Among these, Q_1 , Q_3 and Q_4 contain only inequality join predicates, while Q_2 , Q_5 – Q_7 additionally contain at least one equality join. Queries Q_1 and Q_2 are binary joins, while Q_3 – Q_6 are multi-way join queries over three relations. The inequality predicates in Q_3 are unrelated ($a < d \wedge e < g$) while they form a chain ($a < d < g$) in Q_4 – Q_7 . Note that Q_5 – Q_7 are variants of Q_4 with equijoins added. We have similarly experimented with variants of Q_3 with equijoins added, but the trends are similar to what we obtain for Q_4 – Q_7 . Collectively, Q_3 – Q_7 (and the omitted variants) cover all possible ways in which three relations can be inequality joined in an acyclic manner.

Table 1 Benchmark queries

#	Query expression	Features			
		<	=	π	FC
Q_1	$R(a, b, c), S(d, e, f) \mid a < d$	✓			✓
Q_2	$R(a, b, c, k), S(d, e, f, k) \mid a < d$	✓	✓		✓
Q_3	$R(a, b, c), S(d, e, f), T(g, h, i) \mid a < d \wedge e < g$	✓			✓
Q_4	$R(a, b, c), S(d, e, f), T(g, h, i) \mid a < d \wedge d < g$	✓			✓
Q_5	$R(a, b, c, k), S(d, e, f, k), T(g, h, i) \mid a < d \wedge d < g$	✓	✓		✓
Q_6	$R(a, b, c), S(d, e, f, k), T(g, h, i, k) \mid a < d \wedge d < g$	✓	✓		✓
Q_7	$R(a, b, c, k), S(d, e, f, k), T(g, h, i, k) \mid a < d \wedge d < g$	✓	✓		✓
Q_8	$\pi_{a,b,d,e,f,g,h}(Q_4)$	✓		✓	✓
Q_9	$\pi_{a,d,e,f,g,h,k}(Q_5)$	✓	✓	✓	✓
Q_{10}	$\pi_{d,e,f,g,h,k}(Q_6)$	✓	✓	✓	✓
Q_{11}	$\pi_{a,b,d,e,g,h,k}(Q_7)$	✓	✓	✓	✓
Q_{12}	$\pi_{b,c,e,f,h,i}(Q_4)$	✓		✓	
Q_{13}	$\pi_{b,c,e,f,h,i}(Q_5)$	✓	✓	✓	
Q_{14}	$\pi_{b,c,e,f,h,i}(Q_6)$	✓	✓	✓	
Q_{15}	$\pi_{b,c,e,f,h,i}(Q_7)$	✓	✓	✓	

FC Free-connex acyclic

Queries Q_8 – Q_{15} project over the result of queries Q_4 – Q_7 . Among these, Q_8 – Q_{11} are free-connex acyclic while Q_{12} – Q_{15} are acyclic but not free-connex.

Updates We evaluate on streams of synthetic updates where each update consists of a single-tuple insertion. We focus on the setting where updates are single-tuple insertions for the following reasons. First, single-tuple updates stress-test dynamic query processing since the query results must be kept up-to-date after each and every single tuple, in contrast to the setting for batch updates, where results can be out-of-sync for the duration of the batch. Second, since batch updates can always be processed by executing all updates in the batch individually (using the single-tuple update triggers), performance measurements for single-tuple updates yield an upper bound for the performance of batch updates. Third, while it is true that, in principle, the processing of batch updates can be done faster than that of single tuple updates (e.g., by amortizing index access and data structure updates over the entire batch instead of per tuple), both our implementation and most of the competing systems that we describe below do not implement such optimization.⁶ As such, we restrict attention to single-tuple updates. We focus on insertions because this is supported by all of the systems that we consider whereas explicit deletions are not. While we have experimented with mixed (insert and delete) update streams for IEDYN, performance is similar to that for insert-only streams. This is expected, since IEDYN treats insertions and deletions uniformly.

⁶ In the sense that batch updates are only supported by treating each update tuple in the batch individually.

The database is always empty when we start processing the update stream. We synthetically generate two kinds of update streams: *randomly ordered* and *temporally ordered* update streams. In *randomly ordered* update streams, insertions can occur in any order. In contrast, *temporally ordered* update streams guarantee that any attribute that participates in an inequality in the query has a larger value than the same attribute in any of the previously inserted tuples. Randomly ordered update streams are useful for comparing against systems that allow processing of out-of-order tuples; temporally ordered update streams are useful for comparison against systems that assume events arrive always with increasing timestamp values. Examples of systems that process temporally ordered streams are automaton-based CER systems.

Competitors We compare IEDYN with the following state-of-the-art HIVM and CER engines: DBToaster (DBT) [30], Esper (E) [19], SASE (SE) [3,46,49], Tesla (T) [14,15], and ZStream (Z) [31]. The competing systems differ in their mode of operation (push based vs pull based, cf. Sect. 6) and some of them only support temporally ordered streams. The capabilities of each system is summarized in Table 2 and discussed in detail in “Appendix D”.

Setup Our experiments are run on an 8-core 3.07 GHz machine running Ubuntu with GNU/Linux 3.13.0-57-generic. To compile the different systems or generated trigger programs, we have used GCC version 4.8.2, Java 1.8.0_101, and Scala version 2.12.4. Each query is evaluated ten times to measure update processing delay and two times to measure memory footprint. We present the average over those runs. Each time a query is evaluated, 20 GB of main memory are freshly allocated to the program. To measure the memory

Table 2 Competing systems capabilities overview

	Mode	Stream	Unsupported queries
DBT	Pull	Random	
E	Push	Random	
SE	Push	Ordered	Q_3, Q_5-Q_{15}
T	Push	Ordered	Q_3
Z	Push	Ordered	Q_8-Q_{15}

footprint for Scala/Java-based systems, we invoke the JVM system calls every 10 updates and consider the maximum value. For C/C++-based systems we use the GNU/Linux *time* command to measure memory usage. Experiments that measure memory footprint are always run separately of the experiments that measure processing time.

7.1 Results

We will compare using *total update processing time*, *memory footprint*, and *enumeration delay* as metrics. Here, the total update processing time is the time required to process the entire update stream, where updates are fed into the systems as fast as they can process them. This hence measures the maximum system performance. While measuring the update processing time, we take care to compare fairly with the competing systems by consistently running IEDYN in the same operation mode as the one supported by the competitor. Concretely, for push-based systems, we report the time required to process the entire update stream, while generating the delta result $\Delta Q(db, u)$ after each update (cf. Sect. 6). When comparing against a pull-based system, the measured time includes only processing the entire update stream. For these systems, we later separately report the enumeration delay, i.e., the speed with which the result can be generated from the underlying representation of the output (a *T*-representation in the case of IEDYN). SASE, ZStream, and Tesla only support temporally ordered streams and hence we only compare with them on such streams. While DBToaster and Esper support both random and temporally ordered streams, we only report comparisons using randomly ordered streams. We have also looked at temporally ordered streams for these systems, but their throughput is similar (fluctuating between 3% and 12%) while that of IEDYN significantly improves (fluctuating between 35% and 50%) because insertions to sorted lists become constant instead of logarithmic. We omit these experiments due to lack of space.

Some executions of the competing systems failed either because they required more than 20GB of main memory or they took more than 1500 s. If an execution requires more than 20GB, we report the processing time elapsed until the exception was raised. If an execution is still running after

1500 s, we stop it and report its maximum memory usage while running.

Full join queries We first analyze the performance of the full join queries Q_1-Q_7 in the setting where the join selectivity is relatively large. Here, the selectivity of $R \bowtie S$ is defined as $\frac{|R \bowtie S|}{|R| \times |S|}$. High-selectivity joins are common in CER scenarios. We refer to Table 3(Left) for output sizes. Figure 10 compares the update processing time of IEDYN against the competing systems for full join queries Q_1-Q_7 , grouped per system capability (push/pull, random/temporal). We observe that all of the competing systems have large processing times even for very small update stream sizes, indicating poor scalability. In particular, *DBT* runs out of memory on streams of sizes $\geq 18k$ for query Q_5 , and on streams of size $\geq 15k$ for query Q_6 . *Z* runs out of memory for Q_5 and Q_6 on streams of size $21k$. Similarly, *T* took more than 1500 s for Q_5 on streams of size $\geq 12k$, for Q_6 on streams of size $\geq 18k$, and for Q_7 on streams of size $\geq 9k$ and was aborted at that time. All of these behaviors are due to the large selectivity of joins on this dataset. Note that in all cases, IEDYN scales satisfactorily with increasing stream sizes and significantly outperforms the competitors, often by orders of magnitude. This is confirmed in Table 3 (Right) where we show the processing time and memory footprint used by IEDYN as a percentage of the corresponding usage in the competing systems (for the largest stream size of each query). Writing “(x, y) oom”: to indicate x orders of magnitude improvement in processing time and y orders in memory consumption, we see that IEDYN improves up to (1, 2) oom w.r.t. *Z*; almost (3, 1) oom w.r.t. *T*; almost (2, 2) oom w.r.t. *SE*, and up to (4, 3) oom w.r.t. *DBT*. Moreover, for these queries, even in push-based mode IEDYN can support the enumeration of query results from its data structures at any time while competing push-based systems have no such support. Hence, IEDYN is not only more efficient but also provides more functionality.

Drilling deeper into the specific queries, we see from Fig. 10 that while existing systems already perform poorly on the inequality-only binary join query Q_2 , this is further aggravated when moving to the corresponding three-way join query Q_3 and Q_4 : note that the plot for Q_1 shows streams sizes up to $12k$, whereas the plots for Q_3 and Q_4 only go to $2.7k$ while having much larger runtimes. Because Q_2 and Q_5-Q_7 modify Q_1 , resp. Q_4 by adding equality join predicates, they have a smaller join selectivity. As a result, performance for these queries is much better, across all systems. (To see this, note that, for Q_5-Q_6 the first data point plotted has a significantly lower processing time while processing a data stream three times the size.) Similarly, we note that, while not visually apparent from Fig. 10, the performance of all systems on Q_3 is 10%–100% slower compared to the same systems on Q_4 . This is due to the fact that the two unrelated inequality predicates ($a < d \wedge e < g$) of Q_3

Table 3 (Left) Maximum update stream and result sizes, $k = 10^3$. (Right) Relative performance of IEDYN, expressed as a percentage of the corresponding resources used by competing systems

	Data size			Z		T		SE		E		DBT	
	Stream	Output		Time	Mem	Time	Mem	Time	Mem	Time	Mem	Time	Mem
Q1	12k	18,017k	Q1	45.33	3.4	12.06	96.15	2.74	25.23	29.33	24.51	0.43	2.74
Q2	12k	3.8k	Q2	19.34	87.1	0.52	168.75	97.8	1	22.62	16.67	52.1	0.2
Q3	2.7k	178,847k	Q3	11.3	0.87	N/A	N/A	N/A	N/A	24.49	16.57	0.25	0.19
Q4	2.7k	90,425k	Q4	10.69	0.69	4.64	3.45	13.94	55.56	24.5	14.2	0.01	0.2
Q5	21k	411,669k	Q5	10.85	0.3	13.72	19.11	N/A	N/A	25.5	23.44	0.04	0.15
Q6	21k	297,873k	Q6	10.58	0.31	14.61	11.03	N/A	N/A	23.28	17.37	0.86	0.15
Q7	21k	64,603k	Q7	0.03	2.38	0.05	24.1	N/A	N/A	19.18	67.14	0.32	1.18
Q8	2.7k	114,561k	Q8	N/A	N/A	9.01	6.36	N/A	N/A	24.91	10.24	0.02	0.35
Q9	21k	411,669k	Q9	N/A	N/A	13.96	40	N/A	N/A	34.16	94.34	0.07	0.25
Q10	21k	99,043k	Q10	N/A	N/A	0.07	38.52	N/A	N/A	17.49	23.38	24.22	61.84
Q11	21k	43,564k	Q11	N/A	N/A	0.04	10.23	N/A	N/A	4.95	25.43	0.01	0.45
Q12	2.7k	114,561k	Q12	N/A	N/A	7.11	10.49	N/A	N/A	21.42	14.17	71.19	54.07
Q13	21k	294,139k	Q13	N/A	N/A	9.99	47.58	N/A	N/A	29.59	109.26	10.14	68.29
Q14	21k	297,873k	Q14	N/A	N/A	13.96	47.2	N/A	N/A	34.68	28.92	91.75	77.9
Q15	21k	50,468k	Q15	N/A	N/A	0.16	25.26	N/A	N/A	95.8	70	0.17	15.02

