

Article

Not peer-reviewed version

Matching and Rewriting Rules in Object-Oriented Databases

[Giacomo Bergami](#)^{*}, [Oliver Robert Fox](#), Graham Morgan

Posted Date: 8 August 2024

doi: 10.20944/preprints202408.0536.v1

Keywords: direct acyclic graphs; generalised semistructured model; graph grammars; graph query languages; operator algebras



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Matching and Rewriting Rules in Object-Oriented Databases

Giacomo Bergami , Oliver Robert Fox  and Graham Morgan 

School of Computing, Faculty of Science, Agriculture and Engineering, Newcastle University, Newcastle Upon Tyne NE4 5TG, UK

* Correspondence: giacomo.bergami@newcastle.ac.uk

Abstract: Graph query languages such as Cypher are widely adopted to match and retrieve data in a graph representation, due to their ability of retrieving and transforming information. Despite the most natural way to match and transform information is through rewriting rules, those are scarcely or partially adopted in graph query languages. Their inability of doing so has a major impact on the subsequent way the information is structured, as it might then appear more natural to provide major constraint over the data representation so to consequently constraint the way the information should be represented. On the other hand, recent works are starting to move towards an opposite direction, as the provision of a truly general semistructured model (GSM) allows to both represent all the available data formats (Network-Based, Relational, and Semistructured) as well as support an holistic query language expressing all major queries in such languages. In this paper, we show that the usage of GSM enables the definition of a general rewriting mechanism which can be expressed in current graph query languages only at the cost of adhering the query to the specificity of the underlying data representation. We formalise the proposed query language in terms graph rewriting mechanisms described as a set of production rules $L \rightarrow R$ while providing restriction to the characterization of L , while extending it to support structural graph nesting operations, useful to aggregate similar information around an entry-point of interest. We discuss how GSM, by fully supporting index-based data representation, allows for a better physical model implementation leveraging the benefits of columnar database storages. Preliminary benchmarks shows the scalability of this proposed implementation in comparison with state-of-the-art implementations.

Keywords: direct acyclic graphs; generalised semistructured model; graph grammars; graph query languages; algorithms; operator algebras

1. Introduction

Query languages [1] fulfil the aim of retrieving and manipulating data after being adequately processed according to a physical model requiring preliminary data loading and indexing operations. Good declarative languages for manipulating information, such as SQL [2], are agnostic from the underlying physical model while expressing the information need of combining (JOIN), filtering (WHERE) and grouping (e.g., COUNT(*) over GROUP BY-s) data over its logical representation. This language can also be truly declarative as it does not require a user to explicitly convey this in terms of operations to be performed over the data rather than instructing the machine which information should be used and which types of transformations should be applied. Due to the intrinsic nature of graph data representation, graph query languages such as Gremlin [3], Cypher [4], or SPARQL [5] are procedural (i.e., navigational) due to the navigational structure of the data, thus heavily requiring the user to inform the language how to match and transform the data. These considerations extend to how edge and vertex data should be treated, thus adding further complexity [6,7].

On the other hand, graph grammars [8] provide a high-level declarative way to match specific sub-graphs of interest from vertex-labelled graphs for then rewriting them into another subgraph of interest, thus updating the original graph of interest. These consist in a set of rules $\{L_i \rightarrow R_i\}_{1 \leq i \leq N}$, where each rule $L_i \rightarrow R_i$ consists of two graphs L_i and R_i possibly sharing a common subgraph K : while $L \supseteq K$ specifies a potential removal of either vertices or edges, $K \subseteq R$ specifies their addition. To automate the application of such rules, this requires explicitly following a graph navigational order for

applying each matched subgraph to the structure [9] to guarantee the generation of a unique acceptable graph representation resulting from the querying data. As a by-product of exploiting topological sorts for scheduling visiting and rewriting operations, we then require to precisely identify an *entry match* vertex for each L_i from a rule, thus precisely determining from which position within the graph the rewriting rule should be applied, and in which order.

At the time of the writing, the aforementioned graph query languages are not able to express graph grammars natively without resorting to explicitly telling the query language how to restructure the underlying graph data representation (Lemma 7); furthermore, the replacement of previously matched vertices with new ones invalidates previous matches, thus forcing the user to pipeline different queries until convergence is reached. On the other hand, we would expect any graph query language to directly express the rewriting mechanism by directly expressing the set of rules, without necessarily instructing the querying language in which order such operations shall be performed. Furthermore, when a property-graph model is chosen, and morphisms are expressed through the properties associated with the vertices and edges rather than their associated IDs as per the current Cypher and Neo4J specifications, this makes the deletion of vertices and their update through the resulting morphisms quite impractical, as such data model provides no clear way to refer to both vertices and edges via unique identifiers. As a result of these considerations, we then derive that, at the time of the writing, both the graph query languages and their underlying representational model are insufficient to adequately express rewriting rules as general as the ones postulated by graph grammars over graphs containing property-value associations, independently from their representation of choice [10].

To overcome the aforementioned limitations, we propose for the first time a query language directly expressing such rewriting desiderata: we restrict the set of all the possible matching graphs L_i into ego-nets containing one entry point while incorporating nesting operations; as updating graph vertices' properties breaks the declarative assumption as the user should specify the order in which the operations are performed, we relax the language declarativity for the sole rewritings. To better support this query language, we completely shift the data model of interest to object-oriented databases, where relationships across objects are expressed through object "containment" relationships, and where both objects and such containments are uniquely identified. By also requiring that any object shall never contain itself at any nesting level [6], we obtain cycle-free containment relationships.

The paper is structured as follows: after providing some preliminary notation used throughout the paper (Section 2), we outline the related works supporting the novelty of the proposed approach outlined in the following sections (Section 3). After describing our proposed logical data model with its associated physical counterpart (Section 4), in Section 6 we characterize the formal semantics of such novel graph query language in pseudocode notation (Algorithm 4) characterized in terms of both algorithmic and algebraic notation (Section 6.3 and 6.4) for the matching part, as well as in terms of Structured Operational Semantics (SOS) [11] for describing the rewriting steps (Section 6.5). For implementing this, we design a novel physical model (Section 4.2) leveraging state-of-the-art columnar relational database representations [12]: preliminary benchmarks against Cypher running over Neo4j show the extreme inefficiency of property graph computational model (Section 9), which are fairly restricted due to the impossibility of conveying a novel single query for any possible graph schema (Lemmas in Section 7). Preliminary scalability tests show that our solution outperforms Cypher and Neo4J, providing the query language standard nearer to the recently proposed GQL, by 2 orders of magnitude (Section 9.1) while providing a computational throughput being 600 times faster than Nel4J by solving more queries in the same comparable amount of time (Section 9.1).

These main contributions are then obtained in terms of the following ancillary results:

- Definition of a novel nested relational algebra natural equi-join operator for composing nested morphisms, also supporting left outer joins by parameter tuning (Section 5).
- Definition of an object-oriented database view for updating the databases on the fly without the need for heavy restructuring the loaded and indexed physical model (Section 4.3).

- As the latter view relies on the definition of a logical model extending GSM (Section 4.1), we show that the physical model is isomorphic to an indexed set of GSM databases expressed within the logical model (Lemma 1).

The paper is structured as follows: after providing some preliminary notation being used throughout the paper (Section 2), we introduce the relational and the graph data models from current literature, while focussing on both their data representation and associated query languages (Section 3). Their juxtaposition motivates the proposal of our Generalised Semistructured Model (Section 4), for which we outline both the logical and physical data model; we also introduce the concept of a view for a generalised semistructured model, as well as introduce some morphism notation being used throughout the paper. After introducing the designated operator for instantiating the morphisms by joining matched containments stored in distinct set of tables (Section 5), we finally propose our query language for matching and rewriting object-oriented databases expressed in the aforementioned GSM model (Section 6). The remaining sections discuss some of the expressiveness properties of such query language if compared to Cypher (Section 7), discuss its time complexity (Section 8), and provide benchmarks remarking the efficiency of the proposed query language if compared to Cypher (Section 9). Last, we draw our final conclusions where we propose some future works (Section 10). To improve the paper's readability, we move some definitions (Appendix 10) and the full set of proofs for the Lemmas and Corollaries (Appendix 10) to the Appendix.

2. Preliminary Notation

We denote **sets** $S = \{x_1, \dots, x_n\}$ of cardinality $|S| = n$ as standard. The power set $\wp(S)$ of any set S is the set of all subsets of S , including the empty set \emptyset and S itself. Formally, $\wp(S) = \{S' | S' \subseteq S\}$. A **tuple** or indexed set $t = \langle t_1, \dots, t_n \rangle$ of length $|t| = n$ is defined as a finite function in $\mathbb{N} \rightarrow \mathcal{V}$ where t_i is a shorthand for $t(i)$.

Given two functions f and f' , we denote $f \oplus f'$ as the *overriding* of f by f' returning $f'(x)$ for each $x \in \text{dom}(f')$ and $f(x)$ otherwise. We define a *finite function* $f: A \rightarrow B$ via its *graph* $[(x_1, f(x_1)), \dots, (x_n, f(x_n))] = \oplus_{x_i \in \text{dom}(f)} [(x_i, f(x_i))]$ explicating the association between a value from the domain of f ($\{x_i, \dots, x_i\} = \text{dom}(f)$) and a non-NULL codomain value. Using an abuse of notation, we denote the *restriction* $f|_X$ as the evaluation of f over a domain $X \cap \text{dom}(f)$, i.e. $f|_X = \oplus_{x \in X \cap \text{dom}(f)} [(x, f(x))]$. We can also denote $f(x) := C$ where C is the definition of the function over x as $x \mapsto C$. With an abuse of notation, we denote $|f|$ as the cardinality of its domain, i.e. $|f| = |\text{dom}(f)|$.

We say that two functions f and f' are *equivalent*, $f \doteq f'$, iff. they both share the same domain and, for each element of their domain, both function return the same codomain value, i.e. $f \doteq f' \Leftrightarrow \text{dom}(f) = \text{dom}(f') \wedge \forall x \in \text{dom}(f). f(x) = f'(x)$. Given a set A , an **equivalence class** $[x]_{\mathfrak{R}} \subseteq A$ for $x \in A$ is the set of all the elements in A that are equivalent to x , i.e. $[x]_{\mathfrak{R}} = \{y \in A | x \mathfrak{R} y\}$. Given a set A and an equivalence relationship \mathfrak{R} , the **quotient set** A/\mathfrak{R} is the set of all the equivalence classes of A , i.e. $A/\mathfrak{R} = \{[x]_{\mathfrak{R}} | x \in A\}$. We denote \doteq_X as the equivalence relationship denoting two functions as equivalent if they are equivalent after the same restriction over X , i.e., $f \doteq_X f' \Leftrightarrow f|_X \doteq f'|_X$.

Given a set of all the possible string attributes Σ^* and a set of all the possible values \mathcal{V} , a **record** is defined as a finite function $f: \Sigma^* \rightarrow \mathcal{V}$ with domain $\text{dom}(f) \subseteq \Sigma^*$. We denote as NULL a void value not representing a sensible value in \mathcal{V} : using the usual notation in the database area at the cost of committing an abuse of notation, we say that $f(x)$ returns NULL iff. f is not defined over x (i.e., $x \notin \text{dom}(f)$). We say that a record ϵ is empty if no attribute is associated to a value in \mathcal{V} (i.e., $\text{dom}(\epsilon) = \emptyset$ and $\forall x \in \mathcal{V}. \epsilon(x) = \text{NULL}$).

2.1. Higher-Order Functions

Higher-Order Functions (HOFs) are functions that either take one or more functions as arguments (i.e. procedural parameters of procedures), or return a function as a result. We provide the definition of some HOFs being used in this paper:

- The **zipping** operator maps n tuples (or records) t^1, \dots, t^n to a record of tuples (or records) r defined as $\langle t_i^1, \dots, t_i^n \rangle$ over i iff. all n tuples are defined over i :

$$\zeta(t^1, \dots, t^n) = \bigoplus_{\substack{i \leq \min(|t^1|, \dots, |t^n|) \\ \text{s.t. } \forall 1 \leq j \leq n. i \in \text{dom}(t^j)}} [\langle t_i^1, \dots, t_i^n \rangle]$$

- Given a function $f: A \rightarrow B$ and a generic collection C , the **mapping** operator returns a new collection by applying f to each component of C :

$$\mu(f, C) = \begin{cases} \{f(x) | x \in C\} & C \text{ is a set} \\ [f(C(i))]_{i \in \text{dom}(C)} & C \text{ is a record} \\ \langle f(C_1), \dots, f(C_n) \rangle & C \text{ is a tuple} \end{cases}$$

- Given a binary predicate p and a collection C , the **filter** function trims C by restricting it to its values satisfying p :

$$F(p, C) = \begin{cases} \{x \in C | p(x)\} & C \text{ is a set} \\ \bigoplus_{\substack{x \in \text{dom}(C) \\ \text{s.t. } p(C(x))}} [(x, C(x))] & C \text{ is a record} \\ x \mapsto \arg \min_{\substack{|C| \geq i \geq x \\ \text{s.t. } p(C_i)}} C_i & C \text{ is a tuple} \end{cases}$$

- Given a binary function $f: A \times \mathcal{V} \rightarrow A$, an initial value $\alpha \in A$ (accumulator), and a tuple C , the **(left) fold** operator is a tail-recursive function returning either α for an empty tuple, or $f(\dots f(\alpha, t_1), t_n)$ for a tuple $t = \langle t_1, \dots, t_n \rangle$:

$$\Lambda(f, \alpha, C) = \begin{cases} \alpha & |C| = 0 \\ \Lambda(f, f(\alpha, C(m)), C|_{\text{dom}(C) \setminus \{m\}}) & |C| \neq 0 \wedge m := \min \text{dom}(C) \end{cases}$$

- Given a collection of strings C and a separator s , **collapse** (also referred to as *join* in programming languages such as Javascript or Python) returns a single string where all the strings in C are separated by s . Given “^” the usual string concatenation operator, this can be expressed in terms of Λ as follows:

$$\lambda(C, s) := \Lambda(^, s, C)$$

When s is the empty string “”, then we can use $\lambda(C)$ as a shorthand.

- Given a function $f: A \rightarrow B$ and two values $a \in A$ and $b \in B$, the update of f so that it will return b for a and will return any other previous value for $f(x)$ otherwise is defined as follows:

$$\text{PUT}_f(a, b) := f \oplus [(a, b)]$$

- Given a function f and an input value, the HOF **optionally getting** the value of $f(y)$ if $y \in \text{dom}(f)$ and returning z otherwise is defined as:

$$\text{OPTGET}_f(y, z) := \begin{cases} f(y) & y \in \text{dom}(f) \\ z & \text{oth.} \end{cases}$$

Please observe that we can use this function in combination with PUT to set multiple nested functions:

$$\text{PUT}_f^2(\langle \langle i, j \rangle, v \rangle) := \text{PUT}_f(i, \text{PUT}_{\text{OPTGET}_f(i, \emptyset)}(j, v)) \quad (1)$$

3. Related Works

This section outlines different logical models and query languages for both graph data (Section 3.1.1 and Section 3.1.2) and the nested relational model (Section 3.2.1 and 3.2.2), while starting from their mathematical foundations. We then discuss an efficient physical model being used for the relational model (Section 3.2.3).

3.1. Graph Data

3.1.1. Logical Model

Direct Acyclic Graphs (DAGs) and Topological Sort

A **digraph** $\gamma = (V, E)$ consists of vertices V and edges E being a subset of V^2 . We say that such graph is *weighted* if its definition is extended as (V, E, ω) where $\omega: E \rightarrow \mathbb{R}$ is a weight function associating each edge in the graph to a real weight. A *closed trail* is a sequence of distinct edges $(s_1, d_1), \dots, (s_n, d_n) \in E$ which edge (s_i, d_i) is such that $s_i = d_{i-1}$ and $d_i = s_{i+1}$ if any, and where $s_1 = d_n$. We say that a digraph is *acyclic* if it contains no closed trails consisting of at least one edge, and we refer to it as a *Direct Acyclic Graph* (DAG).

If we assume that edges in a graph reflect dependency relationships across their vertices, a natural way to determine the visiting order of the vertices is first to sort its vertices topologically [13], thus inducing an operational scheduling order [12]. This requires the underlying graph to have no (directed) cycles, also known as Direct Acyclic Graphs (DAGs). Notwithstanding this, any DAG might come with multiple possible valid topological sorts; among these, the scheduling of tasks operations usually prefers visiting the operations bottom-up in a layered way, thus ensuring the operations are always applied starting from the sub-graphs having less structural-refactoring requirements than the higher ones [12].

We say that a *topological sort* of a DAG (V, E) is a linear ordering of its vertices in a tuple t with $|t| = |V|$ so that for any edge $(u, v) \in E$ we have i and $j > 0$ such that $u = t_i$ and $v = t_{i+j}$. Given this linear ordering of the vertices, we can always define a layering algorithm [14] where vertices sharing no mutual interdependencies are placed in the same layer, and where all the vertices appearing in the shallowest layer will be connected by transitive closure to all the vertices in the deeper layers, while the vertices in the deepest layer will share no interconnection with the other vertices in the graph. To do so, we can use Algorithm 1: we use timestamps for determining the layer id: we associate all vertices with no outgoing edge to the deepest layer 0 (Line 5) and, for all the remaining vertices, we set the vertex to the layer immediately below to the none of the deepest outgoing vertex (Line 9). When the graph is cyclic, we can use heuristics to approximate the vertex ordering, such as using the progressive id associated to the vertices to disambiguate and choose an ordering when required.

Algorithm 1 Layering DAGs from a vertex topological order in t

```

1: function LAYERFROMTOPOLOGICALSORT( $G = (V, E), t$ )
2:    $\text{time} := \langle -1, \dots, -1 \rangle$   $\triangleright |\text{time}| = |t|$ 
3:    $\text{firstVisit} := \text{true}$ 
4:   for all  $p \in \text{Reverse}(t)$  do
5:      $\text{time}[p] := 0$ 
6:     if  $\text{firstVisit}$  then
7:        $\text{firstVisit} := \text{false}$ 
8:     else if  $\text{IN}(p) \neq \emptyset$  then
9:        $\text{time}[p] := \max\{\text{time}[u] | u \in \text{OUT}(p)\} + 1$ 
10:    end if
11:  end for
12:  return  $\text{time}$ 
13: end function

```

Property Graphs

Property graphs [15] represent **multigraphs** (i.e., digraphs allowing for multiple edges among two distinct vertices) expressing both vertices and edges as multi-labelled records. The usefulness of this data model is remarked by its implementation in almost all recent Graph DBMSs, such as **Neo4J** [4].

Definition 1 (Property Graph). A *property graph* [9] is a tuple $(V, E, L, A, U, \ell, \kappa, \lambda)$, where V and E are sets of distinct integer identifiers ($V \subseteq \mathbb{N}$, $E \subseteq \mathbb{N}$, $V \cap E = \emptyset$). L is a set of labels, A is a set of attributes and U a set of values. $\ell: V \cup E \rightarrow \wp(L)$ maps each vertex or edge to a set of labels; $\kappa: V \cup E \rightarrow A \rightarrow U$ maps each vertex or edge within the graph and each attribute within A , to a value in U ; last, $\lambda: E \rightarrow V \times V$ maps each edge $e \in E$ to a pair of vertices $\lambda(e) = (s, t) \in V \times V$, where s is the source vertex and t is the target.

Property graphs do not support aggregated values, as values in U cannot contain either vertices or edges, nor U is made to contain collection of values.

RDF

The RESOURCE DESCRIPTION FRAMEWORK (RDF) [16] distinguishes resources (comparable to vertices) via UNIQUE RESOURCE IDENTIFIERS (URI), be them vertices or edges within a graph; those are linked to their properties or other resources via triples, acting as edges. RDF are commonly used in the semantic web and in the ontology field [17,18]. Thus, modern reasoners such as Jena [19] or Pellet [20] assume such data structure as the default graph data model.

Definition 2 (RDF (Graph Data) Model). An *RDF (Graph data) model* [9] is defined as a set of triples (s, p, o) , where s is called “subject”, p is the “predicate” and o is the “object”. Such triple describes an edge with label p linking the source vertex s to the destination vertex o . Such predicate can also be a source vertex [10]. Each vertex is either identified by a unique URI identifier or by a blank vertex b_i . Each predicate is only described by an URI identifier. \square

Despite this model uses unique resource identifiers for either vertices or edges differently from property graphs, RDF are forced to express attribute-value associations for vertices as additional edges through reification. Thus, property graphs can be entirely expressed as RDF triplestore systems as follows:

Definition 3 (Property Graph over Triplestore). Given a property graph $G = (V, E, A, U, \ell, \kappa, \lambda')$, each vertex $v_i \in V$ induces a set of triples (v_i, α, β) for each $\alpha \in A$ such that $\kappa(v_i, \alpha) = \beta$ having $\beta \neq \text{NULL}$. Each edge $e_j \in E$ induces a set of triples (s, e_j, d) such that $\lambda'(e_j) = (s, d)$ and another set of triples (e_j, α', β') for each $\alpha' \in A$ such that $\kappa(e_j, \alpha') = \beta'$ having $\beta' \neq \text{NULL}$.

The inverse transformation is not always possible as RDF properties as property graphs do not allow the representation of edges departing from other edges. RDF also support the unique identification of distinct databases being loaded within the same physical storage through *named graphs* via resource identifier. Even though it allows named graphs to appear as triplet subjects, such named graphs can appear as neither objects nor properties, thus not overcoming property graphs' limitations on representing nested data.

Notwithstanding the model's innate ability of representing both vertices and edges via URIs, the inability of the data model of directly associating each URI a data record while requiring to express the property-value associations via triplets, requires the expression of property updates via the deletion and subsequent creation of new triplets of the data, which might be quite impractical. Given all the above, we still focus our attention over the Property Graph model, as it better matches our data representation desiderata by associating to vertices labels, property-value associations, as well as binary relationships without any data duplications.

3.1.2. Query Languages

Despite the recent adoption of a novel graph query language standard, GQL, (<https://www.iso.org/obp/ui/en/#iso:std:76120:en> Accessed on 19 April, 2024), its definition has still to be implemented yet in existing systems. This motivates to briefly survey currently-available languages. Differently from more common characterization of graph query languages in term of their potential of

expressing traversal queries [21], a better analysis of such languages involves their ability to generate new data. Our previous work [9] proposed the following characterization:

Graph Traversal and Pattern Matching: these are mainly navigational languages performing the graph visit through “tractable” algorithms through polynomial time visits with respect to the graph size [17,22,23]. Consequently, such solutions do not necessarily involve to run a subgraph isomorphism problem, except when expressly requested by specific semantics [21,24].

(Simple) Graph Grammars: as discussed in the forthcoming paragraph, they can add and remove new vertices and edges which do not necessarily depend on previously matched data, but they are unable to express full data transformation operations.

Graph Algebras: these are mainly designed either to change the structure of property graphs through unary operators, or to combine them through n-ary (often binary) ones. These are not to be confused to the path-algebras for expressing graph traversal and pattern matching constructs, as they allow to completely transform graphs alongside the data associated to them as well as dealing with graph data collections [25–28].

“Proper” Graph Query Languages: We say that a graph query language is “proper” when its expressive power includes all the aforementioned query languages, and possibly expressing the graph algebraic operators while being able of expressing, to some extent, graph grammar rewriting rules, independently from their ability of expressing them in a fully-declarative way. This is achieved to some extent to commonly-available languages, such as SPARQL and Cypher [4].

Graph Grammars

Graph grammars [29] are the theoretical foundations of current graph query languages, as they express the capability of *matching* specific patterns L [30] within the data through reachability queries while applying modifications to the underlying graph database structure (*graph rewriting*) R , thus producing a single graph grammar production rule $L \xrightarrow{f} R$, where there is an implicit morphism between some of the vertices (and edges) matched in L and the ones appearing in R : the vertices (and edges) only appearing in R are considered as newly inserted vertices, while the vertices (and edges) only appearing in L are considered as removed edges; we preserve the remaining matched vertices. Each rule is then considered as a function f , taking a graph database as an input and returning a transformed graph.

The process of matching L is usually expressed in terms of subgraph isomorphism: given two graphs G and L , we determine whether G contains a subgraph G_i that is isomorphic to L , i.e. there is a bijective correspondence $L \xleftrightarrow{\mu_i^L} G_0$ between the vertices and edges of L and G_i . In graph query languages, we consider G as our graph database and return $f(G_i)$ for each matched subgraph G_i . When no rewriting is considered, each possible match G_0 for L is usually represented in a tabular form [30,31], where the column header provides the vertex and edge identifiers (e.g., variables) j from L , each row reflects each matched graph G_i , and each cell corresponding to the column j represents $\mu_i(j)$. Figure 1 shows the process of querying a graph g (Figure 1a) through a pattern L (Figure 1b), for which all the subgraphs matching in there former could be reported as morphisms listed as rows within a table, which column headers reflect the vertex and edge variables occurring in the graph pattern (Figure 1c).




Figure 1. Listing all the subgraphs of g being a solution of the subgraph isomorphism problem of g over L .

Figure 2 illustrates graph grammar rules as defined in GraphLog [32] for both matching and transforming any graph: we can first create the new vertices required in R , while updating or removing x as determined by the vertex or edge $f(\mu_i^{-1}(x))$ occurring in R . Deletions can be performed as the last operations from R . GraphLog still allows running of one single grammar rule at a time, while authors assume to have a graph where vertices are associated to data values and edges are labelled. Then, the rewriting operations derived from R will be applied to every subgraph being matched via a previously identified morphism in $MT[L, g]$ for each graph g of interest. Still, GraphLog considered neither the possibility of simultaneously applying multiple rewriting rules to the same graph, nor a formal definition of which order the rules should be applied. The latter is required to ensure that any update on the graph can be performed incrementally while ensuring that any update to a vertex u via the information stored in its neighbours will always rely on the assumption that each neighbour will not be updated in any other subsequent step, thus guaranteeing that the information in u will never become stale.

