



HAL
open science

Condensed Representation for Snapshot-Based RDF Graphs

Jey Puget Gil, Emmanuel Coquery, John Samuel, Gilles Gesquière

► **To cite this version:**

Jey Puget Gil, Emmanuel Coquery, John Samuel, Gilles Gesquière. Condensed Representation for Snapshot-Based RDF Graphs. 2025. ⟨hal-05175454v3⟩

HAL Id: hal-05175454

<https://hal.science/hal-05175454v3>

Preprint submitted on 20 Mar 2026

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

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Condensed Representation for Snapshot-Based RDF Graphs

Jey Puget Gil

jey.puget-gil@liris.cnrs.fr

Universite Claude Bernard Lyon 1, CNRS, INSA Lyon,
LIRIS, UMR 5205
Villeurbanne, FRANCE

John Samuel

john.samuel@cpe.fr

CPE Lyon, CNRS, INSA Lyon, Universite Claude Bernard
Lyon 1, Université Lumière Lyon 2, Ecole Centrale de
Lyon, LIRIS, UMR 5205
Villeurbanne, FRANCE

Emmanuel Coquery

emmanuel.coquery@univ-lyon1.fr

Universite Claude Bernard Lyon 1, CNRS, INSA Lyon,
LIRIS, UMR 5205
Villeurbanne, FRANCE

Gilles Gesquière

gilles.gesquiere@univ-lyon2.fr

Université Lumière Lyon 2, CNRS, Université Claude
Bernard Lyon 1, INSA Lyon, Ecole Centrale de Lyon, LIRIS,
UMR 5205
Villeurbanne, FRANCE

ABSTRACT

Evolving phenomena, often complex, can be represented using knowledge graphs, which have the capability to model heterogeneous data from multiple sources. Nowadays, a considerable amount of sources delivering periodic updates to knowledge graphs in various domains is openly available. The evolution of data is of interest to knowledge graph management systems, and therefore it is crucial to organize these constantly evolving data to make them easily accessible and exploitable for analysis. In this article, we will present and formalize the condensed representation of these evolving graphs and propose a new solution called QuaQue that allows querying across multiple versions of graphs and we also present the results of our benchmark comparing our solution against existing approaches.

Paper Artifact Availability:

The source code has been made available at <https://github.com/VCityTeam/ConVer-G>. The detailed results are available at <https://zenodo.org/records/17198832>.

1 INTRODUCTION AND MOTIVATIONS

1.1 Data evolution and versioning

Time is a fundamental aspect of data, and the need for temporal data management[15] has been recognized for a long time. Temporal snapshots of data help to understand the evolution of data over time and are essential for various applications, such as historical data analysis, trend analysis, and data quality assessment.

While our requirement emerged from urban data management, where temporal data management is highly significant. The challenge of managing evolving data is present in many domains, including but not limited to urban studies, biomedical research, and software engineering. In these fields, data is frequently updated, extended, or corrected, and it is crucial to track and analyze these changes effectively. For example, in biomedical research, tracking changes in ontologies or patient records over time is important; in software engineering, managing the evolution of code or documentation as RDF graphs is increasingly common. The growing availability of open and periodically released data enables researchers to study complex phenomena with incremental changes[6, 10, 18, 23].

A comprehensive understanding of dynamic elements—such as population growth, infrastructure development, or evolving scientific knowledge—is crucial for effective management and analysis in these domains.

To address these challenges, we propose the use of versioning to refer to data at different points or periods of time. The versioning of scientific data enables the exploitation and reproducibility of experiments and analyses[16]. Versioning captures the state of the data at a given point or period in time, facilitating reproducibility, integrity, and the ability to restore previous versions in case of errors. It also supports experimentation with different analytical approaches by allowing easy switching between data versions.

Knowledge graphs, and in particular RDF graphs, are increasingly used to represent evolving data in a variety of domains[32]. By incorporating versioning into these knowledge graphs, it becomes possible to capture concurrent viewpoints and the evolution of entities over time[23]. This enables the representation of different states of the data, providing a comprehensive view of its evolution. Our approach leverages knowledge graphs from heterogeneous sources (concurrent viewpoints) at different periods of time (hereafter, referred to as versions).

1.2 Improving the comprehension of data evolution

In many domains, data is continuously evolving as new information becomes available, corrections are made, or entities change over time. This constant evolution creates challenges for data management systems that need to track and analyze these changes effectively. Additionally, the growing volume of data collected from multiple sources and at different temporal granularities adds complexity to the task of managing and understanding data evolution.

To achieve effective management of evolving data and ensure reproducibility of experiments and analyses, we identified the following requirements for our system:

- **R1: Managing and querying multiple graphs and their different versions:** Our system should allow users to query multiple knowledge graphs from one or more sources and access their different versions simultaneously. This capability enables comprehensive analysis and comparison of data evolution over time.

- **R2: Linking versioned data with metadata:** It is crucial to establish a connection between versioned data and their associated metadata. This linking provides valuable context and additional information about the data, enhancing its usability and interpretability. For example, when tracking changes to a biomedical record or a software artifact, metadata could include the source, method, and date of observation, allowing users to better understand and validate the data’s provenance and reliability.
- **R3: Isolating and querying individual graph versions:** Our system must provide a mechanism to isolate and query specific versions of a knowledge graph independently. This capability allows users to focus their analysis on particular versions without interference from other versions, enabling precise examination of the data at specific snapshots.

By addressing these requirements, our system aims to empower researchers and practitioners in diverse fields to gain deeper insights into data dynamics, facilitate informed decision-making, and ensure the reproducibility and reliability of scientific experiments and analyses. Although our focus in this article is on the formalization of versioned RDF graphs, the principles and mechanisms we present are generic and can be applied to any domain where RDF data evolves over time.

In this article, to address the above requirements, we present the formalization of flat and condensed representations and querying of evolving knowledge graphs. We also present our proposed solution that uses and compares both representations using benchmark data. The state of the art in entity versioning is detailed in Section 2. Section 3 presents the necessary preliminaries and notations. Section 4 presents the contributions of our work: datasets using flat and condensed models, algebra for handling query operators and the resulting data in flat and condensed representations. In Section 5, we briefly present our implementation—**QuaQue** (part of the ConVer-G system)[12]—and compare the results obtained for flat and condensed representations. Section 6 concludes the article with a discussion of the limitations and possible enhancements to our proposed work.

2 STATE OF THE ART - VERSIONING

Resources evolve over time and hence the statements associated to describe them. Version control systems [39], for example, handle this evolution of data (or code) by associating a tag and a timestamp to a snapshot of the data at a given instant of time. This supplementary information associated to a snapshot, referred to as a version can be used to not only represent the state of the data at any given point of time, but also to understand the evolution of data over a period of time. As mentioned by Bayoudhi et al. [3], versioning is a critical aspect of management in a project, allowing developers to track changes to the elements, manage different versions, and collaborate with other team members. For the sake of familiarity, we will use the word *versioning* to this process of adding information like timestamp, tag to snapshots of one or more resources and their associated statements at any given moment of time. In this section, we will take a look at different approaches of versioning. Table 1 summarizes these results.

Table 1: Comparison of code and data versioning

Characteristics	Code versioning	Data versioning
Object	Files	Datasets (structured, unstructured, etc.)
Granularity	Lines of code, files	Data, entire datasets
Primary goals	History, collaboration, branching	Reproducibility, provenance, experimentation,
Tools	Git, Mercurial, SVN	DVC, Pachyderm, Machine Learning platforms

2.1 Code versioning

Version control systems [39] such as Git, Mercurial, Subversion, CVS, and Bazaar are commonly used for managing the evolution of source code. However, these solutions are not specifically designed for data versioning, especially for structured data [24] and have certain limitations for efficient querying [4, 22, 38]. Apart from some exceptions, like Git LFS, they don’t provide adequate support for large files and versioned data queries. Therefore, they may not be the ideal choice for managing data versioning in a project.

2.2 Dataset versioning

2.2.1 Independent copies (IC). This approach allows querying data as it existed in a specific version. *Example: Which buildings were present in the city at a particular version?*

A common approach to dataset versioning is to maintain independent copies (snapshots) of the data for each version. Tools such as Qri [29], SemVersion [33], GeoGig [28], and UrbanCo2Fab [22] follow this strategy, often leveraging Git-based mechanisms. While this method simplifies version management and enables straightforward retrieval of any version by checking out its corresponding snapshot, it has notable limitations. In particular, querying across multiple versions simultaneously is inefficient, as each query requires materializing the full state of the relevant versions. This restricts the ability to perform cross-version analysis or answer queries that span several versions. Similarly, machine learning data management tools such as DVC, DagsHub, and Weights and Biases provide robust tracking and collaboration features for structured data, but do not support efficient querying across multiple snapshots, making them unsuitable for use cases requiring multi-version queries.

2.2.2 Timestamp-based versioning (TB). This representation allows querying data as it existed at a specific point in time. *Example: Which buildings were present in the city at a particular version?*

A common approach to data versioning is to associate each record with a timestamp or a validity interval, as seen in temporal or bitemporal tables [17]. Tools such as ConVer-G [12], Drydra [1], RDF-TX [11], v-RDFCSA [5], and x-RDF-3X [19] adopt this strategy. This method enables querying data as it existed at a specific point in time or within a given interval. While effective for linear histories, timestamp-based versioning is limited when handling concurrent changes, such as branches or forks, which are typical in collaborative environments and version control systems. As a

result, it may not fully support scenarios requiring representation of multiple, diverging timelines.

2.2.3 Change-based versioning (CB). This representation enables retrieval of the differences between two versions. *Example: Which buildings were added or removed between two versions?*

Change-based versioning systems, such as Delta Lake and Apache Hudi store changes made to data as changes, allowing for efficient retrieval of any version or the differences between versions. In the context of RDF, tools such as R43ples [13], R&WBase [31], and Stardog follow this paradigm. A prominent system for versioning RDF data is OSTRICH [26].

However, OSTRICH has notable limitations for our needs. It only supports linear version histories, lacking support for branching or concurrent versions, which are essential for modeling multiple viewpoints in urban evolution. Additionally, its aggregated change-sets grow with each new version, resulting in slow ingestion for archives with long histories.

2.2.4 Fragment-based versioning (FB). This approach enables random access to different versions by managing indexes of file fragments for each version. This archiving policy stores snapshots of each changed fragment of an archive. These fragments can be defined at any level of granularity, such as individual graphs within a dataset. An index is maintained to reference which specific fragments belong to any given version of the dataset. This approach avoids the need to store full, independent copies for each version, and in contrast to change-based systems, it does not require the reconstruction of versions by applying a series of change operations. The Quit Store implements this policy by treating modified files within its Git repository as the fragments to be versioned.

2.3 Querying versioned data

This subsection reviews how different systems implement versioning, the types of queries they support, and the consequences for managing knowledge graphs. The following analysis compares existing solutions and highlights the need for advanced versioning models that address complex, multi-version scenarios. We adopt the query types defined by Ruben Taelman in the OSTRICH system [26], which offer a comprehensive framework for querying versioned data.

2.3.1 Version materialization (VM). We consider version materialization as a means to create and manage distinct versions of knowledge graphs. This approach allows for the reconstruction of any specific version. This is particularly linked with the concept of (IC).

2.3.2 Delta materialization (DM). Delta materialization is a technique used in databases and data management systems to track and retrieve changes in query results between two versions of data.

Instead of recomputing the entire result of a query for both versions, DM focuses on computing the delta—the set of changes (additions and deletions) in the query result that occurred between two versions. This is especially useful when data changes incrementally and you want to know only what has changed, not the full result.

2.3.3 Version query (VQ). A version query annotates each result of query Q with the specific versions of the RDF archive A in which that result appears. This allows users to see not only the query results, but also the history of when each result was valid across different versions of the data.

2.3.4 Cross-version query (CV). A cross-version query typically refers to a database or information retrieval operation that compares or joins the results of two queries executed on different versions of a dataset. This allows you to analyze changes, differences, or relationships between data across versions.

2.3.5 Cross-delta (CD). A cross-delta query combines the results of multiple changesets of a dataset. It is similar to a delta materialization (DM) query but supports more complex comparisons than just identifying additions and deletions. Instead of only showing what has changed, a CD query can analyze how specific values have evolved between versions. For instance, it can be used to compare the value of a particular attribute for a given entity across two versions, enabling a more detailed analysis of data evolution.

OSTRICH supports all query types, but its native capabilities are limited to single triple pattern queries. To execute full SPARQL queries, an additional query engine such as GLENDA [20] is required. Furthermore, OSTRICH relies on a linear sequence of integer revision numbers for versioning, whereas our use case demands a more flexible model capable of representing arbitrary relationships between versions, such as those found in a version graph.