N/A Not applicable

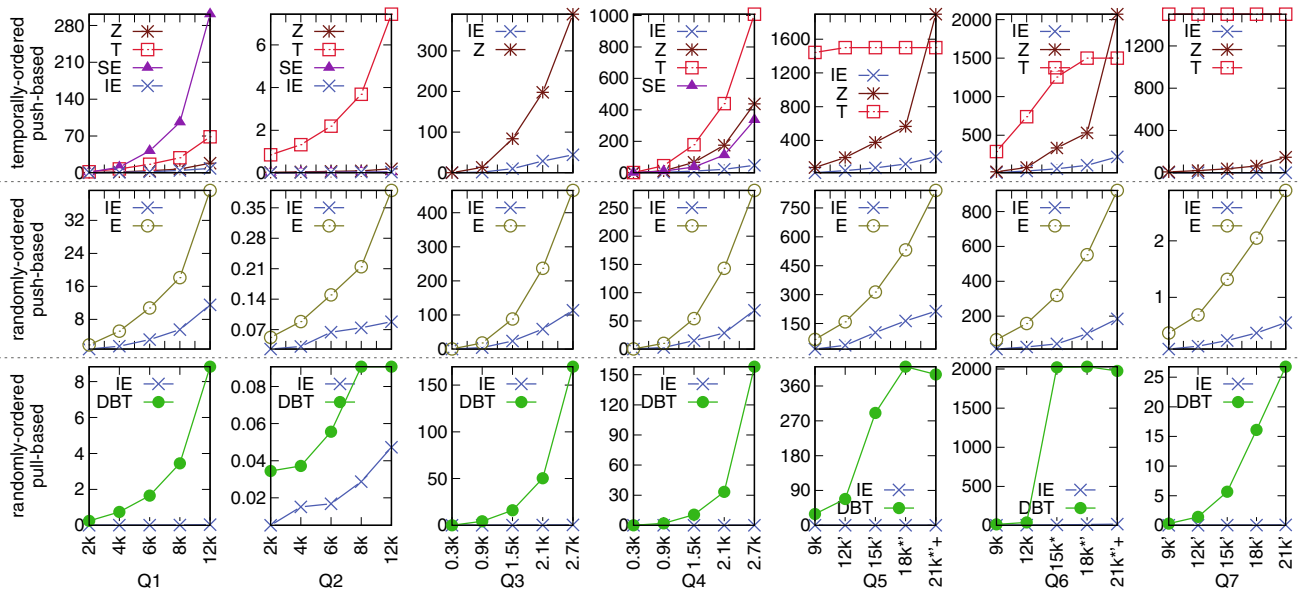


Fig. 10 Processing time in seconds on full join queries for various stream sizes. (*: DBT out of memory, +: Z out of memory, †: T was stopped after 1500 s)

are linearly ordered in Q_4 ($a < d \wedge d < g$), which causes Q_4 to have higher selectivity, improving performance. From these observations we may conclude that the performance of all systems increases as the join selectivity decreases. Never-theless, in contrast to existing systems, IEDYN continues to scale and perform satisfactorily even with large selectivities. To confirm this trend in the setting of very low selectivities, where processing is hence less output-size dependent, we have

also generated datasets with probability distributions that are parametrized by a *selectivity* s , such that the expected number of output tuples is s percent of the Cartesian product of all relations in the query. The results in Fig. 11 show that IEDYN consistently continues to perform better also on very selective inequality joins. For super selective inequality joins the measurements come similar to what we observe for equality joins, which we investigated in detail in [25].




Fig. 11 IEDYN as percentage of (E , DBT , SE , Z , T) for higher join selectivities. X-axis shows queries with tuples per relation and selectivities




Fig. 12 Enumeration of query results: IEDYN vs DBT for various input stream sizes. Different bars for Q_4 , Q_5 , Q_6 show their projected versions

Projections The trends observed for full join queries are confirmed for the queries Q_8 – Q_{15} with projections, as shown in Table 3 (Right). In particular, IEDYN improves up to almost (4, 1) oom w.r.t. T ; up to almost (1, 1) w.r.t. E . and up to (4, 3) oom w.r.t. DBT . Again, we observe that, while IEDYN consistently outperforms the competitors, when the selectivity of the queries decreases, the performance gain of IEDYN increases. In addition, we note that the performance gain of IEDYN is bigger on queries that are free-connex acyclic (namely: Q_8 – Q_{10}) as opposed to those that are not (Q_{11} – Q_{14}). This is to be expected since our implementation evaluates nonfree-connex queries by approximating them by free-connex variants, requiring additional materialization (cf. Sect. 6).

Result enumeration CDE is theoretical notion that hides a constant factor which could decrease performance in practice when compared to enumerating from a fully materialized representation. In Fig. 12, we show the practical application of CDE in IEDYN and compare against DBT which materializes the full query results. We plot the time required to enumerate the result from IEDYN’s T -rep as a fraction of the time required to enumerate the result from DBT ’s materialized views. As can be seen from the figure, both enumeration times are comparable on average. Note that we do not compare enumeration time for push-based systems, since for these systems the time required for delta enumeration is already included in the update processing time.

8 Conclusions

Traditional techniques for dynamic query evaluation are based on a trade-off between materialization of join subresults (to avoid recomputing these subresults) and their recomputation (to avoid the space overhead induced by materialization). We have shown that instead of materializing full join subresults, it suffices to maintain and index semi-join subresults. This methodology, called General Dynamic Yannakakis, allows us to maintain a data structure that, like recomputation, has low memory overhead; and yet supports all operations one commonly expects from full materialization: enumeration with bounded delay and efficient processing of updates. In addition, the framework supports bounded-delay delta enumeration under single-tuple update, hence allowing to operate in push-based mode similar to streaming systems. Our experiments against state-of-the-art engines in the IVM and CER domains show that Dynamic Yannakakis can improve performance by orders of magnitude in both time and space. The performance gap with existing systems is the largest for output-dominated queries (i.e., queries where the join result is large). While the gap decreases for more selective joins, GDYN continues to consistently outperform existing systems. In addition, while GDYN’s theory is developed only for free-connex acyclic GCQs, our experiments show that GDYN’s adaptation to acyclic GCQs that are not free connex (by means of free-connex approximation followed by post-processing) is equally effective compared to existing systems.

From a theoretical viewpoint, it would be satisfying to establish lower bounds on the complexity of processing inequality joins dynamically. For equijoins, such lower bounds were recently established by Berkholz et al. [8]. Since GDYN meets these lower bounds on queries with equijoins only [25], it would be interesting to know whether it is similarly optimal for inequality joins.

A Proofs of Sect. 4

Lemma 1 $\rho_n = Q[T_n, n](db)$, for every node $n \in T$.

Proof We proceed by induction on the number of descendants of n . If n has no descendants, then T_n is a single atom $r(\bar{x})$ with $\bar{x} = \text{var}(n) = \text{out}(Q[T_n, n])$. Then, $Q[T_n, n](db) = (\pi_{\text{var}(n)} r(\bar{x}))(db) = r(\bar{x})(db) = db_{r(\bar{x})} = \rho_n$, concluding the basic case. Now, for the inductive case, we distinguish whether n has one or two children.

Assume n has a single child c . Then, $\text{at}(T_n) = \text{at}(T_c)$ and $\text{pred}(T_n) = \text{pred}(T_c) \cup \text{pred}(n)$. Therefore, by definition of $Q[\cdot]$, we have $Q[T_n] \equiv \sigma_{\text{pred}(n)} Q[T_c]$, which implies that $Q[T_n, n] = \pi_{\text{var}(n)} Q[T_n] \equiv \pi_{\text{var}(n)} \sigma_{\text{pred}(n)} Q[T_c]$. Furthermore, since $\text{pred}(n)$ only mentions variables in $\text{var}(c) \cup$

$var(n)$ and $var(n) \subseteq var(c)$, as c is a guard of n , this is equivalent to

$$\begin{aligned}\pi_{var(n)}\sigma_{pred(n)}\mathcal{Q}[T_c] &\equiv \pi_{var(n)}\sigma_{pred(n)}\pi_{var(c)}\mathcal{Q}[T_c] \\ &= \pi_{var(n)}\sigma_{pred(n)}\mathcal{Q}[T_c, c].\end{aligned}$$

By induction, $\pi_{var(n)}\sigma_{pred(n)}\mathcal{Q}[T_c, c](db) = \pi_{var(n)}\sigma_{pred(n)}\rho_c = \rho_n$, showing that $\mathcal{Q}[T_n, n](db) = \rho_n$.

Assume now that n has two children c_1 and c_2 . We assume w.l.o.g. that c_1 is a guard for n . Note that $at(T_n) = at(T_{c_1}) \cup at(T_{c_2})$ and $pred(T_n) = pred(T_{c_1}) \cup pred(T_{c_2}) \cup pred(n)$. Therefore,

$$\mathcal{Q}[T_n] \equiv \sigma_{pred(n)}\sigma_{pred(T_{c_1})}\sigma_{pred(T_{c_2})} (at(T_{c_1}) \bowtie at(T_{c_2})).$$

Here, we abuse notation and write $at(T_i)$ for the natural join of all atoms in T_{c_i} . Since $pred(T_{c_i})$ only mentions variables of atoms in T_{c_i} (for $i \in \{1, 2\}$), we can push the selections:

$$\begin{aligned}\mathcal{Q}[T_n] &\equiv \sigma_{pred(n)} (\sigma_{pred(T_{c_1})}at(T_{c_1}) \bowtie \sigma_{pred(T_{c_2})}at(T_{c_2})) \\ &\equiv \sigma_{pred(n)} (\mathcal{Q}[T_{c_1}] \bowtie \mathcal{Q}[T_{c_2}]).\end{aligned}$$

Therefore,

$$\mathcal{Q}[T_n, n] = \pi_{var(n)}\mathcal{Q}[T_n] \equiv \pi_{var(n)}\sigma_{pred(n)} (\mathcal{Q}[T_{c_1}] \bowtie \mathcal{Q}[T_{c_2}]).$$

Since $var(pred(n)) \subseteq var(c_1) \cup var(c_2) \cup var(n)$ and $var(n) \subseteq var(c_1)$, we have $var(pred(n)) \subseteq var(c_1) \cup var(c_2)$. This is combined with the fact that, due to the connectedness property of T , we have $var(\mathcal{Q}[T_{c_1}]) \cap var(\mathcal{Q}[T_{c_2}]) \subseteq var(c_i)$ for $i \in \{1, 2\}$, we can add the following projections

$$\begin{aligned}\mathcal{Q}[T_n, n] &\equiv \pi_{var(n)}\sigma_{pred(n)} (\pi_{var(c_1)}\mathcal{Q}[T_{c_1}] \bowtie \pi_{var(c_2)}\mathcal{Q}[T_{c_2}]) \\ &\equiv \pi_{var(n)}\sigma_{pred(n)} (\mathcal{Q}[T_{c_1}, c_1] \bowtie \mathcal{Q}[T_{c_2}, c_2]).\end{aligned}$$

Hence, by induction hypothesis, we have

$$\mathcal{Q}[T_n, n](db) = \pi_{var(n)}\sigma_{pred(n)} (\rho_{c_1} \bowtie \rho_{c_2}) = \rho_n,$$

concluding our proof. \square

Lemma 3

1. $\mathcal{Q}(db)$ is a positive GMR, for any GCQ \mathcal{Q} and any database db .
2. If R is a positive GMR over \bar{x} and $\bar{y} \subseteq \bar{x}$, then $\mathbf{t}[\bar{y}] \in \pi_{\bar{y}}R$ for every tuple $\mathbf{t} \in R$.

Proof (1) Follows by straightforward induction on \mathcal{Q} , using the fact that the GMRs in db are themselves positive by definition. (2) Is a standard result in relational algebra, which hence transfers to the case of positive GMRs. \square

Lemma 10 Let R be a positive GMR over \bar{x} , S a positive GMR over \bar{y} and \mathbf{t} a tuple over \bar{z} . If $\bar{z} \subseteq \bar{y} \subseteq \bar{x}$, then $R \ltimes (S \ltimes \mathbf{t}) = (R \ltimes S) \ltimes \mathbf{t}$.

Proof This results well know in standard relational algebra, and its proof transfers to the case of positive GMRs. \square

Lemma 2 For every node $n \in N$ and every tuple \mathbf{t} in ρ_n , $\text{ENUM}_{T,N}(n, \mathbf{t}, \rho)$ enumerates $\mathcal{Q}[T_n, N_n](db) \ltimes \mathbf{t}$.

Proof. Let $n \in N$ and $\mathbf{t} \in \rho_n$. We need to show that executing $\text{ENUM}_{T,N}(n, \mathbf{t}, \rho)$ outputs all (tuple, multiplicity) pairs of $\mathcal{Q}[T_n, N_n](db) \ltimes \mathbf{t}$ exactly once. We proceed by induction on the number of nodes in N_n . If $N_n = \{n\}$, then $\mathcal{Q}[T_n, N_n] = \mathcal{Q}[T_n, n]$. Therefore, by Lemma 1, $\mathcal{Q}[T_n, N_n](db) = \rho_n$. Since $\mathbf{t} \in \rho_n$, this implies that the only tuple in $\mathcal{Q}[T_n, N_n](db)$ that is compatible with \mathbf{t} is \mathbf{t} itself. Furthermore, since $N_n = \{n\}$, n must be in the frontier of n . Therefore, $\text{ENUM}_{T,N}(n, \mathbf{t}, \rho)$ will output precisely $\{(\mathbf{t}, \rho_n(\mathbf{t}))\}$ (line 4), which concludes the base case.

For the inductive step we need to consider two cases depending on the number of children of n .