This paper solves the ordering issue by applying the changes on the vertices according to their inverse topological order, thus updating the vertices sharing the least dependencies with their direct descendants first.

Furthermore, as GraphLog authors are considering no property graph model, authors do not consider the possibility of updating multiple property value associations over a vertex, while not providing a formal definition of how aggregation functions over vertex data should be defined. While the first problem can be solved only by properly exploiting a suitable data model, the latter is solved by the combined provision of nested morphism, explicitly nesting the vertices associated to a grouped variable, while exploiting a scripting language for expressing how level data manipulations when strictly required. Given these considerations, we will also discuss graph data models (Section 3.1) and their respective query languages (Section 3.1.2).

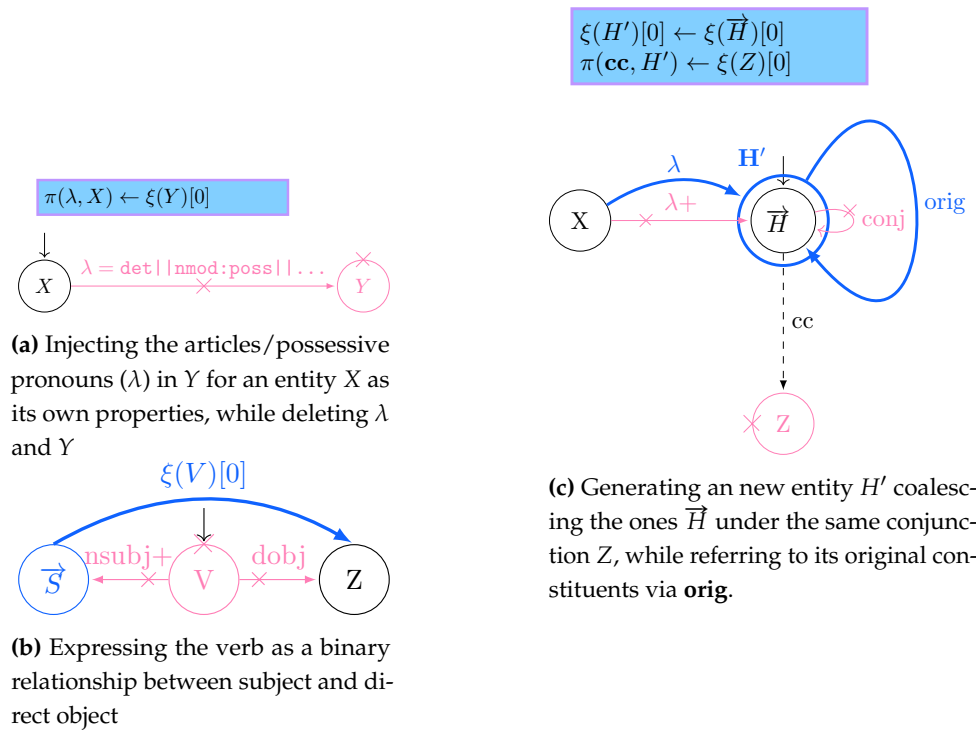


Figure 2. Graph grammar production rules *à la* GraphLog [32] in this paper's use case scenario: thick denotes insertions, crosses deletions, and optional matches are dashed. We extended it with multiple optional edge label matches (\parallel), key-value association $\pi(\lambda, X)$ for property λ and vertex X , and multiple vertex values $\xi(X)$.

Proper Graph Query Languages

Given the above, we will mainly focus our attention to the last type of languages. Even though these languages can be closed under either property graphs or RDF, graphs must not be considered as their main output result, since specific keywords like RETURN for Cypher and CONSTRUCT for SPARQL must be used to force the query result to return graphs. Given also the fact that such languages have not been formalized from the graph returning point of view, such languages prove to be quite slow in producing new graph outputs [6,7].

GQL is largely inspired by **Cypher**, for which this standard might be considered its natural extension. Such query language enables the partial definition of graph grammar operations by supporting the following operators restructuring vertices and edges within the graph: SET, for setting new values within vertices and edges, MERGE, for merging set of attributes within a single vertex or edges, REMOVE for removing labels and properties from vertices and edges, and CREATE for the creating of new vertices, edges and paths.

Notwithstanding the former, such language cannot express all the possible data transformation operations, thus requiring an extension via its APOC library (<https://neo4j.com/developer/neo4j-apoc/>, Accessed on November 13 2023) for defining User-Defined Functions (UDF) in a controlled way, like concatenating the properties associated to vertices' multiple matches:

```
MATCH (a)-[b:cc]->(c)
WITH Collect(a.name) as names, c
CREATE (x {name: apoc.text.join(names, ' ')})
RETURN x
```

A further limitation of Cypher is the inability to create a relationship with a variable as the name, as standard Cypher syntax only allows a string parameter. Using the APOC library we can use

apoc.create.relationship to pass a variable name from an existing vertex for example. Given we our last query creating the vertex *x* we can continue the aforementioned query query as so:

```
MATCH (a)-[:dobj]->(b)
CREATE (y {name: b.name})
WITH a, x, y
CALL apoc.create.relationship(x, a.name, {}, y) YIELD rel
RETURN x, y, rel
```

As there is no associated formal semantics for the entirety of this language's operators except from its fragment related to graph traversal and matching [33], for our proofs regarding this language (e.g. Lemma 7) we are forced to reduce our arguments to common-sense reasoning an experience-driven observation from the usage of Cypher over Neo4J similarly to the considerations in the former paragraph. In fact, such algebra does not involve the creation of new graphs: this is also reflected by its query evaluation plan, which preferred evaluation is a morphism table rather than expressing the rewriting in terms of updated and resulting property graph. As a result, the process of creating or deleting vertices and edges is not optimized.

Overall, Cypher suffers from the limitations posed by the property graph data model which, by having no direct way to refer to the matched vertices or edges by reference, forces the querying user to always refer to the properties associated to them; as a consequence, the resulting morphism tables are carrying out redundant information that cannot reap the efficient data model posed by columnar databases, where entire records can be referenced by their ID. This is evident for **DELETE** statements, voiding objects represented within the morphisms. This limitation of the property graph model, jointly with the need for representing acyclic graphs, motivates us to use the Generalised Semistructured Model (GSM) as an underlying data model for representing graphs, thus allowing us to refer to the vertices and edges by their ID [34]. Consequently, our implementation represents morphisms for acyclic property graphs as per Figure 1c.

Figure 3a provides a possible property graph instantiation of the digraph originally presented in Figure 1a. Notwithstanding the former definition, Neo4J's implemnetation of the Property Graph model substantially differs from the aforementioned mathematical definition, as it does not allow the definition of an explicit resource identifier for both vertices and edges. After expressing the matching query in Figure 1b in Cypher as **MATCH (a)-[b]->(c)RETURN ***, the resulting morphism table from Figure 3b does not explicitly reference the IDs referring to the specific vertices and edges, thus making it quite impractical to update the values associated to the vertices while continuing to restructure the graph, as this will require to re-match the previously updated data to retain it in the match. Despite this issue might be partially solved by exploiting explicit incremental views over the property graph model [31], this solution had no application in current Neo4J implementations, thus making it impossible to fully test its feasibility within the context of current Neo4J implementation. Furthermore, the elimination of an object previously matched within a morphism will update the table by providing an empty object rather than providing a NULL match. This will motivate us to investigate other ID-based graph data models.

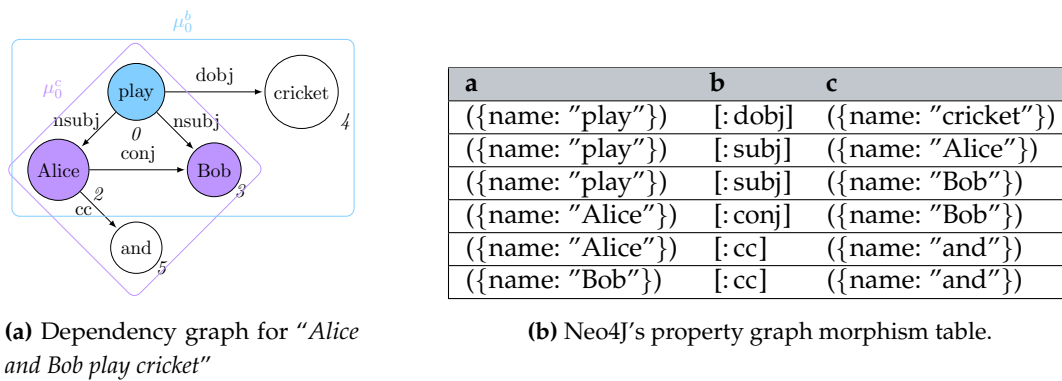


Figure 3. Framing Figure 1 in the context of Neo4J's implementation of the Property Graph model.

Neo4J lists some additional existing limitations beyond APOC and the expressibility of graph grammars in a declarative way with Cypher within their documentation (<https://neo4j.com/docs/operations-manual/current/authentication-authorization/limitations/>, Accessed on November 13, 2023), mainly related to the interplay between data access security requirements and query computations.

At the time of writing, the most studied graph query language both in terms of semantics and expressive power is **SPARQL**, which allows a specific class of queries that can be sensibly optimized [5,30]. The algebraic language used to formally represent SPARQL performs queries' incremental evaluations [35], and hence allows to boost the querying process while data undergoes updates (both incremental and decremental).

While the clauses represented within the WHERE statement are mapped to an optimisable intermediate algebra [30,36], thus including execution of "optional joins" paths [37] for the optional matching of paths, such considerations do not apply for the constraints related to the graph update or return, such as CONSTRUCT, INSERT, and DELETE. While CONSTRUCT is required for returning a graph view as a final outcome, INSERT and DELETE create and remove RDF triplets by chaining them with matching operations. These operations also come with sensible limitations: while the first does not allow to return updated graphs that can be subsequently queried by the same matching algorithm, the two latter statements merely update the underlying data structure and require the re-computation of the overall matching query to retain the updated results. We will partially address these limitations in our query language and data model by associating IDs properties directly via an object-oriented representation while keeping track of the updated information on an intermediate view, which is always accessible within the query evaluation phase.

Last but not least, the usage of so-called *named graphs* allows the selection of over two distinct RDF graphs, which substantially differs from the queries expressible on Cypher, where those can be only computed one graph database at a time. Notwithstanding the former, the latest graph query language standard is very different from SPARQL, hence, for the rest of the paper, we are going to draw our attention to Cypher.

3.2. Nested Relational Model

3.2.1. Logical Model

A nested relational model describes data represented in tabular format, where each table, composed of multiple records, comes with a schema.

Given a set of attributes Σ^* and a set of *datatypes* \mathcal{T} , a **schema** S is a finite function mapping each string attribute in Σ^* to its associated data type ($S: \Sigma^* \rightarrow \mathcal{T}$). A schema is said to be not *nested* if it maps attributes to all basic data types, and *nested* otherwise. This syncretises the traditional function-based notation for schemas within the traditional relational model [38] with the tree-based

characterization of nested schemas [39,40]. While the former focuses on mapping each attribute to a specific data type, the latter focuses on representing attributes associated to other nested attributes.

In this paper, we restrict the *basic datatypes* in $\mathcal{B} \subseteq \mathcal{T}$ to the following ones: vertex-id *ni*, containment-id *ci*, and a label or string *str*. Each of these types are associated to a set of possible values through a \bar{B} function: vertex- and containment-id are associated to natural numbers ($\bar{B}(\text{ni}) = \bar{B}(\text{ci}) = \mathbb{N}$), while the string type is associated to the set of all the possible strings ($\bar{B}(\text{str}) = \Sigma^*$).

A **record** Γ , associated to a schema $\mathcal{S}(\Gamma) = S$, is also a finite function mapping each attribute in $\text{dom}(S)$ to a possible value, either a natural number, a string, or a list of records (**tables**) as specified by the associated schema S ($\forall x \in \text{dom}(\Gamma). \Gamma(x) \in \bar{B}(S(x))$). We define a **table** T with schema $\mathcal{S}(T) = S$ as a list of records all having schema S : i.e., $\forall \Gamma \in T. \mathcal{S}(\Gamma) = S$.

3.2.2. Query Languages

Relational algebra [38] is the de facto standard to decompose relational queries expressed in SQL to its most fundamental operational constituents while providing a well-founded semantics for the SQL implementations. This algebra was later on extended [39,41] to accommodate the extension of the relational model to consider nested relationships. This was also adapted to represent the combination of single edge/triple traversals in SPARQL, so to express the traversal semantics of both required and optional patterns [5].

We now detail a set of operators of interest that will be used across the paper.

We relax the union operation from standard relational algebra by exploiting the notion of **outer union** [42]: given two tables t and s respectively with schema S and U , their outer union is a table $t \sqcup s$ with schema $S \oplus U$ and containing each record from each table where shared attributes are associated to the same types ($\forall x \in \text{dom}(S) \cap \text{dom}(U). S(x) = U(x)$):

$$t \sqcup s = \{\Gamma \mid \Gamma \in t \vee \Gamma \in s\}$$

A **restriction** [38] or **projection** [43] operation $\pi_L(t)$ over a relational table t with schema S returns a new table $\pi_L(t)$ with schema in $S|_L$ where both its schema and its records have a domain restricted to the attributes in L :

$$\pi_L(t) := \langle \Gamma|_L \mid \Gamma \in t \rangle$$

A **renaming** [43] operation $R_{L \rightarrow R}(t)$ over a relational table t with schema S and $L \subseteq \text{dom}(S)$ replaces all the occurrences of attributes in L with ones in R , thus returning a new table with schema $S|_{\text{dom}(S) \setminus L} \oplus \left(\bigoplus_{\langle l_i, r_i \rangle \in \zeta(L, R)} [(r_i, S(l_i))] \right)$:

$$R_{L \rightarrow R}(t) = \left\langle \Gamma|_{\text{dom}(\Gamma) \setminus L} \oplus \left(\bigoplus_{\substack{\langle l_i, r_i \rangle \in \zeta(L, R) \\ l_i \in \text{dom}(\Gamma)}} [(r_i, \Gamma(l_i))] \right) \mid \Gamma \in t \right\rangle$$

A **nesting** [41] operation $\nu_{B \rightarrow A}$ over a table t with schema S returns a new table $\nu_{B \rightarrow A}(t)$ with schema $S|_{\text{dom}(S) \setminus B} \oplus [(A, S|_B)]$, where all the attributes referring to B are nested within the attribute A , which is associated to the type $S|_B$ resulting from the nested attributes. Operatively, it coalesces all the tuples in t sharing the same values not in B as a single equivalence class c : we then restrict one representative of this class to the attributes not in B for then extending it by associating to a novel attribute A the projection of c over the attributes in B $\pi_B(c)$:

$$\nu_{B \rightarrow A}(t) := \langle (\min c)|_{\text{dom}(S) \setminus B} \oplus [(A, \pi_B(c))] \mid c \in t / \cdot|_{\text{dom}(S) \setminus B} \doteq \cdot|_{\text{dom}(S) \setminus B} \rangle \quad (2)$$

A **(natural) join** [44] between two non-nested tables t and s with schemas S and U respectively, if $\text{dom}(S) \cap \text{dom}(U)$ then it combines records from both tables that share the same attributes, and

otherwise extends each record from the left table with a record coming from the right one (cross product):

$$t \bowtie s = \langle \Gamma_i \oplus \Gamma_j \mid \Gamma_i \in t, \Gamma_j \in s, (\Gamma_i \oplus \Gamma_j)|_S = \Gamma_i, (\Gamma_i \oplus \Gamma_j)|_U = \Gamma_j \rangle$$

Given a sequence of attributes $\vec{L} = \langle L_1 \dots L_n \rangle$, this operation can be extended to join the relationship coming from the right at any desired depth level by specifying a suitable path to traverse the nested schema from the left relationship [45]:

$$t \vec{L} \bowtie s = \begin{cases} t \bowtie s & \vec{L} = \emptyset \\ \langle \tilde{t}|_{S \setminus \{L_1\}}, \oplus (\tilde{t}(L_1) \langle L_2, \dots, L_n \rangle \bowtie s) \mid \tilde{t} \in t \rangle & \vec{L} = L_1, L_2, \dots, L_n \end{cases} \quad (3)$$

This paper will automate the determination of \vec{L} given the schemas of the two tables (Section 5). Despite it can be shown that the nested relational model can be easily represented in terms of the traditional not-nested relational model [46], this paper uses the nested relational model for compactly representing graph nesting operations over morphisms. Last, the **left (outer) join** $t \bowtie s$ extends the results from $t \bowtie s$ by also adding all the records from t having no matching tuples in s :

$$t \bowtie s = (t \bowtie s) \cup (t - \pi_S(t \bowtie s)) \quad (4)$$

As per SPARQL semantics, the left outer join represents patterns that might optionally appear.

3.2.3. Columnar Physical Model

Columnar physical models offer a fast and efficient way to store and retrieve data where each table \mathcal{R} [47] with a schema having a domain $\{\text{id}, A_1, \dots, A_n\}$ is decomposed into distinct binary relations \mathcal{R}_{A_i} with a schema with domain $\{\text{id}, A_i\}$ for each attribute A_i in $\text{dom}(\mathcal{R})$, thus requiring to only refer to one record by its id while representing data boolean conditions through algebraic operations. As this decomposition guarantees that the full-outer natural join $\bowtie_{1 \leq i \leq n} \mathcal{R}_{A_i}$ of the decomposed tables is equivalent to the initial relation \mathcal{R} , we can avoid listing NULL values in each \mathcal{R}_{A_i} , thus limiting our space allocation to the values effectively present in our original table \mathcal{R} . Another reason for adopting a logical model compatible to this columnar physical model is to keep provenance information [48] while querying and manipulating the data while carrying out data transformations. By exploiting the features of the physical representation, we are no more required to use the logical model for representing both data and provenance information as per previous attempts for RDF graph data [49], as the columnar physical model allows to natively supporting *id* information for both objects (i.e., vertices) and containments (i.e., edges), thus further extending the RDF model by extensively providing unique ID similarly to what was originally postulated by the EPGM data model for allowing efficient distributed computation [50]. These considerations remarks the generality of our proposed model relying upon this representation.

We now discuss a specific instantiation of this columnar relational model for representing temporal logs: KnoBAB [12]. Despite this representation might sound distant from the aim of supporting an Object-Oriented database, Section 4.2 will outline how this might be achieved. KnoBAB stores each temporal log into three distinct types of tables: (i) a *CountingTable* storing the number of occurrences n of a specific activity label $a \in \Sigma$ in a trace $\sigma^i \in \mathcal{L}$ as a record $\langle a, i, n \rangle$. Such a table, created in almost linear time while scanning the log, comes at no significant cost at data loading. (ii) An *ActivityTable* preserving the traces' temporal information through records $\langle \text{id}, a, i, j, p, x \rangle$ asserting that the j -th event σ_j^i of the i -th log trace σ^i comes with an activity label a and it is stored as the id -th record of such a table. p (and x) points to the records containing the immediately preceding (and following) event of the trace, thus allowing linear scans of the traces. This is of the uttermost importance as the records are sorted by activity label, trace id, and event id for enhancing query run times. (iii) The model also instantiate as many *AttributeTable* $^\kappa$ as the keys $\kappa \in K$ in the data payload associated to temporal events, where each record $\langle a, v, \text{id} \rangle$ remarks that the event occurring as the id -th element within the *ActivityTable* $_\mathcal{L}$ table

with activity label a associates a key κ to a non-null value v . This data model also come with primary and secondary indices further enhancing the access to the underlying data model; further details are provided in [12].

At the time of the writing, no property graph database exploits such a model for efficiently querying data: in particular, Neo4J stores property graphs using a simple data model representation where the whole vertices, relationships, or properties are stored in distinct tables. Notwithstanding the former, from the documentation being available, they do not provide further details on their latest version, so most of the details are missing. On the other hand, triple stores for RDF data have a similar idea for storing edges as in tables. These do not exploit query caching mechanisms as the ones proposed in [51] for also enhancing the execution of multiple relational queries. This approach, on the other hand, was recently implemented in KnoBAB, thus reinforcing our previous claims for extending such previous implementation so to adopt it to a novel logical model.

4. Generalised Semistructured Model v2.0

We continue the discussion of this paper's methodology by discussing the logical model (Section 4.1 on the following page) as a further extension of the Generalised Semistructured Model [34] to explicitly support property-value associations for vertices via π . Despite we propose no property-value extension for containments, we argue that this is minor and that does not substantially change the theoretical and implementation results discussed in this paper for Generalized Graph Grammars. We also introduce novel columnar-oriented physical model (Section 4.2 on the next page) being a direct application of our previous work on temporal databases [12] already supporting collections of databases (log of traces): this will be revised for loading and indexing collection of GSM databases defining our overall physical storage. As the external data to be loaded into the physical model directly represents the logical model, we show that these two representations are isomorphic (Lemma 1) by defining explicitly the loading and serialization operations (Algorithm 2).

Algorithm 2 Loading a Logical Model into the Physical Model and serialising it back

```

1: function LOADINGANDINDEXING( $G = \{g_1, \dots, g_n\}$ )
2:   for all  $g_i = \langle O_i, \ell_i, \xi_i, \epsilon_i, \pi_i, \phi_i, t_{i,\phi} \rangle \in G$  do
3:     for all  $j \in O_i$  do
4:        $L_i(j) := \ell_i(j)$ ;  $X_i(j) := \xi_i(j)$ ;  $C_i(j) := \epsilon_i(j)$ 
5:       ActivityTable.add( $\langle \ell_i(j)[0], i, j, \text{NULL}, \text{NULL} \rangle$ )
6:     end for
7:   end for
8:   INDEX(ActivityTable) ▷ Also sorting the table, [12]
9:   for all  $\kappa \in \Sigma^*$   $j \in O_i$  s.t. ActivityTable[r] =  $\langle \ell_i(j)[0], i, j, p, x \rangle$  do
10:    for all  $t \in \phi_i(j, \kappa)$  s.t.  $t_{i,\phi}(t) = \langle t, w \rangle$  do
11:      PhiTable $^\kappa$ .add( $\langle \ell_i(j)[0], i, j, w, t, t \rangle$ )
12:    end for
13:    if  $\pi_i(j, \kappa) \neq \text{NULL}$  then, AttributeTable $^\kappa$ .add( $\langle \ell_i(j)[0], \pi(j, \kappa), r \rangle$ )
14:  end for
15:  INDEX(ActivityTable $^\kappa$ , PhiTable $^\kappa \mid \kappa \in \Sigma^*$ )
16:  return  $db := \langle L, X, C, \text{ActivityTable}, \bigoplus_{\kappa \in \Sigma^*} [(\kappa, \text{AttributeTable}^\kappa)], \bigoplus_{\kappa \in \Sigma^*} [(\kappa, \text{PhiTable}^\kappa)] \rangle$  ▷ As in [12]
17: end function

18: function SERIALIZATION( $db$ )
19:    $n := \max_{r \in \text{AttributeTable}} r(1)$  (1) Determining the maximum number of GSM databases
20:   for all  $i := 0$  to  $n$  do
21:      $O_i := \{j \mid \exists l, p, x. \langle l, i, j, p, x \rangle \in \text{ActivityTable}\}$ 
22:      $\ell_i := \bigoplus_{j \in O_i} [(l, L_i(j))]$ ;  $\xi_i := \bigoplus_{j \in O_i} [(l, X_i(j))]$ ;  $\epsilon_i := \bigoplus_{j \in O_i} [(i, C_i(j))]$ 
23:      $\pi_i := \bigoplus_{\langle l, v, r \rangle \in \text{AttributeTable}^\kappa} [((j, \kappa), v)]$ 
24:      $\phi_i := \bigoplus_{\langle l, i, s, w, d, t \rangle \in \text{PhiTable}^\kappa} [((s, \kappa), \{t' \mid \exists l', w', d', \langle l', i, j, w', d', t' \rangle \in \text{PhiTable}^k\})]$ 
25:      $t_{i,\phi} := \bigoplus_{\kappa \in \Sigma^*} [(\kappa, \{d, w\})]$ 
26:     yield  $\langle O_i, \ell_i, \xi_i, \epsilon_i, \pi_i, \phi_i, t_{i,\phi} \rangle$ 
27:   end for
28: end function

```

As directly updating the physical model with the changes specified in the rewriting steps of the graph grammar rules might require massive restructuring costs, we instead keep track of such changes in a separate non-indexed representation acting as an incremental view to the entire database. We then refer to this additional structure as a GSM view $\Delta(g)$ for each GSM database g loaded in the physical model (Section 4.3). We then characterise the semantics of $\Delta(g)$ by defining the updating function for g

as a materialisation function considering the incremental changes recorded in $\Delta(g)$ (Section 4.3.2 on page 18).

Last, as we consider the execution of the rewiring steps for each graph as a transformation of the vertices and edges as referenced within each morphism, modulo the updates tracked in $\Delta(g)$, it becomes necessary to introduce some preliminary notation for resolving vertex and edge variables from such morphism (Section 4.4).

4.1. Logical Model

The logical model for GSM describes a single object-oriented database g as a tuple $\langle O, \ell, \xi, \epsilon, \pi, \phi, t_\phi \rangle$, where $O \subseteq \mathbb{N}$ is a collection of object (references). Each object is associated to a list of possible types $\ell(o)$ and to a list of string-representations $\xi(o)$ providing its human-readable descriptions. As an object might be the result of an automated entity-relationship extraction process, each object is associated with a list of confidence values $\epsilon: O \rightarrow \wp(\mathbb{R})$ describing the trustworthiness of the provided information. Differently from the previous definition [34], we extend the previous model to also associate each object to an explicit property-value association for each object $o \in O$ through a finite function $\pi: O \times \Sigma^* \rightarrow \mathcal{V}$.