This work builds on concepts from version control systems [2, 25, 26, 29, 39] and extends them to knowledge graphs. We introduce a "condensed representation" for querying with SPARQL, where the graph is modeled as a set of quads annotated with the versions in which they are valid. Our main contributions include the formalization of this condensed representation for evolving knowledge graphs and the development of an algebra for query operators and the resulting data in both flat and condensed forms.

Table 9 in Appendix is an extension of Olivier Paul Pelgrin's survey of versioning tools for RDF data [21] which summarizes the capabilities of different versioning tools with respect to the types of queries they support.

3 PRELIMINARIES AND NOTATIONS

In this section, we will present the key concepts and notations related to representing knowledge graphs in the form of RDF and querying them using the SPARQL query language.

3.1 RDF context representation

For the sake of simplicity, we considered the following RDF context representation based on the W3C definitions [34]:

3.1.1 Literal. A literal denotes resources of the real world.

- $l, l_1, l_2 \dots$ a literal
- $L, L_1, L_2 \dots$ a finite set of literals
- \mathcal{L} the infinite set of all literals

3.1.2 IRI. Internationalized Resource Identifier (IRI) is a Unicode string, analogous to Uniform Resource Identifiers (URI), and used to identify resources (e.g., `http://example.org/resource/Paris` is an IRI identifying the resource "Paris").

- $i, i_1, i_2 \dots$ an IRI
- $I, I_1, I_2 \dots$ a finite set of IRI
- \mathcal{I} the infinite set of all IRIs

3.1.3 Blank Node. Not all resources and the relationships can be explicitly named and are represented as blank nodes indicating the existence of something. To handle blank nodes, we use Skolemization [14], a technique that replaces blank nodes with generated IRIs. This allows us to treat blank nodes as regular resources with globally unique identifiers, enabling better integration and querying of RDF data.

3.1.4 Terms. Terms are used to refer to IRIs or literals.

- $t, t_1, t_2 \dots$ a term
- $T, T_1, T_2 \dots$ a finite set of terms
- \mathcal{T} the infinite set of all terms
- $\mathcal{T} = \{\mathcal{L}, \mathcal{I}\}$ where:
 - $\mathcal{L} \subseteq \mathcal{T}$
 - $\mathcal{I} \subseteq \mathcal{T}$
 - $\mathcal{L} \cap \mathcal{I} = \emptyset$

3.1.5 Triples. An RDF triple is a statement about a resource that consists of three parts: a subject, an object, and a predicate. The subject is an IRI that represents the resource being described. The object can be an IRI or a literal that represents the value of the property. The predicate is an IRI that represents the relationship between the subject and the object.

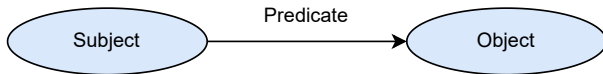


Figure 1: Example of a RDF triple

- $tr = (t_s, t_p, t_o)$ an RDF triple where:
 - $t_s \in \mathcal{I}$ is a subject
 - $t_p \in \mathcal{I}$ is a predicate
 - $t_o \in \mathcal{T}$ is an object
- $TR, TR_1, TR_2 \dots$ a finite set of triples
- \mathcal{TR} the infinite set of all triples

3.1.6 Graphs. An RDF graph is a set of RDF triples. It represents a collection of resources and the relationships between them. An RDF graph can be named or unnamed by a resource. An unnamed graph is called the default graph.

3.1.7 Quads. A quad is a statement that consists of four parts: all triple elements (the subject, the predicate, the object), and the graph name. The graph name is an IRI that represents the context in which the triple is defined.

- $q = (t_s, t_p, t_o, t_g)$ an RDF quad
- $Q, Q_1, Q_2 \dots$ a finite set of quads
- \mathcal{Q} the infinite set of all quads
- $t_g \in \mathcal{T}_g$ where $\mathcal{T}_g \subseteq \mathcal{I}$, t_g is the graph name and \mathcal{T}_g is the set of all graph names

3.1.8 Datasets. A dataset d is a collection of graphs. A dataset can contain zero or more named graphs, each of which contains a set of RDF triples. A dataset can also contain a default graph, which is a set of triples that are not associated with any specific graph name.

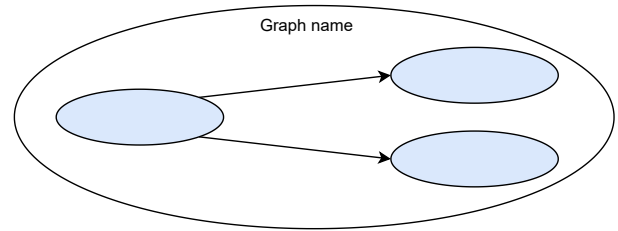


Figure 2: Example of a named graph

3.1.9 Versioning. A version is a snapshot of the RDF data at a given point in time. We define the following versioning representation:

- $v, v_1, v_2 \dots$ a version
- $V, V_1, V_2 \dots$ a finite set of versions
- \mathcal{V} the infinite set of all versions

Because the RDF data model is atemporal as explained in the "RDF and Change over Time" paragraph of the W3C RDF 1.1 Concepts and Abstract Syntax [34], we consider the versioning of RDF data as a set of versions.

3.2 SPARQL queries

SPARQL [35] is a declarative programming language used to query data stores containing a set of triples or quads.

3.2.1 Variables. A SPARQL variable is a placeholder for a value that is not known beforehand. It allows querying for information without having to specify the exact values the user is looking for. Variables are denoted by a question mark "?" or a dollar sign "\$", immediately followed by a string without any white space characters.

- $var, var_1, var_2 \dots$ a SPARQL variable
- $VAR, VAR_1, VAR_2 \dots$ a finite set of SPARQL variables
- \mathcal{VAR} the infinite set of all SPARQL variables

3.2.2 Query. A SPARQL query is a formal language expression used to retrieve and manipulate data stored in RDF format, according to the W3C definition [35].

- $qu, qu_1, qu_2 \dots$ a SPARQL query
- $QU, QU_1, QU_2 \dots$ a finite set of SPARQL queries
- \mathcal{QU} the infinite set of all SPARQL queries

SPARQL expression. A SPARQL expression refers to a part of a SPARQL query that computes or filters values within the query. We note:

- $e, e_1, e_2 \dots$ a SPARQL expression
- $E, E_1, E_2 \dots$ a finite set of SPARQL expressions
- \mathcal{E} the infinite set of all SPARQL expressions

SPARQL expression can be:

- A term t (e.g., 10) is an expression if $t \in \mathcal{T}$
- A variable var (e.g., ?x) is an expression if $var \in \mathcal{VAR}$
- A function $f(e_1, \dots, e_n)$ is an expression, if f is a function and $\forall i \in \{1, \dots, n\}, e_i$ is a SPARQL expression (e.g., $AVG(e)$ where e is a SPARQL expression)
- $e_1 \neq e_2$ is an expression if e_1 and e_2 are expressions

SPARQL expressions are used to:

- Combine values from different RDF triples to form more complex patterns (e.g., triple patterns, basic graph patterns, or subqueries)
- Filter results based on conditions (e.g., using the FILTER keyword)
- Transform data using functions like mathematical operations, string manipulation, or logical operators (e.g., CONCAT, STRLEN...)

Solution mapping. A solution mapping is a mapping between variables and RDF terms. It is used to represent the results of a SPARQL query. We note $m : \mathcal{VAR} \rightarrow \mathcal{T}$ a solution mapping that maps a set of variables to a set of RDF terms and \mathcal{M} an infinite set of solution mappings.

Solution sequence. A solution sequence is a sequence of solution mappings. We denote a solution sequence as $seq = [m_1, \dots, m_n]$ where each m_i is a solution mapping.

4 CONTRIBUTIONS

In this section, we present our contributions (Figure 3). Firstly, we present the key extensions to the RDF model discussed in Section 3 and the resulting dataset models. We discuss in detail the flat and condensed models of representation with examples. We then show how these datasets can be queried with the existing operators of the SPARQL query. Finally, we present the necessary formalisms and theorems to show how such queries can be evaluated.

4.1 RDF versioning definitions

4.1.1 Versioned IRI. In order to represent the versioning of IRI, we introduce the concept of a versioned IRI. A versioned IRI is an IRI that is associated with a version. We define the following versioned IRI representation:

- Let $VI : \mathcal{V} \times \mathcal{I} \rightarrow \mathcal{VI}$ be a bijective function.
- We note $vi, vi_1, vi_2 \dots$ the elements of \mathcal{VI} , referred to as versioned IRI
- We note $VI, VI_1, VI_2 \dots$ a finite set of versioned IRI

Let $v \in \mathcal{V}$, $t_g \in \mathcal{I}$ and $vi \in \mathcal{VI}$. We define the reverse functions of VI :

$$\begin{aligned} VI_v^{-1}(vi) &= v \Leftrightarrow \exists t_g VI(v, t_g) = vi \\ VI_{t_g}^{-1}(vi) &= t_g \Leftrightarrow \exists v VI(v, t_g) = vi \end{aligned} \quad (1)$$

4.1.2 Versioned quad. A versioned quad $vq = (t_s, t_p, t_o, t_g, v)$ is a quintuplet where the quad (t_s, t_p, t_o, t_g) is valid at the version v . We denote $VQ, VQ_1, VQ_2 \dots$ a finite set of versioned quads and \mathcal{VQ} an infinite set of versioned quads.

4.1.3 Versioned dataset. A versioned dataset is a dataset that contains a set of versioned graphs. A versioned graph is a graph that contains a set of RDF triples associated to a versioned named graph built with a version and a graph. By convention, default graphs are assigned a special name, such as `default`, to distinguish them from named graphs.

4.1.4 Dataset representation. Consider a RDF context containing data and its associated metadata. We define the dataset function such that holds all the versioned quads stored in the database: $d = 2^{\mathcal{VQ}}$. We note $V(d)$ the set of versions in the dataset.

The figure 3 provides an overview of the different levels of representation and processing within our framework. It highlights the relationships between the dataset, query, algebra, evaluation, and results, as well as the possible transformations between flat and condensed forms at each stage. The main aspects illustrated in the figure are as follows:

- **Dataset level:** the figure illustrates that the dataset can be represented in either a flat or condensed form, and that it is possible to convert between these two representations using the `condensation` and `flattening` functions.
- **Query:** A SPARQL query is represented in a flat algebra, which is the original query written by the user. The `compile` function returns the flat algebra of the query.
- **Algebra level:** it shows the representation of the algebra in both flat and condensed forms, as well as the translation process from the flat algebra (derived from the original SPARQL query) to the condensed algebra using the `translate` function.
- **Evaluation level:** it illustrates the evaluation of the algebra. Taking the associated dataset and algebra, it shows the evaluation of the algebra in both flat and condensed forms. The evaluation process is represented by the `eval` relationship, which indicates that the evaluation can be performed on either representation.
- **Results level:** it depicts the representation of results in both flat and condensed forms, and the `represent` relationship that enables conversion between these representations, again using the `condense` and `flatten` functions.

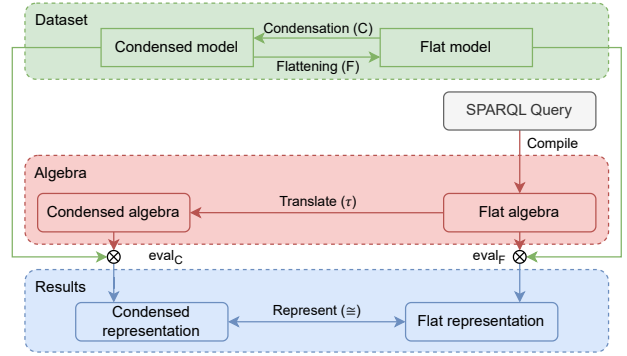


Figure 3: Dataset, algebra and results representation

4.2 Dataset models

Let us explore the different dataset models that we will use in our work. The following sections define a flat model and a condensed model for dataset versioning. To illustrate these concepts, we begin with an example that will be referenced throughout the paper for demonstration purposes.

Example 4.1. Let's assume that we have a dataset representing concurrent points of view about the height of some buildings. In Table 2, we present two versions ($v:1$ and $v:2$) with quads. The column *Named graph* consists of two named graphs corresponding to the data coming from two sources named, Gr-Lyon and IGN.

Table 2: Flat model of the versioned dataset

Subject	Predicate	Object	Named graph	Version
ex:bldg#1	height	10.5	ng:Gr-Lyon	v:1
ex:bldg#2	height	9.1		
ex:bldg#1	height	11	ng:IGN	v:2
ex:bldg#1	height	10.5	ng:IGN	
ex:bldg#1	height	10.5	ng:Gr-Lyon	
ex:bldg#3	height	15		

Table 3 represents all the versioned IRI of the dataset. In this example, each IRI starting with vi : are images (results) of the pair (version, named graph) produced by the function VI (Versioned IRI).