Case 1 If n has a single child c , then necessarily c is a guard of n , i.e., $var(n) \subseteq var(c)$. In this case, Algorithm 1 will call $\text{ENUM}_{T,N}(c, \mathbf{s}, \rho)$ for each tuple $\mathbf{s} \in (\rho_c \ltimes_{pred(n)} \mathbf{t})$. By induction hypothesis and Lemma 1, this will correctly enumerate and output the elements of $\mathcal{Q}[T_c, N_c](db) \ltimes \mathbf{s}$, for every \mathbf{s} in $\mathcal{Q}[T_c, c](db) \ltimes_{pred(n)} \mathbf{t}$. Note that the sets $\mathcal{Q}[T_c, N_c](db) \ltimes \mathbf{s}$ are disjoint for different values of \mathbf{s} . Thus, no element is output twice. Hence, $\text{ENUM}_{T,N}(n, \mathbf{t}, \rho)$ enumerates the GMR

$$\mathcal{Q}[T_c, N_c](db) \ltimes (\mathcal{Q}[T_c, c](db) \ltimes_{pred(n)} \mathbf{t}). \quad (2)$$

Since $var(pred(n)) \subseteq var(c) \cup var(n) = var(c) = out(\mathcal{Q}[T_c, c])$, we can pull out the selection:

$$(2) = \mathcal{Q}[T_c, N_c](db) \ltimes \sigma_{pred(n)}(\mathcal{Q}[T_c, c](db) \ltimes \mathbf{t}). \quad (3)$$

Subsequently, because $var(pred(n)) = var(c) \subseteq out(\mathcal{Q}[T_c, N_c])$, we can pull out the selection again:

$$(3) = \sigma_{pred(n)}(\mathcal{Q}[T_c, N_c](db) \ltimes (\mathcal{Q}[T_c, c](db) \ltimes \mathbf{t})). \quad (4)$$

Because the variables in \mathbf{t} are a subset of $var(c)$, because $var(c) \subseteq var(N_c)$, and because $\mathcal{Q}[T_c, N_c](db)$ and $\mathcal{Q}[T_c, c](db)$ are positive (Lemma 3(1)), we can apply Lemma 10:

$$(4) = \sigma_{pred(n)}((\mathcal{Q}[T_c, N_c](db) \ltimes \mathcal{Q}[T_c, c](db)) \ltimes \mathbf{t}). \quad (5)$$

Next, observe that, since $\text{var}(n_c) \subseteq \text{var}(N_c)$ as $c \in N_c$, we have

$$\begin{aligned}\mathcal{Q}[T_c, n_c] &= \pi_{\text{var}(c)} \mathcal{Q}[T_c] \\ &\equiv \pi_{\text{var}(c)} \pi_{\text{var}(N_c)} \mathcal{Q}[T_c] \\ &\equiv \pi_{\text{var}(c)} \mathcal{Q}[T_c, N_c]\end{aligned}$$

Then, because $\mathcal{Q}[T_c, N_c](db)$ is positive, we obtain from Lemma 3(2) that

$$(5) = \sigma_{\text{pred}(n)}(\mathcal{Q}[T_c, N_c](db) \bowtie \mathbf{t}). \quad (6)$$

Finally, because $\text{pred}(n) \subseteq \text{var}(n) \subseteq \text{var}(N_c)$, we push the selection again and obtain

$$(6) = (\sigma_{\text{pred}(n)} \mathcal{Q}[T_c, N_c](db)) \bowtie \mathbf{t} \quad (7)$$

$$= (\pi_{\text{var}(N_n)} \sigma_{\text{pred}(n)} \mathcal{Q}[T_c, N_c](db)) \bowtie \mathbf{t}. \quad (8)$$

Here, the last equality is due to the fact that $\text{var}(N_n) = \text{var}(n) \cup \text{var}(N_c) = \text{var}(N_c)$, as $\text{var}(n) \subseteq \text{var}(c)$ and $c \in N_c$, which implies that projecting on $\text{var}(N_n)$ does not modify the result. The result then follows from the observation that $\mathcal{Q}[T_n, N_n] \equiv \pi_{\text{var}(N_n)} \sigma_{\text{pred}(n)} \mathcal{Q}[T_c, N_c]$.

Case 2 Otherwise, n has two children c_1 and c_2 . We assume w.l.o.g. that c_1 is a guard of n , i.e., $\text{var}(n) \subseteq \text{var}(c_1)$. Since $|N_n| > 1$ and N is sibling closed, we have $\{c_1, c_2\} \subset N$. In this case, Algorithm 1 will first enumerate $\mathbf{t}_i \in \rho_{c_i} \bowtie_{\text{pred}(n \rightarrow c_1)} \mathbf{t}$ for $i \in \{1, 2\}$. By Lemma 1, this is equivalent to enumerate every \mathbf{t}_i in $\mathcal{Q}[T_{c_i}, c_i](db) \bowtie_{\text{pred}(n \rightarrow c_1)} \mathbf{t}$. Then, for each such \mathbf{t}_i the algorithm will enumerate every pair (\mathbf{s}_i, μ_i) generated by $\text{ENUM}_{T, N}(c_i, \mathbf{t}_i, \rho)$, which by induction is the same as enumerating every (\mathbf{s}_i, μ_i) in $\mathcal{Q}[T_{c_i}, N_{c_i}](db) \bowtie \mathbf{t}_i$. Note that the sets $\mathcal{Q}[T_{c_i}, N_{c_i}](db) \bowtie \mathbf{t}_i$ are disjoint for distinct \mathbf{t}_i . Therefore, no (\mathbf{s}_i, μ_i) is generated twice. the algorithm is hence enumerating

$$\mathcal{Q}[T_{c_i}, N_{c_i}](db) \bowtie (\mathcal{Q}[T_{c_i}, c_i](db) \bowtie_{\text{pred}(n \rightarrow c_i)} \mathbf{t})$$

By the same reasoning as in Case (1), this is equivalent to enumerating every (\mathbf{s}_i, μ_i) in $(\sigma_{\text{pred}(n \rightarrow c_i)} \mathcal{Q}[T_{c_i}](db)) \bowtie \mathbf{t}$. From the connectedness property of T , it follows that $\text{var}(\mathcal{Q}[T_{c_1}]) \cap \text{var}(\mathcal{Q}[T_{c_2}]) \subseteq \text{var}(n)$. Thus, $\text{var}(\mathcal{Q}[T_{c_1}]) \cap \text{var}(\mathcal{Q}[T_{c_2}])$ is a subset of the variables of \mathbf{t} . Hence, every tuple \mathbf{s}_1 will be compatible with every tuple \mathbf{s}_2 , and therefore, enumeration of every pair $(\mathbf{s}_1 \cup \mathbf{s}_2, \mu_1 \times \mu_2)$ is the same as the enumeration of

$$\begin{aligned}&((\sigma_{\text{pred}(n \rightarrow c_1)} \mathcal{Q}[T_{c_1}, N_{c_1}](db)) \bowtie \mathbf{t}) \\ &\bowtie ((\sigma_{\text{pred}(n \rightarrow c_2)} \mathcal{Q}[T_{c_2}, N_{c_2}](db)) \bowtie \mathbf{t}).\end{aligned} \quad (9)$$

The semijoin with \mathbf{t} factors out of the join:

$$(9) = (\sigma_{\text{pred}(n \rightarrow c_1)} \mathcal{Q}[T_{c_1}, N_{c_1}] \bowtie \sigma_{\text{pred}(n \rightarrow c_2)} \mathcal{Q}[T_{c_2}, N_{c_2}](db)) \bowtie \mathbf{t} \quad (10)$$

We can now pull out the selections and obtain

$$\begin{aligned}(10) &= (\sigma_{\text{pred}(n \rightarrow c_1)} \sigma_{\text{pred}(n \rightarrow c_2)} (\mathcal{Q}[T_{c_1}, N_{c_1}] \bowtie \mathcal{Q}[T_{c_2}, N_{c_2}](db))) \bowtie \mathbf{t} \\ &= (\sigma_{\text{pred}(n)} (\mathcal{Q}[T_{c_1}, N_{c_1}] \bowtie \mathcal{Q}[T_{c_2}, N_{c_2}](db))) \bowtie \mathbf{t} \\ &= (\pi_{\text{var}(N_n)} \sigma_{\text{pred}(n)} (\mathcal{Q}[T_{c_1}, N_{c_1}] \bowtie \mathcal{Q}[T_{c_2}, N_{c_2}](db))) \bowtie \mathbf{t}\end{aligned}$$

Here, the last equality is due to the fact that $\text{var}(N_n) = \text{var}(n) \cup \text{var}(N_{c_1}) \cup \text{var}(N_{c_2}) = \text{var}(N_{c_1}) \cup \text{var}(N_{c_2})$, as $\text{var}(n) \subseteq \text{var}(c_1) \subseteq \text{var}(N_{c_1})$. This implies that

$$\text{var}(N_n) = \text{out}(\mathcal{Q}[T_{c_1}, N_{c_1}]) \cup \text{out}(\mathcal{Q}[T_{c_2}, N_{c_2}])$$

Hence, projecting the join result on $\text{var}(N_n)$ does not modify the result. The result then follows from the observation that $\mathcal{Q}[T_n, N_n] \equiv \pi_{\text{var}(N_n)} \sigma_{\text{pred}(n)} (\mathcal{Q}[T_{c_1}, N_{c_1}] \bowtie \mathcal{Q}[T_{c_2}, N_{c_2}])$. \square

Proposition 4 Assume that all join indexes in the (T, N) -rep have access time g , and that all indexes (join and enumeration) have update time h , where g and h are monotone functions. Further assume that, during the entire execution of UPDATE, K and U bound the size of ρ_n and Δ_n , respectively, for all n . Then, $\text{UPDATE}_{T, N}(\rho, u)$ runs in time $\mathcal{O}(|T| \cdot (U + h(K, U) + g(K, U)))$.

Proof First, note that the initialization of Δ_n in line 15 can be done in $\mathcal{O}(U)$ time (by copying $u_{r(\bar{x})}$ to Δ_n tuple by tuple) and the initialization of Δ_n in line 17 in $\mathcal{O}(1)$ time. Therefore, lines 14–17 run in $\mathcal{O}(|T| \cdot U)$ time, which falls within the claimed bounds. We next show that the for loop of lines 18–23 also runs within the claimed bounds. Since the body of this for loop is executed $|T|$ times, it suffices to show that each of the lines 19–23 run in time $\mathcal{O}(U + h(K, U) + g(K, U))$. Since $|\Delta_n| \leq U$ by assumption, the statement $\rho_n + = \Delta_n$ of line 19 can be executed in $\mathcal{O}(U)$ time by iterating over the tuples $\mathbf{t} \in \Delta_n$, and updating $\rho_n(\mathbf{t})$ for each such tuple. (Recall that multiplicity lookup and modification in a GMR are $\mathcal{O}(1)$ operations). The indexes associated with ρ_n (if any) are updated in time $h(K, U)$. Therefore, the total time require to execute line 19 is $\mathcal{O}(U + h(K, U))$. We next bound the complexity of line 21. Computing $\pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}p} \Delta_n)$ using the join index on ρ_m takes $\mathcal{O}(g(K, U))$ time. Furthermore, the number of tuples in $\pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}p} \Delta_n)$ can be at most $2U$. This is because $|\Delta_p| \leq U$ at any time during the execution. In the worst case, therefore, $\pi_{\text{var}(p)}(\rho_m \bowtie_{\text{pred}p} \Delta_n)$ can at most delete the tuples already present in Δ_p (which requires U tuples) and subsequently insert U new tuples

(requiring another U tuples), for at most $2U$ tuples in total. For each of the $2U$ resulting tuples, we update Δ_p accordingly in $\mathcal{O}(1)$ time. The total time to execute line 21 is hence $\mathcal{O}(2 \cdot U + g(K, U))$. Finally, using similar reasoning, the complexity of line 23 can be shown to be $\mathcal{O}(U)$. \square

B Proofs of Sect. 5.1

B.1 Proof of Proposition 7

Because no infinite sequences of reduction steps are possible, it suffices to demonstrate local confluence:

Proposition 14 *If $\mathcal{H} \rightsquigarrow \mathcal{I}_1$ and $\mathcal{H} \rightsquigarrow \mathcal{I}_2$, then there exists \mathcal{J} such that both $\mathcal{I}_1 \rightsquigarrow^* \mathcal{J}$ and $\mathcal{I}_2 \rightsquigarrow^* \mathcal{J}$.*

Indeed, it is a standard result in the theory of rewriting systems that confluence (Lemma 7) and local confluence (Lemma 14) coincide when infinite sequences of reductions steps are impossible [5].

Before proving Lemma 14, we observe that the property of being isolated or being a conditional subset is preserved under reductions, in the following sense.