Differently from our previous definition of our GSM model, we express object relationships through uniquely identified vertex containments similarly to edges as in Figure 1a, so to explicitly reference such containments in morphism tables similarly to Figure 1c. We associate each object o with a containment attribute κ referring to multiple containment IDs via $\phi(o, \kappa) \in \wp(\mathbb{N})$. An explicit index t_ϕ maps each of these IDs to a concrete containment $\langle o_j, w \rangle$ denoting that o_j is contained in o through the containment relationship κ with a confidence score of w . This separation between indices and values is necessary to allow the removal of containment values when computing queries efficiently. This also guarantees the correct correspondence between each value ι in the domain of t_ϕ to the label associated to the ι -th containment requiring each ι will be associated to one solve containment relationship (e.g., $\arg \min_{\kappa \in \Sigma^*} \exists j \in g. \iota \in \phi(j, \kappa)$).

We avoid objects containing themselves at any nesting level by imposing a *recursion constraint* [6] to be checked at indexing time: this allows both the definition of sound structural aggregations, which can be then conveniently used to represent multi-dimensional data-warehouses [34]. Thus, we freely assume that no object shall store the same containment reference across containment attributes [6]: $\forall o, \kappa. \forall o', \kappa'. ((o = o' \wedge \kappa \neq \kappa') \vee o \neq o') \Rightarrow \phi(o, \kappa) \cap \phi(o', \kappa') = \emptyset$. This property can also be adapted to conveniently represent Direct Acyclic Graphs, by representing each vertex $v \in V$ of a weighted property-graph $G = (V, E)$ as a GSM object, and each edge $u \rightarrow v \in E$ with weight w and label κ reflects a containment $\langle v, w \rangle \in t_\phi(\phi(i, \kappa))$. Given this isomorphism $g \simeq G_g$, we can also apply a reverse topological ordering of the vertices of g and denote it as $O_{\text{rtopo}}(g)$.

This also achieves the representation aim of EPGM by natively supporting unique identifiers for both vertices and edges, so to better operate on those by directly referring to their IDs without the need to necessarily carrying out all of its associated payload information [50]. As such, this model leverages all the pros and cons of the previous graph data models by framing them within an object-oriented semistructured model.

4.2. Physical Model

We now describe how the former model can be represented in primary memory for fastening up the matching and rewriting mechanism.

First, we would like to support the loading of multiple GSM databases while being able to operate over them simultaneously similarly to the named graphs in the RDF model. Differently from Neo4J and similarly to RDF's named graphs, we ensure an unique and progressive id across different databases.

We directly exploit our ActivityTable for listing all the objects appearing within each loaded GSM: as such, the i -th record $\langle id, a, g, i, p, x \rangle$ will refer to the i -th object in the g -th GSM database, while a will refer to the first label of $\ell_j(i)$, that is $\ell_j(i)[1]$.

We extend KnobAB's ActivityTable to represent the ϕ containment relationship; the result is a PhiTable^κ for each containment attribute κ : each record $\langle \ell_0, g, o_{src}, w, o_{dst}, \iota \rangle$ refers to a specific GSM database g through which object o_i associated to the first label $\ell_0 = \ell_g(o_i)[0]$ contains o_j with an uncertainty score of w and is associated to an index ι : this expresses the idea that $\iota \in \phi(o_{src}, \kappa)$ with $t_\phi(\iota) = \langle o_{dst}, w \rangle$. At the indexing phase, the table is sorted by lexicographical order over the record's constituents. We extend this table with two indices: a primary index P_I^1 mapping each first occurring ℓ value to the first and the last object within the collection, and a secondary index P_I^2 mapping such each graph id g and object ID o_i to a collection of containment records expressing $\mu(t_\phi, \phi(o_i, \kappa))$ for each o_i and κ such that $\phi(o_i, \kappa) \neq \emptyset$.

We retain the $\text{AttributeTable}^\kappa$ for expressing the properties associated to the GSM vertices, for which we keep the same interpretation from Section 3.2.3 on page 13: thus, each record $\langle a, v, id \rangle$ refers to $\phi(o_i, \kappa) = v$ where o_i appears as the id -th record within the $\text{ActivityTable}_\kappa$, now being used to list all the objects occurring across GSM models.

Last, ℓ and ξ properties are stored straightforwardly into a one-column table and a secondary index mapping each graph id and object ID to a list of record referring to the strings associated to the object ID. The same approach is also used to store the confidence values ϵ associated to each GSM object.

Algorithm 2 shows the algorithms being used for loading all of the GSM databases g_i to be loaded of interest to a single columnar database representation db , as well as providing the algorithm being used to serialize back the stored data into a GSM model for data visualization and materialisation purposes. Given this, we can easily prove the isomorphism between the two shared data structures, thus also providing a proof of correctness of the two transformations which, as a result, explains the formal characterization of such algorithm.

Lemma 1. *A collection of logical model GSM databases $\langle g_i \rangle_{i \leq n}$ is isomorphic to the ones loaded and indexed physical model.*

The proof is given in Appendix B.1. We refer to Section 8 for proofs related to the time complexity for loading, indexing, and serializing operations.

4.3. GSM View $\Delta(g)$

To avoid massive restructuring costs while updating the information indexed in the physical model, we use a direct extension of the logical model to keep track of which vertices were newly generated, removed, or updated during the application of the rule rewriting mechanisms. At query time, we instantiate a *view* $\Delta(g_i)$ for each g_i being loaded within the physical model (Section 6.5). We want our view to support the following operations: (i) creation of new objects, (ii) update of the type/labelling information ℓ , (iii) updating the human-readable value characterization ξ , (iv) update of the containment values, (v) removal of specific objects, and (vi) substitution of previously matched vertices with newly-created or other previously matched ones. While we are not explicitly supporting the removal of specific properties or values, these can be easily simulated by setting specific fields to empty strings or values. A view for g_i is defined as follows:

$$\Delta(g_i) = \langle g_i^\Delta, \Gamma_i, \Gamma_i^v, O_i^+, O_i^-, E_i^-, \rho_i \rangle \quad (5)$$

where $g_i^\Delta = \langle O_i^\Delta, \ell_i^\Delta, \xi_i^\Delta, \epsilon_i^\Delta, \pi_i^\Delta, \phi_i^\Delta, t_{\phi,i}^\Delta \rangle$ is a GSM database holding all the objects being newly inserted alongside with the properties as well as the updated properties for the objects within the graph g (i-iv), Γ refers to the nested morphism being considered while evaluating the query, Γ^v denotes the extension of such morphism with the newly inserted objects through variable declaration, which resulting objects are then collected in O^+ (i). O^- (and E^-) tracks all the removed objects (and specific containment objects, resp.) through their id (v). Last, ρ is a replacement map $\oplus_i[(o_i, o_{\alpha(i)})]$ to be used when evaluating the transformations over morphisms occurring at a higher topological sort

layer, stating to refer to an object $o_{\alpha(i)}$ when o_i occurs (vi). Γ^v and O^+ are used to retain the update information locally to each evaluated morphism, while the rest are shared across the evaluation of each distinct morphism.

We equip $\Delta(g)$ with update operations reflecting the insertion, deletion, update, and replacement operations as per rewriting semantics associated to each graph grammar production rule. Such operations are the following:

START: re-initialises the view to evaluate a new morphism Γ' by discarding any information being local to each specific morphism:

$$\text{START}_{\Delta(g)}(\Gamma') = \langle g^\Delta, \Gamma', \emptyset, \emptyset, O^-, E^-, \rho \rangle \quad (6)$$

DELCONT: We remove the i -th containment relationship from the database:

$$\text{DELCONT}_{\Delta(g)}(i) := \langle g^\Delta, \Gamma, \Gamma^v, O^+, O^-, E^- \cup \{i\}, \rho \rangle$$

NEWOBJ: Generates a new empty object associated to a variable j and with a new unique object ID $|g| + |g^\Delta| + 1$:

$$\begin{aligned} \text{NEWOBJ}_{\Delta(g)}(j) &:= \text{let fresh} := |g| + |g^\Delta| + 1 \text{ in} \\ &\quad \langle g^\Delta, \Gamma, \text{PUT}_{\Gamma^v}(j, \Gamma^v(j) \cup \{\text{fresh}\}), O^+ \cup \{\text{fresh}\}, O^-, E^-, \rho \rangle \end{aligned}$$

REPLOBJ: replaces o_i with o_j iff. o_j was not removed ($o_j \notin O^-$):

$$\text{REPLOBJ}_{\Delta(g)}((o_i, o_j)) := \begin{cases} \Delta(g) & o_j \in O^- \\ \langle g^\Delta, \Gamma, \Gamma^v, O^+, O^-, E^-, \rho \circ [(o_i, o_j)] \rangle & \text{oth.} \end{cases}$$

DELOBJ: We remove an object o_i only if this was already in g or if this was inserted in a previous evaluation of a morphism and, within the evaluation of the current morphism, we remove the original object o_i being replaced by $\tilde{o} = \rho(o_i)$:

$$\begin{aligned} \text{DELOBJ}_{\Delta(g)}(o_i) &:= \text{let } \tilde{o} := \text{OPTGET}_\rho(o_i) \text{ in} \\ &\quad \begin{cases} \langle g^\Delta, \Gamma', \emptyset, O^+, O^- \cup \{o_i\}, E^-, \rho \rangle & \tilde{o} \in O^+ \\ \langle g^\Delta, \Gamma', \emptyset, O^+, O^- \cup \{\tilde{o}\}, E^-, \rho \rangle & \text{oth.} \end{cases} \end{aligned}$$

UPDATE: updates one of the object property functions by specifying which of those should be updated. In other words, this is the extension of PUT^2 (Eq. 1) as a higher-function for updating the view alongside one of this components:

$$\text{UPDATE}_{\Delta(g)}^f(\langle i, j \rangle, u) := \begin{cases} \langle O_i^\Delta, \text{PUT}_{\ell_i}^2(\langle i, j \rangle, u), \xi_i^\Delta, \epsilon_i^\Delta, \pi_i^\Delta, \phi_i^\Delta, t_{\phi,i}^\Delta \rangle & f \equiv \ell \\ \langle O_i^\Delta, \ell_i^\Delta, \text{PUT}_{\xi_i}^2(\langle i, j \rangle, u), \epsilon_i^\Delta, \pi_i^\Delta, \phi_i^\Delta, t_{\phi,i}^\Delta \rangle & f \equiv \xi \\ \langle O_i^\Delta, \ell_i^\Delta, \xi_i^\Delta, \epsilon_i^\Delta, \text{PUT}_{\pi_i}^2(\langle i, j \rangle, u), \phi_i^\Delta, t_{\phi,i}^\Delta \rangle & f \equiv \pi \\ \text{let } n := |\text{dom}(t_{i,\phi}^\Delta)| \text{ in} & f \equiv \phi \\ \text{let } \tilde{\phi} := \text{PUT}_{\phi_i^\Delta}(\langle i, j \rangle, \{n, \dots, n + |u|\}) \text{ in} \\ \text{let } \tilde{t} := t_{\phi,i}^\Delta \oplus \bigoplus_{0 \leq j < |u|} [(j + n, u(j))] \text{ in} \\ \langle O_i^\Delta, \ell_i^\Delta, \xi_i^\Delta, \epsilon_i^\Delta, \pi_i^\Delta, \tilde{\phi}, \tilde{t} \rangle \end{cases}$$

Concerning the time complexity, we can easily see that all operations take $O(1)$ time to compute by assuming efficient hash-based sets and dictionaries. Concerning UPDATE^ϕ , this operation takes $O(1)$ time, as we are merely inserting one pair at a time.

4.3.1. Object Replacement and Resolution

As our query language will have variables to be resolved via matched morphisms and view updates (Appendix A.1), we focus on specific variable resolution operations. Replacement operations should be interpreted as a reflexive and transitive closure over the step-wise replacement operations performed while running the rewriting query (Section 6.5 on page 31).

Definition 4 (Active Replacement). *The **active replacement** function resolves any object ID x into its final replacement vertex following the chain of subsequent unique substitutions of each single vertex in ρ , or otherwise returns x :*

$$\rho_{\Delta(g)}^*(x) := \begin{cases} \rho^n(x) & n' = \max_n . x \in \text{dom}(\rho^n) \wedge \rho^n(x) \neq \rho^{n+1}(x) \\ x & \text{oth.} \end{cases}$$

During an evaluation of a morphism to be rewritten, such replacements and changes should be effective from the next morphism while we would like to preserve the original information while evaluating the current morphism.

Definition 5 (Local Replacement). *The **local replacement** function blocks any notion of replacement while evaluating the original data matched by the current morphism, while activates the changes from the evaluation of any subsequent morphism where such newly-created vertices from the current morphism will not be considered:*

$$\downarrow \rho_{\Delta(g)}(x) := \begin{cases} \rho^*(x) & \rho^*(x) \notin O^+ \\ x & \text{oth.} \end{cases}$$

We consider objects as removed if they have no effective replacements to be enacted in any subsequent morphism evaluation: $x \in \text{dom}(\rho) \wedge x \in O^-$. Thus, we also need to resolve objects' properties (such as ℓ , ξ , π , and ϕ) by considering the changes registered in $\Delta(g_i)$. We want to define a HOF property extraction that is independent from the specific function of choice. We exploit the notion of local replacement (Definition Definition 5) as follows:

Definition 6 (Property Resolution). *Given any property access function ℓ, ξ, ϕ, π , a GSM database g_i and a corresponding view $\Delta(g_i)$ we define the following **property resolution** high-order function:*

$$\downarrow \rho_{\Delta(g)}^f(o) = \begin{cases} \emptyset & \downarrow \rho_{\Delta(g)}(o) \in O_i^- \\ f_{\Delta(g_i)}(\downarrow \rho_{\Delta(g)}(o)) & \downarrow \rho_{\Delta(g)}(o) \in O_i^+ \wedge f_{\Delta(g_i)}(\downarrow \rho_{\Delta(g)}(o)) \neq \emptyset \\ f_{g_i}(\downarrow \rho_{g_i}(o)) & \text{oth.} \end{cases}$$

where we ignore any value associated to a removed vertex in O_i^- (first case), we consider any value stored in $\Delta(g_i)$ as overriding any other value originally in the loaded graph (second case), while returning the original value if the object underwent no updates (last case).

4.3.2. View Materialisation

Last, we define a materialization function as a function updating a GSM database g_i with the updates stored in the incremental view $\Delta(g_i)$. We consider all the objects being inserted (implicitly associated with a 1.0 ε score) and removed, as well as extending all the properties as per the view.

$$\begin{aligned} \text{MATERIALIZE}'(g_i, \Delta(g_i)) = & \langle O_i \cup O_i^+ \setminus O_i^-, \\ & \ell \oplus \ell_i^\Delta, \\ & \xi \oplus \xi_i^\Delta, \\ & \epsilon \oplus \left(\bigoplus_{o \in O_i^\Delta \setminus O_i^-} [(o, 1.0)] \right), \\ & \pi \oplus \pi_i^\Delta, \\ & \bigoplus_{\langle p, k \rangle \in \text{dom}(\phi)} [(\langle p, k \rangle, F(y \mapsto y \notin E_i^-, \phi(p, k)))] \oplus \phi_i^\Delta, \\ & (t_\phi \oplus t_{\phi, i}^\Delta) \rangle \end{aligned}$$

As a rewriting mechanism might add edges violating the recursion constraint, we prune the containments loading to its violation by adopting the following heuristic: after approximating the topological sort by prioritizing the object ID, we remove all the containments generating a cycle where the contained object has an id with a lower value than its container id. From this definition, we then derive the definition of the update of all the GSM databases loaded in the physical model G with their corresponding updates in Δ via the following expression:

$$\text{MATERIALIZE}(G, \Delta) = \mu(\text{MATERIALIZE}', \zeta(G, \Delta)) \quad (7)$$

4.4. Morphism Notation

We consider nested relationships mapping attributes to either basic data types or to not nested schemas, as our query language will syntactically avoid the possibility of arbitrary nesting depths. Given this, any attribute A_i can nest at a maximum level 1 of depth. This will then motivate a similar requirement for the envisioned operator for composing matched containments (collected in relational tables) into nested morphisms (Section 5).

As our query language requires to resolve variables by associating each variable A_i to the values stored in a specific morphism Γ , we need a dedicated function enabling this. We can define a value extraction function for each morphism Γ and attribute $A_i \in \text{dom}(\Gamma)$, returning directly the value associated to A_i in Γ if A_i directly appears in the schema of Γ ($\text{dom}(\Gamma)$), and otherwise return the list of all the possible values associate to it within a nested relationship A_j having A_i in its domain:

$$\text{IDX}_\Gamma(A_i) := \text{let } S := \mathcal{S}(\Gamma) \text{ in } \begin{cases} \langle \Gamma(A_i) \rangle & S(A_i) \in \mathcal{B} \\ \langle \gamma_i(A_i) \mid \gamma_i \in \Gamma(A_j) \rangle & \exists! A_j. A_i \in \text{dom}(S(A_j)) \\ \emptyset & \text{oth.} \end{cases} \quad (8)$$

When resolving a variable, we need to determine whether this refers to a containment or to an object, thus selecting to remove the most appropriate type of constituent indicated within a morphism. So, we can define a function similar to the former for extracting the basic datatypes associated to a given attribute:

$$\text{TIDX}_\Gamma(A_i) := \text{let } S := \mathcal{S}(\Gamma) \text{ in } \begin{cases} S(A_i) & S(A_i) \in \mathcal{B} \\ (S(A_j))(A_i) & \exists! A_j. A_i \in \text{dom}(S(A_j)) \wedge S(A_j)(A_i) \in \mathcal{B} \\ \emptyset & \text{oth.} \end{cases} \quad (9)$$

We also need a function determining the occurrence of an attribute x nested in one of the attributes of S . This will be used for both automating the discovery of the path \vec{L} for joining nested tables from

our recently designed operator (Section 5) or for determining whether two variables belong to the same nested cell of the morphism while updating the GSM view. This boils down to defining a function returning A_j if A_i is an attribute of a table nested in A_j , and NULL otherwise.

$$\text{IDNEST}_S(A_i) := \arg \min_{\substack{A_j \in \text{dom}(S) \\ \text{s.t. } S(A_j) \neq \mathcal{B}}} A_i \in \text{dom}(S(A_j)) \quad (10)$$

Last, we need a function for returning all the object and containment IDs under the circumstance that these contribute to the satisfaction of a boolean expression. We then define such a function returning such IDs at any level of depth of a nested morphism:

$$SE(\Gamma) = \text{let } S = S(\Gamma) \text{ in } \{x \in \text{dom}(\Gamma) \mid S(x) \in \mathcal{B}\} \cup \bigcup_{x \in \text{dom}(S), S(x) \notin \mathcal{B}} SE(\Gamma(x)) \quad (11)$$

5. Nested Natural Equi-Join

Despite previous literature provides a definition for nested natural join, no known algorithmic implementation is available. As our query language will return nested morphisms by gradually composing intermediate tables through natural or left joins, is therefore important to provide an implementation for such an operator. This will be required to combine tables derived from the containment matching (Section 6.3) into nested morphisms, where it is required to join via attributes appearing within nested tables (Section 6.4). Our lemmas show the necessity of this operator by demonstrating the impossibility of expressing it via Eq. 3 directly, while capturing the desired features for generating nested morphisms.

We propose for the first time Algorithm 3 for computing the nested (left outer) equi-join with a path of depth at most 1. The only parameter provided to the algorithm is whether we want a left outer equi-join or a natural one otherwise (`isLeft`) and, given that the determination of the nesting path will depend on the schema of both the left and right operand, we automate (Line 9) the determination of the $\tilde{L} = \langle N \rangle$ path along which compute the nested join for which, we freely assume that we navigate on the nested schema of the left operand similarly to Eq. 3: this assumption comes from our practical use case scenario that we are gradually composing the morphisms provided as a left operand argument with the containment relationships provided as the right argument. Furthermore, to apply the definition from Equation 3 while automating the discovery of the path to navigate to nest the relationship, we require that each attribute appearing from the right table schema might appear as nested in one single attribute from the left table or, otherwise, we cannot automatically determine which left attribute to choose to nest the graph visit (Line 8). Otherwise, we determine an unique attribute from the left table alongside which apply the path descent (Line 9).

The algorithm also takes into account whether no nesting path $\tilde{L} = \langle N \rangle$ is derivable, thus resorting to traditional relational algebra operations: if there are no shared attributes, we boil down to the execution of a cross product (Line 4) and, if none of the attributes from the right table appears within a nested attribute from the left table, then we boil down to a classical left-outer or natural equijoin depending on the `isLeft` parameter (Line 6).

Otherwise, we know that some attributes from the right table appear as nested within the same attribute N of the left table and that the two tables share same non-nested attributes. Then, we initialize the join across the two tables by first identifying the nested attribute N from the left (Line 9). Given IR the attributes appearing as non nested attributes from the left table also appearing in the right one, we partition the tables by \doteq_{IR} , thus identifying the records having the same values for the same attributes in IR (Line 10-11). Then, we start producing the results for the nested join by iterating over the values k appearing in either of the two tables (Line 12): if k appears only over the left table and we want to compute a left nested join (Line 13), we reconstruct the original rows appearing from such table and return them (Line 14). Last, we only consider values k for IR appearing on both tables and, for each

row y from the left table having values in k , we compute the left (or natural equi-)join of $y(N)$ with each row z from the right table and combine the results with k (Line 18).

Algorithm 3 Nested Natural Equi-Join

```

1: function NESTEDNATURALEQUIJOINisLeft( $L, R$ )
2:    $S_L := S(L); S_R := S(R)$ 
3:    $IR := (\text{dom}(S_L) \setminus \{x \in \text{dom}(S_L) \mid S_L(x) \notin \mathcal{B}\}) \cap \text{dom}(S_R)$ 
4:   if  $IR = \emptyset$  then return  $L \times R$  ▷ Cross Product
5:   if  $\bigcup \text{dom}(A_i) \mid A_i \in \text{dom}(S_L) \wedge S_L(A_i) \notin \mathcal{B} \cap \text{dom}(S_R) = \emptyset$  then
6:     if  $\text{isLeft}$  then return  $L \bowtie R$  else return  $L \bowtie R$ 
7:   end if
8:   assert  $\left| \bigcup_{x \in \text{dom}(S_R)} \text{IDNEST}_{S_L}(x) \right| = 1$  ▷ Equation 10
9:    $N := \min \bigcup_{x \in \text{dom}(S_R)} \text{IDNEST}_{S_L}(x)$ 
10:   $LM := \bigoplus_{c \in L / \neq_{IR}} [(c(IR), \pi_{S_L \setminus IR}(c))]$ 
11:   $RM := \bigoplus_{c \in R / \neq_{IR}} [(c(IR), \pi_{S_R \setminus IR}(c))]$ 
12:  for all  $k \in \text{dom}(LM) \cup \text{dom}(RM)$  do
13:    if  $k \notin \text{dom}(RM)$  and  $\text{isLeft}$  then
14:      for  $y \in LM(k)$  yield  $k \oplus y$ 
15:    else if  $k \in \text{dom}(LM)$  then
16:      for all  $y \in LM(k), z \in RM(k)$  do
17:         $y' := \text{copy of } y$ 
18:         $y'(N) := \text{if } \text{isLeft} \text{ then } y'(N) \bowtie z \text{ else } y'(N) \bowtie z$ 
19:        yield  $k \oplus y'$ 
20:      end for
21:    end if
22:  end for
23: end function

```

5.1. Properties

We proof that $\vec{L} \bowtie$ cannot trivially boil down to \bowtie unless $\vec{L} = \emptyset$; otherwise, if A_i is in \vec{L} not appearing as an attribute for the to-be-joined table schemas, we will be left out with the left table and not a classic un-nested natural join. Proofs are postponed to Appendix B.2.

Lemma 2. Given $\mathcal{S}(t) = S$ and $\vec{L} = \langle A_1, A_2, \dots, A_n \rangle$, if $A_1 \notin \text{dom}(S)$, then $t^{\vec{L} \bowtie S} = t$

As this is not a desired feature for an operator which application should be automated, this justifies the need for a novel nested algebra operator for composing nested morphisms, which should shift to left joins [37] for composing optional patterns provided within the right operand (isLeft), while also backtracking to natural joins or cross products if no nested attribute is shared across matched containments. The following lemma discards the possibility of the aforementioned limitation to occur from our operator, by instead capturing the notion of cross-products when tables are not sharing non-nested attributes.

Lemma 3. Given tables L and R respectively with schema S and U with non shared attributes ($\text{dom}(S) \cap \text{dom}(U) = \emptyset$), either $\text{NESTEDNATURALEQUIJOIN}_{\text{False}}(L, R)$ or $\text{NESTEDNATURALEQUIJOIN}_{\text{True}}(L, R)$ compute $L \times R$.

We also demonstrate that the proposed operator falls back to the natural join when no attribute nested in the left operand appears in the right one, while also capturing the notion of left join by changing the isLeft

Lemma 4. Given tables L and R respectively with schema S and U where no nested attribute appearing in the left table appears in the schema of the second, then $\text{NESTEDNATURALEQUIJOIN}_{\text{false}}(L, R) = L \bowtie R$ and $\text{NESTEDNATURALEQUIJOIN}_{\text{true}}(L, R) = L \bowtie R$.

The next lemma observes that our proposed operator not only nests the computation of the join operator within a table, but also the implementation of an equijoin doing a value match across the table fields that are shared within the shallowest level. This is a *desideratum* to guarantee the composition of nested morphisms within the same GSM database id, thus requiring to share at least the same dbid

field (Section 6.3). Still, these operations cannot be expressed through the nested join operator available from current literature (Eq. 3).