Table 3: Versioned IRI of the dataset

Version	Named graph	Versioned IRI
v:1	ng:Gr-Lyon	vi:1
	ng:IGN	vi:2
v:2	ng:Gr-Lyon	vi:3
	ng:IGN	vi:4

4.2.1 Flat model. We define the *flat model* as a representation of a dataset as a set of quintuplets. In this model, each triple is augmented with two additional elements: the name of the graph corresponding to the data source, and the version information.

Example 4.2. Table 2 represents the flat model of the dataset. $d_F : I \times I \times \mathcal{T} \times I \times \mathcal{V}$ designates the relation such as the quad (t_s, t_p, t_o, t_g) is valid in the version v .

The supplementary information related to the named graphs and versions is stored separately in a metadata table. Table 4 shows this metadata for the flat model in Table 2. For example, in Table 3 the first row indicate that vi:1 is a versioned IRI of the named graph Gr-Lyon and belongs to version 1. It could also be used to store supplementary information such as provenance, authorship, or other metadata relevant to the versioned dataset.

Table 4: Metadata of the flat model

Subject	Predicate	Object	N. graph	Version
vi:1	version-of	ng:Gr-Lyon	ng:Metadata	v:0
	in-version	1		
vi:2	version-of	ng:IGN		
	in-version	1		
vi:3	version-of	ng:Gr-Lyon		
	in-version	2		
vi:4	version-of	ng:IGN		
	in-version	2		

4.2.2 Condensed model. $d_C : \mathbb{N} \times \mathbb{N} \times \mathbb{N} \times \mathbb{N} \rightarrow 2^{\mathcal{V}}$ is a function that takes four identifiers that are bound to an RDF quad and returns a set of finite elements of \mathcal{V} such that $2^{\mathcal{V}}$ represents a set of finite versions where the quad is valid.

Example 4.3. The table 5 represents the condensed model of the versioned dataset (simplified by replacing the identifiers by their values). The first row, for example states that the triple present in the named graph Gr-Lyon is valid in versions v:1 and v:2. Each entry of this table can be replaced by the identifiers listed in Appendix 10.

Table 5: Condensed model of the versioned dataset

Subject	Predicate	Object	N. Graph	Versions
ex:bldg#1	height	10.5	ng:Gr-Lyon	{v:1, v:2}
ex:bldg#2	height	9.1	ng:Gr-Lyon	{v:1}
ex:bldg#1	height	11	ng:IGN	{v:1}
ex:bldg#3	height	15	ng:Gr-Lyon	{v:2}
ex:bldg#1	height	10.5	ng:IGN	{v:2}

Metadata plays a role in our approach, particularly for managing concurrent snapshots or concurrent viewpoints. By associating a constant version $v : \emptyset$ to all metadata, we ensure that version information and relationships between versions remain consistent and accessible across the entire dataset. This is especially important when dealing with branching version histories or multiple concurrent viewpoints, as the metadata helps track and organize these complex relationships. While this means metadata is not condensed and can only be represented in the flat model, this design choice provides a stable foundation for managing versioned data across different temporal and conceptual dimensions.

4.2.3 Model relationships.

Definition 4.4 (Condensation). Consider a flat dataset $d_F = I \times I \times \mathcal{T} \times I \times \mathcal{V}$. We can condense the dataset d_F into the dataset $d_C = I \times I \times \mathcal{T} \times I \rightarrow 2^{\mathcal{V}}$. The condensation function C takes a quad and returns the set of versions in which that quad exists. This function is defined as:

$$C = (t_s, t_p, t_o, t_g) \mapsto \{v \mid (t_s, t_p, t_o, t_g, v) \in d_F\}$$

The domain of C is the set of quads in the dataset d_F where the quad occurs at a set of versions such that: $\mathcal{D}(C) = \{(t_s, t_p, t_o, t_g) \mid \exists v (t_s, t_p, t_o, t_g, v) \in d_F\}$.

Definition 4.5 (Flattening). Let $d_C = I \times I \times \mathcal{T} \times I \rightarrow 2^{\mathcal{V}}$ be a dataset. We can flatten the dataset d_C into the dataset: $d_F = I \times I \times \mathcal{T} \times I \times \mathcal{V}$. The flattening function F is defined as:

$$F = \{(t_s, t_p, t_o, t_g, v) \mid v \in d_C(t_s, t_p, t_o, t_g)\}$$

The domain of F is the set of versioned quads in the dataset d_C where the quad exists in each version according to the condensed model: $\mathcal{D}(F) = \{(t_s, t_p, t_o, t_g, v) \mid \exists v \in d_C(t_s, t_p, t_o, t_g)\}$.

THEOREM 4.6 ($F \circ C(d_F) = d_F$). *Given a flat model d_F , its image by the composition of the F and the C functions is equal to the flat model.*

PROOF. We want to prove the following equivalence: $F \circ C(d_F) = d_F$. This means we need to prove:

- On the one hand, that $F \circ C(d_F) \subseteq d_F$;
- On the other hand, that $d_F \subseteq F \circ C(d_F)$.

In other words, we need to show that every element of $F \circ C(d_F)$ belongs to d_F , and conversely, every element of d_F belongs to $F \circ C(d_F)$.

Let the versioned quad $(t_s, t_p, t_o, t_g, v) \in d_F$ such that $v \in C(d_F)(t_s, t_p, t_o, t_g)$, then $(t_s, t_p, t_o, t_g, v) \in F(C(d_F))$.

Let $(t_s, t_p, t_o, t_g, v) \in F(C(d_F))$, then $v \in C(d_F)(t_s, t_p, t_o, t_g) = \{v' \mid (t_s, t_p, t_o, t_g, v') \in d_F\}$. This proves that $F \circ C(d_F) = d_F$. \square

THEOREM 4.7 ($C \circ F(d_C) = d_C$). *Given a condensed model d_C , its image by the composition of the C and the F functions is equal to the condensed model.*

PROOF. We want to prove the following equivalence: $C \circ F(d_C) = d_C$. This means we must prove:

- For every element of $C \circ F(d_C)$, there exists a corresponding element in d_C ;
- For every element of d_C , there exists a corresponding element in $C \circ F(d_C)$.

In other words, we show equality by proving both inclusions.

Let $V^C = C(F(d_C))(t_s, t_p, t_o, t_g)$, so $V^C = \{v \mid (t_s, t_p, t_o, t_g, v) \in F(d_C)\} = \{v \mid v \in d_C(t_s, t_p, t_o, t_g)\} = d_C(t_s, t_p, t_o, t_g)$. \square

4.3 SPARQL algebra

A SPARQL query algebra provides a formal framework for representing SPARQL queries [9]. It also defines the execution semantics by translating the high-level query syntax into a series of operations that an RDF query engine can process. We use \mathcal{A} to denote the SPARQL algebras set. This paper will present two types of algebra: a condensed algebra and a flat algebra that will be respectively noted as \mathcal{A}_C and \mathcal{A}_F . Based on the SPARQL algebra proposed by [9, Cyganiak, Richard, A relational algebra for SPARQL], we modified the SPARQL Algebra to take into account graph versioning to allow us to query them. An initial proposition was made by [8, Cuevas, Hogan] to represent queries over RDF Archives. We have taken this proposition as a basis to define what we call *the flat SPARQL algebra*.

4.3.1 Flat SPARQL algebra. The flat SPARQL algebra contains operators that define its expressive power. For example, the quad pattern $QP_F(?s, ?p, ?o, ?g)$ retrieves triples matching the triple pattern $TP_F(?s, ?p, ?o)$ within a versioned graph $?g$, and is considered syntactic sugar for the graph operator, as shown by [8, Cuevas, Hogan]. The JOIN operator $\bowtie_F(e_1, e_2)$ combines results from two expressions based on common variable bindings, while the FILTER operator $\sigma_F(\mathcal{E}, e)$ restricts results according to filter functions over variable bindings. The UNION operator $\cup_F(e_1, e_2)$ merges results from two expressions into a single set. The DIFF operator $\setminus_F(e_1, e_2, e_3)$ computes the difference between two sets of solutions, conditioned by an additional expression, and is essential for defining the OPTIONAL operator. The OPTIONAL operator $\bowtie_{\mathcal{A}}(e_1, e_2, e_3)$ extends the join by including all results from the first expression and, where possible, matching results from the second, returning nulls otherwise. Finally, the aggregation operator groups results based on variables and supports operations such as count, sum, average, min, and max, typically used with the GROUP BY clause in SPARQL queries.

4.3.2 Condensed algebra. The condensed algebra provides operators tailored for querying over all versions of a graph in a compact representation. The quad pattern $QP_C(?s, ?p, ?o, ?g)$ retrieves triples matching the triple pattern $TP_F(?s, ?p, ?o)$ across all versioned graphs $?g$. The condensed algebra includes optimized aggregation operators that can process data across versions. For example, aggregation operators can leverage optimizations such as count-views (as defined by Sarah Cohen [7]) to efficiently compute aggregate values over versioned data without having to materialize each version separately.

4.3.3 Variable representation. SPARQL queries are composed of operators and contain variables whose representations need to be defined. The variable representation is used to determine its behavior in the query and how it is processed in the algebra.

It is important to note that in our current formalization, we do not handle cases where there are variable clashes. Variable clashes occur when two variables with the same name but different scopes or contexts are used within the same query, leading to ambiguity and potential errors in query evaluation. By not addressing these clashes, we assume that all variables are uniquely named within their respective scopes. This simplification allows us to focus on the core aspects of our model and its algebraic semantics without the added complexity of variable disambiguation. However, in a more robust formalization, mechanisms to detect and resolve such clashes would be necessary to ensure the correctness and reliability of query results. In our formalization, each variable in a SPARQL query is assigned a *variable representation* that determines how its value is interpreted and manipulated during query evaluation. Figure 4 illustrates the hierarchy and relationships between the different representations.

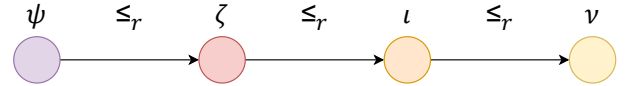


Figure 4: Hierarchy of variable representations in SPARQL algebra

A variable representation is a type that is assigned to a variable in a SPARQL query.

- We note:
 - r, r_1, r_2, \dots a variable representation
 - R, R_1, R_2, \dots a set of variable representations
 - \mathcal{R} an infinite set of variable representations
- We define four variable representations:
 - **value** ν , where variable's bound to a term
 - **id** ι , where variable's bound to a term identifier
 - **condensed** ζ , where variable's bound to a list of versioned named graph identifiers
 - **unbound graph** ψ , where variable's bound to the list of all versioned named graph identifiers

Definition 4.8 (Variable representation). Let a bounded total order of variable representations defined as follows: $(\{\nu, \iota, \zeta, \psi\}, \leq_r)$ where:

- $\{\nu, \iota, \zeta, \psi\}$ is the set of elements

- \leq_r is the total order relation with:
 - $\psi \leq_r \zeta \leq_r \iota \leq_r \nu$
 - ν is the greatest element
 - ψ is the least element

This total order serves as a guide to determine the variable representation of the output of a SPARQL operator. For any two representations r_1, r_2 , we write $r_1 <_r r_2$ to denote their minimum according to \leq_r .

Definition 4.9 (Environment). Let the environment $\Gamma : \mathcal{VAR} \rightarrow \mathcal{R}$ be a partial function that maps a variable to a variable representation. It is represented as a set of pairs (var, r) where var is a variable and r is a variable representation. We note $\mathcal{D}(\Gamma)$ the domain of the function Γ .

Definition 4.10 (Domain of a set of environments). Let Δ be a set of environments. We note:

- $\Delta, \Delta_1, \Delta_2, \dots$ a finite set of environments
- $\mathcal{D}(\Delta) = \{\mathcal{D}(\Gamma) \mid \Gamma \in \Delta\}$ the set of domains of all environments in Δ

Definition 4.11 (Lower environment of a variable). $\sqcap_\Delta : 2^\Delta \rightarrow \mathcal{VAR} \rightarrow \Delta$ is a function that takes a set of environments and a variable and returns the environment where the variable is reduced to the lowest representation. It is defined as follows:

$$\begin{aligned} \mathcal{D}(\sqcap_\Delta(\Delta)) &= \cap \{\mathcal{D}(\Gamma) \mid \Gamma \in \Delta\} \\ \sqcap_\Delta(\Delta)(var) &= r \mid \forall \Gamma \in \Delta, \\ & r \leq_r \Gamma(var) \wedge \forall r', (\forall \Gamma \in \Delta, r' \leq_r \Gamma(var)) \Rightarrow r' \leq_r r \end{aligned} \quad (2)$$

Definition 4.12 (Logical provability of an environment). Logical provability, denoted by \vdash , is a relation between an environment Γ and an expression e such that $\Gamma \vdash e$ holds if and only if, under the assumptions and variable representations specified by Γ , the expression e is well-typed and can be evaluated according to the typing rules of the algebra.