Lemma 11 *Assume that $\mathcal{H} \rightsquigarrow \mathcal{I}$. Then, $\text{pred}(\mathcal{I}) \subseteq \text{pred}(\mathcal{H})$ and for every hyperedge e , we have $\text{ext}_{\mathcal{I}}(e) \subseteq \text{ext}_{\mathcal{H}}(e)$, $\text{fv}_{\mathcal{I}}(e) \subseteq \text{fv}_{\mathcal{H}}(e)$, and $\text{isol}_{\mathcal{H}}(e) \subseteq \text{isol}_{\mathcal{I}}(e)$. Furthermore, if $e \sqsubseteq_{\mathcal{H}} f$ then also $e \sqsubseteq_{\mathcal{I}} f$.*

Proof First, observe that $\text{pred}(\mathcal{I}) \subseteq \text{pred}(\mathcal{H})$, since reduction operators only remove predicates. This implies that $\text{ext}_{\mathcal{I}}(e) \subseteq \text{ext}_{\mathcal{H}}(e)$ for every hyperedge e . Furthermore, because reduction operators only remove hyperedges and never add them, it is easy to see that $\text{fv}_{\mathcal{H}}(e) \subseteq \text{fv}_{\mathcal{I}}(e)$. Hence, if $x \in \text{isol}_{\mathcal{H}}(e)$ then $x \notin \text{fv}_{\mathcal{H}}(e) \supseteq \text{fv}_{\mathcal{I}}(e)$ and $x \notin \text{var}(\text{pred}(\mathcal{H})) \supseteq \text{var}(\text{pred}(\mathcal{I}))$. Therefore, $x \in \text{isol}_{\mathcal{I}}(e)$. As such, $\text{isol}_{\mathcal{I}}(e) \subseteq \text{isol}_{\mathcal{H}}(e)$.

Next, assume that $e \sqsubseteq_{\mathcal{H}} f$. We need to show that $\text{fv}_{\mathcal{I}}(e) \subseteq f$ and $\text{ext}_{\mathcal{I}}(e \setminus f) \subseteq f$. The first condition follows since $\text{fv}_{\mathcal{I}}(e) \subseteq \text{fv}_{\mathcal{H}}(e) \subseteq f$ where the last inclusion is due to $e \sqsubseteq_{\mathcal{H}} f$. The second also follows since $\text{ext}_{\mathcal{I}}(e \setminus f) \subseteq \text{ext}_{\mathcal{H}}(e \setminus f) \subseteq f$ where the last inclusion is due to $e \sqsubseteq_{\mathcal{H}} f$. \square

Proof of Proposition 14 If $\mathcal{I}_1 = \mathcal{I}_2$, then it suffices to take $\mathcal{J} = \mathcal{I}_1 = \mathcal{I}_2$. Therefore, assume in the following that $\mathcal{I}_1 \neq \mathcal{I}_2$. Then, necessarily \mathcal{I}_1 and \mathcal{I}_2 are obtained by applying two different reduction operations on \mathcal{H} . We make a case analysis on the types of reductions applied.

(1) *Case (ISO, ISO)* assume that \mathcal{I}_1 is obtained by removing the nonempty set $X_1 \subseteq \text{isol}_{\mathcal{H}}(e_1)$ from hyperedge e_1 , while \mathcal{I}_2 is obtained by removing nonempty $X_2 \subseteq \text{isol}_{\mathcal{H}}(e_2)$ from e_2 with $X_1 \neq X_2$. There are two possibilities.

(1a) $e_1 \neq e_2$. Then, e_2 is still a hyperedge in \mathcal{I}_2 and e_1 is still a hyperedge in \mathcal{I}_1 . By Lemma 11, $\text{isol}_{\mathcal{H}}(e_1) \subseteq \text{isol}_{\mathcal{I}_2}(e_1)$

and $\text{isol}_{\mathcal{H}}(e_2) \subseteq \text{isol}_{\mathcal{I}_1}(e_2)$. Therefore, we can still remove X_2 from \mathcal{I}_1 by means of rule ISO and similarly remove X_1 from \mathcal{I}_2 . Let \mathcal{J}_1 (resp. \mathcal{J}_2) be the result of removing X_2 from \mathcal{I}_1 (resp. \mathcal{I}_2). Then, $\mathcal{J}_1 = \mathcal{J}_2$ (and hence equals triplet \mathcal{J}):

$$\begin{aligned} \text{hyp}(\mathcal{J}_1) &= \text{hyp}(\mathcal{H}) \setminus \{e_1, e_2\} \cup \{e_1 \setminus X_1 \mid e_1 \setminus X_1 \neq \emptyset\} \\ &\quad \cup \{e_2 \setminus X_2 \mid e_2 \setminus X_2 \neq \emptyset\} \\ &= \text{hyp}(\mathcal{J}_2) \end{aligned}$$

$$\text{pred}(\mathcal{J}_1) = \text{pred}(\mathcal{H}) = \text{pred}(\mathcal{J}_2)$$

(1b) $e_1 = e_2$. We show that $X_2 \setminus X_1 \subseteq \text{isol}_{\mathcal{I}_1}(e_1 \setminus X_1)$ and similarly $X_1 \setminus X_2 \subseteq \text{isol}_{\mathcal{I}_2}(e_2 \setminus X_1)$. This suffices because we can then apply ISO to remove $X_2 \setminus X_1$ from \mathcal{I}_1 and $X_1 \setminus X_2$ from \mathcal{I}_2 . In both cases, we reach the same triplet as removing $X_1 \cup X_2 \subseteq \text{isol}_{\mathcal{H}}(e_1)$ from \mathcal{H} .⁷

To see that $X_2 \setminus X_1 \subseteq \text{isol}_{\mathcal{I}_1}(e_1 \setminus X_1)$, let $x \in X_2 \setminus X_1$. We need to show $x \notin \text{fv}_{\mathcal{I}_1}(e_1 \setminus X_1)$ and $x \notin \text{var}(\text{pred}(\mathcal{I}_1))$. Because $x \in X_2 \subseteq \text{isol}_{\mathcal{H}}(e_1)$, we know $x \notin \text{fv}_{\mathcal{H}}(e_1)$. Then, since $x \notin X_1$, also $x \notin \text{fv}_{\mathcal{H}}(e_1 \setminus X_1)$. By Lemma 11, $\text{fv}_{\mathcal{I}_1}(e_1 \setminus X_1) \subseteq \text{fv}_{\mathcal{H}}(e_1 \setminus X_1)$. Therefore, $x \notin \text{fv}_{\mathcal{I}_1}(e_1 \setminus X_1)$. Furthermore, because $x \in \text{isol}_{\mathcal{H}}(e_1)$, we know $x \notin \text{var}(\text{pred}(\mathcal{H}))$. Since $\text{var}(\text{pred}(\mathcal{I}_1)) \subseteq \text{var}(\text{pred}(\mathcal{H}))$ by Lemma 11, also $x \notin \text{var}(\text{pred}(\mathcal{I}_1))$.

$X_1 \setminus X_2 \subseteq \text{isol}_{\mathcal{I}_2}(e_2 \setminus X_1)$ is shown similarly.

(2) *Case (CSE, CSE)* assume that \mathcal{I}_1 is obtained by removing hyperedge e_1 because it is a conditional subset of hyperedge f_1 , while \mathcal{I}_2 is obtained by removing e_2 , conditional subset of f_2 . Since $\mathcal{I}_1 \neq \mathcal{I}_2$, it must be $e_1 \neq e_2$. We need to further distinguish the following cases.

(2a) $e_1 \neq f_2$ and $e_2 \neq f_1$. In this case, e_2 and f_2 remain hyperedges in \mathcal{I}_1 while e_1 and f_1 remain hyperedges in \mathcal{I}_2 . Then, by Lemma 11, $e_2 \sqsubseteq_{\mathcal{I}_1} f_2$ and $e_1 \sqsubseteq_{\mathcal{I}_2} f_2$. Let \mathcal{J}_1 (resp. \mathcal{J}_2) be the triplet obtained by removing e_2 from \mathcal{I}_1 (resp. e_1 from \mathcal{I}_2). Then, $\mathcal{J}_1 = \mathcal{J}_2$ since clearly $\text{out}(\mathcal{J}_1) = \text{out}(\mathcal{J}_2)$ and

$$\begin{aligned} \text{hyp}(\mathcal{J}_1) &= \text{hyp}(\mathcal{H}) \setminus \{e_1, e_2\} = \text{hyp}(\mathcal{J}_2) \\ \text{pred}(\mathcal{J}_1) &= \{\theta \in \text{pred}(\mathcal{H}) \mid \text{var}(\theta) \cap (e_1 \setminus f_1) = \emptyset, \\ &\quad \text{var}(\theta) \cap (e_2 \setminus f_2) = \emptyset\} \\ &= \text{pred}(\mathcal{J}_2) \end{aligned}$$

From this the result follows by taking $\mathcal{J} = \mathcal{J}_1 = \mathcal{J}_2$.

(2b) $e_1 \neq f_2$ but $e_2 = f_1$. Then, $e_1 \sqsubseteq_{\mathcal{H}} e_2$ and $e_2 \sqsubseteq_{\mathcal{H}} f_2$ with $f_2 \neq e_1$. It suffices to show that $e_1 \sqsubseteq_{\mathcal{H}} f_2$ and $e_1 \setminus f_2 = e_1 \setminus f_1$, because then (CSE) due to $e_1 \sqsubseteq_{\mathcal{H}} f_1$ has the same effect as CSE on $e_1 \sqsubseteq_{\mathcal{H}} f_2$, and we can apply the reasoning of case (2a) because $e_1 \neq f_2$ and $e_2 \neq f_2$.

⁷ Should $X_2 \setminus X_1$ be empty, we don't actually need to do anything on \mathcal{I}_1 : $X_1 \cup X_2$ is already removed from it. A similar remark holds for \mathcal{I}_2 when $X_1 \setminus X_2$ is empty.

We first show $e_1 \setminus f_2 = e_1 \setminus f_1$. Let $x \in e_1 \setminus f_2$ and suppose for the purpose of contradiction that $x \in e_2 = f_1$. Then, since $e_1 \neq e_2$, $x \in jv_{\mathcal{H}}(e_2) \subseteq f_2$ where the last inclusion is due to $e_2 \sqsubseteq_{\mathcal{H}} f_2$. Hence, $e_1 \setminus f_2 \subseteq e_1 \setminus f_1$. Conversely, let $x \in e_1 \setminus f_1$. Since $f_1 = e_2$, $x \notin e_2$. Suppose for the purpose of contradiction that $x \in f_2$. Because $e_1 \neq f_2$, $x \in jv_{\mathcal{H}}(e_1) \subseteq e_2$ where the last inclusion is due to $e_1 \sqsubseteq_{\mathcal{H}} e_2$. Therefore, $e_2 \setminus f_1 = e_1 \setminus f_2$.

To show that $e_1 \sqsubseteq_{\mathcal{H}} f_2$, let $x \in jv_{\mathcal{H}}(e_1)$. Because $e_1 \sqsubseteq_{\mathcal{H}} e_2$, $x \in e_2$. Because x occurs in two distinct hyperedges in \mathcal{H} , also $x \in jv_{\mathcal{H}}(e_2)$. Then, because $e_2 \sqsubseteq_{\mathcal{H}} f_2$, $x \in f_2$. Hence, $jv_{\mathcal{H}}(e_1) \subseteq f_2$. It remains to show $ext_{\mathcal{H}}(e_1 \setminus f_2) \subseteq f_2$. To this end, let $x \in ext_{\mathcal{H}}(e_1 \setminus f_2)$ and suppose for the purpose of contradiction that $x \notin f_2$. By definition of ext there exists $\theta \in pred(\mathcal{H})$ and $y \in var(\theta) \cap (e_1 \setminus f_2)$ such that $x \in var(\theta) \setminus (e_1 \setminus f_2)$. In particular, $y \notin f_2$. Since $e_1 \setminus f_2 = e_1 \setminus e_2$, $y \in var(\theta) \cap (e_1 \setminus e_2)$ and $x \in var(\theta) \setminus (e_1 \setminus e_2)$. Thus, $x \in ext_{\mathcal{H}}(e_1 \setminus e_2)$. Then, since $e_1 \sqsubseteq_{\mathcal{H}} e_2$, $x \in e_2$. Thus, $x \in e_2 \setminus f_2$ since $x \notin f_2$. Hence, $x \in var(\theta) \cap (e_2 \setminus f_2)$. Furthermore, since $y \notin e_2$ also $y \notin e_2 \setminus f_2$. Hence, $y \in var(\theta) \setminus (e_2 \setminus f_2)$. But then θ shows that $y \in ext_{\mathcal{H}}(e_2 \setminus f_2)$. Then, by because $e_2 \sqsubseteq_{\mathcal{H}} f_2$, also $y \in f_2$ which yields the desired contradiction.

(2c) $e_1 = f_2$ but $e_2 \neq f_1$. Similar to case (2b).