Lemma 5. *Given tables L and R respectively with schema S and U , that is $\mathcal{S}(L) = S$ and $\mathcal{S}(R) = U$, where the left table has a column N ($N \in \text{dom}(S)$) being nested ($S(N) \notin \mathcal{B}$) and also appearing in the right table ($N \in \text{dom}(U)$), $\text{NESTEDNATURALEQUIJOIN}_{\text{false}}(L, R)$ cannot be expressed in terms of $L \langle^N \bowtie R$ for $N \in \text{dom}(S) \cap \text{dom}(U)$, $N \in \text{dom}(S(N))$, and $\text{dom}(S(N)) \cap \text{dom}(U) \neq \emptyset$.*

6. Generalised Graph Grammar

After a preliminary and example-driven representation of the proposed query language (Section 6.1), we characterize the semantics of the proposed query language in terms of procedural semantics being subsumed in Algorithm 4. This is defined by the following phases: after determining the order of application of the matching and rewriting rules (Line 2), we match and cache the traversal of each containment relationship to reduce the number of accesses to the physical model (Line 3), from which we then proceed to the instantiation of the morphisms, so to produce the $MT[\cdot, \cdot]$ table (Line 4). This fulfils the matching phase. Finally, by visiting the objects from each GSM database in a reverse topological order, we then access each morphism stored in $MT[\cdot, \cdot]$ for then applying (Line 5) the rewriting rules according to the sorting in Line 2. As this last phase produces the views for each g_i GSM database we then materialize this view and store the resulting logical model on disk (Line 6). Each of the forthcoming sections discusses each of the aforementioned phases.

Algorithm 4 Generalized Graph Grammar (gg) evaluation

1: function GENERALISEDGRAPHGRAMMARS(gg, $G = \{g_1, \dots, g_n\}$)	$\triangleright G \simeq db$ (Lemma 1)
2: SORTRULES(gg)	\triangleright Algorithm 5
3: CACHEINTERMEDIATERESULTS(gg, db)	\triangleright Algorithm 6
4: INSTANTIATEMORPHISMS(gg, db)	\triangleright Algorithm 7
5: $\Delta := \text{GENERATEGRAPHVIEWS}(gg, G)$	\triangleright Algorithm 8
6: return MATERIALIZE(G, Δ)	\triangleright Equation 7
7: end function	

6.1. Syntax and Informal Semantics

We now discuss our novel proposed matching and rewriting language by taking inspiration from graph grammars. To achieve language declarativeness, we not force the user to specify the order of application of the rules as in graph rewriting.

Figure 4 provides the Backus-Naur Form (BNF) [52] for the proposed query language for matching and rewriting object-oriented databases by extending the original definition of graph grammars. Each query (gg) is a list of semi-colon separated rules (rule), where each of those describes a matching and rewriting rule, $L_i \rightarrow_{\Theta} R_i$. For each uniquely-identified rule (pt), we identify a match L_i (obj cont* joining*) and an optional (?) rewrite R_i (op* obj). Those are separated by an optional condition predicate Θ (Appendix A.2 on page 42), providing the conditions upon which the rewriting needs to be applied to the database view, and \leftrightarrow .

L_i is characterized by one single entry point similarly to GraphQL [53], thus defining the main GSM object through which we relativize the graphs' structural changes or updates around its neighbouring vertices, as defined by its ego-network cont of objects being either contained ($-- \text{clabel} \rightarrow \text{obj}$) or containing ($<- \text{clabel} -- \text{obj}$) the entry-point obj. While objects should be always referred through variables, containment relationships might be optionally referred through a variable. Edge traversal beyond the ego-net are expressed through additional edges (joining). We require that at least one of the edges should be a mandatory one. Differently from Cypher, we can match containments by providing more possible alternatives for the containment label rather than just considering one single alternative: this acts as the SPARQL's union operator across differently matched edges, each for a distinct edge label. Please observe that this boils down to a potentially polynomial sub-problem of the usual subgraph isomorphism problem, being still in NP despite the results in Section 8 proposed for the present query language.


String	str	\in	Σ^*
Integer	int	\in	\mathbb{N}
Variables	vars	$::=$	str
Labels	labels	$::=$	(str \parallel)* str
Query	gg	$::=$	(rule ;)* rule
Rule	rule	$::=$	(pt:str) = (entrypoint:obj) cont+ joining* (WHERE Θ)? (\hookleftarrow op* (res:obj))?
Object	obj	$::=$	($\gg?$ vars)
Contain	cont	$::=$	-- clabel \rightarrow obj \leftarrow clabel -- obj -- clabel hook
Join-Contain	joining	$::=$	obj -- clabel \rightarrow obj obj \leftarrow clabel -- obj
CLabel	clabel	$::=$	[$\forall?$?? (var:str :)? labels]
Pred	Θ	$::=$	arg ₁ = arg ₂ arg ₁ \neq arg ₂ arg ₁ \leq arg ₂ arg ₁ \geq arg ₂ $\Theta_1 \vee \Theta_2$ $\Theta_1 \wedge \Theta_2$ TEST <i>script</i> str ₁ unmatched str ₂ .str ₃ str ₁ matched str ₂ .str ₃ FILL Θ
Rewrite	op	$::=$	del str new str set expr ₁ as expr ₂
Val	expr	$::=$	str SCRIPT <i>script</i> label expr src expr dst expr ζ init @ expr ℓ init @ expr π expr ₁ @ expr ₂ ϕ expr ₁ , expr ₂ if Θ over str then expr ₂ else expr ₃
PredArg	arg	$::=$	str expr

Figure 4. Proposed language for Graph Grammars over the GSM expressed in ANTLR4-flavoured BNF notation [54]. Terminal symbols are expressed in green. *script* refers to a double-quoted string representing a program in the Script v2.0 language [34]. Similarly to Regular Expression syntax [52], ? refer to optional sub-expressions, + (or *) indicate one (or zero) or more occurrences of a given sub-expression.

Up to this point, all these features are shared with graph query languages. We now start discussing features extending those: we generate nested embeddings by structurally grouping entry point matches sharing a same containing vertex: this is achieved by specifying the need of grouping an object (\gg) along one of its containment vertices via a containment relationship remarked with \forall . Last, we use return statements in the rewritings to specify which entry-point objects should be replaced by any other matched or newly created objects.

Example 1. Listing 1 expresses the graph grammar rules from Figure 2 in our proposed language: we achieve declarativeness when associating multiple string values coming from nested vertices (and therefore associated to a single variable) to one single vertex, as strings will be normally space-joined (Line 11) with a syntax equivalent to

setting such properties where no nestings are in a morphism (Line 2). A return statement at Line 14 guarantees that, while considering matching a GSM database from Figure 3a, objects 2 and 3 for Alice and Bob in X will be replaced by the newly instantiated “Alice Bob” object h when considering the subsequent creation of the edge “plays” (Line 18). This remarks the need for visiting the GSM database in topological order so to minimise the rewriting costs when the updates are applied. This is also guaranteed by the matching assumption only considering objects within the entry-point’s ego-net, as we ensure to pass-on the information by layers via return statements. Figure 5a shows the result of this rewrite.



(a) Generating a binary relationship between the subject as a single entity and the direct object.

graph	V	n_label	S		l_label	Z	
0	0	nsubj	n	S	dobj	I	Z
			↑2	2		↑1	4
			↑3	3			

(b) Morphisms $M[p_3, g_0]$

graph	X	*					
0	0	H	l_label	I	c_label	c	Z
		2	nsubj	↑2	cc	↑4	5

(c) Morphisms $M[p_2, g_0]$

Figure 5. Applying the rewriting rules expressed in Figure 2 to the graph originally presented in Figure 3a: different colours refer to different matching rules. Filled vertices in the left (and right) graph refer to the distinct vertex entry points (and newly generated components), while uparrows \uparrow are used to differentiate containment ids from the ones for the objects.

Listing 1: Expressing the graph grammar rules represented visually in Figure 2 in our proposed language.

```

1  p1 = (X) --[l:det || nmod_poss || amod || mark || case || punct]-->(Y)
2      ↪ set (π (label l) @ X) as (ξ 0 @ Y)
3      del Y
4      (X);
5
6  p2 = (>> H) <-[∀l:]--(X)
7      --[conj] hook
8      --[? c : cc]-->(Z)
9      ↪ new h // H
10     set (φ orig , h) as H
11     set (π cc @ h) as (ξ 0 @ Z)
12     del Z
13     del l
14     (h);
15

```

```

16 p3 = (V) --[∀n: nsubj]-->(>>S)
17      --[∀l: dobj || ccomp || nmod]-->(Z)
18      ↪ set(φ (SCRIPT "(_[(ξ {\"A\"}) 0])" ), S) as Z
19      del V
20      (S)

```

When grouping of entry-points, we require those to be grouped over one same containing object, to unambiguously refer the nested entry points to one single morphism. This allows the query language to coalesce morphisms.

Example 2. The usefulness of a nested morphism representation can be promptly showed with the example in Figure 5 while focussing on the morphism tables referring to the matching of the subject-verb.object structure of a sentence (Figure 5b). Each morphism can contain two distinct nested relationships, one referring to the subject (S) and one to the object (Z). The possibility of representing such values in a nested morphism allows us to better group vertices to be considered while referring to those with the same variable while keeping unique entry-point instances.

Example 3. With reference to the morphism resulting from matching (con-)/dis-junctions with a sentence (Figure 5c), entry-point grouping allows the creation of one single vertex matching as a single subject for the sentence, thus ensuring the creation of one final single vertex per group of matched entry-points.

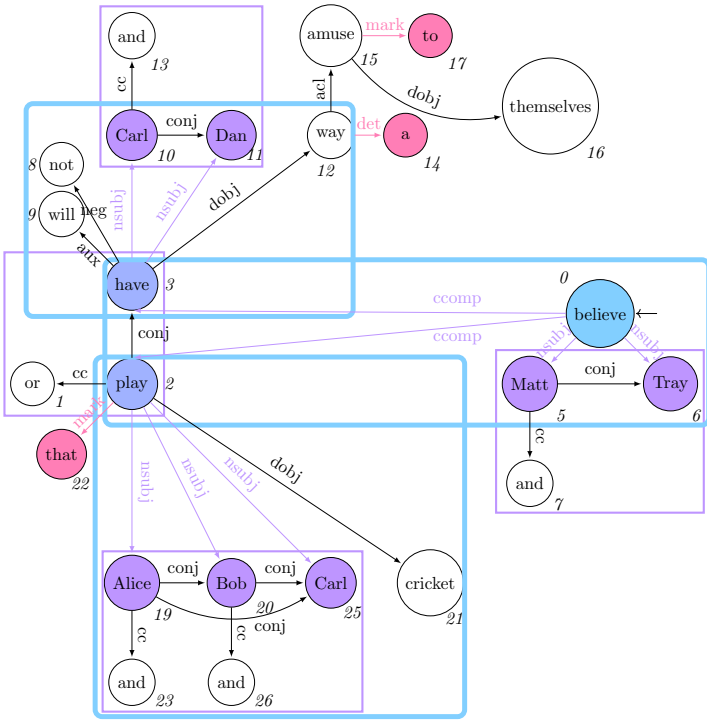
We also show how the language allows to break the declarativeness assumption when we want to specifically compose values according to a specific value composition function:

Example 4. The user is free to break this declarative assumption by directly specifying the order of combination when it is required to combine the values from different variables. This can be observed in a longer query considering more morphosyntactic language features, which is provided online (https://github.com/datagram-db/datagram-db/blob/v2.0/data/test/einstein/einstein_query.txt, Accessed on 18 July, 2024). This can be used to fully rewrite the database as per Figure 6. As the creation of the will-not-have containment in this will require to combine values from vertices 3, 8, and 9 and respectively associated to variables V, A, and N, we can use a scripting language being a direct extension of Script v2.0 [34] for determining the order of strings' composition. Please observe that this formulation, contrary to Neo4j's APOC, also supports optional object matches, where the values associated to non-existing NULL objects are resolved into empty strings (see Proof for Lemma 7 in Appendix B.3).

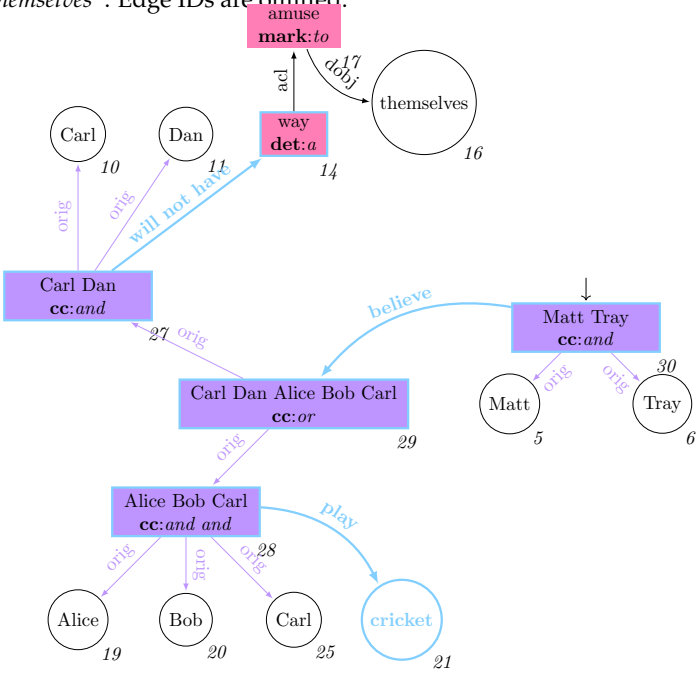
By explicitly expressing a containment relationship across the nested entry-point X via so called hook containments defining equivalence classes, we split the nested morphism Γ into as many new morphisms as the equivalence classes in $\frac{\text{Idx}_{\Gamma}(X)}{\mathfrak{R}}$ (see Section 6.3).

Example 5. We appreciate the usefulness of such morphism splitting while looking at a more convoluted example, as the one considering the rewriting in the sentence depicted in Figure 6a: vertices Matt and Tray as well as play and have from conjunctions, but at different branches of the sentence structure. Furthermore, all four constituent vertices have the same containing vertex believe for which, if no hook relationship was considered, they would have been added within the same morphism as per the previous example.

Still, given that those vertices are associated with different conj as they appear in different coordinating conjunctions, we can use this as a hook relationship to distinguish those, for then obtaining two separate morphisms as illustrated by the first two rows in Figure 7b. Thus, hooks help in splitting nested entry points structurally by identifying similar elements through structural information.



(a) Dependency graph for “Matt and Tray believe that either Alice and Bob and Carl play cricket or Carl and Dan will not have a way to amuse themselves”. Edge IDs are omitted.



(b) Generating a binary relationship between the subject as a single entity and the direct object.

Figure 6. Applying the rewriting rules expressed in Figure 2: different colours refer to different graph grammar rules (b and c), filled vertices in the left (and right) graph refer to the distinct vertex entry points (and newly generated components).

graph	X	l_label	l	Y
1	2	mark	↑12	22
1	12	det	↑1	14
1	15	mark	↑13	17

(a) Morphisms $M[p_1, g_1]$

graph	X	*				
1	0	H	l_label	l	c_label	c
		2	ccomp	↑29	cc	↑22
		3	ccomp	↑30	NULL	NULL
1	0	H	l_label	l	c_label	c
		5	nsubj	↑14	cc	↑23
		6	nsubj	↑15	NULL	NULL
1	2	H	l_label	l	c_label	c
		19	nsubj	↑16	cc	↑25
		20	nsubj	↑17	cc	↑26
		25	nsubj	↑18	NULL	NULL
1	3	H	l_label	l	c_label	c
		10	nsubj	↑19	cc	↑24
		11	nsubj	↑20	NULL	NULL

(b) Morphisms $M[p_2, g_1]$

graph	V	n_label	S	l_label	Z	A	N	M
1	0	nsubj	n	ccomp	l	NULL	NULL	NULL
			S		Z			
			↑14 5		↑29 2			
1	2	nsubj	n	dobj	l	NULL	NULL	22
			S		Z			
			↑16 19		↑3 21			
			↑17 20					
1	3	nsubj	n	dobj	l	9	8	NULL
			S		Z			
			↑19 10		↑4 12			
			↑20 11					

(c) Morphisms $M[p_3, g_1]$

Figure 7. Resulting morphisms from the application of the graph grammar rules from Listing 1 over the GSM database in Figure 6a, from which the resulting rewritten database Figure 6b is then obtained.

Last, these examples provide an intuitive motivation why the current matching expressible within our query language can express distances of at most traversing one containment relationship from the entry-point match. We want to guarantee that, given two objects matching the query entry points, located at different distances from the vertex appearing last within the reverse topological order, these are still reachable within the distance of crossing a single containment. Similarly to semistructured data literature, we refer to one as (direct) ancestor and to the other as its (direct) descendant. If the entry-point considering in its match the aforementioned direct descendant object is replaced with another

object, be it recently created during the application of the rewriting rules or already existing within the database, we want this information to be passed directly during the execution of the rewriting associated with the match of the object of the direct descendant. For this, we need both an explicit return mechanism, which allows the possibility of explicitly telling the objects appearing at the higher layers induced by the topological order that the previous object has been replaced, and to keep the match size compact, so that we can guarantee that any entry-point value updated at a lower level is obtained immediately.

6.2. Determining the Order of Application of the Rules

We determine the application order of our language rules for each entry point vertex of interest (Algorithm 5). This boils down to solving a scheduling problem, which requires first determining the interdependencies across the graph grammar rules. All the matching constructs L_i for each rule $L_i \rightarrow_{\Theta} R_i$ in our query gg have variables which might be shared across morphisms.

Algorithm 5 Sorting the Graph Grammar rules by application order

```

1: procedure SORTRULES( $gg$ )
2:    $V := \{p_i \mid p_i = L_i \rightarrow_{\Theta} R_i \in gg\}$ 
3:    $E := \{p_i \rightarrow p_j \mid p_i.res \neq NULL \wedge p_i.res = p_j.entrypoint \vee$ 
       $p_i.entrypoint = p_j.entrypoint, p_i \in V, p_j \in V\}$ 
4:    $G := (V, E)$ 
5:    $time := LAYERFROMTOPOLOGICALSORT(G, APPROXV_{TOPO}(G))$ 
6:   sort each  $x$  in  $gg$  by  $time[x]$  in ascending order
7: end procedure

```

▷ Algorithm 1

As per Example 1, each rewriting R_i might replace the entry-points with a single new object, or we preserve the previously matched ones otherwise. These are then an input to any later morphism being considered while applying the rewritings. For this, we might consider the variables across patterns as hints to the query language on how the updated or matched objects are going to influence their updates, thus declaring their interdependencies. By reducing this to a dependency ordering, we consider the dependency graphs for the matching and rewriting rules, (Line 4), where each vertex represents a rule (Line 2). Each edge connecting two vertices (or patterns) represents a connection between the entry point or returned variable from the source pattern and any other non-entrypoint variable occurring in the target pattern (Line 3). As the resulting graph might contain loops as some patterns might exhibit mutual dependencies, we are then forced to run an approximated topological sorting algorithm (Line 5) to determine an approximated scheduling order across such tasks through a DFS visit of the graph [55]: we start traversing the graph from the first rule appearing in the graph grammar while avoiding visiting edges leading to already-visited vertices; if we visited all the vertices reachable from such initial vertex while still having unvisited ones, we recommence the same visit starting from the earliest vertex that was not already visited. We add each visited vertex inside a stack after each of its children has been fully visited. By doing so, we prioritise the rules' declaration order which then act as an heuristic for guiding the algorithm to decide upon a specific visiting order.

6.3. Containment Matching

With Algorithm 6, we define the steps realizing containment matching for each L_i from a rule p_i for later on generating a morphism table $MT[L_i, g_j]$ per GSM database g_j , as discussed Section 6.4. The algorithm works as follows: (i) after caching the $\Phi Table^x$ referenced by the containments in the matching patterns so to minimize the tables' access in an uniform schema specification, (ii) we specialize such tables to the specific schema induced by the variables names and nesting occurring in the matching L_i . Last, (iii) we collect the matching containments by separating them between required or optional ones.

Algorithm 6 Intermediate Edge Result Caching

```

1: procedure UNIQUECACHEID( $c = (q, isOut), L$ )
2:   global queryCache, queryMap, emptySet
3:   if  $L = \emptyset$  then emptySet = emptySet  $\cup \{c\}$  else
4:     for all  $l \in L$  do
5:       queryCache = queryCache  $\cup \{l\}$ 
6:     end for
7:   end procedure
8: function FROMDB( $\kappa, db$ )
9:    $t := \langle \rangle$ ;  $\delta(t) := \{ (dbid, N), (src, ni), (edge, ci), (edgeLabel, \Sigma^*), (dst, ni) \}$ 
10:  if  $\Phi Table \in db$  then
11:     $t := \langle (dbid, t), (src, i), (edge, r), (edgeLabel, \kappa), (dst, d) \rangle$  |
12:     $r < |\Phi Table|$ ,  $\Phi Table^{\kappa}(r) = \langle l, l_x, w, d, t \rangle$ 
13:  end if
14:  return  $t$ 
15: function TOTABLE( $t, x, y, l_x, nestCont$ )
16:  if  $l_x = \text{NULL}$  then  $t := R_{src, dst \rightarrow x, y}(\pi_{dbid, src, dst}(t))$ 
17:  else  $t := R_{src, edge, edgeLabel, dst \rightarrow x, l_x + \text{"\_label"}, y}(t)$ 
18:  if nestCont then  $t := \pi_{l_x, y \rightarrow y}(t)$ 
19:  return  $t$ 
20: end function
21: procedure CACHEINTERMEDIATERESULTS( $gg, db$ )
22:  global queryCache, queryMap, emptySet
23:  for all  $p_i = L \rightarrow \Theta R \in gg$  s.t.  $L = \langle ep, in, out, join, hook \rangle$  do
24:    for all  $q \in in$  s.t.  $q = \langle u, l, l_x, all, opt \rangle$  do UNIQUECACHEID( $(q, \text{false}), l$ )
25:    for all  $q \in out$  s.t.  $q = \langle u, l, l_x, all, opt \rangle$  do UNIQUECACHEID( $(q, \text{true}), l$ )
26:    for all  $q \in join$  s.t.  $q = \langle u, l, l_x, all, opt, v \rangle$  do UNIQUECACHEID( $(q, \text{true}), l$ )
27:    queryCache = queryCache  $\cup hook$ 
28:  end for
29:  if emptySet  $\neq \emptyset$  then queryCache =  $\{ \kappa | \Phi Table^{\kappa} \in db \}$ 
30:  cache :=  $\bigoplus_{x \in queryCache} [x, FROMDB(x, db)]$ 
31:  for all  $p_i = L \rightarrow \Theta R \in gg$  s.t.  $L = \langle ep, agg_{ep}, in, out, join, hook \rangle$  do
32:    hook :=  $\pi_{dbid, src, dst}(\bigcup_{i \in hook} TOTABLE(cache(i), src, dst, \text{NULL}, \text{false}))$ 
33:    for all  $q \in out, l \in queryCache(q, \text{true})$  s.t.  $q = \langle u, agg_u, l, l_x, all, opt \rangle$  do
34:       $t := TOTABLE(cache(i), ep, u, l_x, all \text{ and } agg_u)$ 
35:      if opt then  $opt_i := opt_i \cup \{t\}$  else  $req_i := req_i \cup \{t\}$ 
36:    end for
37:    for all  $q \in in, l \in queryCache(q, \text{false})$  s.t.  $q = \langle u, agg_u, l, l_x, all, opt \rangle$  do
38:       $t := TOTABLE(cache(i), u, ep, l_x, \text{false})$ 
39:      if opt then  $opt_i := opt_i \cup \{t\}$  else  $req_i := req_i \cup \{t\}$ 
40:    end for
41:    for all  $q \in join$  s.t.  $q = \langle u, \text{false}, l, l_x, all, opt, v, agg_v \rangle$  do
42:       $t := TOTABLE(cache(i), u, v, l_x, all \text{ and } agg_u)$ 
43:      if opt then  $opt_i := opt_i \cup \{t\}$  else  $req_i := req_i \cup \{t\}$ 
44:    end for
45:  end for
46: end procedure

```

6.3.1. Pseudocode Notation for L_i

We describe the notation used in our pseudocode being the machine readable representation of the query syntax discussed in Section 6.1.

We define each object variable as a pair $\langle x, \mathbf{agg} \rangle \in \mathcal{N} = \Sigma^* \times \{0, 1\}$, where x is the variable name and \mathbf{agg} denotes whether the vertex should be aggregated (1) or not (0). In our pseudocode, each matching L_i is represented as the tuple $\langle ep, in, out, join, hook \rangle$, where $ep \in \mathcal{N}$ specifies the pattern entry-point, and each ingoing (or outgoing) edge is represented by a pair $\langle u, l, l_x, all, opt, v \rangle$, where $u \in \mathcal{N}$ remarks the variable associated to the container (or contained) object alongside the containment relationship through ϕ and t_ϕ , $l \in \wp(\Sigma^*)$ provides a potentially empty-set of containment relationships, l_x is an optional containment variable, $all \in \{0, 1\}$ determines whether the edges should be considered in the aggregation or not, and $opt \in \{0, 1\}$ determines whether the match should be considered as optional or not. The *join* edges extend such records as $\langle u, l, l_x, all, opt \rangle$ by specifying both the containment ($v \in \mathcal{N}$) and container ($u \in \mathcal{N}$) variable explicitly, and $hook \in \wp(\Sigma^*)$ determines whether the aggregated entry-points over the single incoming container should be subdivided in equivalence classes according to the containment labels in *hook*, and we perform no clustering otherwise.

6.3.2. Procedural Semantics for Matching and Caching Containments

We now narrate the operations required to match each containment occurring across patterns while representing those as relational tables expressing either required (req_i), optional (opt_i), or hook-driven equivalence relationships ($hook_i$) per pattern p_i .

First, we define the semantics associated to the matching of each object ego-net as described in TOTABLE from Algorithm 6 (Line 15): either containment $(src) \rightarrow [l_x : L] \rightarrow (dst)$ or $(dst) \leftarrow [l_x :$

$L]$ - (src) are represented as records with fixed schema (Line 9) where each record refers to a single containment ι (edge) in $\phi(\text{src}, \ell)$ for a container src, where $\langle \text{dst}, w \rangle = t_\phi(\iota)$ and dst refers to the contained object; in this, we also retain the containing db as dbid as well as the containment $\ell \in L$ (edgeLabel). At this stage, all containments are associated with the same schema and are not specialized to abide by a specific schema induced by a matching specification. This allows to easily cache PhiTable^κ containments (Line 30).