LEMMA 4.13 (GREATEST ENVIRONMENT TYPING). *Let Δ a set of environments and e is an expression. If $\forall \Gamma \in \Delta, \Gamma \vdash e$, then $\geq_r(\Delta) \vdash e$.*

In our approach, SPARQL operators impose constraints on variable representations. Let's take the FILTER operator as an example. When using comparison operators (like $<$, $>$, $=$, $!=$) in a FILTER clause, we can only compare actual values - we cannot compare condensed graph identifiers (ζ) or internal identifiers (ι) directly. Therefore, when using FILTER, we must promote any variables of type ζ or ι to the ν representation by resolving them to their actual values. This promotion ensures the variables can be properly compared in filter conditions. In the case of the JOIN operator, it is necessary for the joined variables to be of the same type in order to apply the appropriate join operation. If the variables are of different types, they need to be promoted to a common type.

Definition 4.14 (Optimal representation). We note that $R^*(e)$ is the optimal environment for the expression e . It is the environment where all the variables are promoted to the lowest representation. We define the function $R^* : \mathcal{E} \rightarrow (\mathcal{VAR} \rightarrow \mathcal{R})$ such that:

$$R^*(e) = \sqcap_\Delta(\{\Gamma \mid \Gamma \vdash e\}) \quad (3)$$

Example 4.15. Let's consider the SPARQL query 1 below.

Query 1: Get all the building heights for a specific named graph

```
SELECT ?height WHERE {
  GRAPH ?vng {
    ?s building:height ?height
  }
  ?vng v:is-version-of ng:Gr-Lyon .
}
```

The optimal environment at the quad pattern operator is $\{(?s, \iota), (?height, \iota), (?vng, \zeta)\}$.

Definition 4.16 (Value translation). Let $\tau_{r \rightsquigarrow r'}(value_R(r))$ be a function that returns the set of values in the representation r' that are equivalent to the value in the representation r . The function $\tau_{r \rightsquigarrow r'}(value_R(r))$ is defined as:

$$\begin{aligned} \tau_{\psi \rightsquigarrow \zeta}(\bullet) &= \{(n, V) \mid \text{itov}(n) \in \mathcal{T}_g \wedge V \subseteq V(d)\} \\ \tau_{\zeta \rightsquigarrow \iota}((n, V)) &= \{n' \mid \exists v \in V : VI(v, \text{itov}(n)) = \text{itov}(n')\} \quad (4) \\ \tau_{\iota \rightsquigarrow \nu}(n) &= \{\text{itov}(n)\} \end{aligned}$$

Definition 4.17 (Type downgrade of a variable's mapping). Let $\downarrow_{r \rightarrow r'}^{var}(m)$ be a function that returns the condensed mappings where the variable var is reduced to the representation r' . The function $\downarrow_{r \rightarrow r'}^{var}(m)$ is defined as:

$$\begin{aligned} \downarrow_{\Gamma \rightarrow \Gamma'}^{var}(m) &= [m' \mid \\ & \begin{cases} m'(var') = m(var) & \text{if } var \neq var' \\ m'(var) \in \tau_{\Gamma(var) \rightsquigarrow \Gamma'(var)}(m(var)) & \text{otherwise} \end{cases} \\ &] \end{aligned} \quad (5)$$

Definition 4.18 (Type reduction of a solution sequence). Let $\downarrow_{r \rightarrow r'}^{var}(seq)$ be a function that returns the condensed solution sequence where the variable var is reduced to the representation r' . The function $\downarrow_{r \rightarrow r'}^{var}(seq)$ is defined as:

$$\downarrow_{\Gamma \rightarrow \Gamma'}^{var}(seq) = \parallel_{m \in seq} (\downarrow_{r \rightarrow r'}^{var}(m)) \quad (6)$$

Definition 4.19 (Condensed solution sequence's variable translation). The translation function $\downarrow_{\Gamma \rightarrow \Gamma'}^{var}$ is used to translate the variable representation of a solution sequence from an environment Γ to another environment Γ' . It is defined as follows:

$$\begin{aligned} \downarrow_{\Gamma \rightarrow \Gamma'}^{var}(seq) &= \downarrow_{\Gamma \rightarrow \Gamma'}^{var}([m_1, \dots, m_n]) \\ &= \downarrow_{\Gamma(var) \rightarrow \Gamma'(var)}^{var}(m_1) \parallel \dots \parallel \downarrow_{\Gamma(var) \rightarrow \Gamma'(var)}^{var}(m_n) \\ &= \parallel_{m \in seq} (\downarrow_{\Gamma(var) \rightarrow \Gamma'(var)}^{var}(m)) \end{aligned} \quad (7)$$

Where \parallel is the concatenation operator that concatenates two solution sequences, and \parallel is the function that concatenates all the solution sequences in the list.

Definition 4.20 (Environment transformation). The environment transformation function $\downarrow_r : ((\mathcal{VAR} \rightarrow \mathcal{R}) \times (\mathcal{VAR} \rightarrow \mathcal{R})) \rightarrow (\mathcal{AC} \rightarrow \mathcal{AC})$ is used to transform the variable representation to a target representation. This function takes two environments (i.e., two ways of associating variables to representations) and constructs a transformation that converts all variables from the old environment to the new environment. For each variable whose representation changes, a specific translation operation is applied to that

variable. All these operations are then composed to obtain the overall transformation. It is defined as follows:

$$\downarrow_r(\Gamma, \Gamma') = \bigcirc_{\{var \mid \Gamma(var) \neq \Gamma'(var)\}} \downarrow_{\Gamma \rightarrow \Gamma'}^{var} \text{ if } \Gamma <_r \Gamma' \quad (8)$$

4.3.4 Metadata operators. The metadata operators are used to manipulate and query the metadata associated with versioned named graphs. These operators work together to provide fine-grained control over which versions and graphs are included in query results. They are particularly useful when working with versioned RDF datasets where we need to access data from specific versions or named graphs. A detailed ontology of the versioning operators is shown in Figure 5 in the appendix.

Versioned IRI restriction operator. Given a SPARQL GRAPH operator, we can restrict the versioned graph to match a specific versioned IRI. The versioned IRI restriction operator *vi*-restriction is used in conjunction with the GRAPH keyword in SPARQL queries.

vi-restr is a function that takes a versioned IRI and returns the singleton that contains the couple (n, V) where n is the identifier of the named graph, V the singleton containing the variable associated to the versioned named graph, and *itov* the function that converts an RDF identifier to its value.

The expression of this operator is defined as follows:

$$\begin{aligned} vi\text{-restr} : VI &\rightarrow 2^{\mathbb{N} \times 2^{\mathcal{V}}} \\ vi\text{-restr} : vi &\mapsto \{(n, \{v\}) \mid VI(v, t_g) = vi \wedge itov(n) = t_g\} \end{aligned} \quad (9)$$

Example 4.21. We can restrict the versioned graph using a versioned IRI shown in the SPARQL query 2 defined below. This query retrieves all the building heights for a specific versioned named graph `<vng:Gr-Lyon-v1>`.

Query 2: Get all the building heights for a specific versioned graph

```
SELECT ?height WHERE {
  GRAPH <vng:Gr-Lyon-v1> {
    ?s building:height ?height
  }
}
```

Graph restriction operator. Given a SPARQL GRAPH operator, we can restrict the graph by setting the named graph. The graph restriction operator *vngs-of* is used in conjunction with the triple pattern inside the metadata graph in SPARQL queries.

Example 4.22. We can restrict the graph by setting the named graph linked to the versioned graph shown in the SPARQL query 1 previously defined. This query retrieves all the building heights for a specific named graph `ng:Gr-Lyon`.

Given the flat dataset in Table 2, the result of this query is: $\{\{height \mapsto 10.5\}, \{height \mapsto 9.1\}, \{height \mapsto 10.5\}, \{height \mapsto 15\}\}$.

$$\begin{aligned} vngs\text{-of} : \mathcal{T}_g &\rightarrow \mathbb{N} \times 2^{\mathcal{V}} \\ vngs\text{-of} : t_g &\mapsto \{(n, V_{t_g}) \mid itov(n) = t_g \wedge V_{t_g} = \{v \in V(d) \mid t_g \text{ present in } v\}\} \end{aligned} \quad (10)$$

Version restriction operator. Given a SPARQL GRAPH operator, we can restrict the version by setting the version. The version restriction operator *vngs-in* is used in conjunction with the triple pattern inside the metadata graph in SPARQL queries.

$$\begin{aligned} vngs\text{-in} : \mathcal{V} &\rightarrow 2^{\mathbb{N} \times 2^{\mathcal{V}}} \\ vngs\text{-in} : v &\mapsto \{(n, \{v\}) \mid itov(n) = t_g \wedge t_g \in \mathcal{T}_g\} \end{aligned} \quad (11)$$

Example 4.23. We can restrict the version by setting the version linked to the versioned graph described in the SPARQL query 3 below. This query retrieves all the building heights for a specific version `v:1`.

Query 3: Get all building heights for a specific version

```
SELECT ?height WHERE {
  GRAPH ?vng {
    ?s building:height ?height
  }
  ?vng v:in-version v:1 .
}
```

The result of this query is: $\{\{height \mapsto 10.5\}, \{height \mapsto 9.1\}\}$.

4.3.5 Typing rules.

Expression typing. $\Gamma \vdash e : r$, if assuming that the variables are representable as defined in Γ , e can evaluate to a value that can be represented by r .

The expression typing rules determine the variable representation of a SPARQL expression. We define the SPARQL expressions typing rules. A premise (judgment) is a statement or an assertion on a given abstract syntax tree. An inference rule denoted by (*Rule*) is given by $\frac{P_1, \dots, P_n}{Conc}$ where P_1, \dots, P_n are the premises and *Conc* is the conclusion (composed of one judgment).

Terms : $\frac{}{\Gamma \vdash e : v}$ where $e \in \mathcal{T}$

Variables : $\frac{}{\Gamma \vdash x : r}$ where $x \in VAR$ and $\Gamma(x) = r$

Functions : $\frac{\Gamma_1 \vdash e_1 : v, \dots, \Gamma_n \vdash e_n : v}{\Gamma \vdash f(e_1, \dots, e_n) : v}$ where f is a function (12)

Equality : $\frac{\Gamma \vdash e_1 : r, \Gamma \vdash e_2 : r}{\Gamma \vdash e_1 = e_2 : v}$ where $\iota \leq_r r$

The equality rule is used to determine the variable representation of the equality and inequality operators. Those operators take two expressions as input and returns a boolean value (considered as a v). The two expressions needs to be of the same type and at least an ι to be able to compare them.

Operators typing. The typing rules are used to determine the variable representation of the output of a SPARQL operator. In this paper we define typing rules for selected SPARQL operators.

$$\begin{aligned} (QP \text{ with IRI}) : & \frac{}{\Gamma \vdash QP(s, p, o, x)} \text{ and } x \in \mathcal{T} \\ (QP \text{ with variable}) : & \frac{}{\Gamma \vdash QP(s, p, o, x)} \text{ if } x \in \mathcal{D}(\Gamma) \end{aligned} \quad (13)$$

Query 4: Concatenation of the building names by height where the height is greater than 10

```
Example 4.24. SELECT ?height (GROUP_CONCAT(CONCAT("B.",
  ?b); separator=",")) AS ?concatbuilding) WHERE {
  GRAPH ?vng {
```

```

    ?b height ?height .
  }
  FILTER(?height > 10)
} GROUP BY ?height

```

For certain aggregations, we need to convert the aggregated variables into ν to be able to perform the aggregation operations. For example, for a sum operation, we need to convert the ι variables into ν to be able to perform the operation. Let's consider the SPARQL query in Listing 4. The grouping attributes are `?height` and the aggregate is `GROUP_CONCAT(CONCAT("B.", ?b); separator=",")`. The variable `?height` is of type ι and the variable `?b` is of type ν .