(2d) $e_1 = f_2$ and $e_2 = f_1$. Then, $e_1 \sqsubseteq_{\mathcal{H}} e_2$ and $e_2 \sqsubseteq_{\mathcal{H}} e_1$ and $e_1 \neq e_2$. Let \mathcal{K}_1 (resp. \mathcal{K}_2) be the triplet obtained by applying (FLT) to remove all $\theta \in pred(\mathcal{I}_1)$ (resp. $\theta \in pred(\mathcal{I}_2)$) for which $var(\theta) \subseteq var(e_2)$ (resp. $var(\theta) \subseteq var(e_1)$). Furthermore, let \mathcal{J}_1 (resp. \mathcal{J}_2) be the triplet obtained by applying ISO to removing $isol_{\mathcal{I}_1}(e_2)$ from \mathcal{K}_1 (resp. removing $isol_{\mathcal{I}_2}(e_1)$ from \mathcal{K}_2). Here, we take $\mathcal{J}_1 = \mathcal{K}_1$ if $isol_{\mathcal{K}_1}(e_2)$ is empty (and similarly for \mathcal{J}_2). Then, clearly $\mathcal{H} \rightsquigarrow \mathcal{I}_1 \rightsquigarrow^* \mathcal{K}_1 \rightsquigarrow^* \mathcal{J}_1$ and $\mathcal{H} \rightsquigarrow \mathcal{I}_2 \rightsquigarrow^* \mathcal{K}_2 \rightsquigarrow^* \mathcal{J}_2$. The result then follows by showing that $\mathcal{J}_1 = \mathcal{J}_2$. Toward this end, first observe that $out(\mathcal{J}_1) = out(\mathcal{K}_1) = out(\mathcal{I}_1) = out(\mathcal{H}) = out(\mathcal{I}_2) = out(\mathcal{K}_2) = out(\mathcal{J}_2)$. Next, we show that $pred(\mathcal{J}_1) = pred(\mathcal{J}_2)$. We first observe that $pred(\mathcal{J}_1) = pred(\mathcal{K}_1)$ and $pred(\mathcal{J}_2) = pred(\mathcal{K}_2)$ since the ISO operation does not remove predicates. Then, observe that

$$\begin{aligned} pred(\mathcal{K}_1) &= \{\theta \in pred(\mathcal{I}_1) \mid var(\theta) \not\subseteq var(e_2)\} \\ &= \{\theta \in pred(\mathcal{H}) \mid var(\theta) \cap (e_1 \setminus e_2) = \emptyset \text{ and } var(\theta) \not\subseteq e_2\}, \\ pred(\mathcal{K}_2) &= \{\theta \in pred(\mathcal{I}_2) \mid var(\theta) \not\subseteq e_1\} \\ &= \{\theta \in pred(\mathcal{H}) \mid var(\theta) \cap (e_2 \setminus e_1) = \emptyset \text{ and } var(\theta) \not\subseteq e_1\}. \end{aligned}$$

We only show the reasoning for $pred(\mathcal{K}_1) \subseteq pred(\mathcal{K}_2)$, the other direction being similar. Let $\theta \in pred(\mathcal{K}_1)$. Then, $var(\theta) \cap (e_1 \setminus e_2) = \emptyset$ and $var(\theta) \not\subseteq e_2$. Since $var(\theta) \not\subseteq e_2$ there exists $y \in var(\theta) \setminus e_2$. Then, because $var(\theta) \cap (e_1 \setminus e_2) = \emptyset$, $y \notin e_1$. Thus, $var(\theta) \not\subseteq e_1$. Now, suppose for the purpose

of obtaining a contradiction, that $var(\theta) \cap (e_2 \setminus e_1) \neq \emptyset$. Then take $z \in var(\theta) \cap (e_2 \setminus e_1)$. But then $y \in ext_{\mathcal{H}}(e_2 \setminus e_1)$. Hence, $y \in e_1$ because $e_2 \sqsubseteq_{\mathcal{H}} e_1$, which yields the desired contradiction with $y \notin e_2$. Therefore, $var(\theta) \cap (e_2 \setminus e_1) = \emptyset$, as desired. Hence, $\theta \in pred(\mathcal{K}_2)$.

It remains to show that $hyp(\mathcal{J}_1) = hyp(\mathcal{J}_2)$. To this end, first observe

$$\begin{aligned} hyp(\mathcal{J}_1) &= hyp(\mathcal{K}_1) \setminus \{e_2\} \cup \{e_2 \setminus isol_{\mathcal{K}_1}(e_2)\}, \\ &= hyp(\mathcal{H}) \setminus \{e_1\} \setminus \{e_2\} \cup \{e_2 \setminus isol_{\mathcal{K}_1}(e_2)\}, \\ hyp(\mathcal{J}_2) &= hyp(\mathcal{K}_2) \setminus \{e_1\} \cup \{e_1 \setminus isol_{\mathcal{K}_2}(e_1)\} \\ &= hyp(\mathcal{H}) \setminus \{e_2\} \setminus \{e_1\} \cup \{e_1 \setminus isol_{\mathcal{K}_2}(e_1)\}. \end{aligned}$$

Clearly, $hyp(\mathcal{J}_1) = hyp(\mathcal{J}_2)$ if $e_2 \setminus isol_{\mathcal{K}_1}(e_2) = e_1 \setminus isol_{\mathcal{K}_2}(e_1)$.

We only show $e_2 \setminus isol_{\mathcal{K}_1}(e_2) \subseteq e_1 \setminus isol_{\mathcal{K}_2}(e_1)$, the other inclusion being similar. Let $x \in e_2 \setminus isol_{\mathcal{K}_1}(e_2)$. Since $x \notin isol_{\mathcal{K}_1}(e_2)$ one of the following hold.

- $x \in out(\mathcal{K}_1)$. But then, $x \in out(\mathcal{K}_1) = out(\mathcal{I}_1) = out(\mathcal{H}) = out(\mathcal{I}_2) = out(\mathcal{K}_2)$. In particular, x is an equi-join variable in \mathcal{H} and \mathcal{K}_2 . Then, $x \in jv_{\mathcal{H}}(e_2) \subseteq e_1$ because $e_2 \sqsubseteq_{\mathcal{H}} e_1$. From this and the fact that x remains an equi-join variable in \mathcal{K}_2 , we obtain $x \in e_1 \setminus isol_{\mathcal{K}_2}(e_1)$.
- x occurs in e_2 and in some hyperedge g in \mathcal{K}_1 with $g \neq e_2$. Since e_1 is not in \mathcal{K}_1 also $g \neq e_1$. Since every hyperedge in \mathcal{K}_1 is in \mathcal{I}_1 and every hyperedge in \mathcal{I}_1 is in \mathcal{H} , also g is in \mathcal{H} . But then, x occurs in two distinct hyperedges in \mathcal{H} , namely e_2 and g , and hence $x \in jv_{\mathcal{H}}(e_2) \subseteq e_1$ because $e_2 \sqsubseteq_{\mathcal{H}} e_1$. However, because x also occurs in g which must also be in \mathcal{I}_2 and therefore also in \mathcal{K}_2 , x also occurs in two distinct hyperedges in \mathcal{K}_2 , namely e_1 and g . Therefore, $x \in jv_{\mathcal{K}_2}(e_1)$ and hence $x \in e_1 \setminus isol_{\mathcal{K}_2}(e_1)$, as desired.
- $x \in var(pred(\mathcal{K}_1))$. Then, there exists $\theta \in pred(\mathcal{K}_1)$ such that $x \in var(\theta)$. Since $pred(\mathcal{K}_1) = pred(\mathcal{K}_2)$, $\theta \in pred(\mathcal{K}_2)$. As such, $\theta \in pred(\mathcal{H})$, $var(\theta) \cap (e_2 \setminus e_1) = \emptyset$, and $var(\theta) \not\subseteq e_1$. But then, since $x \in var(\theta)$; $x \in e_2$; and $var(\theta) \cap (e_2 \setminus e_1) = \emptyset$, it must be the case that $x \in e_1$. As such, $x \in e_1$ and $x \in var(\mathcal{K}_2)$. Hence, $x \in e_1 \setminus isol_{\mathcal{K}_2}(e_1)$.

(3) *Case (ISO, CSE)* assume that \mathcal{I}_1 is obtained by removing the nonempty set of isolated variables $X_1 \subseteq isol_{\mathcal{H}}(e_1)$ from e_1 , while \mathcal{I}_2 is obtained by removing hyperedge e_2 , conditional subset of hyperedge f_2 . We may assume w.l.o.g. that $e_1 \neq isol_{\mathcal{H}}(e_1)$: if $e_1 = isol_{\mathcal{H}}(e_1)$, then the ISO operation removes the complete hyperedge e_1 . However, because no predicate in \mathcal{H} shares any variable with e_1 , it is readily verified that $e_1 \sqsubseteq_{\mathcal{H}} e_2$ and thus the removal of e_1 can also

be seen as an application of CSE on e_1 ,⁸ and we are hence back in case (2).

Now reason as follows. Because $e_2 \sqsubseteq_{\mathcal{H}} f_2$ and because isolated variables of e_1 occur in no other hyperedge in \mathcal{H} , it must be the case that $e_2 \cap X_1 = \emptyset$. In particular, e_1 and e_2 must hence be distinct. Therefore, $e_1 \in \text{hyp}(\mathcal{I}_2)$ and $e_2 \in \text{hyp}(\mathcal{I}_1)$. By Lemma 11, we can apply ISO on \mathcal{I}_2 to remove X_1 from e_1 . It then suffices to show that e_2 remains a conditional subset of some hyperedge f'_2 in \mathcal{I}_1 with $e_2 \setminus f_2 = e_2 \setminus f'_2$. Indeed, we can then use ECQ to remove e_2 from $\text{hyp}(\mathcal{I}_1)$ as well as predicates θ with $\text{var}(\theta) \cap (e_2 \setminus f_2) \neq \emptyset$ from $\text{pred}(\mathcal{I}_1)$. This clearly yields the same triplet as the one obtained by removing X_1 from e_1 in \mathcal{I}_2 . We need to distinguish two cases.

(3a) $f_2 \neq e_1$. Then, $f_2 \in \text{hyp}(\mathcal{I}_1)$ and hence $e_2 \sqsubseteq_{\mathcal{I}_1} f_2$ by Lemma 11. We hence take $f'_2 = f_2$.

(3b) $f_2 = e_1$. Then, we take $f'_2 = e_1 \setminus X$. Since $e_1 \neq \text{isol}_{\mathcal{H}}(e_1)$, it follows that $e_1 \setminus X_1 \neq \emptyset$. Therefore, $f'_2 = e_1 \setminus X_1 \in \text{hyp}(\mathcal{I}_1)$. Furthermore, since $X \subseteq \text{isol}_{\mathcal{H}}(e_1)$, no variable in X is in any other hyperedge in \mathcal{H} . In particular $X \cap e_2 = \emptyset$. Therefore, $e_2 \setminus f'_2 = e_2 \setminus (e_1 \setminus X) = (e_2 \setminus e_1) \cup (e_2 \cap X) = e_2 \setminus e_1 \setminus e_1 = e_2 \setminus f_2$. It remains to show that $e_2 \sqsubseteq_{\mathcal{I}_1} e_1 \setminus X_1$.

- $jv_{\mathcal{I}_1}(e_2) \subseteq e_1 \setminus X_1$. Let $x \in jv_{\mathcal{I}_1}(e_2)$. By Lemma 11, $x \in jv_{\mathcal{I}_1}(e_2) \subseteq jv_{\mathcal{H}}(e_2) \subseteq e_1$ where the last inclusion is due to $e_2 \sqsubseteq_{\mathcal{H}} e_1$. In particular, x is an equijoin variable in \mathcal{H} . But then it cannot be an isolated variable in any hyperedge. Therefore, $x \notin X_1$.
- $ext_{\mathcal{I}_1}(e_2 \setminus e_1) \subseteq e_1 \setminus X$. Let $x \in ext_{\mathcal{I}_1}(e_2 \setminus e_1)$. Then, $x \in ext_{\mathcal{I}_1}(e_2 \setminus e_1) \subseteq ext_{\mathcal{H}}(e_2 \setminus e_1) \subseteq e_1$ where the first inclusion is by Lemma 11 and the second by $e_2 \sqsubseteq_{\mathcal{H}} e_1$. Then, because $x \in ext_{\mathcal{H}}(e_2 \setminus e_1)$, it follows from the definition of ext that x occurs in some predicate in $pred(\mathcal{H})$. However, X is disjoint with $var(pred(\mathcal{H}))$ since it consists only of isolated variables. Therefore, $x \notin X$.

(4) *Case (ISO, FLT)* Assume that \mathcal{I}_1 is obtained by removing the nonempty set $X_1 \subseteq \text{isol}_{\mathcal{H}}(e_1)$ from hyperedge e_1 , while \mathcal{I}_2 is obtained by removing all predicates in the nonempty set $\Theta \subseteq \text{pred}(\mathcal{H})$ with $\text{var}(\Theta) \subseteq e_2$ for some hyperedge e_2 in $\text{hyp}(\mathcal{H})$. Observe that $e_1 \in \text{hyp}(\mathcal{I}_2)$. By Lemma 11, $X \subseteq \text{isol}_{\mathcal{H}}(e_1) \subseteq \text{isol}_{\mathcal{I}_2}(e_1)$. Therefore, we may apply reduction operation (ISO) on \mathcal{I}_2 to remove X_1 from e_1 . We will now show that, similarly, we may still apply (FLT) on \mathcal{I}_1 to remove all predicates in Θ from $\text{pred}(\mathcal{I}_1) = \text{pred}(\mathcal{H})$. The two operations hence commute, and clearly, the resulting triplets in both cases is the same. We distinguish two possibilities. (i) $e_1 \neq e_2$. Then, $e_2 \in \mathcal{I}_1$ and, $\text{var}(\Theta) \subseteq e_2$ and, since (ISO) does not remove predicates, $\Theta \subseteq \text{pred}(\mathcal{H}) =$

⁸ Note that, since e_1 does not share variables with any predicate, the CSE operation also does not remove any predicates from \mathcal{H}_1 , similar to the ISO operation and hence yields \mathcal{I}_1 .