Next, we discuss how we specialize the results from the cache according to the schema induced by the variables occurring in each matching L_i . This is carried out by renaming generic containing/-containment/contained object labels (src/edge/dst) with the variable names in L_i associated to them; if the patterns also contain references to the edge variable (l_x), we also retain each id in $\phi_x(u, L)$ as l_x , as well as l_x 's label (Line 17) and we discard such information otherwise (Line 16). If the edge expresses an aggregation from the container to the content (e.g., (src)-[$\forall L$]-(\gg dst)), we nest l_x (if available) and dst from the table obtained at the previous step (Line 18). This makes the major difference with existing graph query languages, as we consider containment identifiers as a separate class from object ones (thus differently from SPARQL) and we also produce nested morphisms according to the query annotations.

Last, we collect the tables while differentiating where those are associated to a required or optional pattern (Lines 35, 39, and 43), acting as the basic building step for defining nested morphisms as in the subsequent section.

6.3.3. Algorithmic Choices for Optimization

After discussing the procedural semantics of the matching of containments occurring across all L_i -s, we describe the algorithmic choices for optimization purposes. First, to minimize the access to a relational database, we ensure that each PhiTable^κ is accessed at most once across all the edges by collecting all the labels of interest across table matched containments in queryCache (Line 5). If some containments have no containment attribute specified, we remember the existence of such containment (Line 3) for which we will then require considering all the PhiTable^κ , as a containment for which no label is specified forces to consider all the containments (Line 29). Only after this, we access the physical model to transform each PhiTable^κ per $\kappa \in \text{queryCache}$ to a containment table to be then composed into the final nested morphism table, thus ensuring minimal memory access (Line 30).

6.4. Morphism Instantiation and Indexing

Algorithm 7 combines each table produced from the previous phase to generate the morphisms describing the result of matching L_i through the generation of morphisms being records within a $MT[L_i, g_i]$ table for each GSM database g_i . Similarly to SPARQL's triple navigation semantics, we generate the morphisms P_i for all GSM databases by natural joining all the tables associated to the mandatory containments (Line 10), while left-joining P_i (the resulting table from the natural join of the required containments) against the optional patterns set and updating P_i with such a result (Line 13).

Algorithm 7 Morphism instantiation and indexing

```

1: procedure INSTANTIATEMORPHISMS( $gg, G = \{g_1, \dots, g_n\}$ )  $\triangleright$  global  $MT$ 
2:   for all  $p_i = L_i \rightarrow_{\Theta} R_i \in gg$  do  $\triangleright$  Algorithm 6
3:     global  $req_i, opt_i, hook_i$ 
4:      $\bar{e} := L_i.entrypoint$ 
5:     if  $|req_i| = 0$  then continue
6:     sort each  $t$  in  $req_i$  by  $|t|$  in ascending order
7:      $p_i := req_i(0)$ 
8:     for  $j = 1$  to  $|req_i| - 1$  do
9:       if  $|p_i| = 0$  then break
10:       $p_i := NESTEDNATURALEQUIJOIN_{false}(P_i, req_i(j))$ 
11:    end for
12:    if  $|p_i| = 0$  then continue
13:     $p_i := \Lambda(NESTEDNATURALEQUIJOIN_{true}, P_i, opt_i)$ 
14:    if  $L_i.entrypoint = (\gg z)$  and  $\exists x, obj. \langle \neg \forall x: \dots \rangle_i - obj \in cont$  then
15:       $\bar{e} := obj$ 
16:       $p_i := \nu_{\bar{e}(p_i)}(\{obj, x\_label\} \star (P_i))$ 
17:       $hook_i := TRANS(SYMMETRIC(hook_i))$ 
18:       $\mathfrak{R}_i := \{\mathfrak{R}_i | g_i \in GF \wedge t_{\mathfrak{R}_i} s \Leftrightarrow (t, t(z), s(z)) \in hook_i\}$ 
19:      for all  $\Gamma \in p_i$  and  $\gamma_i \in if(\star) / \mathfrak{R}_i^{\Gamma(dbid)}$  do
20:         $MT[L_i, \Gamma(graph)].add(\Gamma |_{\{graph, x, x\_label\}} \oplus [(\star, \gamma_i)])$ 
21:      end for
22:    else
23:      for all  $\Gamma \in p_i$  do  $MT[L_i, \Gamma(graph)].add(\Gamma)$ 
24:    end if
25:    for all  $g_j \in G$  do
26:      sort each  $\Gamma$  in  $MT[L_i, g_j]$  by  $O_{rtopo}(\Gamma(\bar{e}))$  in ascending order
27:    end for
28:  end for
29: end procedure

```

We further nest the morphism if and only if the entry point is aggregated via a single containment object obj (Line 14), for which we then nest in a fresh attribute \star all the attributes except the database where obj is contained, obj itself, and optionally the edge labels for the containment if the pattern exhibits its variable (Line 16).

Hooks derive an equivalence relationship \mathfrak{R}_i per GSM database g_i having a \star -nested morphism through which optionally split the morphisms. We retain only the container and containment relationship as well as their containing database id (Line 32 from Algorithm 6), for then obtaining a suitable equivalence relationship \mathfrak{R}_i by computing the reflective, symmetric, and transitive closure of that relationship (Line 17). Then, we potentially split the table nested in \star according to the equivalence classes associated with each equivalence relationship \mathfrak{R}_i obtained from hooks (Line 19) and update the morphism table accordingly.

Concerning the composition of cached tables, to reduce the equi-join time we first sort the tables by increasing size (Line 6) for then stopping the joins as soon as we find or compute an empty morphism, thus entailing that no collection of objects across all dbs can match L_i (Line 5 and 9). As a last optimization, we populate MT so collecting the morphisms by database id and rule id (Line 20 and 23). Last, we sort each morphism in $MT[L_i, g_j]$ by entry point (Line 4) reverse topological order (Line 26) or, if these were nested in \star , by their container object (Line 15). This induces a primary block indexing mapping each elected vertex \bar{e} to a set of morphisms containing it.

6.5. Graph Rewriting Operations (op from R_i)

Finally, we apply the transformation operations required by the rewriting side of the rule for each instantiated morphism in MT across all GSM databases. This works by keeping track of the desired changes within a GSM view per loaded GSM database.

We now discuss Algorithm 8. For updating GSM views, we apply the rewriting query for each database g_j as described by the rewriting patterns R_i in gg (Line 2): we visit per GSM database its objects according to the reverse topological order from $O_{rtopo}(g_i)$ (Line 4) while retaining the objects v appearing in the aforementioned primary block index of a non-empty morphism table $M[L_i, g_j]$ for each production rule $p_i = L_i \rightarrow_{\Theta} R_i \in gg$ (Line 5): we skip the morphisms Γ associated with v if either a previously matched vertex was deleted and not replaced with a new one (Line 6), or if the current morphism Γ does not satisfy a possible WHERE condition Θ associated to L_{Θ} (Line 8). For the remaining morphisms, we run the operations listed in R_i in order of appearance (Line 9).

Algorithm 8 Rewriting phase

```

1: function GENERATEGRAPHVIEWS( $gg, G = g_1, \dots, g_n$ )
2:   for all  $g_i \in G$  do
3:      $\Delta(g_i) := \text{new GraphView}(|V(g_i)|)$ 
4:     for all  $v \in \text{Ortopp}(g_i)$  s.t.  $v \notin O_i^-$  do ▷ Section 4.1
5:       for all  $p_i \in L_i \rightarrow O_i$  s.t.  $MT[L_i, g_i] \neq \emptyset$  do
6:         for all  $\Gamma \in MT[L_i, g_i]$  s.t.  $\neg \text{Optional}_L(\text{col}) \Rightarrow \Gamma(\text{col}) \notin O_i^-$  do ▷ Eq. 6
7:            $\Delta(g_i) := \text{START}_{\Delta(g_i)}(\Gamma)$ 
8:           if  $\Theta \neq \text{true}$  and  $|\llbracket \Theta \rrbracket_{\Gamma}^{1, MT}| = 0$  then continue ▷ Figure A2
9:            $\llbracket \Theta \rrbracket_{\Gamma}^{1, MT} := \nu(R, T, p, \Delta(g_i), MT)$  ▷ Figure 8
10:          end for
11:        end for
12:      end for
13:    yield  $\Delta(g_i)$ 
14:  end for
15: end function

```

We now discuss the definition of ν through SOS semantics enclosed within the evaluation of Line 9 in detail. Figure 8 describes the interpretation of all rewriting rules R_j updating the GSM view $\Delta(g_i)$, where the first three update and exploit the functions from Section 4.3. All the y -s are interpreted as evaluated expressions without intermediate assignments. Rule NEWOBJECT creates a new object and refers it to the variable j . Rule DELOBJECT deletes a pre-existing or a newly-created object, and Rule DELOBJECT deletes a single container-containment relationship that was defined at loading time. We can easily distinguish between variables referring to objects or containments by the simple type associated with the attribute. For now, we do not allow the explicit removal of containments at run-time unless the containment is explicitly specified via containment update.

We discuss the `set` update for the vertices' label values; we distinguish the following cases: if both the number of variable and values are in the same number, we assign for each i -th variable the u -th occurring resolved value (LABELZIP); otherwise, we assign to each object associated to the resolved variable the collapsed string values (LABELVALFLAT). We can obtain similar rules for ξ which are omitted here for conciseness, but that can be retrieved in the appendix (Figure A4).

We treat the property π and the containment ϕ updates differently, as they deal with the resolution of three variables. We are also interested whether different resolved variables belong to a same nested table within the morphism or not, with the rationale being that we can freely associate each value within the same nesting row-by-row (P2ZIP, P2'ZIP, and P2''ZIP), while we need to compute the cross-product between the assignments if the value belongs to distinct netsings (Clearly in P3ZIP). This is determined via the second parameter of expression evaluation \rightarrow_E (Appendix A.3) by transferring the attribute information originated by the variable resolution \rightarrow_ρ (Appendix A.1). In all the remaining occasions, we arbitrarily decide to flatten the associated expressions. Please observe that, if the querying user wants more control on the precise value to be associated, they can always refer to SCRIPT expressions, thus breaking the declarative assumptions for a more controlled output.

$$\begin{array}{c}
\frac{\langle \text{"}\iota\text{"}, \Gamma, p, \text{NEWOBJ}_j(\Delta(g)), MT \rangle \rightarrow_v R}{\langle \text{"new } j; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R} \text{NEWOBJECT} \\
\\
\frac{\boxed{\text{TID}_\Gamma(j) = \text{ni}} \quad \langle \text{"}\iota\text{"}, \Gamma, p, F(\text{DELOBJ}, \Delta(g), \text{IDX}_\Gamma(j)), MT \rangle \rightarrow_v R}{\langle \text{"del } j; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R} \text{DELOBJECT} \\
\\
\frac{\boxed{\text{TID}_\Gamma(j) = \text{ci}} \quad \langle \text{"}\iota\text{"}, \Gamma, p, F(\text{DELCONT}, \Delta(g), \text{IDX}_\Gamma(j)), MT \rangle \rightarrow_v R}{\langle \text{"del } j; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R} \text{DELCONT} \\
\\
\frac{\langle x, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, J \rangle \quad \boxed{|\text{Var}| = |\text{Val}|} \quad \text{idx} \in \mathbb{N}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle i, u \rangle) \mapsto \text{UPDATE}_x^\ell(\langle i, \text{idx} \rangle, u), \Delta(g'), \zeta(\text{Var}, \text{Val})), MT \rangle \rightarrow_v R} \text{LABELZIP} \\
\langle \text{"set } \ell \text{ idx @ } x \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle x, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, J \rangle \quad \boxed{|\text{Val}| \neq |\text{Var}|} \quad \text{idx} \in \mathbb{N}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, i) \mapsto \text{UPDATE}_x^\ell(\langle i, \text{idx} \rangle, \lambda(\text{Val}))), \Delta(g'), \text{Var}), MT \rangle \rightarrow_v R} \text{LABELVALFLAT} \\
\langle \text{"set } \ell \text{ idx @ } x \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_N = I_L}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle a, b, c \rangle) \mapsto \text{UPDATE}_x^\pi(\langle a, b \rangle, c), \Delta(g), \zeta(\text{Var}, \text{Name}, \text{Val}))), MT \rangle \rightarrow_v R} \text{PI3ZIP} \\
\langle \text{"set } \pi t @ z \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_L}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle b, a, c \rangle) \mapsto \text{UPDATE}_x^\pi(\langle a, b \rangle, c), \Delta(g), \text{Name} \times \zeta(\text{Var}, \text{Val}))), MT \rangle \rightarrow_v R} \text{PI2ZIP} \\
\langle \text{"set } \pi t @ z \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_N = -1 \vee I_R = I_N \wedge I_L = -1}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle a, b \rangle) \mapsto \text{UPDATE}_x^\pi(\langle a, b \rangle, \lambda(\text{Val})), \Delta(g), \zeta(\text{Name}, \text{Var}))), MT \rangle \rightarrow_v R} \text{PI2'ZIP} \\
\langle \text{"set } \pi t @ z \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_L = I_N \wedge I_R = -1}}{\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle b, c \rangle) \mapsto \text{UPDATE}_x^\pi(\langle \text{Var}(0), b \rangle, c), \Delta(g), \zeta(\text{Name}, \text{Val}))), MT \rangle \rightarrow_v R} \text{PI2''ZIP} \\
\langle \text{"set } \pi t @ z \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \\
\\
\frac{\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \quad \langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_L = I_N \wedge I_R = -1}}{\langle \text{"}\iota\text{"}, \Gamma, p, \text{UPDATE}_{\Delta(g)}^\pi(\langle \text{Var}(0), \lambda(\text{Name}) \rangle, \lambda(\text{Val})), MT \rangle \rightarrow_v R} \text{PIALLFLAT} \\
\langle \text{"set } \pi t @ z \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R
\end{array}$$

Figure 8. Cont.

$$\begin{array}{c}
\frac{
\begin{array}{l}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad I_L = I_N \neq -1 \wedge I_R = -1 \\
\langle \text{"}\epsilon\text{"}, \Gamma, p, F((x, y) \mapsto \text{UPDATE}_x^\pi(\langle \text{Var}(0), y \rangle, \lambda(\text{Val})), \Delta(g), \text{Name}), MT \rangle \rightarrow_v R
\end{array}
}{
\langle \text{"set } \pi t @ z \text{ as } y; \epsilon\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R
} \text{PIVAREXT}
\\[10pt]
\frac{
\begin{array}{l}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_q \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad I_N = -1 \vee I_V = -1 \\
\langle \text{"}\epsilon\text{"}, \Gamma, p, F((x, (a, b)) \mapsto \text{UPDATE}_x^\pi(\langle a, b \rangle, \lambda(\text{Val})), \Delta(g), \text{Id} \times \text{Name}), MT \rangle \rightarrow_v R
\end{array}
}{
\langle \text{"set } \pi t @ z \text{ as } y; \epsilon\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R
} \text{PIEXTOTH}
\\[10pt]
\frac{
p = \langle \langle p', \text{agg} \rangle, \text{in}, \text{out}, \text{join}, \text{hook} \rangle
}{
\langle \text{"}(p')\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v \langle \text{"}\text{"}, \Gamma, p, \Delta(g), MT \rangle
} \text{NOREWR}
\\[10pt]
\frac{
\begin{array}{l}
p = \langle \langle q, \text{false} \rangle, \text{in}, \text{out}, \text{join}, \text{hook} \rangle \quad q \neq p' \\
\langle p', \Delta(g), \text{false} \rangle \rightarrow_o \langle \text{Repl}, \Delta(g), I \rangle \quad \langle q, \Delta(g), \text{false} \rangle \rightarrow_o \langle \text{Orig}, \Delta(g), J \rangle \\
\Delta' := F((x, (y, z)) \mapsto \text{REPLOBJ}_x(y, z), \Delta(g), \text{Orig} \times \text{Repl})
\end{array}
}{
\langle \text{"}(p')\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v \langle \text{"}\text{"}, \Gamma, p, \Delta', MT \rangle
} \text{NOAGGREWR}
\\[10pt]
\frac{
\begin{array}{l}
p = \langle \langle q, \text{false} \rangle, \langle \langle \text{cont}, \text{false} \rangle, l, l_x, \text{all}, \text{opt}, v \rangle, \text{out}, \text{join}, \text{hook} \rangle \quad q \neq p' \\
\langle p', \Delta(g), \text{false} \rangle \rightarrow_q \langle \text{Repl}, \Delta(g), I \rangle \quad \langle q, \Delta(g), \text{false} \rangle \rightarrow_q \langle \text{Orig}, \Delta(g), J \rangle \\
\langle \text{cont}, \Delta(g), \text{false} \rangle \rightarrow_q \langle \text{Cont}, \Delta(g), K \rangle \\
C := \{ \kappa \in \Sigma^* \mid \exists o \in K. \exists l \in \downarrow \rho_{\Delta(g)}^\phi(o). \downarrow \rho_{\Delta(g)}^{t_\phi}(l) \cap \text{Orig} \neq \emptyset \} \\
\Delta' := F((x, (y, z)) \mapsto \text{REPLOBJ}_x(y, z), \Delta(g), \text{Orig} \times \text{Repl}) \\
\Delta'' := F((x, (y, z, t)) \mapsto \text{UPDATE}_x^\phi(\langle y, z \rangle, \text{Repl}), \Delta', \text{Cont} \times C)
\end{array}
}{
\langle \text{"}(p')\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v \langle \text{"}\text{"}, \Gamma, p, \Delta'', MT \rangle
} \text{AGGREWR}
\end{array}$$

Figure 8. Graph Rewriting SOS for view update.

Even in this case, we can formulate the rules for setting the ϕ -s similarly to our previous discussion for the π -s. Therefore, we defer to Figure A5 provided in the Appendix.

We conclude by analysing the semantics associated with the replacement via R_i 's return statement, the last and mandatory operation for declared rewriting operations. We apply no rewriting if the returned variable is the variable entry point (NOREWR). Otherwise, if the entry point variable is not aggregated, we resolve the replacement and entry point (Repl and Orig respectively) and we replace any object in Orig associated with the entry point variable p' with an object in Repl associated with the replacing variable p' (NOAGGREWR). Otherwise (AGGREWR), the rewriting occurs by replacing the objects in Orig, associated with the entry point's container and the objects in Repl, associated with the returned variable p' . Furthermore, given C the containment labels for which the entry-point is contained by its aggregating object in cont, we also update the containing for the cont objects so to also contain via C the replacing objects in Repl. As this provides the final update, we then consider this last resulting GSM view of the resulting view for our rewriting step.

As no SOS rule matches the empty string, no further operation is conducted, and the rewriting program terminates after considering the final rewriting statement.

7. Language Properties

Given the previous language description, we want to characterize its properties by juxtaposing them with Cypher's. Full proofs are provided in Appendix B.3. We start by showing that this, unlike current graph query languages, we propose a rewriting language framed as *generalised graph grammars*:

we relate our proposed language to the graph grammars as, similarly to these, the absence of a matched pattern leads to no view updates. Still, we claim that such language provides a generalization of the former by supporting explicit data-aware update operations over the objects and containments, while also defining an explicit semantics determining the order of the application of such rules, both across rules and within each GSM database.

Lemma 6. *If either we query all the GSM databases with no rules, or all the rules have no rewritings, or none of the matches returned a morphism, or none of the ones being matched pass the associated Θ condition, then we return the same GSM databases as the ones being originally loaded.*

Next, we ensure that the rules are applied in reverse topological order, thus minimising the restructuring cost of the GSM database while achieving declarativeness for rule application, as the user does not specify this within the query formulation, as no matched pattern leads to no view updates.

Property 1. *The rules are applied in (reversed) topological order while visiting each GSM.*

On the other hand, we can show that Cypher forces the user to specify in which order the updates shall be applied, thus breaking the declarative assumption for a good query language.

Lemma 7. *The application of the rewriting in Cypher in lexicographical order requires explicitly determining the order of the morphisms' application and the order of the graph's visit.*

Next, we state the usefulness of an explicit return statement within the interpretation of a rule as it allows to propagate the updates on the subsequently evaluated morphisms.

Lemma 8. *If an entry-point match is deleted and a new object is re-created inheriting all of its properties, this new object will not be connected to the entrypoint's containers unless the newly-returned object was applied.*

To a minor extent, we also show a greater amount of declarativeness if compared to current graph query languages by automatically generating a new object if an update operation occurs with reference to an optionally matched variable that was not matched to an existing object.

Lemma 9. *The setting of properties to variables not associated with a newly-created object variable and not associated with an actual GSM object due to an optional match reduces to the creation of a new object associated with the variable, which then always receives this and any following update associated to the same variable.*

At this point, it could be argued that, although our proposed rewriting language performs queries that cannot be expressed in other graph query languages, this does not return the matched subgraphs as in such other languages, similarly to Cypher's return statement due to the considerations from Lemma 6. The following Lemma shows otherwise. Thanks to this, this language is considered more expressive than Cypher.

Lemma 10. *The proposed graph grammar query language can express traversal queries retaining only the objects and containments being matched and traversed.*

Corollary 1. *The proposed graph query language is more expressive than current graph query languages.*

8. Time Complexity

We now study the computational complexity associated with the algorithms discussed in the previous section and infer this from the implementation details discussed while reasoning on the SOS discussion. Proofs are postponed to Appendix 6.1. Please observe that, as previously noted from previous graph query language literature (Section 3.1.2), the following results do not intend to prove P

vs NP, as we are deliberately expressing a sub-problem of the more generic subgraph isomorphism problem that can be easily captured through algebraic operations. Proofs are postponed to Appendix B.4.

Lemma 11. *The time complexity of sorting the rules within the query is quadratic over the number of query rules.*

Lemma 12. *The intermediate edge result caching time complexity has a polynomial cost being linear in both the loaded databases and the number of available objects.*

Corollary 2. *The cost for generating the nested morphisms is polynomial with the size of the entire data loaded within the physical model.*

Lemma 13. *The rewriting cost is polynomial and linear with the number of rewriting operations and the number of the morphisms.*

Corollary 3. *The time complexity of the whole Generalised Graph Grammars is polynomial with the size of the loaded physical model.*

9. Empirical Evaluation

For our empirical evaluation, we study the use case of graph grammars in the context of rewriting graphs representing the grammatical structure of a natural-language sentence. Universal dependencies [56] capture these syntactical features across languages by exploiting a shared annotation scheme. In this context, the usual approach to graph rewriting boils down to rewriting a PROLOG-like logical program by applying declarative rewriting rules (https://wiki.opencog.org/w/RelEx_Dependency_Relationship_Extractor Accessed on the 22 April 2024) via an unification algorithm [57], where *compound terms* are equivalently expressing binary relationship as well as properties associated to specific vertices. Given the general interest for such an approach within the domain of Natural Language Processing (NLP), the present paper is going to specifically focus on use case scenarios within such domain. This will also give us the opportunity of re-using freely available datasets for sentences for our experiments (https://osf.io/btjqw/?view_only=f31eda86e7b04ac886734a26cd2ce43d and <https://osf.io/rpu37/>, Accessed on the 21 April 2024), which can be then deemed as repeatable.

Notwithstanding the previous approaches, we want to achieve a more data-driven approach to sentence rewriting, where atoms can also be associated with properties and labels, thus motivating the definition of the proposed query language. Furthermore, the extension of the graph grammar language with a script can be used to compose data as well as define boolean conditions allows us to break the declarativeness assumption only when the user wants more control on how data is processed. Thus, we postulate that our proposed query language extends current literature in several aspects way beyond graph database literature, while postulating the possibility of applying concurrently multiple rewriting rules over disparate sentences via a sole declarative query. The proposed approach shares some similarities with programming language research where, after representing a program in terms of its operational semantics, we can apply graph rewriting rules over *abstract semantic graphs* [58], which are usually represented as trees, for which similar considerations like the former can be applied.

We test the implementation of our physical model and associated query language as implemented in our novel object-oriented database, named *Datagram-DB*, which source code associated with the current paper is freely available online (<https://github.com/datagram-db/datagram-db/releases/tag/v2.0>, Accessed on the 22 April 2024). We conducted all our benchmarks on a Dell Precision mobile workstation 5760 running Ubuntu 22.04 LTS. The specifications of the machine include an Intel® Xeon(R) W-11955M CPU @ 2.60GHz × 16, 64GB DDR4 3200MHz RAM.

9.1. Comparing Cypher and Neo4J with Our Proposed Implementation

Given the problems being evidenced by the Cypher query language from Lemma 7, we cannot automate the benchmarking of Cypher for all the possible sentences coming from an on-line available dataset. By recalling the results of this lemma, we observe that, when the query is not able to capture a pattern albeit optionally, this will have a cascade effect on the entire query for which none of the following rewriting operations will be applied. Given this, the same query might not necessarily provide the same output across different sentences having a different sentence structure. Thus, we cannot use one single query for dealing with unpredictable sentence structure that, in the worst case scenario, would require us to write one single query per input sentence. We then preferred to limit our comparison on two sentences, for which we design the minimal Cypher query exactly capturing the syntactical features expressed by these sentences, while using the same query for *Datagram-db* across all the sentences.

We considered two distinct dependency graphs: “*Alice and Bob play cricket*”, the one in Figure 3a, and the one resulting from the dependency parsing of the “*Matt and Tray...*” sentence from Figure 6a. We loaded them in both Neo4j and our proposed GSM database. In Cypher, we then run the query as formulated in the previous section, while we construct a fully declarative query in our proposed graph query language syntax directly representing an extension of the patterns in Figure 2. Examining Table 1, we can see our solution consistently outperforms the Neo4j solution by two orders of magnitude. Furthermore, the data materialisation phase does not significantly impact the overall running time, as its running times are always negligible compared to the other ones. Additionally, Neo4j does not consider a materialisation phase, as the graph resulting from the graph rewriting pattern is immediately returned and stored as a distinct connected component of the previously loaded graph. This then clearly remarks the benefit of the proposed approach for rewriting complex sentences into a more compact machine representation of the dependency graphs.