The equation below presents several typing rules for algebraic operations in formal logic. Each rule describes how an expression can be typed based on its environment and the sub-expressions or values it contains. For each rule, the premises are the conditions that must be met for the rule to apply, and the conclusion is the type of the expression.

$$\begin{aligned}
Join(\bowtie) &: \frac{\Gamma \vdash e_1, \Gamma \vdash e_2}{\Gamma \vdash \bowtie(e_1, e_2)} \\
Filter(\sigma) &: \frac{\Gamma \vdash e, \Gamma \vdash c : \nu}{\Gamma \vdash \sigma(e, c)} \\
Union(\cup) &: \frac{\Gamma \vdash e_1, \Gamma \vdash e_2}{\Gamma \vdash \cup(e_1, e_2)} \\
Diff(\setminus) &: \frac{\Gamma \vdash e_1, \Gamma \vdash e_2, \Gamma \vdash e_3}{\Gamma \vdash \setminus(e_1, e_2, e_3)} \\
Optional(\bowtie) &: \frac{\Gamma \vdash e_1, \Gamma \vdash e_2, \Gamma \vdash e_3}{\Gamma \vdash \bowtie(e_1, e_2, e_3)} \quad (14) \\
Projection(\Pi) &: \frac{\Gamma \vdash \Pi(e, \{var_1, \dots, var_n\})}{\Gamma \vdash e} \\
GroupBy(Gr) &: \frac{\forall 1 \leq i \leq m, \Gamma \vdash e_i : \nu, \Gamma \vdash e}{\Gamma \vdash Gr(\{e'_1, \dots, e'_m\}, e)} \\
Aggregation(Agg) &: \frac{\forall 1 \leq i \leq n, \Gamma \vdash e_i : \nu, \Gamma \vdash Gr}{\Gamma \vdash Agg(\{e_1, \dots, e_n\}, AggFun, S, Gr)} \\
AggregateJoin(AggJ) &: \frac{\forall 1 \leq i \leq k, \Gamma \vdash Agg_i}{\Gamma \vdash AggJ(Agg_1, \dots, Agg_k)}
\end{aligned}$$

Example 4.25. Let's consider the SPARQL query above 4. The term `10` in the `FILTER` expression is of type ν .

Example 4.26. Let's consider the SPARQL query in Listing 4. The function `FILTER` is of type ν because the expression `?height > 10` is of type ν . The variable `?height` is of type ν because it needs to compare it with the term `10`. In this case, the function `>` is a function that takes two terms as input and returns a boolean value (considered as a ν).

4.3.6 Operators translation. As shown in Figure 3, the dataset is stored in a condensed model. In order to evaluate a SPARQL query on this condensed dataset, we need to translate the input SPARQL query to another query language in order to make it compatible with the condensed model. For this purpose, we will now define the translation functions for the SPARQL operators of the flat algebra to the condensed algebra. The translation of these operators are done with the function τ and is defined as:

The translation of the flat quad pattern to the condensed quad pattern is straightforward. It doesn't require any variable type

transformation.

$$\tau(QP_F(s, p, o, g)) = QP_C(s, p, o, g) \quad (15)$$

The translation of the flat join operator is done by promoting the variables to the lowest common representation between the two expressions.

$$\begin{aligned}
\tau(\bowtie_F(e_1, e_2)) &= \bowtie_C(\downarrow_r(R^*(e_1), \Gamma)(\tau(e_1)), \downarrow_r(R^*(e_2), \Gamma)(\tau(e_2))) \\
&\text{where } \Gamma = R^*(\bowtie_F(e_1, e_2))
\end{aligned} \quad (16)$$

The translation of the flat filter operator is straightforward and doesn't require any variable type transformation because they are already in the lowest common representation in the flat algebra.

$$\tau(\sigma_F(\mathcal{E}, e)) = \sigma_C(\mathcal{E}, \tau(e)) \quad (17)$$

The translation of the flat optional operator is done by promoting the variables to the lowest common representation between the two joined expressions and the filter expression.

$$\begin{aligned}
\tau(\bowtie_F(e_1, e_2, e_3)) &= \bowtie_C(\\
&\downarrow_r(R^*(e_1), \Gamma)(\tau(e_1)), \\
&\downarrow_r(R^*(e_2), \Gamma)(\tau(e_2)), \\
&\downarrow_r(R^*(e_3), \Gamma)(\tau(e_3))) \\
&\text{where } \Gamma = R^*(\bowtie_F(e_1, e_2, e_3))
\end{aligned} \quad (18)$$

The translation of the flat diff operator preserves the three-argument structure without variable promotion, as the conditions are already evaluated on ν -typed variables.

$$\tau(\setminus_F(e_1, e_2, e_3)) = \setminus_C(\tau(e_1), \tau(e_2), \tau(e_3)) \quad (19)$$

The translation of the flat union operator is done by promoting the variables to the lowest common representation such that the variables are of the same type.

$$\begin{aligned}
\tau(\cup_F(e_1, e_2)) &= \cup_C(\downarrow_r(R^*(e_1), \Gamma)(\tau(e_1)), \downarrow_r(R^*(e_2), \Gamma)(\tau(e_2))) \\
&\text{where } \Gamma = R^*(\cup_F(e_1, e_2))
\end{aligned} \quad (20)$$

The translation of the flat group operator requires promoting the grouping variables to their lowest common representation, ensuring compatibility across all group expressions. The translation is defined as follows:

$$\begin{aligned}
\tau(Gr_F(\mathcal{E}_n, e)) &= Gr_C(\\
&\{\downarrow_r(R^*(e_1), \Gamma)(\tau(e_1)), \dots, \downarrow_r(R^*(e_n), \Gamma)(\tau(e_n))\}, \\
&\tau(e)
\end{aligned} \quad (21)$$

The translation of the flat aggregation operator involves promoting each aggregated expression to the appropriate representation in the target environment. The translation is defined as follows:

$$\begin{aligned}
\tau(Agg_F(\mathcal{E}_n, AggFun_F, S, Gr_F)) &= Agg_C([\downarrow_r(R^*(e_1), \Gamma)(\tau(e_1)), \\
&\dots, \downarrow_r(R^*(e_n), \Gamma)(\tau(e_n))], \\
&\tau(AggFun_F), S, \tau(Gr_F))
\end{aligned} \quad (22)$$

The translation of the flat aggregate join operator is performed by translating each aggregated expression individually and then

combining them using the condensed aggregate join operator. The translation is defined as follows:

$$\tau(\text{Agg}_F(\text{Agg}_{F_1}, \dots, \text{Agg}_{F_n})) = \text{Agg}_J(\tau(\text{Agg}_{F_1}), \dots, \tau(\text{Agg}_{F_n})) \quad (23)$$

The translation of the flat projection operator preserves the projected variable set unchanged; the variable representation of each variable is determined by the sub-expression before projection.

$$\tau(\Pi_F(e, \{var_1, \dots, var_n\})) = \Pi_C(\tau(e), \{var_1, \dots, var_n\}) \quad (24)$$

4.4 SPARQL query evaluation

4.4.1 Algebra semantics.

Solution mapping. We define four partial functions:

- $m^v : \mathcal{VAR} \rightarrow \text{value}_R(v)$ that maps a set of variables to a set of RDF terms as defined in the "Solution mappings" of SPARQL Query language for RDF [36].
- $m^i : \mathcal{VAR} \rightarrow \text{value}_R(i)$ that maps a set of variables to a set of term identifiers.
- $m^\zeta : \mathcal{VAR} \rightarrow \text{value}_R(\zeta)$ that maps a set of variables to a set of pairs composed of a term identifier and finite sets of versions. This solution mapping is called versioned solution mapping.
- $m^\psi : \mathcal{VAR} \rightarrow \text{value}_R(\psi)$ that maps a set of variables to the singleton that contains all the versioned named graphs denoted $\{\bullet\}$.

We note m the function that returns the instances of a variable in a solution mapping in the representation r . The function m is defined as: $m = \text{value}_R(r)$ where value_R is the domain of the solution mapping such that $\text{value}_R(v) = \mathcal{T}$, $\text{value}_R(i) = \mathbb{N}$, $\text{value}_R(\zeta) = \mathbb{N} \times 2^V$ and $\text{value}_R(\psi) = \{\bullet\}$. We consider a mapping valid when all variables appear only in one representation. The domain of each representation is disjoint such that $\text{value}_R(v) \cap \text{value}_R(i) = \emptyset$, $\text{value}_R(v) \cap \text{value}_R(\zeta) = \emptyset$, $\text{value}_R(v) \cap \text{value}_R(\psi) = \emptyset$, $\text{value}_R(i) \cap \text{value}_R(\zeta) = \emptyset$, $\text{value}_R(i) \cap \text{value}_R(\psi) = \emptyset$ and $\text{value}_R(\zeta) \cap \text{value}_R(\psi) = \emptyset$.

Definition 4.27 (Expression). An expression is composed of terms, variables and filter functions. We inductively define \mathcal{E} as:

$$\left\{ \begin{array}{ll} t \in \mathcal{E} & \text{if } t \in \mathcal{T} \\ var \in \mathcal{E} & \text{if } var \in \text{VAR} \\ g(e_1, \dots, e_n) \cup_i \text{VAR}'_i & \text{if } e_i \in \mathcal{E} \\ & \text{and } g \text{ is a filter function with arity } n. \\ & \text{OR/2, AND/2 and NOT/1} \\ & \text{are filter functions} \end{array} \right. \quad (25)$$

Flat algebra semantics. The function $\text{eval}_F(d_F, e)$ evaluates a SPARQL expression e over a flat dataset d_F and returns a solution sequence $[m_1^v, \dots, m_n^v]$.

Definition 4.28 (Evaluation of an expression). We build *values* according to the interpretation of the "Effective Boolean Value" [35] that is used to calculate the result of a filter expression such that $t \in \text{values}$ if $t \sim_{EBV} \text{true} \vee t \sim_{EBV} \text{false}$. We consider also that $\text{null} \in \text{values}$.

The function $\llbracket t \rrbracket : \text{value}$ is a function that returns the value of a term t and the function $\llbracket g/n \rrbracket : \text{value}^n \rightarrow \text{value}$ is a function that returns the value of a filter function g/n applied to n values. They are defined as follows:

$$\begin{aligned} \text{eval}_e(t, m^v) &= \llbracket t \rrbracket \\ \text{eval}_e(var, m^v) &= \begin{cases} m^v(var) & \text{if } var \in \mathcal{D}(m^v) \\ \text{null} & \text{otherwise} \end{cases} \\ \text{eval}_e(g(e_1, \dots, e_n), m^v) &= \llbracket g/n \rrbracket(\text{eval}_e(e_1, m^v), \dots, \text{eval}_e(e_n, m^v)) \\ \text{eval}_e(e, m) &= \text{eval}_e(e, m^v) \\ \text{eval}_e(\text{card}, m^v) &= 1 \\ \text{eval}_e(\text{card}, m) &= \prod_{var \in m^\zeta} |m^\zeta(var)| \times \prod_{var \in m^\psi} |m^\psi(var)| \end{aligned} \quad (26)$$

Example 4.29. Let's consider the SPARQL query 4. The interpretation of the filter expression is $\llbracket ?height > 10 \rrbracket = \llbracket ?height \rrbracket > 10$. Then, the evaluation of the filter is $\text{eval}_e(?height > 10, m^v) = \llbracket ?height \rrbracket > 10$. The result of the filter is $\text{values} = \{\text{true}, \text{false}\}$.

LEMMA 4.30 (INVARIANCE UNDER MAPPING EQUIVALENCE). Let m_1, m_2 be two mappings such that:

$$\begin{aligned} \forall var \in \text{VAR}, (m_1(var) = m_2(var) \vee \\ (var \notin \mathcal{D}(m_1) \wedge var \notin \mathcal{D}(m_2))) \\ \forall e_i \in \mathcal{E}, \text{eval}_e(m_1) = \text{eval}_e(m_2) \end{aligned} \quad (27)$$

INVARIANCE UNDER MAPPING EQUIVALENCE. We proceed by induction on the structure of the expression \mathcal{E} . First, consider the case where the expression is a variable var and $var \notin \mathcal{D}(m_1)$:

$$var \notin \text{dom}(m_2) \wedge \text{eval}_e(var, m_1) = \text{null} \wedge \text{eval}_e(var, m_2) = \text{null} \quad (28)$$

Next, consider the case where the expression is a variable var and $var \in \mathcal{D}(m_1)$:

$$\begin{aligned} \text{eval}_e(var, m_1) = m_1(var) \wedge \text{eval}_e(var, m_2) = m_2(var) \wedge \\ m_1(var) = m_2(var) \end{aligned} \quad (29)$$

□

Definition 4.31 (Evaluation of a list of expressions). We define the function $\text{eval}_\mathcal{E}(\mathcal{E}_n, m)$ that evaluates a list of expressions \mathcal{E}_n as:

$$\begin{aligned} \text{eval}_\mathcal{E}(\mathcal{E}_n, m) &= \text{eval}_\mathcal{E}(\{e_1, \dots, e_n\}, m) \\ &= [\text{eval}_e(e_1, m), \dots, \text{eval}_e(e_n, m)] \end{aligned} \quad (30)$$

Definition 4.32 (Flat grouping function). The function $\text{Gr}_F(\mathcal{E}_n, \Omega)$ is defined as:

$$\begin{aligned} \text{Gr}_F(\mathcal{E}_n, \Omega) &= \{\text{eval}_\mathcal{E}(\mathcal{E}_n, m^v) \rightarrow [m'^v \mid m'^v \in \Omega, \\ &\text{eval}_\mathcal{E}(\mathcal{E}_n, m'^v) = \text{eval}_\mathcal{E}(\mathcal{E}_n, m^v)] \mid m^v \in \Omega\} \end{aligned} \quad (31)$$

Example 4.33. After compiling query 4 and applying it to the dataset 2, the function $\text{Gr}_F(\mathcal{E}_n, \Omega) = \text{Gr}_F((?height), \Omega)$, where Ω is shown in Table 6.