$pred(I_1)$. As such the (FLT) operation indeed applies to remove all predicates in Θ from $pred(I_1)$. (ii) $e_1 = e_2$. Then, since $X \subseteq isol_{\mathcal{H}}(e_1)$ and isolated variables do not occur in any predicate, $X \cap var(\Theta) = \emptyset$. Then, since $var(\Theta) \subseteq e_2 = e_1$, it follows that also $var(\Theta) \subseteq e_1 \setminus X$. In particular, since we disallow nullary predicates and Θ is nonempty, $e_1 \setminus X \neq \emptyset$. Thus, $e_1 \setminus X \in hyp(I_1)$ and hence operation (FLT) applies indeed applies to remove all predicates in Θ from $pred(I_1)$

(5) *Case (CSE, FLT)* assume that \mathcal{I}_1 is obtained by removing hyperedge e_1 , conditional subset of e_2 in \mathcal{H} , while \mathcal{I}_2 is obtained by removing all predicates in the nonempty set $\Theta \subseteq \text{pred}(\mathcal{H})$ with $\text{var}(\Theta) \subseteq e_3$ for some hyperedge $e_3 \in \text{hyp}(\mathcal{H})$. Since the (FLT) operation does not remove any hyperedges, e_1 and e_2 are in $\text{hyp}(\mathcal{I}_2)$. Then, since $e_1 \sqsubseteq_{\mathcal{H}} e_2$ also $e_1 \sqsubseteq_{\mathcal{I}_2} e_2$ by Lemma 11. Therefore, we may apply reduction operation (CSE) on \mathcal{I}_2 to remove e_1 from $\text{hyp}(\mathcal{I}_2)$ as well as all predicates $\theta \in \text{pred}(\mathcal{I}_2)$ for which $\text{var}(\theta) \cap (e_1 \setminus e_2) \neq \emptyset$. Let \mathcal{J}_2 be the triplet resulting from this operation. We will show that, similarly, we may apply (FLT) on \mathcal{I}_1 to remove all predicates in $\Theta \cap \text{pred}(\mathcal{I}_1)$ from $\text{pred}(\mathcal{I}_1)$, resulting in a triplet \mathcal{J}_1 . Observe that necessarily, $\mathcal{J}_1 = \mathcal{J}_2$ (and hence they form the triplet \mathcal{J}). Indeed, $\text{out}(\mathcal{J}_1) = \text{out}(\mathcal{I}_1) = \text{out}(\mathcal{H}) = \text{out}(\mathcal{I}_2) = \text{out}(\mathcal{J}_2)$ since reduction operations never modify output variables. Moreover,

$$\begin{aligned} \text{hyp}(\mathcal{J}_1) &= \text{hyp}(\mathcal{I}_1) \\ &= \text{hyp}(\mathcal{H}) \setminus \{e_1\} \\ &= \text{hyp}(\mathcal{I}_2) \setminus \{e_1\} \\ &= \text{hyp}(\mathcal{J}_2) \end{aligned}$$

where the first and third equality is due to fact that (FLT) does not modify the hypergraph of the triplet it operates on. Finally, observe

$$\begin{aligned}
pred(\mathcal{J}_1) &= pred(\mathcal{I}_1) \setminus (\Theta \cap pred(\mathcal{I}_1)) \\
&= pred(\mathcal{I}_1) \setminus \Theta \\
&= \{\theta \in pred(\mathcal{H}) \mid var(\theta) \cap (e_1 \setminus e_2) = \emptyset\} \setminus \Theta \\
&= \{\theta \in pred(\mathcal{H}) \setminus \Theta \mid var(\theta) \cap (e_1 \setminus e_2) = \emptyset\} \\
&= \{\theta \in pred(\mathcal{I}_2) \mid var(\theta) \cap (e_1 \setminus e_2) = \emptyset\} \\
&= pred(\mathcal{J}_2)
\end{aligned}$$

It remains to show that we may apply (FLT) on \mathcal{I}_1 to remove all predicates in $\Theta \cap \text{pred}(\mathcal{I}_1)$, resulting in a triplet \mathcal{J}_1 . There are two possibilities.

- $e_3 \neq e_1$. Then, $e_3 \in \mathcal{I}_1$, $\Theta \cap \text{pred}((I_1)) \subseteq \text{pred}(\mathcal{I}_1)$, and $\text{var}(\Theta \cap \text{pred}(\mathcal{I}_1)) \subseteq \text{var}(\Theta) \subseteq e_3$. Hence, the (FLT) operation indeed applies to \mathcal{I}_1 to remove all predicates in $\Theta \cap \text{pred}(\mathcal{I}_1)$.
- $e_3 = e_1$. In this case, we claim that for every $\theta \in \Theta \cap \text{pred}(\mathcal{I}_1)$, we have $\text{var}(\theta) \subseteq e_2$. As such, $\text{var}(\Theta \cap$

$pred(\mathcal{I}_1) \subseteq e_2$. Since $e_2 \in hyp(\mathcal{I}_1)$ and $\Theta \cap pred(\mathcal{I}_1) \subseteq pred(\mathcal{I}_1)$, we may hence apply (FLT) to remove all predicates in $\Theta \cap pred(\mathcal{I}_1)$ from \mathcal{I}_1 . Concretely, let $\theta \in \Theta \cap pred(\mathcal{I}_1)$. Because, in order to obtain \mathcal{I}_1 , (CSE) removes all predicates from \mathcal{H} that share a variable with $e_1 \setminus e_2$, we have $var(\theta) \cap (e_1 \setminus e_2) = \emptyset$. Moreover, because $\theta \in \Theta$, $var(\theta) \subseteq e_1$. Hence, $var(\theta) \subseteq e_2$, as desired.

The remaining cases, (CSE, ISO), (FLT, ISO), and (FLT, CSE), are symmetric to case (3), (4), and (5), respectively. \square

B.2 Proof of Proposition 9

Proposition 9 *For every GJT pair there exists an equivalent canonical pair.*

Proof Let T be a GJT. The proof proceeds in three steps.

Step 1 Let T_1 be the GJT obtained from T by (i) removing all predicates from T , and (ii) creating a new root node r that is labeled by \emptyset and attaching the root of T to it, labeled by the empty set of predicates. T_1 satisfies the first canonicity condition, but is not equivalent to T because it has none of T 's predicates. Now re-add the predicates in T to T_1 as follows. For each edge $m \rightarrow n$ in T and each predicate $\theta \in pred_T(m \rightarrow n)$, if $var(\theta) \cap (var(n) \setminus var(m)) \neq \emptyset$ then add θ to $pred_{T_1}(m \rightarrow n)$. Otherwise, if $var(\theta) \cap (var(n) \setminus var(m)) = \emptyset$, do the following. First, observe that, by definition of GJTs, $var(\theta) \subseteq var(n) \cup var(m)$. Because $var(\theta) \cap (var(n) \setminus var(m)) = \emptyset$ this implies $var(\theta) \subseteq var(m)$. Because we disallow nullary predicates, $var(m) \neq \emptyset$. Let a be the first ancestor of m in T_1 such that $var(\theta) \not\subseteq var(a)$. Such an ancestor exists because the root of T_1 is labeled \emptyset . Let b be the child of a in T_1 . Since a is the first ancestor of m with $var(\theta) \not\subseteq var(a)$, $var(\theta) \subseteq var(b)$. Therefore, $var(\theta) \subseteq var(b) \cup var(a)$ and $var(\theta) \cap (var(b) \setminus var(a)) \neq \emptyset$. As such, add θ to $pred_{T_1}(a \rightarrow b)$. After having done this for all predicates in T , T_1 becomes equivalent to T and satisfies canonicity conditions (1) and (3). Then, take $N_1 = N \cup \{r\}$. Clearly, N_1 is a connex subset of T_1 and $var(N) = var(N')$.

Therefore, (T_1, N_1) is equivalent to (T, N) .

Step 2 Let T_2 be obtained from T_1 by adding, for each leaf node l in T_1 a new interior node n_l labeled by $var(l)$ and inserting it in-between l and its parent in T_1 , i.e., if l has parent p in T_1 , then we have $p \rightarrow n_l \rightarrow l$ in T_2 with $pred_{T_2}(p \rightarrow n_l) = pred_{T_1}(p \rightarrow n)$ and $pred_{T_2}(n_l \rightarrow l) = \emptyset$.⁹ Furthermore, let N_2 be the connex subset of T_2 obtained by replacing every leaf node l in N_1 by its newly inserted node n_l . Clearly, $var(N_2) = var(N_1) = var(N)$ because $var(l) = var(n_l)$ for every leaf l of T_1 . By our construction, (T_2, N_2) is equivalent to (T, N) ; T_2 satisfies canonicity conditions (1), (2), and (4); and N_2 is canonical.

⁹ Note that all leaves have a parent since the root of T_1 is an interior node labeled by \emptyset .

Step 3 It remains to enforce condition (3). To this end, observe that, by the connectedness condition of GJTs, T_2 violates canonicity condition (3) if and only if there exist internal nodes m and n where m is the parent of n such that $var(m) = var(n)$. In this case, we call n a culprit node. We will now show how to obtain an equivalent pair (U, M) that removes a single culprit node; the final result is then obtained by iterating this reasoning until all culprit nodes have been removed.

The culprit removal procedure is essentially the reverse of the binarization procedure of Fig. 9. Concretely, let n be a culprit node with parent m and let n_1, \dots, n_k be the children of n in T_2 . Let U be the GJT obtained from T_2 by removing n and attaching all children n_i of n as children to m with edge label $pred_U(m \rightarrow n_i) = pred_{T_2}(n \rightarrow n_i)$, for $1 \leq i \leq k$. Because $var(n) = var(m)$, the result is still a valid GJT. Moreover, because $var(n) = var(m)$ and T_2 satisfied condition (4), we had $pred_{T_2}(m \rightarrow n) = \emptyset$, so no predicate was lost by the removal of n . Finally, define M as follows. If $n \in N_2$, then set $M = N_2 \setminus \{n\}$; otherwise, set $M = N_2$. In the former case, since N_2 is connex and $n \in N_2$, m must also be in N_2 . It is hence in M . Therefore, in both cases, $var(N) = var(N_2) = var(M)$. Furthermore, it is straightforward to check that M is a connex subset of U . Finally, since N_2 consisted only of interior nodes of T_2 , M consists only of interior nodes of U and hence remains canonical. \square

B.3 Proof of Lemma 5

We first require a number of auxiliary results.

We first make the following observations regarding canonical GJT pairs.

Lemma 12 *Let (T, N) be a canonical GJT pair; let n be a frontier node of N and let m be the parent of n in T .*

1. $x \notin var(N \setminus \{n\})$, for every $x \in var(n) \setminus var(m)$.
2. $hyp(T, N \setminus \{n\}) = hyp(T, N) \setminus \{var(n)\}$.
3. $\theta \notin pred(m \rightarrow n)$, for every $\theta \in pred(T, N \setminus \{n\})$.
4. $pred(T, N \setminus \{n\}) = pred(T, N) \setminus pred(m \rightarrow n)$.
5. $pred(m \rightarrow n) = \{\theta \in pred(T, N) \mid var(\theta) \cap (var(n) \setminus var(m)) \neq \emptyset\}$.
6. $pred(T, N \setminus \{n\}) = \{\theta \in pred(T, N) \mid var(\theta) \cap (var(n) \setminus var(m)) = \emptyset\}$.

Proof (1) Let $x \in var(n) \setminus var(m)$ and let c be a node in $N \setminus \{n\}$. Clearly the unique undirected path between c and n in T must pass through m . Because $x \notin var(m)$, it follows from the connectedness condition of GJTs that also $x \notin var(c)$. As such, $x \notin var(N \setminus \{n\})$.

(2) The \supseteq direction is trivial. For the \subseteq direction, assume that $m \in N \setminus \{n\}$ with $var(m) \neq \emptyset$. Then, clearly $m \in N$ and hence $var(m) \in hyp(T, N)$. Further-

more, because N is canonical, both m and n are interior nodes in T . Then, because T is canonical and $m \neq n$, we have $\text{var}(m) \neq \text{var}(n)$. Therefore, $\text{var}(m) \in \text{hyp}(T, N) \setminus \{\text{var}(n)\}$.