Table 1. Table displaying results from rewriting the aforementioned sentences.

Data Model		Loading/Indexing (avg. ms)	Querying (avg. ms)	Materialisation (avg. ms)	Total (ms)
Neo4J	Simple	1.83×10^0	7.03×10^0	N/A	8.86×10^1
	Complex	9.32×10^0	1.93×10^2	N/A	2.02×10^2
GSM	Simple	9.63×10^{-2}	4.82×10^{-1}	2.40×10^{-2}	6.02×10^{-1}
	Complex	6.91×10^{-1}	9.00×10^{-1}	6.67×10^{-1}	1.13×10^0

9.2. Scalability for the Proposed Implementation

For testing the scalability of our implemented system, we used a corpora of sentences used for natural-language common-sense question answering [59] which we rewrote into dependency graphs using Stanford NLP [60]. As its output is failing systematically at correctly recognising verbs in passive form when at the present time as well as at recognising negations due to its training-based nature [61], we provide a piece of software amending its output, so that all the syntactic and the semantic information retained within the sentence could be pertained in a graph structure. This server also transforms the internal Stanford NLP dependency graph into a collection of GSM databases serialised in textual format. The resulting server is available online (https://github.com/datagram-db/stanfordnlp_dg_server) Available on the 19th of April, 2024).

Given the resulting GSM representation of the two sentences, we provide two distinct scalability tests: in the first one, we sample the dataset into sentences to determine the algorithm’s scalability in terms of both the number of GSM databases as well as the number of vertices, while in the second we will only consider scalability regarding the sole number of GSM databases while maintaining the sample distribution of the sentence’s lengths, thus reflecting the number of GSM objects within the database.

Concerning the first scenario, we choose the sentences containing $|O_i| \in \{5, 10, 15, 18\}$ words, and for each of them we choose 300 sentences each, thus obtaining sample sets $S_{300}^{|O_i|}$. Then, we further sample $S_4^{|O_i|}$ into three distinct subsets $S_i^{|O_i|}$ having cardinality of $S_i^{|O_i|} = 75 \cdot i$ for which $S_i^{|O_i|} \subset S_{i+1}^{|O_i|}$ for $n > 0$ and $1 \leq i < i + n \leq 4$. This will be useful to then plot the rewriting running times using for the x axis either the number of sentences (or gsm dbs) or the sequence length, so to better analyze the overall time complexity. The results for these experiments are reported in Figure 9.

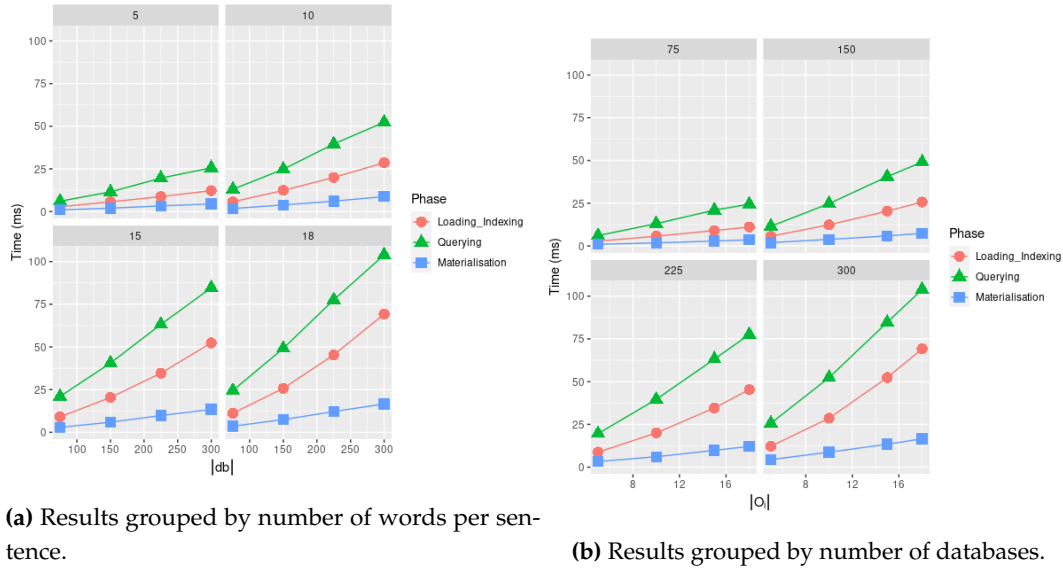


Figure 9. Analysing Datagram-DB

From these plots, we can clearly remark that querying and materialization times are clearly linear over the size of objects being loaded or GSM databases being matched and rewritten, thus remarking the efficiency for the overall envisioned approach: please observe that we cannot achieve a better time complexity than a linear scan without additional heuristics, as we still have to perform the visit over each GSM database by also performing a linear scan of the database objects in reverse topological (layered) order. We can also motivate these results by having graphs in which the branching factor is relatively small compared to the overall number of available vertices, thus $\beta \ll |O_i|$ for each gsm database g_i . We also observe that, for these smaller datasets, the result materialization time is almost negligible compared to the query time, which, across the board, dominates the loading and indexing time.

By comparing such running times with the ones from neo4J, we can easily observe that, while we were able to process 300 GSM databases in 100 milliseconds, Neo4J could rewrite just one single graph in double time. Given this, our approach has a throughput of almost 600 times over Neo4J. This further remarks the impracticality of using the competing solution for analysing more sentences in the future, e.g., while considering sentences crawled from the web.

While it might initially seem that all phases (loading, querying, and materialization) exhibit linear time complexity, we will try to consider a larger set of data to better outline the time complexity associated with our implementation.

Concerning the second scenario, we decided to sample the internet dataset in a subset of sentences S_1, \dots, S_4, S where $|S_i| = 10^i$. S represents the entire dataset, while ensuring that each dataset of a lesser size is always contained in the larger dataset. Furthermore, we ensure that the number of words, and therefore, of objects on each sampled database reflects a similar frequency distribution of the number of objects per resulting GSM database (Figure 10). By doing so, for our final scalability tests in which we consider more data, we make up for the lack of long sentences with the number of sentences reflecting in the number of the GSM databases to be processed.




Figure 10. Sampled probability density function associated with the number of words within the sentences for each subset of traces.

Figure 11 provides the benchmarks for these experiments: a non-linear but polynomial loading time might be possibly related to the data parsing time and on the time to store such data in primary memory, while the remaining running times keep a linear time complexity concerning the increase of the number of the GSM databases to be rewritten, similarly to the previous experiments. Querying time always dominates in indexing time by at least one order of magnitude, thus showing that most of the significant computations occur while matching and rewriting. Materialization times are on the same order of magnitude as indexing time, thus also showing that this cost does not exceed the actual querying phase. Overall, the efficiency of our system is also reflected by a linear querying time for the datasets being considered.




Figure 11. Running time of each algorithm with different sentence samples.

10. Conclusions and Future Works

This paper offers the definition of a matching and rewriting mechanism similarly to the one provided by graph grammars, implemented over object-oriented databases. As our definition supports both data matches and complex data update operations over the objects of interest, which features were not considered in previous graph grammar formulations, we name our proposed language Generalised Graph Grammars. Our theoretical results prove the impossibility of expressing the same query with the same degree of generality of Cypher, which requires to specify a different query for any potential graph to be queried, thus posing a major limitation to automating rewriting operations over graphs. Empirical results prove that our query language offers an implementation being faster

than the ad-hoc query formalised in Cypher and run over Neo4J, in terms of both running time and throughput expressed in number of queries runnable per comparable amount of time.

At this stage, we considered nested morphisms of at most depth of 1, thus requiring that each cell of a morphism table should contain at most one non-nested table. Future works will investigate whether there might be any benefit for further generalising this nesting to arbitrary depth, thus requiring to further extend the definition of the Nested Natural Equi-Join operator to arbitrary depths.

Notwithstanding the possibility of the current model to express uncertainty within the data, the rewriting operations always assume to deal with perfect information, thus generating objects or containments containing no uncertainty. Future works will address this gap by further extending the SOS semantics of these rewriting steps while considering provenance information [48], thus relieving the user from explicitly defining the uncertainty of the generated data by adding further declarativeness to the query language here proposed.

Despite Lemma 10 showed that the proposed graph query language is also able to remove the unmatched objects and contents, our current algorithm is not tailored for effectively matching this case, as the removing pattern will be forced to return all the objects and containments for then removing them. Future works will extend this by optimising testing conditions for our general language while matching the objects and containments. As a direct consequence of this, our returned morphisms are not pruned after testing the Θ condition, which is just evaluated in the rewriting phase due to the fact that any updates to the GSM views will also alter the outcome of Θ . Future works will use static analysis approaches to determine when Θ can be effectively used for pruning morphisms before being returned in the matching phase, as well as condition rewriting strategies to push condition evaluations towards the generation of the morphism in Algorithm 6.4 and 6.3.

Despite the current language supports the update of objects and containments within GSM objects, the provided query semantics does not consider the possibility that such updates can be still matched by the query, thus triggering additional rewriting operations. Future works will also consider further generalising the database matching and rewriting approach by considering the fix-point over the loaded database until convergence is met, thus meeting the former desiderata.

Finally, we will also consider extending our containments with property-values associations similarly to the property graph model, as well as considering updating the objects' and containments' costs while performing the rewriting operations.

Author Contributions: Conceptualization, G.B. and O.R.F.; methodology, G.B.; software, G.B. and O.R.F.; validation, G.B. and O.R.F.; formal analysis, G.B.; investigation, O.R.F.; resources, G.B.; data curation, O.R.F.; writing—original draft preparation, G.B.; writing—review and editing, O.R.F. and G.M.; visualization, G.B. and O.R.F.; supervision, G.B. and G.M.; project administration, G.B.; funding acquisition, G.B. All authors have read and agreed to the published version of the manuscript.

Funding: Not applicable

Institutional Review Board Statement: Not applicable

Informed Consent Statement: Not applicable.

Data Availability Statement: The datasets are available at the following repositories: https://osf.io/btjqw/?view_only=f31eda86e7b04ac886734a26cd2ce43d and <https://osf.io/rpu37/>. The codebase associated to the implementation of the data model and query language interpretation are available on GitHub: <https://github.com/datagram-db/datagram-db/releases/tag/v2.0>. All the URLs were accessed on the 21st of April, 2024.

Conflicts of Interest: The authors declare no conflicts of interest.

Acronyms

BNF Backus-Naur Form. 22, 23

DAG Direct Acyclic Graph. 5

GSM Generalised Semistructured Model. 10, 14, 15, 16, 18, 19, 20, 22, 23, 28, 31, 32, 35, 37, 46

HOF Higher-Order Function. 3, 4, 18

NLP Natural Language Processing. 36, 37

SOS Structured Operational Semantics. 2, 32, 34, 35, 41, 44, 45

Appendix A. Full Definitions

Appendix A.1. Variable Resolution

Before discussing value or boolean expression evaluation, we need to consider their base case first, which is the evaluation of variables occurring in such expressions. We therefore discuss these first before introducing the other steps. We consider variables referring to both objects and containments, while supporting the update operations only for the latter. We therefore resolve the variables while returning the IDs for either matched objects or containments.

We furthermore want to achieve declarativeness while considering variables: if associated to null values as a result of a missed optional match, we want to return empty values, while if we set values to them, we want to create implicitly new vertices. We need to then provide a new parameter to explicitly consider the creation of new objects when the expression context allows to due so over unreferenced and unmatched variables. Furthermore, the single variable might be associated with multiple objects if referring to a morphism attribute within a nested relationship. Since the resolution of variables within our language may require to access previously replaced or substituted values such as from $\Delta(g)$, we also have to refer to $\Delta(g)$ and g to carry out the variable resolution.

$$\begin{array}{c}
 \frac{x \in \text{dom}(\Gamma) \quad \mathcal{S}(\Gamma)(x) \in \mathcal{B}}{\langle x, \Delta(g), \text{new} \rangle \rightarrow_q \langle \downarrow \rho_{\Delta(g)}(\Gamma(x)), \Delta(g), -1 \rangle} \text{EXRES} \\
 \\
 \frac{x \in \text{dom}(\Gamma) \quad \mathcal{S}(\Gamma)(x) \notin \mathcal{B}}{\langle x, \Delta(g), \text{new} \rangle \rightarrow_q \langle \mu(\downarrow \rho_{\Delta(g)}, \text{IDX}_{\Gamma}(x)), \Delta(g), \text{ID}_{\text{NEST}_{\mathcal{S}(\Gamma)}(x)} \rangle} \text{EXRESNEST} \\
 \\
 \frac{x \in \text{dom}(\Gamma^v)}{\langle x, \Delta(g), \text{new} \rangle \rightarrow_q \langle \Gamma^v(x), \Delta(g), -1 \rangle} \text{NEWRES} \quad \frac{x \notin \text{dom}(\Gamma^v) \quad x \notin \text{dom}(\mathcal{S}(\Gamma))}{\langle x, \Delta(g), \text{new} \rangle \rightarrow_q \langle \emptyset, \Delta(g), -1 \rangle} \text{NORES} \\
 \\
 \frac{x \notin \text{dom}(\Gamma) \quad x \in \mathcal{S}(\Gamma) \quad \mathcal{S}(\Gamma)(x) = \text{ni} \quad g' := \text{NEWOBJ}_{\Delta(g)}(x) \quad \tilde{o} := \max g'.O}{\langle x, \Delta(g), \text{true} \rangle \rightarrow_q \langle \tilde{o}, g', -1 \rangle} \text{FORCERES} \\
 \\
 \frac{x \notin \text{dom}(\Gamma) \quad x \in \mathcal{S}(\Gamma) \quad \mathcal{S}(\Gamma)(x) \in \mathcal{B}}{\langle x, \Delta(g), \text{false} \rangle \rightarrow_q \langle \emptyset, \Delta(g), -1 \rangle} \text{NOFRES} \\
 \\
 \frac{x \notin \text{dom}(\Gamma) \quad x \in \mathcal{S}(\Gamma) \quad \mathcal{S}(\Gamma)(x) = \text{ci}}{\langle x, \Delta(g), \text{true} \rangle \rightarrow_q \langle \emptyset, \Delta(g), -1 \rangle} \text{NOFCRES}
 \end{array}$$

Figure A1. Variable resolution with potential view updates (q)

Since these operations can also involve updating the environment, we express these operations via SOS rather than using algebraic notation. Figure A1 shows the SOS for the sole variable resolution, returning a tuple containing (in order of appearance) the resolved variables, the possibly updated view due to object insertion, and the potential nested attribute containing a nested table expressing the variable as an attribute.

If the variable belongs to the morphism's schema and is associated to a non NULL value within the morphism while being associated to a basic type, we return the value stored in it (EXRES). If, otherwise, the variable refers to a nested attribute, we resolve the variable ($\text{IDX}_{\Gamma}(x)$, Eq. 8) and return all the

associated values via $\downarrow \rho_{\Delta(g)}^f$ (Definition 6, EXRESNEST). If, otherwise, the variable was declared within the rewriting pattern alongside the creation of a new object, we return the id associated to the newly created object (NEWRES). If the variable is neither newly declared nor in the morphism's schema, we return no result, as there exists no binding and the query language is not expected to return a value (NORES). Last, if we require the object to be created (`new=true`) as (e.g.,) we set values associated to an object and the morphism did not return an object ID associated to x due to an optional match, we then create a new object in the GSM view \tilde{o} associated to a fresh value, while returning the updated view (FORCERES). This behaviour is completely disabled and no view is updated if the original expression does not force the creation of a new object due to an expression evaluation (NOFRES). In all the other circumstances, we resolve no reference ids (NOFCRES).

Appendix A.2. Predicate Evaluation (Θ)

We now discuss boolean predicate evaluation: given that the variables might refer to nested attributes referring to multiple objects or containment, for which we are interested to compare the labels, values, or properties associated with them, we might be interested in returning not only one single boolean value, but one single boolean value per comparison outcome. Given this, we need a notion of sets explicitly remarking that some elements were not inserted. A **maximum cardinality set** is a pair of a set S and a natural number n denoted as $\wr S, n \wr$ such that $|S| \leq n$. This type of set can be used to efficiently represent how many elements satisfied a given condition if the number of elements is previously known. So, if $|S| < n$, we know that not all the n elements of interest satisfied a given condition. We also constraint such sets to contain at most n items. We can easily override traditional set operations over such sets as follows

$$\wr S, n \wr \cap \wr S', m \wr = \wr S \cap S', \max(n, m) \wr$$

$$\wr S, n \wr \cup \wr S', m \wr = \wr S \cup S', \max(n, m) \wr$$

$$\wr S, n \wr \setminus \wr S', m \wr = \wr S \setminus S', \max(n, m) \wr$$

We say a maximum cardinality set $\wr S, n \wr$ is **full** if the cardinality of S is equal to n : $\text{FULL}(\wr S, n \wr) \Leftrightarrow |S| = n \wedge n \neq 0$. We say that such maximum cardinality set is empty if S is empty: $\wr S, n \wr \simeq \emptyset \Leftrightarrow S = \emptyset$. The cardinality of the maximum cardinality set is the cardinality of its constituent set: $|\wr S, n \wr| = |S|$.

$$\begin{aligned}
\llbracket \text{TEST}_{script} \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \begin{cases} \{S, |S|\} & \text{eval}_{(script)} \text{ holds, } S = \bigcup_{x \in SE(\Gamma)} \text{Idx}_{\Gamma}(x) \\ \{\emptyset, M\} & \text{oth.} \end{cases} \quad (\text{TestScript}) \\
\llbracket \text{arg}_1 \leq \text{arg}_2 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \text{let } t_1 := \eta(\text{arg}_1, \Delta(g), \Gamma, MT) \text{ and } t_2 := \eta(\text{arg}_2, \Delta(g), \Gamma, MT) \text{ in} \quad (\text{A1}) \\
&\quad \begin{cases} \{i \in \mu(\leq, \xi(t_1, t_2)) | i = \text{true}\}, |t_1| & |t_1| = |t_2| \\ \{i \in \mu(x \mapsto t_1(0) \leq x, t_2) | i = \text{true}\}, |t_2| & |t_1| = 1 \\ \{i \in \mu(x \mapsto x \leq t_2(0), t_2) | i = \text{true}\}, |t_1| & |t_2| = 1 \end{cases} \quad (\text{TestComp}) \\
\llbracket \text{FILL } \Theta \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \begin{cases} \{S, |S|\} & |\llbracket \Theta \rrbracket_{\Gamma, g, \Delta(g)}^{MT}| > 0, S = \bigcup_{x \in SE(\Gamma)} \text{Idx}_{\Gamma}(x) \\ \{\emptyset, M\} & \text{oth.} \end{cases} \quad (\text{TestFill}) \\
\llbracket X \text{ matched } L.Y \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \text{IDX}_{\Gamma}(X) \cap \bigcup_{\Gamma' \in M[L, g]} \text{IDX}_{\Gamma'}(Y) \quad (\text{TestMatch}) \\
\llbracket X \text{ unmatched } L.Y \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \text{IDX}_{\Gamma}(X) \setminus \bigcup_{\Gamma' \in M[L, g]} \text{IDX}_{\Gamma'}(Y) \quad (\text{TestUmatch}) \\
\llbracket \Theta_1 \wedge \Theta_2 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \llbracket \Theta_1 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} \cap \llbracket \Theta_2 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} \quad (\text{TestConj}) \\
\llbracket \Theta_1 \vee \Theta_2 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} &:= \llbracket \Theta_1 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} \cup \llbracket \Theta_2 \rrbracket_{\Gamma, g, \Delta(g)}^{MT} \quad (\text{TestDisj})
\end{aligned}$$

Figure A2. Predicate Evaluation semantics $\llbracket \Theta \rrbracket_{\Gamma, g, \Delta(g)}^{MT}$.

Figure A2 provides the definition providing the considered semantics. For this, we require that two original variables refer to the same cardinality of values, or that at least one of them is associated to one single value (TESTCMP). For this, we can then return a maximal cardinality set $\{S, M\}$ where S indicates the resulting tuple indices associated with a true value, and M refers to the maximum size of the morphisms.

Furthermore, we might also be interested to determine whether objects being matched in the current morphism also appear (TESTMATCH) or not (TESTUMATCH) in another morphism L associated to a variable y . Given that we are interested in the objects actually being matched and not in only later changes provided by subsequent transformations, we can simply refer to $\text{IDX}_{\Gamma}(X)$ listing the reference ids associated to an attribute X in Γ . This can be then easily represented as a maximum cardinality set $\{x \in \text{Idx}_{\Gamma}(X) | x \neq \text{NULL}\}, |\text{Idx}_{\Gamma}(X)|$ by stripping the null values, while returning the intersection (TESTMATCH) or the difference (TESTUMATCH) for those IDs.

As the tests in the second last paragraph will return no sets whose indices are related to neither object nor containment id while the former will return such indices, we need an intermediate predicate where these two set-boolean results are computed, so to return for the former a full maximum cardinality set containing all the id references from the morphism Γ if the underlying expression will not return an empty maximum cardinality set (TESTFILL).

We associate a similar behaviour to the typecasting of a script evaluation to a boolean value, for which we return an empty set if this is false, and tall the occurring references within the morphism otherwise (TESTSCRIPT).

Last, we interpret the conjunction and the disjunction of such multivalued boolean semantics as the intersection (TESTCONJ) or the union (TESTDISJ) at the set of reference id satisfying the base case candidates.

Appendix A.3. Expression Evaluation (expr)

Last, we are interested in evaluating expressions exploiting the variables withheld by each morphism. Please also observe that these expressions will have a different operational semantics if compared to the ones associated to assignments, as while the former will only retrieve values associated from the expressions, the latter describe how to update the view with the newly assigned

value. For this, these expressions will only return the final value associated to them. As the morphisms might be also nested, their final representation will be a one column table with an arbitrary attribute, which time is one of the basic types.

$$\begin{array}{c}
\frac{\langle x, \Delta(g), \mathbf{false} \rangle \rightarrow_{\mathcal{Q}} \langle T, \Delta(g), I \rangle}{\langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle T, I \rangle} \text{VARE} \\
\\
\frac{}{\langle \text{"x"}, \Delta(g), \mathbf{false}, \Gamma, MT \rangle \rightarrow_E \langle x, -1 \rangle} \text{STRE} \\
\\
\frac{\text{TIDE}_{\Gamma}(x) = \text{ci} \quad \langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle V, I \rangle}{\langle \text{"label x"}, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle \mu(\ell_{g_i}, V), I \rangle} \text{LABELE} \\
\\
\frac{\text{TIDX}_{\Gamma}(x) = \text{ni} \quad \text{idx} \in \mathbb{N} \quad \langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle V, I \rangle}{\langle \text{"\xi idx@ x"}, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle \mu(\downarrow \rho_{\Delta(g_i)}^{\xi}, V), I \rangle} \text{XIE} \\
\\
\frac{\text{TIDX}_{\Gamma}(x) = \text{ni} \quad \text{idx} \in \mathbb{N} \quad \langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle V, I \rangle}{\langle \text{"\ell idx@ x"}, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle \mu(\downarrow \rho_{\Delta(g_i)}^{\ell}, V), I \rangle} \text{ELLE} \\
\\
\frac{\langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle V, I \rangle \quad \langle \text{"str"}, \Delta(g), \mathbf{false}, \Gamma, MT \rangle \rightarrow_E \langle P, J \rangle \quad |P| = |V|}{\langle \text{"\phi str, x"}, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle \mu(\downarrow \rho_{\Delta(g_i)}^{\phi}, \zeta(P, V)), I \rangle} \text{CONTZIPE} \\
\\
\frac{\langle \text{"x"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle V, I \rangle \quad \langle \text{"str"}, \Delta(g), \mathbf{false}, \Gamma, MT \rangle \rightarrow_E \langle P, J \rangle \quad |P| = 1 \quad |P| < |V|}{\langle \text{"\phi str, x"}, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle \mu(y \mapsto \downarrow \rho_{\Delta(g_i)}^{\phi}(y, P(0)), V), I \rangle} \text{CONTE} \\
\\
\frac{\llbracket \Theta \rrbracket_{\Gamma, g_i}^{1, MT} = \{S, n\} \quad |S| = n \quad \langle y, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle R, I \rangle}{\langle \text{"if } \Theta \text{ over x then y else z"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle R, I \rangle} \text{ALLTRUEE} \\
\\
\frac{\exists M. \llbracket \Theta \rrbracket_{\Gamma, g_i}^{1, MT} = \{\emptyset, M\} \quad \langle z, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle R, I \rangle}{\langle \text{"if } \Theta \text{ over x then y else z"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle R, I \rangle} \text{ALLFALSEE} \\
\\
\frac{\llbracket \Theta \rrbracket_{\Gamma, g_i}^{1, MT} = \{T, n\} \quad \langle x, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle X, I \rangle \quad \langle y, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle L, J \rangle \quad \langle z, \Delta(g), \mathbf{f}, \Gamma, MT \rangle \rightarrow_E \langle R, K \rangle}{\langle \text{"if } \Theta \text{ over x then y else z"}, \Delta(g), \mathbf{true}, \Gamma, MT \rangle \rightarrow_E \langle \mu(i \mapsto \text{ifte}(X(i) \in T, L(i), R(i)), \zeta(\text{dom}(X), L, R)), I \rangle} \text{ZIPE}
\end{array}$$

Figure A3. Expression evaluation (E) with no view update, where $\text{ifte}(x, y, z)$ is an expression returning y if x holds and z otherwise

Figure A3 shows the SOS associated to the expression evaluation. From this, we provide η as a shorthand for the above relationship:

$$\eta(x, \delta, \Gamma, NT) := T \text{ s.t. } \langle x, \delta, \mathbf{false}, \Gamma, MT \rangle \rightarrow_E \langle T, I \rangle$$

In other words, η computes the first projection of \rightarrow_E , being the evaluation of an expression x .