Definition 4.34 (Flat aggregation function). We define the function $\text{Agg}_F(\mathcal{E}_n, \text{AggFun}_F, \mathcal{S}, \text{Gr}_F)$ that creates the results list of an aggregation function. It is defined as:

$$\begin{aligned} \text{Agg}_F(\mathcal{E}_n, \text{AggFun}_F, \mathcal{S}, \text{Gr}_F) &= \\ \text{AggFun}_F([\text{eval}_\mathcal{E}(\mathcal{E}_n, m) \mid m \in \mathcal{D}(\text{Gr}_F)], \mathcal{S}) \end{aligned} \quad (32)$$

Table 6: Value of Ω before evaluating the flat group operation for query 4 on dataset 2.

ψ	ζ	l	v		
			?b	?height	?vng
			ex:bldg#1	10.5	vi:1
			ex:bldg#1	10.5	vi:4
			ex:bldg#1	10.5	vi:3
			ex:bldg#1	11	vi:2
			ex:bldg#3	15	vi:3

Where \mathcal{E} is the list of expressions, $AggFun_F$ is the aggregation function and \mathcal{S} a partial function passed to $AggFun_F$ representing the parameters of the aggregation function.

Definition 4.35 (Flat aggregate join function). The result of the aggregation is a mapping of the variables to the result of the aggregation function. We note $\mathcal{AE} ::= Agg(\mathcal{E}_n, AggFun, \mathcal{S}, Gr)$ the expressions of aggregation. $AggJ_F(\mathcal{AE}_n) = [agg_i \rightarrow val_i \mid k \in \{k' \mid \exists 1 \leq j \leq n, k' \in \mathcal{D}(Agg_{F_j})\}, \forall 1 \leq i \leq n, (k \rightarrow val_i) \in Agg_{F_i}]$.

Example 4.36. The result of the aggregate join ($AggJ_F(\mathcal{AE}_n)$) can be represented as follows:

Table 7: Result of $AggJ_F(\mathcal{AE}_n)$ for the example.

ψ	ζ	l	v	
			?1	?2
			10.5	B.ex:bldg#1,B.ex:bldg#1,B.ex:bldg#1
			11	B.ex:bldg#1
			15	B.ex:bldg#3

Definition 4.37 (Flat compatibility function). $m_1^v \sim_F m_2^v$ is a function that checks if two solution mappings are compatible. It is defined as:

$$m_1^v \sim_F m_2^v = \forall var \in \mathcal{D}(m_1^v) \cap \mathcal{D}(m_2^v), m_1^v(var) = m_2^v(var) \quad (33)$$

In order to evaluate the join of two quad patterns, we need to define the merge function \vee_F of two mappings. The combination function is defined as:

$$\begin{aligned} \vee_F(m_1^v, m_2^v) &= m^v \mid m_1^v \sim_F m_2^v \wedge \forall x \in \mathcal{D}(m_1^v) \cap \mathcal{D}(m_2^v), \\ m^v(x) &= \begin{cases} m_1^v(x), & \text{if } x \in \mathcal{D}(m_1^v) \\ m_2^v(x), & \text{otherwise} \end{cases} \end{aligned} \quad (34)$$

The evaluation functions of those graph patterns are defined in the appendix B.3.1.

Condensed algebra semantics. The function $eval_C(d_C, E)$ evaluates a SPARQL expression E over a condensed dataset d_C and produces a sequence of condensed solution mappings $[m_1, \dots, m_n]$.

Definition 4.38 (Condensed grouping function). The grouping function $Gr_C(\mathcal{E}, \Omega)$ is defined as:

$$\begin{aligned} Gr_C(\mathcal{E}, \Omega) &= \{eval_{\mathcal{E}}(\mathcal{E}_n, m) \rightarrow [m' \mid \\ & m' \in \Omega, eval_{\mathcal{E}}(\mathcal{E}_n, m') = eval_{\mathcal{E}}(\mathcal{E}_n, m)] \mid m \in \Omega\} \end{aligned} \quad (35)$$

Example 4.39. After compiling query 4 and applying it to the dataset 5, the function $Gr_C(\mathcal{E}, \Omega) = Gr_C((?height), \Omega)$, where Ω is shown in Table 8.

Table 8: Value of Ω before evaluating the condensed group operation for query 4 on dataset 5.

ψ	ζ	l		v
		?vng	?b	
	(6, {v:1, v:2})	1	2	
	(7, {v:2})	1	2	
	(7, {v:1})	1	3	
	(6, {v:2})	4	5	

Definition 4.40 (Condensed aggregation function). Let's define the function $Agg_C(\mathcal{E}_n, AggFun_C, \mathcal{S}, Gr_C)$ that creates a list of result of an aggregation function. It is defined as:

$$\begin{aligned} Agg_C(\mathcal{E}_n, AggFun_C, \mathcal{S}, Gr_C) &= \\ AggFun_C([eval_{\mathcal{E}}([e'_1, \dots, e'_n, card], m) \mid m \in \Omega], \mathcal{S}) \end{aligned} \quad (36)$$

Definition 4.41 (Condensed aggregate join function). After the group by operation, we can apply the aggregation function to the results list. The function is defined as:

$$\begin{aligned} AggJ_C(\mathcal{AE}_n) &= [agg_i \rightarrow val_i \mid \\ & k \in \{k' \mid \exists 1 \leq j \leq n, k' \in \mathcal{D}(Agg_{C_j})\}, \\ & \forall 1 \leq i \leq n, (k \rightarrow val_i) \in Agg_{C_i}] \end{aligned} \quad (37)$$

Definition 4.42 (Condensed compatibility function). It checks compatibility by requiring that, for each component of the mappings, the values or identifiers are equal for shared variables, condensed elements have the same graph and overlapping version sets, and unbound graphs are equal for common keys. In summary, two mappings are compatible if all their shared domains agree according to these criteria. It is defined as:

$$\begin{aligned} m_1 \sim_C m_2 &= \forall x \in \mathcal{D}(m_1^v) \cap \mathcal{D}(m_2^v), m_1^v(x) = m_2^v(x) \wedge \\ & \forall x \in \mathcal{D}(m_1^l) \cap \mathcal{D}(m_2^l), m_1^l(x) = m_2^l(x), \\ & \forall x \in \mathcal{D}(m_1^{\zeta}) \cap \mathcal{D}(m_2^{\zeta}), \\ & m_1^{\zeta}(x) = (t_{g1}, V_1), m_2^{\zeta}(x) = (t_{g2}, V_2), t_{g1} = t_{g2}, \\ & V_1 \cap V_2 \neq \emptyset, \forall x \in \mathcal{D}(m_1^{\psi}) \cap \mathcal{D}(m_2^{\psi}), m_1^{\psi}(x) = m_2^{\psi}(x) \end{aligned} \quad (38)$$

A basic graph pattern inside a graph can be interpreted as a join of two quad patterns. In order to evaluate the join of two quad patterns, we need to define the combination of two mappings m_1, m_2 . This equation describes how to merge two mappings by defining the resulting values for each common domain based on the values of the original mappings. The combination function is

defined as:

$$\begin{aligned}
\forall_C (m_1, m_2) &= m \mid m_1 \sim_C m_2 \wedge \forall x \in \mathcal{D}(m_1^v) \cap \mathcal{D}(m_2^v), \\
m^v(x) &= \begin{cases} m_1^v(x), & \text{if } x \in \mathcal{D}(m_1^v) \\ m_2^v(x), & \text{otherwise} \end{cases} \\
\forall x \in \mathcal{D}(m_1^t) \cap \mathcal{D}(m_2^t), \\
m^t(x) &= \begin{cases} m_1^t(x), & \text{if } x \in \mathcal{D}(m_1^t) \\ m_2^t(x), & \text{otherwise} \end{cases} \\
\forall x \in \mathcal{D}(m_1^\zeta) \cap \mathcal{D}(m_2^\zeta), \\
m^\zeta(x) &= \begin{cases} m_1^\zeta(x), & \text{if } x \in \mathcal{D}(m_1^\zeta) \setminus \mathcal{D}(m_2^\zeta) \\ m_2^\zeta(x), & \text{if } x \in \mathcal{D}(m_2^\zeta) \setminus \mathcal{D}(m_1^\zeta) \\ (t_{g1}, V_1 \cap V_2), & \text{otherwise} \\ \text{where } m_1^\zeta(x) = (t_{g1}, V_1) \wedge m_2^\zeta(x) = (t_{g2}, V_2) \end{cases} \\
\forall x \in \mathcal{D}(m_1^\psi) \cap \mathcal{D}(m_2^\psi), m^\psi(x) &= m_1^\psi(x)
\end{aligned} \tag{39}$$

The compatibility function \sim and the merge function \forall_C are used to evaluate the join expression. The evaluation functions of those graph patterns are defined in the appendix B.3.2.

4.4.2 Results equivalence.

Definition 4.43 (Represents). Let $v_m(\text{var}) : \mathcal{VAR} \rightarrow 2^{\mathcal{T}}$ be a function that returns the set of terms for a variable in the mapping m . We consider the injective function itov to be given as the evaluation of a m^t to a m^v . The function $v_m(\text{var})$ is defined as: $v_m(\text{var}) =$

$$\begin{cases} \{m^v(\text{var})\}, & \text{if } \text{var} \in \mathcal{D}(m^v) \\ \{\text{itov}(m^t(\text{var}))\}, & \text{if } \text{var} \in \mathcal{D}(m^t) \\ \{VI(v, t_g) \mid (t_g, V) = m^\zeta(\text{var}) \wedge v \in V\}, & \text{if } \text{var} \in \mathcal{D}(m^\zeta) \\ \{VI(v, t_g) \mid v \in \mathcal{V} \wedge t_g \in \mathcal{T}_g\}, & \text{if } \text{var} \in \mathcal{D}(m^\psi) \end{cases} \tag{40}$$

Example 4.44. Let $m = \{m^v \mapsto (), m^t \mapsto (?b \mapsto 4, ?height \mapsto 5), m^\zeta \mapsto (?vng \mapsto (6, \{v:2\})), m^\psi \mapsto ()\}$.

Then $v_m(?b) = \{\text{ex:bldg\#3}\}$, $v_m(?height) = \{15\}$, $v_m(?vng) = \{VI(v:2, 6) = vi : 3\}$.

Definition 4.45 (Flattening of a condensed solution mapping). We write m_F the mapping that is of the form $m_F \in d_F(t_s, t_p, t_o, t_g, v)$. Let Ξ be a function that returns the sequence of all flat solution mappings for a condensed solution mapping. The function Ξ is defined as:

$$\begin{aligned}
\Xi(m) &= [m_F \mid \mathcal{D}(m_F) = \mathcal{D}(m^v) \cup \mathcal{D}(m^t) \cup \mathcal{D}(m^\zeta) \cup \mathcal{D}(m^\psi) \wedge \\
&\quad \forall \text{var} \in \mathcal{D}(m_F), m_F(\text{var}) \in v_m(\text{var})] \tag{41}
\end{aligned}$$

Definition 4.46 (Permutation of a sequence). Let $\pi(\text{seq})$ the function that returns all the permutations of a sequence where S_n is the symmetric group on n elements, i.e., the set of all permutations from the set $\{1, \dots, n\}$ to itself. The function π is defined as: $\pi(\text{seq}) = \pi([m_1, \dots, m_n]) = \{[m_{\sigma(1)}, \dots, m_{\sigma(n)}] \mid \sigma \in S_n\}$.

LEMMA 4.47 (CANONICAL REPRESENTATION OF AN EVALUATION). *If $\text{seq}_F \cong \text{eval}_C(d_C, e)$ then $\text{seq}_F \cong \text{eval}_C(d_C, \downarrow_r (\Gamma, \Gamma')(e))$.*

The proof is provided in Appendix B.3.3.