- (3) Let $\theta \in \text{pred}(T, N \setminus n)$. Then, θ occurs on the edge between two nodes in $N \setminus n$, say $m' \rightarrow n'$. By definition of GJTs, $\text{var}(\theta) \subseteq \text{var}(n') \cup \text{var}(m') \subseteq \text{var}(N \setminus \{n\})$. Now suppose for the purpose of contradiction that also $\theta \in \text{pred}(m \rightarrow n)$. Because T is nice, there is some $x \in \text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) \neq \emptyset$. Hence, by (1), $x \notin \text{var}(N \setminus \{n\})$, which contradicts $\text{var}(\theta) \subseteq \text{var}(N \setminus \{n\})$.
- (4) Clearly, $\text{pred}(T, N) \setminus \text{pred}(m \rightarrow n) \subseteq \text{pred}(T, N \setminus \{n\})$. The converse inclusion follows from (3).
- (5) The \subseteq direction follows from the fact that m and n are in N , and T is nice. To also see \supseteq , let $\theta \in \text{pred}(T, N)$ with $\text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) \neq \emptyset$. There exists $x \in \text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m))$. By (1), $x \notin \text{var}(N \setminus \{n\})$. Therefore, θ cannot occur between edges in $N \setminus \{n\}$ in T . Since it nevertheless occurs in $\text{pred}(T, N)$, it must hence occur in $\text{pred}(m \rightarrow n)$.
- (6) Follows directly from (4) and (5). \square

Lemma 13 *Let (T, N) be a canonical GJT pair, let n be a frontier node of N and let m be the parent of n in T . Let $\bar{z} \subseteq \text{var}(N \setminus \{n\})$.*

1. $\text{var}(n) \sqsubseteq_{\mathcal{H}(T, N, \bar{z})} \text{var}(m)$.
2. $x \notin \text{fv}(\mathcal{H}(T, N, \bar{z}))$, for every $x \in (\text{var}(n) \setminus \text{var}(m))$.

Proof For reasons of parsimony, let $\mathcal{H} = \mathcal{H}(T, N, \bar{z})$. We first prove (2) and then (1).

(2) Let $x \in \text{var}(n) \setminus \text{var}(m)$. By Lemma 12(1), $x \notin \text{var}(N \setminus \{n\})$. Therefore, x occurs in $\text{var}(n)$ in \mathcal{H} and in no other hyperedge. Furthermore, because $\bar{z} \subseteq \text{var}(N \setminus \{n\})$, also $x \notin \bar{z}$. Hence, $x \notin \text{fv}_{\mathcal{H}}(\text{var}(n))$.

(1) We need to show that $\text{fv}_{\mathcal{H}}(\text{var}(n)) \subseteq \text{var}(m)$ and $\text{ext}_{\mathcal{H}}(\text{var}(n) \setminus \text{var}(m)) \subseteq \text{var}(m)$. Let $x \in \text{fv}_{\mathcal{H}}(\text{var}(n))$. By contraposition of (2), we know that $x \notin (\text{var}(n) \setminus \text{var}(m))$. Therefore, $x \in \text{var}(m)$ and thus $\text{fv}_{\mathcal{H}}(\text{var}(n)) \subseteq \text{var}(m)$. To show $\text{ext}_{\mathcal{H}}(\text{var}(n) \setminus \text{var}(m)) \subseteq \text{var}(m)$, let $y \in \text{ext}_{\mathcal{H}}(\text{var}(n) \setminus \text{var}(m))$. Then, $y \notin \text{var}(n) \setminus \text{var}(m)$ and there exists $\theta \in \text{pred}(T, N)$ with $\text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) \neq \emptyset$ and $y \in \text{var}(\theta)$. By Lemma 12(5), $\theta \in \text{pred}_T(m \rightarrow n)$. Thus, $y \in \text{var}(m) \cup \text{var}(n)$. Since also $y \notin \text{var}(n) \setminus \text{var}(m)$, it follows that $y \in \text{var}(m)$. Therefore, $\text{ext}_{\mathcal{H}}(\text{var}(n) \setminus \text{var}(m)) \subseteq \text{var}(m)$. \square

Lemma 14 *Let (T, N) be a canonical GJT pair and let n be a frontier node of N . Then, $\mathcal{H}(T, N, \bar{z}) \rightsquigarrow^* \mathcal{H}(T, N \setminus \{n\}, \bar{z})$ for every $\bar{z} \subseteq \text{var}(N \setminus \{n\})$.*

Proof For reasons of parsimony, let us abbreviate $\mathcal{H}_1 = \mathcal{H}(T, N, \bar{z})$ and $\mathcal{H}_2 = \mathcal{H}(T, N \setminus \{n\}, \bar{z})$. We make the following case analysis.

Case 1 Node n is the root in N . Because the root of a canonical tree is labeled by \emptyset , we have $\text{var}(n) = \emptyset$. Since n is a frontier node of N , $N = \{n\}$. Thus, $\text{hyp}(T, N) = \emptyset$ and $\text{hyp}(T, N \setminus \{n\}) = \emptyset$. Furthermore, $\text{pred}(T, N) = \text{pred}(T, N \setminus \{n\}) = \emptyset$ and $\bar{z} \subseteq \text{var}(N \setminus \{n\}) = \text{var}(\emptyset) = \emptyset$. As such, both \mathcal{H}_1 and \mathcal{H}_2 are the empty triplet $(\emptyset, \emptyset, \emptyset)$. Therefore, $\mathcal{H}_1 \rightsquigarrow^* \mathcal{H}_2$.

Case 2 n has parent m in N and $\text{var}(m) \neq \emptyset$. Then, $\text{var}(n) \neq \emptyset$ since in a canonical tree the root node is the only interior node that is labeled by the empty hyperedge. Therefore, $\text{var}(n) \in \text{hyp}(T, N)$, $\text{var}(m) \in \text{hyp}(T, N)$, and $\text{var}(n) \sqsubseteq_{\mathcal{H}_1} \text{var}(m)$ by Lemma 13(1). We can hence apply reduction (CSE) to remove $\text{var}(n)$ from $\text{hyp}(\mathcal{H}_1)$ and all predicates that intersect with $\text{var}(n) \setminus \text{var}(m)$ from $\text{pred}(\mathcal{H}_1)$. By Lemma 12(2) and 12(6) the result is exactly \mathcal{H}_2 :

$$\begin{aligned} \text{hyp}(\mathcal{H}_2) &= \text{hyp}(T, N \setminus \{n\}) \\ &= \text{hyp}(T, N) \setminus \{\text{var}(n)\} = \text{hyp}(\mathcal{H}_1) \setminus \{\text{var}(n)\} \\ \text{pred}(\mathcal{H}_2) &= \text{pred}(T, N \setminus \{n\}) \\ &= \{\theta \in \text{pred}(T, N) \mid \text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) = \emptyset\} \\ &= \{\theta \in \text{pred}(\mathcal{H}_1) \mid \text{var}(\theta) \cap (\text{var}(n) \setminus \text{var}(m)) = \emptyset\} \end{aligned}$$

Case 3 n has parent m in N and $\text{var}(m) = \emptyset$. Then, $\text{var}(n) \neq \emptyset$ since in a canonical tree the root node is the only interior node that is labeled by the empty hyperedge. By definition of GJTs, it follows that for every $\theta \in \text{pred}(m \rightarrow n)$, we have $\text{var}(\theta) \subseteq \text{var}(n) \cup \text{var}(m) = \text{var}(n)$. In other words: all $\theta \in \text{pred}(m \rightarrow n)$ are filters. As such, we can use reduction (FLT) to remove all predicates in $\text{pred}(m \rightarrow n)$ from \mathcal{H}_1 . This yields a triplet \mathcal{I} with the same hypergraph as \mathcal{H}_1 , same set of output variables as \mathcal{H}_1 , and

$$\begin{aligned} \text{pred}(\mathcal{I}) &= \text{pred}(\mathcal{H}_1) \setminus \text{pred}_T(m \rightarrow n) \\ &= \text{pred}(T, N) \setminus \text{pred}_T(m \rightarrow n) \\ &= \text{pred}(T, N \setminus \{n\}) = \text{pred}(\mathcal{H}_2), \end{aligned}$$

where the third equality is due to Lemma 12(4). We claim that every variable in e is isolated in \mathcal{I} . From this the result follows, because then we can apply (ISO) to remove the entire hyperedge $\text{var}(e)$ from $\text{hyp}(\mathcal{I}) = \text{hyp}(\mathcal{H}_1)$ while preserving $\text{out}(\mathcal{I})$ and $\text{pred}(\mathcal{I})$. The resulting triplet hence equals \mathcal{H}_2 . To see that $e \subseteq \text{isol}(\mathcal{I})$, observe that no predicate in $\text{pred}(\mathcal{I}) = \text{pred}(T, N \setminus \{n\})$ shares a variable with $\text{var}(n) = (\text{var}(n) \setminus \text{var}(m))$ by Lemma 12(6). Therefore, $\text{var}(n) \cap \text{var}(\text{pred}(\mathcal{I})) = \emptyset$. Furthermore, $\text{var}(n) \cap \text{fv}(\mathcal{I}) = \emptyset$ because $\text{fv}(\mathcal{I}) = \text{fv}(\mathcal{H}_1)$ and no $x \in \text{var}(n) = \text{var}(n) \setminus \text{var}(m)$ is in $\text{fv}(\mathcal{H}_1)$ by Lemma 13(2). \square

Lemma 5 Let (T, N_1) and (T, N_2) be canonical GJTpairs with $N_2 \subseteq N_1$. Then, $\mathcal{H}(T, N_1, \bar{z}) \rightsquigarrow^* \mathcal{H}(T, N_2, \bar{z})$ for every $\bar{z} \subseteq \text{var}(N_2)$.

Proof By induction on k , the number of nodes in $N_1 \setminus N_2$. In the base case where $k = 0$, the result trivially holds since then $N_1 = N_2$ and the two triplets are identical. For the induction step, assume that $k > 0$ and the result holds for $k - 1$. Because both N_1 and N_2 are connex subsets of the same tree T , there exists a node $n \in N_1$ that is a frontier node in N_1 and which is not in N_2 . Then, define $N'_1 = N_1 \setminus \{n\}$. Clearly (T, N'_1) is again canonical, and $|N'_1 \setminus N_2| = k - 1$. Therefore, $\mathcal{H}(T, N'_1, \bar{z}) \rightsquigarrow^* \mathcal{H}(T, N_2, \bar{z})$ by induction hypothesis. Furthermore, by $\mathcal{H}(T, N_1, \bar{z}) \rightsquigarrow^* \mathcal{H}(T, N'_1, \bar{z})$ by Lemma 14, from which the result follows. \square

B.4 Proof of Lemma 6

Lemma 6 Let H_1 and H_2 be two hypergraphs such that for all $e \in H_2$ there exists $\ell \in H_1$ such that $e \subseteq \ell$. Then, $(H_1 \cup H_2, \bar{z}, \Theta) \rightsquigarrow^* (H_1, \bar{z}, \Theta)$, for every hyperedge \bar{z} and set of predicates Θ .

Proof The proof is by induction on k , the number of hyperedges in $H_2 \setminus H_1$. In the base case where $k = 0$, the result trivially holds since $H_1 \cup H_2 = H_1$ and the two triplets are hence identical. For the induction step, assume that $k > 0$ and the result holds for $k - 1$. Fix some $e \in H_2 \setminus H_1$ and define $H'_2 = H_2 \setminus \{e\}$. Then, $|H'_2 \setminus H_1| = k - 1$. We show that $(H_1 \cup H_2, \bar{z}, \Theta) \rightsquigarrow^* (H_1 \cup H'_2, \bar{z}, \Theta)$, from which the result follows since $(H_1 \cup H'_2, \bar{z}, \Theta) \rightsquigarrow^* (H_1, \bar{z}, \Theta)$ by induction hypothesis. To this end, we observe that there exists $\ell \in H_1 \setminus \{e\}$ with $e \subseteq \ell$. Therefore, $\text{fv}_{(H_1 \cup H_2, \bar{z}, \Theta)}(e) \subseteq e \subseteq \ell$. Moreover, $e \setminus \ell = \emptyset$. Therefore, $\text{ext}_{(H_1 \cup H_2, \bar{z}, \Theta)}(e \setminus \ell) = \emptyset \subseteq \ell$. Thus $e \sqsubseteq_{(H_1 \cup H_2, \bar{z}, \Theta)} \ell$. We may, therefore, apply (CSE) to remove e from $H_1 \cup H_2$, yielding $H_1 \cup H'_2$. Since no predicate shares variables with $e \setminus \ell = \emptyset$ this does not modify Θ . Therefore, $(H_1 \cup H_2, \bar{z}, \Theta) \rightsquigarrow^* (H_1 \cup H'_2, \bar{z}, \Theta)$. \square

C Proofs of Section 5.2

Lemma 7 Let n be a violator of type 1 in (T, N) and assume $(T, N) \xrightarrow{1,n} (T', N')$. Then, (T', N') is a GJTpair and it is equivalent to (T, N) . Moreover, the number of violators in (T', N') is strictly smaller than the number of violators in (T, N) .

Proof. The lemma follows from the following observations. (1) It is straightforward to observe that T' is a valid GJT: the construction has left the set of leaf nodes untouched; took care to ensure that all nodes (including the newly added node p)

continue to have a guard child; ensures that the connectedness condition continues to hold also for the relocated children of n because every variable in n is present on the entire path between n and p ; and have ensured that also edge labels remain valid (for the relocated nodes this is because $\text{var}(p) = \text{var}(g) \subseteq \text{var}(n)$).

(2) N' is a connex subset of T' because the subtree of T induced by N equals to subtree of T' induced by N' , modulo the replacement of l by p in case that l was in N and p is hence in N' .