For the evaluation of the attribute associated to a containment id (LABEL), we directly apply λ to all the non-null matches, as containment ids are never updated in this version of the language. For extracting values (XIE) or labels (ELLE) associated to the object x at the numerical index idx not associated to an evaluated expression, we resort instead to the Property Resolution function also encompassing the changes in the GSM view (Definition 6). The interpretation of $\phi(x, \text{str})$ considers both x and str as expressions, where only the former is forced to be interpreted as a variable if is a string: if x and str are associated to a tuple of values V and P of the same cardinality, we then return $\phi(y, z)$ for $(y, z) \in \zeta(V, P)$ (CONTZIPE), and if otherwise $|P| = 1$ we return $\phi(y, P(0))$ for $y \in V$ (ContE).

Last, for conditional expressions, we first evaluate a condition Θ which, as per forthcoming discussion, will return a set of object or containment ids for which the entire expression holds. If such set is full, we return the evaluation of the expression associated to the “then” branch (ALLTRUEE), if it is empty, we return the evaluation of the expression associated to the “else” branch (ALLFALSEE) and, otherwise, we first interpret Θ over the ids referenced by a variable x , and then return the i -th value from the left expression if the MES associated with Θ contains the i -th object reference after x , and the i -th value from the else branch otherwise.

If we need to evaluate the string as a variable as it appears as the leftmost argument of a label, ζ , ℓ , or ϕ , then we use variable resolution ρ (Appendix A.1 on page 41) to evaluate the values associated to the variable (VARE), and we consider this as a simple string otherwise (STRE). Please observe that this invented value is then associated to a unexisting nested morphism -1 .

Appendix A.4. Full SOS Rewriting Specifications

This section describes through Figure A4 and A5 the remaining **set**-update operations that were not reported in Figure 8 for conciseness. These refer to ζ updates (Figure A4), similar to the ones for ℓ , and ϕ (Figure A5), similar to the ones for π .

$$\begin{array}{c}
 \frac{\langle x, \Delta(g), \mathbf{true} \rangle \rightarrow_{\varrho} \langle \text{Var}, \Delta(g') \rangle \quad \langle y, \Delta(g'), \mathbf{false}, \Gamma, MT \rangle \rightarrow \text{Val} \quad \begin{array}{c} |\text{Var}| = |\text{Val}| \\ \text{idx} \in \mathbb{N} \end{array} \quad \langle \text{“}\iota\text{”}, F((x, \langle i, u \rangle) \mapsto \text{UPDATE}_x^{\zeta}(\langle i, \text{idx} \rangle, u), \Delta(g'), \zeta(\text{Var}, \text{Val})), MT \rangle \rightarrow_{\nu} R}{\langle \text{“}\mathbf{set} \ \zeta \ \text{idx}@ \ } x \ \mathbf{as} \ y; \iota\text{”}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_{\nu} R} \quad \text{XIZIP} \\
 \\
 \frac{\langle x, \Delta(g), \mathbf{true} \rangle \rightarrow_{\varrho} \langle \text{Var}, \Delta(g') \rangle \quad \langle y, \Delta(g'), \mathbf{false}, \Gamma, MT \rangle \rightarrow \text{Val} \quad \begin{array}{c} |\text{Val}| \neq |\text{Var}| \\ \text{idx} \in \mathbb{N} \end{array} \quad \langle \text{“}\iota\text{”}, F((x, i) \mapsto \text{UPDATE}_x^{\zeta}(\langle i, \text{idx} \rangle, \lambda(\text{Val}))), \Delta(g'), \text{Var}), MT \rangle \rightarrow_{\nu} R}{\langle \text{“}\mathbf{set} \ \zeta \ \text{idx}@ \ } x \ \mathbf{as} \ y; \iota\text{”}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_{\nu} R} \quad \text{XIVALFLAT}
 \end{array}$$

Figure A4. Remaining setting functions for ζ -s being the extension of the cases covered in Figure 8.

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_N = I_L} \\
\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle a, b, c \rangle) \mapsto \text{UPDATE}_x^\phi(\langle a, b \rangle, c), \Delta(g), \zeta(\text{Var}, \text{Name}, \text{Val})), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHI3ZIP}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_L} \\
\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle b, a, c \rangle) \mapsto \text{UPDATE}_x^\phi(\langle a, b \rangle, c), \Delta(g), \text{Name} \times \zeta(\text{Var}, \text{Val})), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHI2ZIP}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_R = I_N = -1 \vee I_R = I_N \wedge I_L = -1} \\
\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle a, b \rangle) \mapsto \text{UPDATE}_x^\phi(\langle a, b \rangle, \lambda(\text{Val})), \Delta(g), \zeta(\text{Name}, \text{Var})), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHI2'ZIP}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_L = I_N \wedge I_R = -1} \\
\langle \text{"}\iota\text{"}, \Gamma, p, F((x, \langle b, c \rangle) \mapsto \text{UPDATE}_x^\phi(\langle \text{Var}(0), b \rangle, c), \Delta(g), \zeta(\text{Name}, \text{Val})), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHI2''ZIP}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_L = I_N \wedge I_R = -1} \\
\langle \text{"}\iota\text{"}, \Gamma, p, \text{UPDATE}_{\Delta(g)}^\phi(\langle \text{Var}(0), \lambda(\text{Name}) \rangle, \lambda(\text{Val})), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHIALLFLAT}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_L = I_N \neq -1 \wedge I_R = -1} \\
\langle \text{"}\iota\text{"}, \Gamma, p, F((x, y) \mapsto \text{UPDATE}_x^\phi(\langle \text{Var}(0), y \rangle, \lambda(\text{Val})), \Delta(g), \text{Name}), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHIVAREXT}
\end{array}$$

$$\begin{array}{c}
\langle z, \Delta(g), \text{true} \rangle \rightarrow_e \langle \text{Var}, \Delta(g'), I_R \rangle \quad \langle t, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Name}, I_N \rangle \\
\langle y, \Delta(g'), \text{false}, \Gamma, MT \rangle \rightarrow_E \langle \text{Val}, I_L \rangle \quad \boxed{I_N = -1 \vee I_V = -1} \\
\langle \text{"}\iota\text{"}, \Gamma, p, 1F((x, (a, b)) \mapsto \text{UPDATE}_x^\phi(\langle a, b \rangle, \lambda(\text{Val})), \Delta(g), \text{Id} \times \text{Name}), MT \rangle \rightarrow_v R \\
\hline
\langle \text{"}\text{set } (\varphi t, z) \text{ as } y; \iota\text{"}, \Gamma, p, \Delta(g), MT \rangle \rightarrow_v R \quad \text{PHIEXTOTH}
\end{array}$$

Figure A5. Remaining setting functions for ϕ -s being the extension of the cases covered in Figure 8.

Appendix B. Proofs

Appendix B.1. Transformation Isomorphism

Proof (for Lemma 1). Proving this reduces to prove that $\langle g_i \rangle_{i \leq n} = \text{SERIALIZATION}(\text{LOADING}(\langle g_i \rangle_{i \leq n}))$ as well as $db = \text{LOADING}(\text{SERIALIZATION}(db))$. We can prove this by extending the definition of the transformations being given in 2.

$\langle \tilde{g}_i \rangle_{i \leq n} = \text{SERIALIZATION}(\text{LOADING}(\langle g_i \rangle_{i \leq n}))$: We can consider one GSM database \tilde{g}_i being returned at a time, for which we consider its constituents $\langle \tilde{O}_i, \tilde{\ell}_i, \tilde{\xi}_i, \tilde{\epsilon}_i, \tilde{\pi}_i, \tilde{\phi}_i, \tilde{t}_{i,\phi} \rangle$. To do this, we want to show that such returned database is equivalent to the originally loaded g_i . To do this, we need to show that each of the constituents are equivalent.

We can prove that $\tilde{O}_i = O_i$ as follows:

$$j \in \tilde{O}_i \Leftrightarrow \exists l, p, x. \langle l, i, j, p, x \rangle \in \text{ActivityTable} \Leftrightarrow j \in O_i$$

We can prove that $\tilde{\ell}_i \doteq \ell_i$, $\tilde{\zeta}_i \doteq \zeta_i$, and $\tilde{\epsilon}_i \doteq \epsilon_i$ as follows:

$$\tilde{\ell}_i(j) = L_i(j) = \ell_i(j) \quad \tilde{\zeta}_i(j) = X_i(j) = \zeta_i(j) \quad \tilde{\epsilon}_i(j) = C_i(j) = \epsilon_i(j)$$

We can also prove that the properties are preserved:

$$\begin{aligned} \tilde{\pi}_i(j, \kappa) = v &\Leftrightarrow \langle \ell_i(j)[0], v, r \rangle \in \text{AttributeTable} \wedge \exists p, x. \text{ActivityTable}[r] = \langle \ell_i(j)[0], i, j, p, x \rangle \\ &\Leftrightarrow \langle \ell_i(j)[0], v, r \rangle \in \text{AttributeTable} \wedge \text{ActivityTable}[r](2) = j \wedge j \in O_i \\ &\Leftrightarrow \pi(j, \kappa) = v \wedge j \in \tilde{O}_i \end{aligned}$$

For $\tilde{\phi}_i(j, \kappa)$, we can easily show that this is associated to no value only if there are no records referring to a containment for j in the loaded database, which we can easily show that this derives from an originally empty contained from the loaded database g_i . On the other hand, this function returns a set S of identifiers only if there exists at least one containment record in PhiTable^κ , for which we can derive the following:

$$\begin{aligned} \tilde{\phi}_i(j, \kappa) = S &\Leftrightarrow S = \{ \iota \mid \exists l, w, d. \langle l, i, j, w, d, \iota \rangle \in \text{PhiTable}^\kappa \} \\ &\Leftrightarrow S = \{ \iota \mid \iota \in \phi_i(j, \kappa) \} \\ &\Leftrightarrow S = \phi_i(j, \kappa) \end{aligned}$$

Given that $\tilde{\phi}_i(j, \kappa)$ and $S = \phi_i(j, \kappa)$, this sub-goal is closed by transitivity closure over the equality predicate.

We can prove similarly the equivalence $\tilde{t}_{i,\phi} \doteq t_{i,\phi}$ as a corollary of the previous sub-cases:

$$\begin{aligned} \tilde{t}_{i,\phi}(\iota) = \langle d, w \rangle &\Leftrightarrow \exists l, j. \langle l, i, j, w, d, \iota \rangle \in \text{PhiTable}^\kappa \\ &\Leftrightarrow \iota \in \phi_i(j, \kappa) \wedge t_{i,\phi}(\iota) = \langle d, w \rangle \\ &\Leftrightarrow \iota \in \tilde{\phi}_i(j, \kappa) \wedge t_{i,\phi}(\iota) = \langle d, w \rangle \\ &\Leftrightarrow t_{i,\phi}(\iota) = \langle d, w \rangle \end{aligned}$$

$\overline{db} = \text{LOADING}(\text{SERIALIZATION}(db))$: In this case, we need to show that each table in \overline{db} is equivalent to those in db . We can prove similarly to the previous step that $\overline{L} = L$, $\overline{X} = X$, and $\overline{C} = C$.

Next, we need to prove that the ActivityTables begin returned are equivalent. We can achieve this by guaranteeing that both tables should contain the same records. After remembering that the values of p and x are determined through the indexing phase, from which we determine the offset where the objects with immediately preceding and following ids are stored, we can then guarantee that these values will be always the same under the condition that the two tables will always contain the same records, for which these values will be always the same. After remembering from the previous proof that $\tilde{\ell}_i(j) \in \equiv L_i(j) \in db$, for $\tilde{\ell}_i \in \text{SERIALIZATION}(db)$ and $L \in db$, we can also have that $\tilde{\ell}_i(j)(0) \equiv L_i(j)(0)$:

$$\begin{aligned} \langle l, i, j, p, x \rangle \in \overline{\text{ActivityTable}} &\Leftrightarrow \tilde{g}_i \in db \wedge \tilde{\ell}_i(j)[0] = l \wedge j \in \tilde{O}_i \\ &\Leftrightarrow \tilde{g}_i \in db \wedge \tilde{\ell}_i(j)[0] = l \wedge \exists l', p', x'. \langle l', i, j, p', x' \rangle \in \text{ActivityTable} \\ &\quad \wedge L_i(j)[0] = l' \\ &\Leftrightarrow \exists p', x'. \langle l, i, j, p', x' \rangle \in \text{ActivityTable} \end{aligned}$$

Given that the first three components of the record always corresponds, this entails that will preserve the number of records across the board, hence preserving the same record ordering, thus also guaranteeing that $x = x'$ and $x = x'$.

By adopting the equivalence of the previous and next offset for the $\overline{\text{AttributeTable}}$ from the previous sub-proof, we can then also prove that to each record of the $\overline{\text{AttributeTable}}$ always corresponds an equivalent one in AttributeTable .

$$\begin{aligned} \langle l, v, r \rangle \in \overline{\text{AttributeTable}}^k &\Leftrightarrow \overline{\text{ActivityTable}}(r) = \langle l, i, j, p, x \rangle \wedge \tilde{\pi}_i(j, \kappa) = v \\ &\quad \exists l', r'. \langle l', v, r' \rangle \in \text{AttributeTable}^k \wedge \text{ActivityTable}[r'] = \langle l', i, j, p, x \rangle \end{aligned}$$

Given that we can prove that $\tilde{l} = l$ also by the previous sub-case as well as $\tilde{r} = r$ as the records' index will always correspond after sorting and indexing, we can close this sub-goal.

Last, we need to prove that the PhiTable^k tables are equivalent, which can be closed as follows:

$$\begin{aligned} \langle l, i, j, w, d, \iota \rangle \in \overline{\text{PhiTable}}^k &\Leftrightarrow \ell_i(j)(0) = l \wedge \iota \in \phi_i(j, \kappa) \wedge t_{i, \phi}(\iota) = \langle d, w \rangle \\ &\Leftrightarrow \langle l, i, j, w, d, \iota \rangle \in \text{PhiTable}^k \end{aligned}$$

□

Appendix B.2. Nested Equi-Join Properties

Proof (for Lemma 2). If $L \neq \emptyset$ with $\vec{L} = \langle A_1, A_2, \dots, A_n \rangle$, then we can rewrite the definition of $\vec{L} \bowtie$ as follows:

$$t^{\vec{L} \bowtie s} = \left\{ \tilde{t}|_{S \setminus A_1} \oplus (\tilde{t}(A_1) \langle A_2, \dots, A_n \rangle \bowtie s) \mid \tilde{t} \in t \right\}$$

If $A_1 \notin \text{dom}(\tilde{t})$, then $\tilde{t}(A_1) = \emptyset$ by assuming to represent a null value as an empty table by default. While doing so, the former result can be rewritten as:

$$= \left\{ \tilde{t}|_{S \setminus \{A_1\}} \mid \tilde{t} \in t \right\} = t$$

□

Proof (for Lemma 3). Given the hypothesis and with reference to Algorithm 3, we have $IR = \emptyset$, which then yields to $L \times R$ (Line 4). □

Proof (for Lemma 4). Given the hypothesis and with reference to Algorithm 3, this satisfies the condition at Line 5, for which we can then immediately close our goal. □

Proof (for Lemma 5). As $\text{dom}(S) \cap \text{dom}(U) = \{N\}$, this violates the condition for the rewriting Lemma 3, which is not then rewritten accordingly. Furthermore, the condition $\text{dom}(S(N)) \cap \text{dom}(U) \neq \emptyset$ violates the condition for the rewriting Lemma 4, which is then not applied. Given IR as defined in Line 3 of Algorithm 3, we show that this algorithm computes the following:

$$\left\langle \tilde{t}|_{IR} \oplus (\tilde{t}(N) \oplus \tilde{s}|_{U \setminus IR}) \mid \tilde{t} \in L, \tilde{s} \in R, \tilde{t} \dot{=}_{IR} \tilde{s} \right\rangle$$

This equates to equi-joining the tables over the IR records where $N \notin IR$ by construction, and by nesting all the records from the right table sharing the same IR values with the records from the left table. Last, we can easily observe that this cannot be easily expressed in terms of $L \langle N \rangle \bowtie R$, as the rewriting of the former expression in the following way, for which it is evident that the former operator did not consider the recursive natural joining of the records by considering the commonly-shared attributes during its descent:

$$\begin{aligned} L \langle N \rangle \bowtie R &= \langle \tilde{t}|_{S \setminus \{N\}} \oplus (\tilde{t}(N) \diamond \tilde{s}) \mid \tilde{t} \in t \rangle \\ &= \langle \tilde{t}|_{IR} \oplus (\tilde{t}(N) \oplus \tilde{s}) \mid \tilde{t} \in t, \tilde{s} \in R \rangle \end{aligned}$$

This can be easily observed by the lack of $\dot{=}_{IR}$ component that captures the essence of such natural join condition. \square

Appendix B.3. Language Properties

Proof (for Lemma 6). As by the conditions stated in the current Lemma we either generate an empty morphism table $MT[]$ (no rules, no rewritings, no matches) or all the generated morphisms are ignored as all the Θ values are falsified (Line 8 from Algorithm 8), then we will always have an empty view $\Delta(g_i)$ for all the GSM databases g_i loaded in the physical model. Considering this, the proof is an application of Equation 7, by which the materialization of a GSM database g_i with a view $\Delta(g_i)$ containing no changes simply returns g_i . As g_i is stored in a physical model, this result is also validated via Lemma 1. \square

Proof (for Property 1). This condition is generated by Line 4 from Algorithm 8, as we use the reverse topological order index $O_{r_{topo}}(g_i)$ associated to each GSM DB g_i to visit the entry-point objects or, when nested, their containing object associated to the reported morphism in $MT[L_i, G_i]$. \square

Proof (for Lemma 7). As per previous discussions, the proof for this lemma will be given intuitively by analysing the Cypher representation of the graph grammar represented visually in Figure 2 and previously represented by our proposed Generalised Graph Grammar language in Listing 1. We then provide the query required for rewriting the sentence expressed in Figure 6a:

The current Neo4j implementation does not support the theorised graph incremental views for Cypher [31]. As we require to update the graph while querying, it is not possible to entirely create a new graph without restructuring or expanding a previously loaded one as per graph joins [7] or nesting [6] by simply returning some newly-created relationships or nodes; returning a new graph and rewriting a previous match will come at the cost of either restructuring the previously loaded graph, thus requiring additional overhead costs for re-indexing and updating the database while querying, or by creating a new distinct connected component within the loaded graph (CREATE statements from Listing 2). As it is impossible to refer by the vertices and edges through their ID, thus exploiting graph provenance techniques for mapping the newly created vertices to the ones from the previously loaded graph [49], we are forced to join the loaded vertices (e.g., Lines 35-37, 50-52, and 67) with the newly created ones (e.g., Lines 38, 53, and 68 respectively) by their property values (e.g., Lines 39, 54, and 70 respectively). Our proposed approach avoids such cost via the aforementioned objects' and containments' ID-based morphism representation while keeping track of the restructuring operations (property UPDATE, insertion NEWOBJ, deletion DELOBJ and DELCONT, and substitution REPLOBJ) over a graph g within an incremental view $\Delta(g)$ (Section 4.3).

Listing 2: Cypher representation for Figure 2 to rewrite the sentence from Figure 6a.

```

1  // Create the grouped vertices
2  MATCH (a)-[:subj]->(b)-[:cc]->(c)
3  WHERE (b)<-[:conj]-() OR (b)-[:conj]->()
4  WITH a, Collect(b.name) as names, c
5  WHERE size(names) > 1 // To make sure it is only grouping vertices with
   ↳ multiple names
6  OPTIONAL MATCH (a)-[:neg]->(d)
7  CREATE (x {name: apoc.text.join(names, ' '), cc: c.name, dobjRel: a.name
   ↳ })
8  WITH a, c, d, x, names
9  CALL apoc.create.addLabels(x, names) YIELD node
10 CALL apoc.do.when(
11     d IS NOT NULL,
12     "SET x.neg = d.name",

```

```

13     '',
14     {c:c, d:d, a:a, x:x})
15 YIELD value as neg
16
17 // Create the origins
18 WITH neg
19 MATCH (n) // Match only vertices with no relationships
20 WITH LABELS(n) as nameLabels, Collect(DISTINCT n) as nameNodes
21 WHERE size(nameLabels) > 0
22 FOREACH (nNode IN nameNodes |
23     FOREACH (nLabel IN nameLabels |
24         CREATE (y {name: nLabel})<-[:orig]-(nNode)))
25
26 // Create dobj node
27 WITH *
28 MATCH (d)<-[:dobj]-(a)
29 WITH Collect(DISTINCT d.name) as objects, a
30 FOREACH (obj IN objects |
31     CREATE (y {name: obj, dobjRel: a.name}))
32
33 // Check for 'det'
34 WITH *
35 MATCH (n) WHERE NOT (n)-[]->() AND NOT (n)<-[]-()
36 WITH n
37 MATCH (d)<-[:dobj]-(a)
38 MATCH (e)-[:orig]->()
39 WHERE e.dobjRel = a.name AND d.name = n.name
40 OPTIONAL MATCH (d)-[:det]->(f)
41 CALL apoc.do.when(
42     f IS NOT NULL,
43     'SET n.det = f.name',
44     '',
45     {n:n, f:f})
46 YIELD value as det
47
48 // Create relationships
49 WITH det
50 MATCH (m) WHERE NOT (m)-[]->() AND NOT (m)<-[]-()
51 WITH m
52 MATCH (d)<-[:dobj]-(a)
53 MATCH (e)-[:orig]->()
54 WHERE e.dobjRel = m.dobjRel
55 WITH DISTINCT e, m
56 CALL apoc.do.when(
57     e.neg IS NOT NULL,
58     'CALL apoc.create.relationship(e, e.neg + "_" + e.dobjRel, {}, m)
59     → YIELD rel RETURN rel',
60     'CALL apoc.create.relationship(e, e.dobjRel, {}, m) YIELD rel RETURN
61     → rel',
62     {e:e, m:m})

```



```

61 YIELD value as dobjRel
62
63 // Create ? acl node
64 WITH dobjRel
65 MATCH (p) WHERE NOT (p)-[]->() AND NOT (p)<-[]-()
66 WITH p
67 MATCH (a)-[:acl]->(b)
68 MATCH (c)<-[:orig]-(d)-[]->(e)
69 WITH DISTINCT p, a, e
70 WHERE e.name = a.name
71 CREATE (y {name: '?'})<-[:acl]-(e)
72 WITH p, y
73 CALL apoc.create.relationship(y, p.dobjRel, {}, p) YIELD rel as aclRel
74
75 // Group the groups
76 WITH aclRel
77 MATCH (a)-[:conj]->(b)-[:cc]->(c)<-[:cc]-(a)
78 MATCH ()<-[:orig]-(d)-[r]->()
79 WHERE type(r) CONTAINS a.name OR type(r) CONTAINS b.name
80 WITH a, b, c, d, r
81 ORDER BY d.name ASC
82 WITH Collect(DISTINCT d.name) as names, c
83 CREATE (x {name: apoc.text.join(names, ' ' + c.name + ' ')})
84 WITH x, names
85 CALL apoc.create.addLabels(x, names) YIELD node as groupedGroup
86
87 // Add properties to rels
88 // mark
89 WITH groupedGroup
90 MATCH (c)<-[:mark]-(a)-[:dobj]->(b)
91 MATCH (d)-[r]->(e)
92 WHERE type(r) CONTAINS a.name AND b.name = e.name
93 CALL apoc.create.setRelProperty(r, 'mark', c.name) YIELD rel as markRel
94
95 // aux
96 WITH markRel
97 MATCH (d)<-[:aux]-(a)-[:subj]->()-[:cc]->()
98 MATCH ()<-[:orig]-(f)-[r]->()
99 WHERE type(r) CONTAINS a.name
100 CALL apoc.create.setRelProperty(r, 'aux', d.name) YIELD rel as auxRel
101
102 // Create rels to orig groups
103 WITH auxRel
104 MATCH (a)-[:orig]->(b)
105 MATCH (x) WHERE NOT (x)-[]->() AND NOT (x)<-[]-()
106 WITH DISTINCT a, x
107 WHERE apoc.text.join(LABELS(a), ' ') IN LABELS(x)
108 CALL apoc.create.relationship(x, 'orig', {}, a) YIELD rel as origRel
109
110 // Rel from M+T to A+B+C|C+D

```

```

111 WITH x
112 MATCH (d)-[:subj]-(a)-[:ccomp]->(b)-[:subj]->(c)
113 MATCH (e)-[:orig]->(f)
114 WITH a, b, c, d, e, f, x
115 ORDER BY d.name ASC, c.name ASC
116 WITH Collect(DISTINCT c.name) as endNames, x, Collect(DISTINCT d.name)
    → as startNames, e, a
117 WHERE apoc.text.join(startNames, ' ') = e.name AND apoc.text.join(
    → endNames, ' ') IN LABELS(x)
118 CALL apoc.create.relationship(e, a.name, {}, x) YIELD rel as finalRel
119
120 WITH finalRel
121 MATCH q=()-[]->()
122 MATCH (z)
123 CALL apoc.create.removeLabels(z, LABELS(z)) YIELD node // Remove all
    → labels from vertices
124 WITH q, z
125 RETURN q, z

```

Cypher does not ensure to apply the graph rewriting rules as intended in our scenarios: let us consider the dependency graph generated from the recursive sentence “*Matt and Tray believe that either Alice and Bob and Carl play cricket or Carl and Dan will not have a way to amuse themselves*” and let us try to express patterns in Figure 2b and 2c as two distinct **MATCH**-es with their respective update operations as per the following Listing:

```

1 MATCH (a)-[b:cc]->(c)
2 WITH Collect(a.name) as names, Collect(DISTINCT a) as nameNodes, c
3 CREATE (x {name: apoc.text.join(names, ' '), cc: c.name})
4 FOREACH (p IN nameNodes | CREATE (y {name: p.name})<-[:orig]-(x))
5 WITH x
6 MATCH (a)-[:dobj]->(b)
7 CREATE (y {name: b.name})
8 WITH a, x, y
9 CALL apoc.create.relationship(x, a.name, {}, y) YIELD rel
10 WITH y
11 MATCH q=()-[]->()
12 RETURN q

```

Instead of generating one single connected component representing the result, we will generate as many distinct connected components as subgraphs being identified as matching the patterns, while this does not occur with a simple sentence structure (Figure 3a) where we achieve the correct result as in Figure 5. We must **MATCH** elements of the graph multiple times, constantly rejoining on data previously **MATCH**-ed in earlier stages of the query (e.g.) for establishing relationships over previously grouped vertices (Lines 108 and 118 from Listing 2). This then postulates the inability of such language to automatically apply an order of visit for restructuring the loaded graph (e.g., we need to tell the query language to first group-by the vertices – Lines 2-15 – and then establish the orig relationships – Lines 18-24) while not expressing an automated way to merge each distinct transformed graph into one cohesive, connected component. This then forces the expression of a generic graph matching and rewriting mechanism to be dependent on the specific recursive structure of the data. Thus, requiring the creation of a broader query, where we need to explicitly instruct the query language on the correct way to visit the data while instructing how to reconcile each generated subgraph from each morphism within one final graph.