THEOREM 4.48 (EVALUATION EQUIVALENCE). *Let a flat algebra A_F , a condensed algebra A_C such that $A_C = \tau(A_F)$ and a dataset d*

and its associated representations d_F and d_C then $\text{eval}_F(d_F, A_F) \cong \text{eval}_C(d_C, \tau(A_F))$.

The proof is provided in Appendix B.3.3.

5 FROM THEORETICAL FOUNDATIONS TO BENCHMARKS

5.1 Implementation of ConVer-G

The ConVer-G system is composed of several components that work together to provide a comprehensive solution for versioned knowledge representation and querying. The Quads loader (QuaDer) component is responsible for loading versions of graphs into the system. The Quads Query (QuaQue) component handles the querying of versioned graphs using either the condensed or the flat algebra.

In its current implementation, core operators are available, including the condensed quad pattern, join, filter, union, optional and aggregation. The translation function is implemented to convert SPARQL queries involving these patterns from the flat algebra into the condensed algebra for execution.

Rather than storing version sets as explicit lists, we encode them as bitstrings, where each bit position corresponds to a specific version. A set bit indicates the presence of that version in the set. This representation enables set operations—such as intersection (bitwise AND) and union (bitwise OR)—which are useful for the condensed join operator. For example, joining two quads valid in versions $\{v:1, v:5\}$ and $\{v:1, v:10\}$ involves a simple bitwise AND of their bitstrings, instantly yielding the intersection $\{v:1\}$. This approach eliminates the need for iteration over version lists.

To compare different methods for querying multiple versions of knowledge, we benchmarked our condensed representation. This allowed us to evaluate and contrast the approaches for querying multiple versions simultaneously.

The benchmark measured query times, memory usage, and space consumption for both the condensed and flat representations. These tests provided insights into the strengths and weaknesses of each model regarding resource usage and performance.

Experiments were conducted on the LIRIS platform PAGoDA <https://projet.liris.cnrs.fr/pagoda/latest/>, which offers a stable, high-performance, and controlled environment to ensure reliable results. We used Argo Workflows, deployed on the PAGoDA Kubernetes cluster, to automate and manage experiment pipelines efficiently. A dataset with fixed numbers of versions, products, and steps was used to ensure consistency across runs, and a fixed memory limit was set for comparability.

The benchmarking environment provided 156 virtual CPU cores and 896GB of RAM, distributed across three HPE Proliant servers (DL380 and DL385 models), each with two AMD EPYC 7443 processors (24 cores per processor, 2.85GHz, 200W).

The environment was configured as follows:

- **Cluster provider:** RKE2.
- **Kubernetes version:** 1.30.7+rke2r1.
- **Architecture:** Amd64.

The benchmark used these parameters:

- **Number of versions in the dataset (Ve):** [5,20,50,100].

- **Number of steps (St):** [0,5,10]. This determined the number of products (i.e., quads) added to the dataset in each version.

These parameters were chosen to cover a broad range of scenarios and ensure the results reflect real-world use cases. A cartesian product of the parameters generated the configurations. For each configuration, an experiment was run and results recorded. The benchmark results are available in the following sections, and for reproducibility, the source code used to conduct the benchmark is available at <https://github.com/VCityTeam/UD-knowledge-evolution-experiment>.

Each experiment deployed the components in a Kubernetes-controlled environment, allowing precise management of memory and CPU resources. Resource allocations for each component were:

- **Processor:** Virtual CPU core request: 1, Virtual CPU core limit: 2.
- **Memory allocation:** Request - 4GB, Limit - 8GB.

5.2 Model - Space usage

The condensed representation may use less space than the flat model. We conducted experiments to compare the space usage of both models. Three plots illustrate performance and resource usage for different types of knowledge graphs: non-evolutive, evolutive, and highly evolutive. These visualizations provide further insight into the comparative analysis. To ensure validity, only experiments completed without errors were included.

Results show that the Postgres-condensed approach consistently uses the least space across all configurations. For instance, with 50 versions and 10 steps, Postgres-condensed uses about 841MB, compared to 2333MB for Postgres-flat, 3091MB for Blazegraph, and 13148MB for Jena. This trend holds for smaller datasets: with 5 versions and 10 steps, Postgres-condensed uses 34MB, while Postgres-flat uses 74MB, Blazegraph 86MB, and Jena 195MB.

As the number of versions and steps increases, space usage grows for all systems, but much more slowly for Postgres-condensed. Blazegraph and Jena show the steepest increases, especially for highly evolutive datasets.

5.3 Dataset query time

For the same query, the condensed representation is expected to be faster than the flat one. We conducted experiments to compare query times for both models.

A set of 17 queries was designed to evaluate system performance, covering both non-aggregative and aggregative use cases. Each query was executed 200 times, with the first 20 runs discarded to account for system warm-up, ensuring only stabilized measurements were analyzed. Each experiment ran every query multiple times on each component to ensure statistical significance and account for variability.

Benchmark results for both non-aggregative and aggregative queries reveal clear trends. For non-aggregative queries, QuaQue-flat consistently achieves the lowest median query times in all configurations, often outperforming other systems by a significant margin. Blazegraph and Jena TDB generally have higher query times, with performance degrading as the number of versions and

steps increases. QuaQue-condensed, while not as fast as QuaQue-flat for non-aggregative queries, remains competitive and scales better than Blazegraph and Jena TDB.

For aggregative queries, performance is more nuanced. QuaQue-flat often provides the best median times, but in some configurations, QuaQue-condensed approaches or surpasses its performance. This suggests that the choice between condensed and flat representations depends on the query type and dataset characteristics. As datasets become more complex, the performance gap between QuaQue-flat and other systems widens, especially for non-aggregative queries.

In summary, the flat representation is generally optimal for non-aggregative queries, while the condensed representation can be advantageous for certain aggregative queries and scales better in terms of space usage. Blazegraph and Jena TDB are less efficient in both query time and scalability compared to the specialized QuaQue implementations. QuaQue-flat queries are faster because they operate on a simple, pre-processed table and use **Hash Joins**, ideal for bulk processing. QuaQue-condensed queries are slower due to complex data transformations performed on the fly, which can mislead the query planner and result in inefficient **Nested Loop** plans, leading to suboptimal performance under real workloads.

6 DISCUSSION

Our formalization provides a solid foundation for understanding and handling RDF versioning in a condensed representation. While our primary motivation and examples have focused on urban data management, the approach is generalizable to a wide range of domains where versioned knowledge graphs are relevant, such as biomedical informatics, digital humanities, software engineering, and scientific data management. Any application that requires querying, and provenance tracking of evolving RDF datasets with different POVs can benefit from the condensed representation and associated algebra. Furthermore, several theoretical limitations need to be addressed in future works.

6.1 Integration with Existing RDF Infrastructure

Our formalization is designed to be compatible with existing RDF standards and infrastructure. The condensed algebra can be implemented as an extension to current RDF stores and SPARQL engines, enabling backward compatibility and incremental adoption. Integration with standards such as PROV [37] would facilitate interoperability and provenance tracking. Furthermore, the translation layer from standard SPARQL to the condensed algebra allows users to leverage familiar query languages while benefiting from the performance and storage optimizations of the condensed model. By aligning our ontology and implementation with established RDF tools and protocols, our approach can be seamlessly incorporated into existing workflows and infrastructures, broadening its impact across multiple domains.

6.2 RDF standards and expressiveness

6.2.1 Integration with Existing RDF Standards. Our formalization has not been fully integrated with existing RDF standards and tools. This would involve aligning our versioning ontology with standards

like PROV [37] and implementing our model in popular RDF query engines.

6.2.2 Handling of Variable Clashes. Our current formalization does not address variable clashes, which can lead to ambiguity and errors in query evaluation. A more robust approach would involve mechanisms to detect and resolve such clashes.

6.2.3 Blank node handling. Our formalization does not address the handling of blank nodes, which are a common feature of RDF data. A more comprehensive model would need to include mechanisms for dealing with blank nodes in versioned data.

6.2.4 Limited Expressiveness of the Versioning Model. Our model is based on a snapshot-based versioning approach, which may not be suitable for all use cases. For example, it cannot handle continuous changes or temporal intervals effectively.

6.3 Solution Sequences and modifiers needs

Projection modifier has been considered in our formalization. However, other modifiers such as Offset, Limit, Order, Distinct, and Reduced are not yet supported. These modifiers are essential for many query scenarios and will be considered in future work.

6.4 Condensed representation limitations

6.4.1 Complexity of the Condensed Representation. While the condensed representation offers potential storage and query optimization benefits, it also introduces complexity in terms of model manipulation and query translation.

6.4.2 Abstraction of the level of compression.

Finite set of compression levels. Our formalization uses a finite set of compression levels. This means that we have defined a limited number of compression levels that our system can handle. While this approach is sufficient for many use cases, it has certain limitations. A formalization that generalizes the typing of variables for any representation could be more advantageous. By allowing more flexible typing, we could support a wider range of compression levels, including intermediate levels or levels defined dynamically based on the specific needs of the application.

From a bounded total order to a bounded lattice. Our formalization uses a bounded total order to represent the variable representations. While this approach is sufficient for many use cases, it has certain limitations. A formalization that uses a bounded lattice could be more advantageous. By allowing more flexible typing, we could support a wider range of compression levels, including intermediate levels or levels defined dynamically based on the specific needs of the application.

ACKNOWLEDGMENTS

This work *Condensed representation of RDF and its application on graph versioning* is supported and funded by the IADoc@UDL (Université de Lyon, Université Claude Bernard Lyon 1) and LIRIS UMR 5205. We also acknowledge the BD team and the Virtual City Project [30] members for their invaluable advice and assistance. We thank the anonymous reviewers for their constructive comments that helped improve the quality of this paper.