(3) (T, N) is equivalent to (T', N') because the construction leaves leaf atoms untouched, preserves edge labels, and $\text{var}(N) = \text{var}(N')$. The latter is clear if $l \notin N$ because then $N = N'$. It follows from the fact that $\text{var}(l) = \text{var}(p)$ if $l \in N$, in which case $N' = N \setminus \{l\} \cup \{p\}$.

(4) All nodes in $\text{ch}_T(n) \setminus N$ (and their descendants) are relocated to p in T' . Therefore, n is no longer a violator in (T', N') . Because we do not introduce new violators, the number of violators of (T', N') is strictly smaller than the number of violators of (T, N) . \square

Lemma 8 Let n be a violator of type 2 in (T, N) and assume $(T, N) \xrightarrow{2,n} (T', N')$. Then, (T', N') is a GJTpair and it is equivalent to (T, N) . Moreover, the number of violators in (T', N') is strictly smaller than the number of violators in (T, N) .

Proof. The lemma follows from the following observations. (1) It is straightforward to observe that T' is a valid GJT: the construction has left the set of leaf nodes untouched; took care to ensure that all nodes (including the newly added node p) continue to have a guard child; ensures that the connectedness condition continues to hold also for the relocated children of n because every variable in n is also present in p , their new parent; and have ensured that also edge labels remain valid (for the relocated nodes this is because $\text{var}(p) = \text{var}(n)$).

(2) N' is a connex subset of T' because (i) the subtree of T induced by N equals to subtree of T' induced by $N' \setminus \{p\}$, (ii) $n \in N$, and (iii) p is a child of n in T' . Therefore, N' must be connex.

(3) (T, N) is equivalent to (T', N') because the construction leaves leaf atoms untouched, preserves edge labels, and $\text{var}(N) = \text{var}(N')$. The latter follows because $\text{var}(N') = \text{var}(N \cup \{p\})$ and because $\text{var}(p) = \text{var}(n) \subseteq \text{var}(N)$ since $n \in N$.

(4) All nodes in $\text{ch}_T(n) \setminus N$ (and their descendants) are relocated to p in T' . Therefore, n is no longer a violator in (T', N') . Because we do not introduce new violators, the number of violators of (T', N') is strictly smaller than the number of violators of (T, N) . \square

D Description of competing systems

DBToaster DBToaster (henceforth denoted DBT) is a state-of-the-art implementation of HIVM. It operates in pull-based mode and can deal with randomly ordered update streams. DBT is particularly meticulous in that it materializes only useful views, and therefore, it is an interesting implementation for comparison. It has been extensively tested on equijoin queries and has proven to be more efficient than a commercial database management system, a commercial stream processing system and an IVM implementation [30]. DBT compiles given SQL statements into executable trigger programs in different programming languages. We compare against those generated in Scala from the DBToaster Release 2.2,¹⁰ and it uses actors¹¹ to generate events from the input files. During our experiments, however, we have found that this creates unnecessary memory overhead. For a fair memory-wise comparison, we have, therefore, removed these actors.

Esper Esper (E) is a CER engine with a relational model based on Stanford STREAM [4]. It is push based and can deal with randomly ordered update streams. We use the Java-based open source¹² for our comparisons. Esper processes queries expressed in the Esper event processing language (EPL).

SASE SASE (SE) is an automaton-based CER system. It operates in push-based mode and can deal with temporally ordered update streams only. We use the publicly available Java-based implementation of SASE.¹³ This implementation does not support projections. Furthermore, since SASE requires queries to specify a match semantics (any match, next match, partition contiguity) but does not allow combinations of such semantics, we can only express queries Q_1 , Q_2 , and Q_4 in SASE. Hence, we compare against SASE for these queries only. To be coherent with our semantics, the corresponding SASE expressions use the any match semantics [3].

Tesla/T-Rex Tesla/T-Rex (T) is also an automaton-based CER system. It operates in push-based mode only, and supports temporally ordered update streams only. We use the publicly available C-based implementation.¹⁴ This implementation operates in a publish-subscribe model where events are published by clients to the server, known as TRexServer. Clients can subscribe to receive recognized composite events. Tesla cannot deal with queries involving inequalities on multiple attributes, e.g., Q_3 , therefore, we do not show results for Q_3 .

Since Tesla works in a decentralized manner, we measure the update processing time by logging the time at the Tesla TRexServer from the stream start until the end.

ZStream ZStream (Z) is a CER system based on a relational internal architecture. It operates in push-based mode and can deal with temporally ordered update streams only. ZStream is not available publicly. Hence, we have created our own implementation following the lazy evaluation algorithm described in the original paper [31]. This paper does not describe how to treat projections, and as such we compare against ZStream only for full join queries Q_1 – Q_8 .

References

1. Abiteboul, S., Hull, R., Vianu, V.: Foundations of Databases. Addison-Wesley Longman Publishing Co., Inc., Boston (1995)
2. Abo Khamis, M., Ngo, H.Q., Rudra, A.: FAQ: questions asked frequently. In: Proceedings of PODS, pp. 13–28 (2016)
3. Agrawal, J., Diao, Y., Gyllstrom, D., Immerman, N.: Efficient pattern matching over event streams. Proc. SIGMOD **2008**, 147–160 (2008)
4. Arasu, A., Babcock, B., Babu, S., Cieslewicz, J., Datar, M., Ito, K., Motwani, R., Srivastava, U., Widom, J.: STREAM: the stanford data stream management system. In: Data Stream Management—Processing High-Speed Data Streams, pp. 317–336 (2016)
5. Baader, F., Nipkow, T.: Term Rewriting and All That. Cambridge University Press, Cambridge (1998)
6. Bagan, G., Durand, A., Grandjean, E.: On acyclic conjunctive queries and constant delay enumeration. In: Proceedings of CSL, pp. 208–222 (2007)
7. Bakibayev, N., Kočíský, T., Olteanu, D., Závodný, J.: Aggregation and ordering in factorised databases. Proc. VLDB **6**(14), 1990–2001 (2013)
8. Berkholz, C., Keppeler, J., Schweikardt, N.: Answering conjunctive queries under updates. In: Proceedings of PODS, pp. 303–318 (2017)
9. Bernstein, P.A., Goodman, N.: The power of inequality semijoins. Inf. Syst. **6**(4), 255–265 (1981)
10. Brault-Baron, J.: De la pertinence de l'énumération: complexité en logiques. Ph.D. thesis, Université de Caen (2013)
11. Brenna, L., Demers, A.J., Gehrke, J., Hong, M., Ossher, J., Panda, B., Riedewald, M., Thatte, M., White, W.M.: Cayuga: a high-performance event processing engine. Proc. SIGMOD **2007**, 1100–1102 (2007)
12. Chirkova, R., Yang, J.: Materialized views. Found. Trends Databases **4**(4), 295–405 (2012)
13. Cormen, T.: Introduction to Algorithms, 3rd edn. MIT Press, Cambridge (2009)
14. Cugola, G., Margara, A.: TESLA: a formally defined event specification language. Proc. DEBS **2010**, 50–61 (2010)
15. Cugola, G., Margara, A.: Complex event processing with T-REX. J. Syst. Softw. **85**(8), 1709–1728 (2012)
16. Cugola, G., Margara, A.: Processing flows of information: from data stream to complex event processing. ACM Comput. Surv. **44**(3), 15:1–15:62 (2012)
17. DeWitt, D.J., Naughton, J.F., Schneider, D.A.: An evaluation of non-equijoin algorithms. VLDB **1991**, 443–452 (1991)
18. Enderle, J., Hampel, M., Seidl, T.: Joining interval data in relational databases. Proc. SIGMOD **2004**, 683–694 (2004)
19. EsperTech. Esper complex event processing engine. <http://www.espertech.com/>

¹⁰ <https://dbtoaster.github.io/>.

¹¹ <https://doc.akka.io/docs/akka/2.5/>.

¹² <http://www.espertech.com/esper/esper-downloads/>.

¹³ <https://github.com/haopeng/sase>.

¹⁴ <https://github.com/deib-polimi/TRex>.

20. Golab, L., Özsu, M.T.: Processing sliding window multi-joins in continuous queries over data streams. In: *Proceedings of VLDB*, pp. 500–511 (2003)
21. Gupata, A., Mumick, I.S. (eds.): *Materialized Views: Techniques, Implementations, and Applications*. MIT Press, Cambridge (1999)
22. Gupta, A., Mumick, I.S., Subrahmanian, V.S.: Maintaining views incrementally. In: *Proceedings of SIGMOD*, pp. 157–166 (1993)
23. Hellerstein, J.M., Naughton, J.F., Pfeffer, A.: Generalized search trees for database systems. In: *VLDB'95*, pp. 562–573 (1995)
24. Henzinger, M., Krinninger, S., Nanongkai, D., Saranurak, T.: Unifying and strengthening hardness for dynamic problems via the online matrix-vector multiplication conjecture. In: *Proceedings of STOC*, pp. 21–30 (2015)
25. Idris, M., Ugarte, M., Vansummeren, S.: The dynamic Yannakakis algorithm: compact and efficient query processing under updates. In: *Proceedings of SIGMOD 2017* (2017)
26. Idris, M., Ugarte, M., Vansummeren, S., Voigt, H., Lehner, W.: Conjunctive queries with inequalities under updates. *PVLDB* **11**(7), 733–745 (2018)
27. Kang, J., Naughton, J.F., Viglas, S.: Evaluating window joins over unbounded streams. In: *Proceedings of ICDE*, pp. 341–352 (2003)
28. Khayyat, Z., Lucia, W., Singh, M., Ouzzani, M., Papotti, P., Quiané-Ruiz, J., Tang, N., Kalnis, P.: Fast and scalable inequality joins. *VLDB J.* **26**(1), 125–150 (2017)
29. Koch, C.: Incremental query evaluation in a ring of databases. In: *Proceedings of PODS*, pp. 87–98 (2010)
30. Koch, C., Ahmad, Y., Kennedy, O., Nikolic, M., Nötzli, A., Lupei, D., Shaikhha, A.: Dbtoaster: higher-order delta processing for dynamic, frequently fresh views. *VLDB J.* **23**, 253–278 (2014)
31. Mei, Y., Madden, S.: Zstream: a cost-based query processor for adaptively detecting composite events. *Proc. SIGMOD* **2009**, 193–206 (2009)
32. Nikolic, M., Olteanu, D.: Incremental view maintenance with triple lock factorization benefits. *Proc. SIGMOD* **2018**, 365–380 (2018)
33. Olteanu, D., Závodný, J.: Size bounds for factorised representations of query results. *ACM TODS* **40**(1), 2:1–2:44 (2015)
34. Roy, P., Teubner, J., Gemulla, R.: Low-latency handshake join. *PVLDB* **7**(9), 709–720 (2014)
35. Sahay, B., Ranjan, J.: Real time business intelligence in supply chain analytics. *Inf. Manage. Comput. Secur.* **16**(1), 28–48 (2008)
36. Schleich, M., Olteanu, D., Ciucanu, R.: Learning linear regression models over factorized joins. In: *Proceedings of SIGMOD*, pp. 3–18 (2016)
37. Schultz-Møller, N.P., Migliavacca, M., Pietzuch, P.R.: Distributed complex event processing with query rewriting. In: *Proceedings of DEBS 2009* (2009)
38. Segoufin, L.: Constant delay enumeration for conjunctive queries. *SIGMOD Rec.* **44**(1), 10–17 (2015)
39. Stonebraker, M., Çetintemel, U., Zdonik, S.: The 8 requirements of real-time stream processing. *SIGMOD Rec.* **4**, 42–47 (2005)
40. Teubner, J., Müller, R.: How soccer players would do stream joins. In: *Proceedings of SIGMOD*, pp. 625–636 (2011)
41. Urhan, T., Franklin, M.J.: Xjoin: a reactively-scheduled pipelined join operator. *IEEE Data Eng. Bull.* **23**(2), 27–33 (2000)
42. Vardi, M.Y.: The complexity of relational query languages (extended abstract). In: *Proceedings of STOC*, pp. 137–146 (1982)
43. Viglas, S., Naughton, J. F., Burger, J.: Maximizing the output rate of multi-way join queries over streaming information sources. In: *Proceedings of VLDB*, pp. 285–296 (2003)
44. Wang, W., Gao, J., Zhang, M., Wang, S., Chen, G., Ng, T.K., Ooi, B.C., Shao, J., Reyad, M.: Rafiki: machine learning as an analytics service system. *PVLDB* **12**(2), 128–140 (2018)
45. Wilschut, A.N., Apers, P.M.G.: Dataflow query execution in a parallel main-memory environment. In: *Proceedings of the First International Conference on Parallel and Distributed Information Systems (PDIS 1991)*, pp. 68–77. IEEE Computer Society (1991)
46. Wu, E., Diao, Y., Rizvi, S.: High-performance complex event processing over streams. *Proc. SIGMOD* **2006**, 407–418 (2006)
47. Yannakakis, M.: Algorithms for acyclic database schemes. In: *Proceedings of VLDB*, pp. 82–94 (1981)
48. Yoshikawa, M., Kambayashi, Y.: Processing inequality queries based on generalized semi-joins. In: *VLDB*, pp. 416–428 (1984)
49. Zhang, H., Diao, Y., Immerman, N.: On complexity and optimization of expensive queries in complex event processing. In: *Proceedings of SIGMOD* (2014)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.