During the delineation of the final Cypher query succeeding in obtaining the correct rewritten graph (also Listing 2), we also highlighted the impossibility of Cypher to propagate the temporary result generated by a rewriting rule and propagate it to another rule to be applied upstream: this requires carrying out intermediate sub-queries establishing connections across patterns sharing intermediate vertices, as well as the re-computation of the same intermediate solutions, such as vertex grouping (cfr. Line 4 and Line 116). Since Cypher also does not support the explicit grouping of vertices based on a pattern as in [50], this required us to identify the vertices satisfying each specific pattern, label them appropriately in a unique way (e.g. Line 9), and then compare the result obtained (e.g. Line 20 for generating orig relationships). This limitation can be overcome by providing two innovations: first, using nested relational tables for representing morphisms, where each nest will contain the sub-pattern of interest possibly to be grouped. Second, we track any vertex substitution for entry-point vertex matches via incremental views. This substitution can be easily propagated at any level by considering the transitive closure of the substitution function (Definition 5), while the order of visit in the graph guarantees the correctness of the application of such substitution (Algorithm 8).

Listing 2, constructed for the specific matches referring to the sentence "*Matt and Tray...*", will not fully execute on a different sentence without the given dependencies, as no match is found, and therefore no rewriting can occur. Current graph query languages are meant to return a subgraph from the given patterns. In Cypher, you must abide with what is contained within the data, if the data is not there we need to remove the match from the query, which we cannot forecast in advance. This results in constant analysis of the data. For us the intention is to have graph grammar rewriting rules whereby if a match is not made, no rewriting occurs.

By leveraging such limitations of Cypher while juxtaposing the desired behaviour of the language, we derive a declarative graph query language where patterns can be expressed similarly to Figure 5. \square

Proof (for Lemma 8). This is a direct application of the SOS rules from Figure 8: any removed vertex will not be replaced by a newly inserted vertex within the matched entry-point ego-net if not explicitly updating the containment to also add the newly created object. If an entry point was removed, the only way to preserve the connectivity of the GSM objects is to exploit the replacement, through which we will explicitly state that, for any explicitly declared container within the matched pattern, we will insert the created object or any other previously matched object of choice within the container's containment relationships also containing the object. \square

Proof (for Lemma 9). This requires discussing the SOS rules from Figure 8, from which we set or update values, labels, containments, and properties associated with objects. Concerning label updates, such updates occur over variable x , which variable resolution ρ is always in the form $\langle x, \Delta(g), \mathbf{true} \rangle$: if the variable does not appear in the morphism, we expand the first two cases from Figure A1. We need also to exclude the case it was declared with a new statement from Figure 8, as we will have otherwise x in Γ^v from $\Delta(g)$. As we have the parameter **true**, this also excludes the rule NOFRES (Figure A1). We can then see that we do not create an optimally matched containment, as expected by the intended semantics. Thus, we restrict our case to FORCERES (Figure A1), on which we see that a new object is created, thus updating $\Delta(g)$, and that we know the ID of this object as it will be naturally the next incrementally number being available. Then, the label update will always occur, which will then preserve the update in $\Delta(g)$. These choices are reflected in the materialization phase by extending each database g and prioritizing the changes described in the view $\Delta(g)$. \square

Proof. Given the possibility of visiting several patterns L_1, \dots, L_n , we can express the matching of those in our proposed query language as rules $p_i = L_i \rightarrow \emptyset$ for $1 \leq i \leq n$, where both objects and containments must be explicitly referenced with a variable. Still, this formulation will not explicitly remove any object or containment not being visited. Enabling this requires the extension of the former query with two additional rules, one for removing all the vertices not visited in the different pattern (om), and the other for explicitly removing unmatched containments (cm). Given the variables

A^o, \dots, W^o referring to matched objects and A^c, \dots, W^c to matched containments in L_1, \dots, L_n , we can then express om and cm as the following rules being defined immediately after p_n :

```
om = (Z)
  where (Z unmatched $p_1$ .  $A^n$   $\wedge$  ( . . . Z unmatched $p_n$ .  $W^n$ )
         $\hookrightarrow$  del Z (Z);

cm = (X)--[Y:]->(Z)
  where (Y unmatched $p_1$ .  $A^c$   $\wedge$  ( . . . Y unmatched $p_n$ .  $W^c$ )
         $\hookrightarrow$  del Z (Z);
```

As we rewrite the same matching object, no replacement will occur and given that the matching (Z) and (X)--[Y:]->(Z) will return all the objects and containments across the databases, we have to further where-test those so to delete only the ones being not matched in L_1, \dots, L_n .

□

Proof (for Corollary 1). This follows from our previous proof, for which we clearly showed that our proposed language can match and rewrite graphs declaratively while considering optional rewrites. Cypher has some limitations in this regard, as it forces the user to specify the order in which the matching and rewriting rules should be applied. Furthermore, our language can return the matched morphisms similarly to SPARQL while allowing the generation of multiple morphism tables rather than just one, as well as selecting just the objects and containments being matched while removing the remaining ones similarly to Cypher. Therefore, the proposed language over GSM generalizes over current graph query languages over a novel generalised semistructured model enabling this. □

Appendix B.4. Time Complexity

Proof (for Lemma 11). Regarding Algorithm 5, as we defined a graph connecting each rule appearing in the query, which will be then represented as a vertex, in the worst case scenario, we will have a fully connected graph with $E = V \times V$. Thus, the cost of creating this graph is quadratic, as $O(|V| + |V|^2)$ is in $O(|V|^2)$. Given that the approximate topological sort uses a DFA visit of the resulting graph and that the layering is linear over the size of the vertices, and given that $|V| = |gg|$ by construction, we therefore obtain an overall global quadratic time complexity over the size of the query when expressed via the number of available rules. □

Proof (for Lemma 12). Suppose that each rule has at most c containment relationships, which will be provided by disjunction reference to all the k containment labels recorded in the physical model. Thus, the caching phase will take at most $ck|gg| + k$ time, as we might still consider all of these labels if we have containments for which the containment relationship is not specified.

Thus, the caching mechanism will guarantee to solely access each PhiTable^k once. By estimating an average branching factor β across the loaded GSM in the physical model and by assuming that, in the worst case scenario, all the objects will contain containments for all the k labels, then the cost of caching the tables to make them ready to the morphism representation takes $k|db|\bar{O}\beta$ time, where \bar{O} is the average number of objects stored across GSM databases.

Now, we consider the cost associated with a table of size $|db|\bar{O}\beta$. we can freely assume that rewriting operations are in $O(1)$, as in our implementation, morphisms are striped by schema information and are merely represented as tuples while associating the schema to the sole table. Similarly, the projection costs are linear over the size of the table, while the nesting operation can be performed in linear time while reducing the size of the table to $|db|\bar{O}$ due to the ego-net assumptions enforced within the structure of the matching pattern. Overall, this comes at $O(|db|\bar{O}\beta)$ time.

In the worst case scenario, the association of containment to a table will take cost $k|db|\bar{O}\beta$, thus totalling to an overall cost in $O(ck|db|\bar{O}\beta)$, which also dominates the time complexity of the other phases. □

Proof (for Corollary 2). This proof refers to the time complexity of Algorithm 7, that can be seen as a corollary of the previous lemma, for which we already derived the the estimated table size $m = |db|\overline{O}\beta$ for each required table composing the final morphism. Let us denote r (and o) as the maximum number of required (and optional) matches appearing across all L_i from gg -s. As the nested relational algebra can be always expressed in terms of “flat” relational algebra, we can upper-bound the time complexity of Algorithm 3 by $O(m^2)$. This gives us a worst-case scenario time complexity of $O(m^{r+o})$ for computing P_i for each rule $p_i \in gg$, which is a linear time complexity over the size of the worst-case scenario yielded table.

The nesting operator $\nu_{B \rightarrow A}(t)$ from Equation 2 being optionally used to nest morphisms if entry-point vertices are grouped by a direct ancestor can be easily implemented with a linear time complexity over the size of the table that we want to nest: this boils down to computing the equivalence class of \doteq_R over the fields $R = \text{dom}(\mathcal{S}(t)) \setminus \{A, B\}$ and holding such information as a map from the values $\tilde{t}(R)$ for each $\tilde{t} \in t$ to the collection of rows $\tilde{t}|_B$ for which $\tilde{t}|_R = \tilde{t}|_R$. Thus, Line 16 comes with a linear cost over the size of the table P_i .

Given that the time complexity of computing symmetric closure of a relationship is trivially linear while the time complexity of computing the transitive closure for a relationship is upper-bounded by the Floyd-Warshall algorithm with a cubic time, this leads to an worst-case time complexity of $O(|db|\overline{O}^3)$ time for computing each \mathbb{R}_i (Lines 17 and 18).

We can freely assume that the insertion cost of each morphism within the $MT[\cdot, \cdot]$ table comes at a linear cost, while the sorting of each $MT[L_i, g_j]$ comes with a cost of $O((r+o)m^{r+o} \log(m))$ with $m = |db|\overline{O}\beta$. This phase clearly dominates over all the previous ones, and thus we can freely assume that the time complexity of computing each morphism is in $O((r+o)m^{r+o} \log(m))$. This leads to an overall time complexity of $O(|gg|(r+o)m^{r+o} \log(m))$ for generating all the morphisms, which can be still upper-bounded by a polynomial time complexity. \square

Proof (for Lemma 13). In section 4.3, we observed that all the functions that were there introduced can be computed in $O(1)$ time via the GSM view update; these operations also occur with the definition of \rightarrow_ν at the basis of the graph rewriting operations outlined in Section 6.5 and Figure 8: the worst case scenario for the evaluation of such expressions refers to the evaluation of variables associated to nested relationships, thus referring to at most β object per morphism Γ . Given that each rewriting operation considers one single morphism at a time and that, within the worst case scenario, we consider the cross-product of the objects associated to both variables, the computation of each operation shall take at most $O(\beta^2)$ time.

Given $|gg|m^{r+o}$ the number of possible nested morphisms as determined from the previous corollary, this overall leads to an overall time complexity of $O(|gg|m^{r+o}\beta^2)$ for overall computing Algorithm 6.5. \square

Proof (for Corollary 3). This is a corollary for all the previous lemmas, as the composition of polynomial-time algorithms leads to an overall algorithm of polynomial time complexity. \square

References

1. Schmitt, I. QQL: A DB&IR Query Language. *VLDB J.* **2008**, 17, 39–56. doi:10.1007/S00778-007-0070-1.
2. Chamberlin, D.D.; Boyce, R.F. SEQUEL: A Structured English Query Language. Proceedings of 1974 ACM-SIGMOD Workshop on Data Description, Access and Control, Ann Arbor, Michigan, USA, May 1-3, 1974, 2 Volumes; Altshuler, G.; Rustin, R.; Plagman, B.D., Eds. ACM, 1974, pp. 249–264. doi:10.1145/800296.811515.
3. Rodriguez, M.A. The Gremlin graph traversal machine and language (invited talk). Proceedings of the 15th Symposium on Database Programming Languages; Association for Computing Machinery: New York, NY, USA, 2015; DBPL 2015, p. 1–10. doi:10.1145/2815072.2815073.
4. Robinson, I.; Webber, J.; Eifrem, E. *Graph Databases*; O'Reilly Media, Inc., 2013.
5. Angles, R.; Gutierrez, C. The Expressive Power of SPARQL; 2008; pp. 114–129.

6. Bergami, G.; Petermann, A.; Montesi, D. THoSP: an algorithm for nesting property graphs. Proceedings of the 1st ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA); Association for Computing Machinery: New York, NY, USA, 2018; GRADES-NDA '18. doi:10.1145/3210259.3210267.
7. Bergami, G. On Efficiently Equi-Joining Graphs. Proceedings of the 25th International Database Engineering & Applications Symposium; Association for Computing Machinery: New York, NY, USA, 2021; IDEAS '21, p. 222–231. doi:10.1145/3472163.3472269.
8. Ehrig, H.; Habel, A.; Kreowski, H.J. Introduction to graph grammars with applications to semantic networks. *Computers & Mathematics with Applications* **1992**, *23*, 557–572. doi:https://doi.org/10.1016/0898-1221(92)90124-Z.
9. Bergami, G. A new Nested Graph Model for Data Integration. PhD thesis, alma, 2018.
10. Das, S.; Srinivasan, J.; Perry, M.; Chong, E.I.; Banerjee, J. A Tale of Two Graphs: Property Graphs as RDF in Oracle. Proceedings of the 17th International Conference on Extending Database Technology, EDBT 2014, Athens, Greece, March 24–28, 2014., 2014, pp. 762–773. doi:10.5441/002/edbt.2014.82.
11. Turi, D.; Plotkin, G.D. Towards a Mathematical Operational Semantics. Proceedings, 12th Annual IEEE Symposium on Logic in Computer Science, Warsaw, Poland, June 29 - July 2, 1997. IEEE Computer Society, 1997, pp. 280–291. doi:10.1109/LICS.1997.614955.
12. Bergami, G.; Appleby, S.; Morgan, G. Quickening Data-Aware Conformance Checking through Temporal Algebras. *Information* **2023**, *14*. doi:10.3390/info14030173.
13. Kahn, A.B. Topological sorting of large networks. *Commun. ACM* **1962**, *5*, 558–562. doi:10.1145/368996.369025.
14. Sugiyama, K.; Tagawa, S.; Toda, M. Methods for Visual Understanding of Hierarchical System Structures. *IEEE Trans. Syst. Man Cybern.* **1981**, *11*, 109–125. doi:10.1109/TSMC.1981.4308636.
15. Hölsch, J.; Grossniklaus, M. An Algebra and Equivalences to Transform Graph Patterns in Neo4j. Proceedings of the Workshops of the EDBT/ICDT 2016 Joint Conference, EDBT/ICDT Workshops 2016, Bordeaux, France, March 15, 2016.; Palpanas, T.; Stefanidis, K., Eds., 2016, Vol. 1558, *CEUR Workshop Proceedings*.
16. Gutierrez, C.; Hurtado, C.A.; Mendelzon, A.O.; Pérez, J. Foundations of Semantic Web databases. *Journal of Computer and System Sciences* **2011**, *77*, 520 – 541. Database Theory.
17. Fionda, V.; Pirrò, G.; Gutierrez, C. NautiLOD: A Formal Language for the Web of Data Graph. *TWEB* **2015**, *9*, 5:1–5:43.
18. Hartig, O.; Pérez, J., The Semantic Web - ISWC 2015: 14th International Semantic Web Conference, Bethlehem, PA, USA, October 11–15, 2015, Proceedings, Part I; Springer International Publishing: Cham, 2015; chapter LDQL: A Query Language for the Web of Linked Data, pp. 73–91. doi:10.1007/978-3-319-25007-6_5.
19. Carroll, J.J.; Dickinson, I.; Dollin, C.; Reynolds, D.; Seaborne, A.; Wilkinson, K. Jena: Implementing the Semantic Web Recommendations. Proceedings of the 13th International World Wide Web Conference on Alternate Track Papers & Posters; ACM: New York, NY, USA, 2004; WWW Alt. '04, pp. 74–83. doi:10.1145/1013367.1013381.
20. Sirin, E.; Parsia, B.; Grau, B.C.; Kalyanpur, A.; Katz, Y. Pellet: A practical OWL-DL reasoner. *Web Semantics: Science, Services and Agents on the World Wide Web* **2007**, *5*, 51 – 53. Software Engineering and the Semantic Web.
21. Angles, R.; Arenas, M.; Barceló, P.; Hogan, A.; Reutter, J.L.; Vrgoc, D. Foundations of Modern Query Languages for Graph Databases. *ACM Comput. Surv.* **2017**, *50*, 68:1–68:40. doi:10.1145/3104031.
22. Fan, W.; Li, J.; Ma, S.; Tang, N.; Wu, Y. Adding regular expressions to graph reachability and pattern queries. *Frontiers of Computer Science* **2012**, *6*, 313–338.
23. Barceló, P.; Fontaine, G.; Lin, A.W., Expressive Path Queries on Graphs with Data. In *Logic for Programming, Artificial Intelligence, and Reasoning: 19th International Conference, LPAR-19, Stellenbosch, South Africa, December 14–19, 2013. Proceedings*; McMillan, K.; Middeldorp, A.; Voronkov, A., Eds.; Springer Berlin Heidelberg: Berlin, Heidelberg, 2013; pp. 71–85. doi:10.1007/978-3-642-45221-5_5.
24. Junghanns, M.; Kießling, M.; Averbuch, A.; Petermann, A.; Rahm, E. Cypher-based Graph Pattern Matching in Gradoop. Proceedings of the Fifth International Workshop on Graph Data-management Experiences & Systems, GRADES@SIGMOD/PODS 2017, Chicago, IL, USA, May 14 - 19, 2017, 2017, pp. 3:1–3:8. doi:10.1145/3078447.3078450.

25. Ghrab, A.; Romero, O.; Skhiri, S.; Vaisman, A.A.; Zimányi, E. GRAD: On Graph Database Modeling. *CoRR* **2016**, *abs/1602.00503*.
26. Ghrab, A.; Romero, O.; Skhiri, S.; Vaisman, A.; Zimányi, E., Advances in Databases and Information Systems: 19th East European Conference, ADBIS 2015, Poitiers, France, September 8-11, 2015, Proceedings; Springer International Publishing: Cham, 2015; chapter A Framework for Building OLAP Cubes on Graphs, pp. 92–105. doi:10.1007/978-3-319-23135-8_7.
27. Junghanns, M.; Petermann, A.; Teichmann, N.; Gomez, K.; Rahm, E. Analyzing Extended Property Graphs with Apache Flink. *SIGMOD workshop on Network Data Analytics (NDA)* **2016**.
28. Wolter, U.; Truong, T.T. Graph Algebras and Derived Graph Operations. *Logics* **2023**, *1*, 182–239. doi:10.3390/logics1040010.
29. Rozenberg, G., Ed. *Handbook of Graph Grammars and Computing by Graph Transformation: Volume I. Foundations*; WSP, 1997.
30. Pérez, J.; Arenas, M.; Gutierrez, C. Semantics and Complexity of SPARQL **2009**. *34*, 16:1–16:45.
31. Szárnyas, G. Incremental View Maintenance for Property Graph Queries. *Proc. of SIGMOD*. ACM, 2018, p. 1843–1845.
32. Consens, M.P.; Mendelzon, A.O. GraphLog: A Visual Formalism for Real Life Recursion. *Proc. of PODS*. ACM, 1990, pp. 404–416.
33. Hölsch, J.; Grossniklaus, M. An Algebra and Equivalences to Transform Graph Patterns in Neo4j. *Fifth International Workshop on Querying Graph Structured Data* **2016**.
34. Bergami, G.; Zegadło, W. Towards a Generalised Semistructured Data Model and Query Language. *SIGWEB Newsl.* **2023**, 2023. doi:10.1145/3609429.3609433.
35. Shmedding, F. Incremental SPARQL Evaluation for Query Answering on Linked Data. *Second International Workshop on Consuming Linked Data, 2011, COLD2011*.
36. Huang, J.; Abadi, D.J.; Ren, K. Scalable SPARQL Querying of Large RDF Graphs. *PVLDB* **2011**, *4*, 1123–1134.
37. Atre, M. Left Bit Right: For SPARQL Join Queries with OPTIONAL Patterns (Left-outer-joins). *SIGMOD Conference*. ACM, 2015, pp. 1793–1808.
38. Codd, E.F. A relational model of data for large shared data banks. *Commun. ACM* **1970**, *13*, 377–387. doi:10.1145/362384.362685.
39. Liu, H.C.; Ramamohanarao, K. Algebraic equivalences among nested relational expressions. *Proceedings of the Third International Conference on Information and Knowledge Management; Association for Computing Machinery: New York, NY, USA, 1994; CIKM '94*, p. 234–243. doi:10.1145/191246.191287.
40. Levene, M.; Loizou, G. The nested universal relation data model. *Journal of Computer and System Sciences* **1994**, *49*, 683–717. 30th IEEE Conference on Foundations of Computer Science, doi:https://doi.org/10.1016/S0022-0000(05)80076-5.
41. Colby, L.S. A recursive algebra and query optimization for nested relations. *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data; Association for Computing Machinery: New York, NY, USA, 1989; SIGMOD '89*, p. 273–283. doi:10.1145/67544.66952.
42. Leser, U.; Naumann, F. *Informationsintegration: Architekturen und Methoden zur Integration verteilter und heterogener Datenquellen*; dpunkt, 2006.
43. Elmasri, R.A.; Navathe, S.B. *Fundamentals of Database Systems*, 7th ed.; Pearson, 2016.
44. Atzeni, P.; Ceri, S.; Paraboschi, S.; Torlone, R. *Database Systems - Concepts, Languages and Architectures*; McGraw-Hill Book Company, 1999.
45. Colby, L.S. A recursive algebra and query optimization for nested relations. *SIGMOD Rec.* **1989**, *18*, 273–283. doi:10.1145/66926.66952.
46. den Bussche, J.V. Simulation of the nested relational algebra by the flat relational algebra, with an application to the complexity of evaluating powerset algebra expressions. *Theor. Comput. Sci.* **2001**, *254*, 363–377. doi:10.1016/S0304-3975(99)00301-1.
47. Idreos, S.; Groffen, F.; Nes, N.; Manegold, S.; Mullender, K.S.; Kersten, M.L. MonetDB: Two Decades of Research in Column-oriented Database Architectures. *IEEE Data Eng. Bull.* **2012**, *35*, 40–45.
48. Green, T.J.; Karvounarakis, G.; Tannen, V. Provenance semirings. *Proceedings of the Twenty-Sixth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems; Association for Computing Machinery: New York, NY, USA, 2007; PODS '07*, p. 31–40. doi:10.1145/1265530.1265535.

49. Chapman, A.; Missier, P.; Simonelli, G.; Torlone, R. Capturing and querying fine-grained provenance of preprocessing pipelines in data science. *VLDB* **2020**, *14*.
50. Junghanns, M.; Petermann, A.; Rahm, E. Distributed Grouping of Property Graphs with GRADOOP. In *Datenbanksysteme für Business, Technologie und Web (BTW 2017)*; Gesellschaft für Informatik, Bonn, 2017; pp. 103–122.
51. Bellatreche, L.; Kechar, M.; Bahloul, S.N. Bringing Common Subexpression Problem from the Dark to Light: Towards Large-Scale Workload Optimizations. *IDEAS*. ACM, 2021.
52. Aho, A.V.; Lam, M.S.; Sethi, R.; Ullman, J.D. *Compilers: Principles, Techniques, and Tools (2nd Edition)*; Addison-Wesley Longman Publishing Co., Inc.: USA, 2006.
53. Ulrich, H.; Kern, J.; Tas, D.; Kock-Schoppenhauer, A.; Ückert, F.; Ingenerf, J.; Lablans, M. QL4MDR: a GraphQL query language for ISO 11179-based metadata repositories. *BMC Medical Informatics Decis. Mak.* **2019**, *19*, 45:1–45:7. doi:10.1186/S12911-019-0794-Z.
54. Parr, T. *The Definitive ANTLR 4 Reference*, 2nd ed.; Pragmatic Bookshelf, 2013.
55. Tarjan, R.E. Edge-Disjoint Spanning Trees and Depth-First Search. *Acta Informatica* **1976**, *6*, 171–185. doi:10.1007/BF00268499.
56. de Marneffe, M.C.; Manning, C.D.; Nivre, J.; Zeman, D. Universal Dependencies. *Computational Linguistics* **2021**, *47*, 255–308, [https://direct.mit.edu/coli/article-pdf/47/2/255/1938138/coli_a_00402.pdf]. doi:10.1162/coli_a_00402.
57. Martelli, A.; Montanari, U. Unification in linear time and space: A structured presentation. Technical Report Vol. IEI-B76-16, Consiglio Nazionale delle Ricerche, Pisa.
58. Rozenberg, G., Ed. *Handbook of Graph Grammars and Computing by Graph Transformation: Volume II.*; WSP, 1999.
59. Talmor, A.; Herzig, J.; Lourie, N.; Berant, J. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers); Burstein, J.; Doran, C.; Solorio, T., Eds. Association for Computational Linguistics, 2019, pp. 4149–4158. doi:10.18653/V1/N19-1421.
60. de Marneffe, M.; MacCartney, B.; Manning, C.D. Generating Typed Dependency Parses from Phrase Structure Parses. Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006; Calzolari, N.; Choukri, K.; Gangemi, A.; Maegaard, B.; Mariani, J.; Odijk, J.; Tapias, D., Eds. European Language Resources Association (ELRA), 2006, pp. 449–454.
61. Chen, D.; Manning, C.D. A Fast and Accurate Dependency Parser using Neural Networks. Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL; Moschitti, A.; Pang, B.; Daelemans, W., Eds. ACL, 2014, pp. 740–750. doi:10.3115/V1/D14-1082.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.