REFERENCES

- [1] James Anderson and Arto Bendiken. 2016. Transaction-Time Queries in Dydra. *MEPDAW/LDQ@ESWC 1585* (2016), 11–19.
- [2] Natanael Arndt. 2020. *Distributed Collaboration on Versioned Decentralized RDF Knowledge Bases*. Ph.D. Dissertation. Universität Leipzig. <https://doi.org/10.33968/9783966270205-00>
- [3] Leila Bayoudhi, Najla Sassi, and Wassim Jaziri. 2020. A survey on versioning approaches and tools. In *International Conference on Intelligent Systems Design and Applications*. Springer, 1155–1164.
- [4] Laurent Bulteau, Pierre-Yves David, and Florian Horn. 2023. The Problem of Discovery in Version Control Systems. In *Proceedings of the XII Latin-American Algorithms, Graphs and Optimization Symposium (LAGOS 2023)*. Elsevier, Huatulco, Oaxaca, Mexico, 209–216. <https://doi.org/10.1016/j.procs.2023.08.231>
- [5] Ana Cerdeira-Pena, Antonio Fariña, Javier D. Fernández, and Miguel A. Martínez-Prieto. 2016. Self-Indexing RDF Archives. In *2016 Data Compression Conference (DCC)*. 526–535. <https://doi.org/10.1109/DCC.2016.40>
- [6] Kanishk Chaturvedi, Carl Stephen Smyth, Gilles Gesquière, Tatjana Kutzner, and Thomas H. Kolbe. 2017. Managing Versions and History Within Semantic 3D City Models for the Next Generation of CityGML. In *Advances in 3D Geoinformation*, Alias Abdul-Rahman (Ed.). Springer International Publishing, Cham, 191–206. https://doi.org/10.1007/978-3-319-25691-7_11
- [7] Sara Cohen, Werner Nutt, and Yehoshua Sagiv. 2006. Rewriting queries with arbitrary aggregation functions using views. *ACM Transactions on Database Systems (TODS)* 31, 2 (2006), 672–715.
- [8] Ignacio Cueva and Aidan Hogan. 2020. Versioned Queries over RDF Archives: All You Need is SPARQL?. In *MEPDAW@ISWC*. 43–52.
- [9] Richard Cyganiak. 2005. A relational algebra for SPARQL. *Digital Media Systems Laboratory HP Laboratories Bristol. HPL-2005-170* 35, 9 (2005).
- [10] Helen Eriksson, Jing Sun, Väino Tarandi, and Lars Harrie. 2021. Comparison of versioning methods to improve the information flow in the planning and building processes. *Transactions in GIS* 25, 1 (Feb. 2021), 134–163. <https://doi.org/10.1111/tgis.12672>
- [11] Shi Gao, Jiaqi Gu, and Carlo Zaniolo. 2016. RDF-TX: A Fast, User-Friendly System for Querying the History of RDF Knowledge Bases. In *EDBT*. 269–280.
- [12] Jey Puget Gil, Emmanuel Coquery, John Samuel, and Gilles Gesquière. 2024. ConVer-G: Concurrent versioning of knowledge graphs. arXiv:2409.04499 [cs.DB] <https://arxiv.org/abs/2409.04499>
- [13] Markus Graube, Stephan Hensel, and Leon Urbas. 2014. R43ples: Revisions for triples. In *Proceedings of the 1st Workshop on Linked Data Quality co-located with 10th International Conference on Semantic Systems (SEMANTiCS 2014)*.
- [14] Aidan Hogan, Marcelo Arenas, Alejandro Mallea, and Axel Polleres. 2014. Everything you always wanted to know about blank nodes. *Journal of Web Semantics* 27–28 (Aug. 2014), 42–69. <https://doi.org/10.1016/j.websem.2014.06.004>
- [15] Christian Søndergaard Jensen and Richard T Snodgrass. 2009. Temporal Database. In *Encyclopedia of Database Systems*. Springer, 2957–2960.
- [16] Jens Klump, Lesley Wyborn, Mingfang Wu, Julia Martin, Robert R Downs, and Ari Asmi. 2021. Versioning data is about more than revisions: A conceptual framework and proposed principles. *Data Science Journal* 20, 1 (2021), 12.
- [17] Krishna Kulkarni and Jan-Eike Michels. 2012. Temporal features in SQL:2011. *SIGMOD Rec.* 41, 3 (oct 2012), 34–43. <https://doi.org/10.1145/2380776.2380786>
- [18] Tatjana Kutzner, Kanishk Chaturvedi, and Thomas H. Kolbe. 2020. CityGML 3.0: New Functions Open Up New Applications. *PFG – Journal of Photogrammetry, Remote Sensing and Geoinformation Science* 88, 1 (Feb. 2020), 43–61. <https://doi.org/10.1007/s41064-020-00095-z>
- [19] Thomas Neumann and Gerhard Weikum. 2010. x-RDF-3X: fast querying, high update rates, and consistency for RDF databases. *Proc. VLDB Endow.* 3, 1–2 (Sept. 2010), 256–263. <https://doi.org/10.14778/1920841.1920877>
- [20] Olivier Pelgrin, Ruben Taelman, Luis Galárraga, and Katja Hose. 2023. GLENDA: querying RDF archives with full SPARQL. In *European Semantic Web Conference*. Springer, 75–80.
- [21] Olivier Paul Pelgrin. [n.d.]. Efficient Management of Large RDF Archives. ([n. d.]).
- [22] John Samuel, Sylvie Servigne, and Gilles Gesquière. 2018. Urbano2fab: comprehension of concurrent viewpoints of urban fabric based on git. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* 4 (2018), 65–72.
- [23] John Samuel, Sylvie Servigne, and Gilles Gesquière. 2020. Representation of concurrent points of view of urban changes for city models. *Journal of Geographical Systems* (Feb. 2020). <https://doi.org/10.1007/s10109-020-00319-1>
- [24] Atle Frenvik Sveen. 2020. GeomDiff – an algorithm for differential geospatial vector data comparison. *Open Geospatial Data, Software and Standards* 5, 1 (Dec. 2020), 3. <https://doi.org/10.1186/s40965-020-00076-4>
- [25] Wouter Swierstra and Andres Löb. 2014. The Semantics of Version Control. In *Proceedings of the 2014 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software* (Portland, Oregon, USA) (*Onward!* 2014). Association for Computing Machinery, New York, NY, USA, 43–54. <https://doi.org/10.1145/2661136.2661137>

- [26] Ruben Taelman, Miel Vander Sande, Joachim Van Herwegen, Erik Mannens, and Ruben Verborgh. 2019. Triple storage for random-access versioned querying of RDF archives. *Journal of Web Semantics* 54 (2019), 4–28.
- [27] Ruben Taelman, Miel Vander Sande, and Ruben Verborgh. 2018. OSTRICH: versioned random-access triple store. In *Companion Proceedings of the The Web Conference 2018*. 127–130.
- [28] GeoGig team. 2024. GeoGig - Geospatial Distributed Version Control System. <https://github.com/locationtech/geogig>
- [29] Qri team. 2024. Qri CLI. <https://github.com/qri-io/qri>
- [30] VCity Team. 2024. VCity project website. <https://projet.liris.cnrs.fr/vcity/>
- [31] Miel Vander Sande, Pieter Colpaert, Ruben Verborgh, Sam Coppens, Erik Mannens, and Rik Van de Walle. 2013. R&Wbase: git for triples. *LDOW* 996 (2013).
- [32] Diego Vinasco-Alvarez, John Samuel, Sylvie Servigne, and Gilles Gesquière. 2021. Towards a semantic web representation from a 3D geospatial urban data model. In *SAGEO 2021, 16ème Conférence Internationale de la Géomatique, de l'Analyse Spatiale et des Sciences de l'Information Géographique*. 227–238. https://hal.science/hal-03240567/file/SAGEO_2021.pdf
- [33] Max Volkkel, Wolf Winkler, York Sure, S Ryszard Kruk, and Marcin Synak. 2005. Semversion: A versioning system for rdf and ontologies. In *Proc. of ESWC*.
- [34] W3C. 2024. RDF 1.1 Concepts and Abstract Syntax. <https://www.w3.org/TR/rdf11-concepts/>
- [35] W3C. 2024. SPARQL 1.1 Query Language. <https://www.w3.org/TR/sparql11-query/>
- [36] W3C. 2024. SPARQL Query Language for RDF. <https://www.w3.org/TR/rdf-sparql-query/>
- [37] W3C. 2025. PROV-Overview. <https://www.w3.org/TR/prov-overview/>
- [38] Gunce Su Yilmaz and Jens Dittrich. 2025. Generic Version Control: Configurable Versioning for Application-Specific Requirements. In *Proceedings of the 15th Annual Conference on Innovative Data Systems Research (CIDR 2025)*. CIDR, Amsterdam, The Netherlands, 24. <https://vidb.org/cidrdb/papers/2025/p24-yilmaz.pdf>
- [39] Nazatul Nurlisa Zolkifli, Amir Ngah, and Aziz Deraman. 2018. Version Control System: A Review. *Procedia Computer Science* 135 (2018), 408–415. <https://doi.org/10.1016/j.procs.2018.08.191> The 3rd International Conference on Computer Science and Computational Intelligence (ICCS CI 2018) : Empowering Smart Technology in Digital Era for a Better Life.

A STATE OF THE ART - VERSIONING

A.1 Dataset versioning

Table 9: Comparison of versioning tools by paradigms and query capabilities and area of application adapted from [21]

Tool	Paradigms	Queries	Operators
ConVer-G[12]	TB	all	all
Drydra[1]	TB	all	all
OSTRICH[27]	IC/CB/TB	VM/DM/VQ	TP/BGP/Agg
Quit Store[2]	FB	all	all
RDF-TX[11]	TB	all	TP/BGP/Agg
R43ples[13]	CB	all	all
R&WBase[31]	CB	all	TP/BGP
RBDMS	CB	all	TP/BGP
SemVersion[33]	IC	VM/DM	TP
stardog	CB	all	all
v-RDFCSA[5]	TB	VM/DM/VQ	TP
x-RDF-3X[19]	TB	VM/VQ	TP/BGP/Agg

Operators: [TP] Triple pattern, [BGP] Basic graph pattern, [Graph] Graph name filter, [Agg] Aggregation.

B CONTRIBUTIONS

B.1 Dataset models

Table 10: Association of the identifiers with their values

identifier	Value
1	ex:bldg#1
2	10.5
3	11
4	ex:bldg#3
5	15
6	ng:Gr-Lyon
7	ng:IGN
8	ex:bldg#2
9	height

B.2 Versioning ontology

B.3 Algebra semantics

B.3.1 Flat algebra semantics.

Quad pattern.

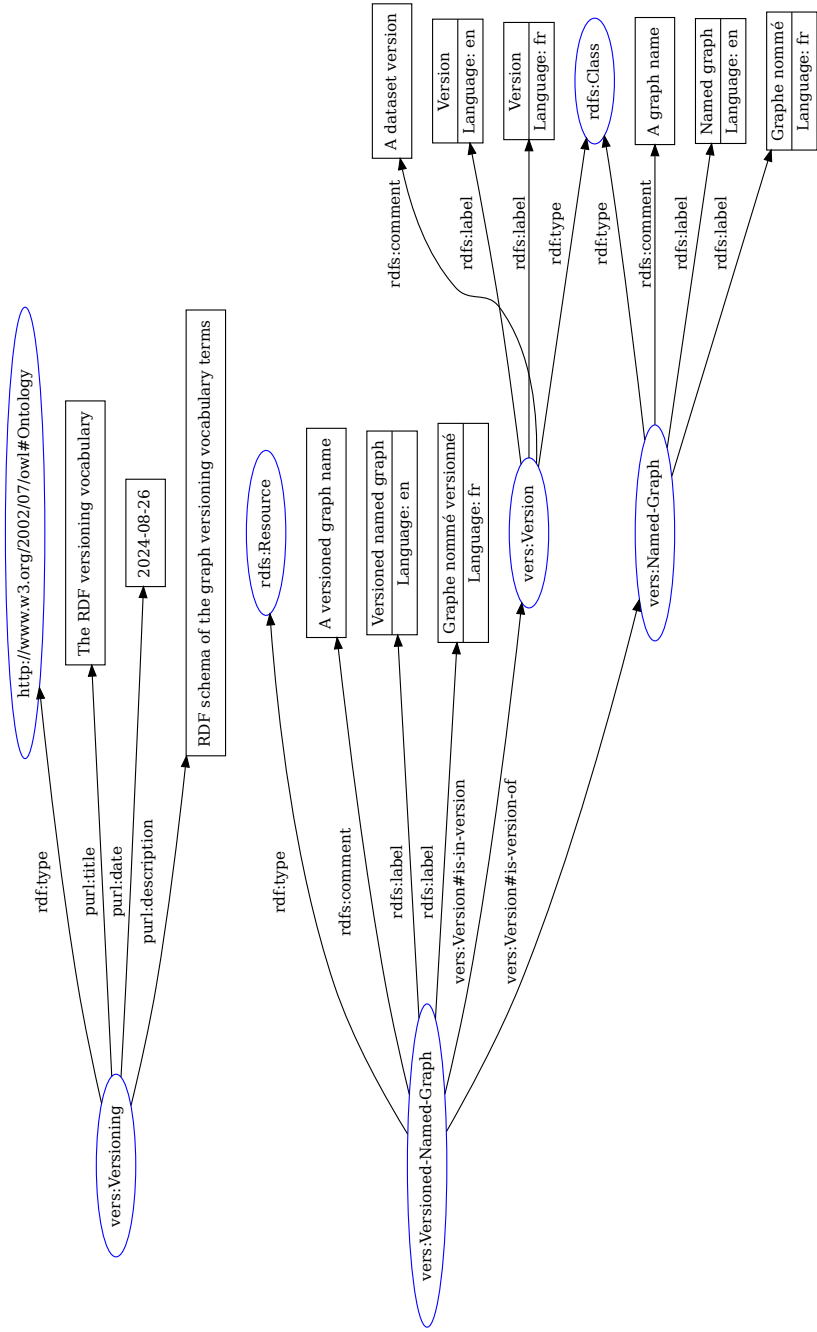
$$eval_F(d_F, QP) = \begin{cases} eval_F(d_F, QP_F(?s, ?p, ?o, ?ag)) = [m \mid \exists t_g \exists v, m(?ag) = VI(v, t_g) \wedge (m(?s), m(?p), m(?o), t_g, v) \in d_F] & \text{if } ?ag \text{ is a variable} \\ eval_F(d_F, QP_F(?s, ?p, ?o, ag)) = [m \mid \exists t_g \exists v, VI(v, t_g) = ag \wedge (m(?s), m(?p), m(?o), t_g, v) \in d_F] & \text{if } ag \text{ is an IRI} \end{cases} \quad (42)$$

Join.

$$eval_F(d_F, \bowtie) = eval_F(d_F, \bowtie_F(e_1, e_2)) = \Omega \quad (43)$$

$$\mathcal{D}(\Omega) = [m \mid \exists m_1, \exists m_2 m_1 \in \Omega_1, m_2 \in \Omega_2, m = \vee_F(m_1, m_2)]$$

$$\Omega(m) = \sum_{m_1 \in \mathcal{D}(\Omega_1), m_2 \in \mathcal{D}(\Omega_2), m = \vee_F(m_1, m_2)} \Omega_1(m_1) \times \Omega_2(m_2)$$



Namespaces:
 vers: <https://github.com/ACityTeam/ConVer-G/>
 rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
 rdfs: <http://www.w3.org/2000/01/rdf-schema#>
 purl: <http://purl.org/dc/elements/1.1>

Figure 5: Versioning ontology

Union.

$$eval_F(d_F, \cup) = eval_F(d_F, \cup_F(e_1, e_2)) = eval_F(d_F, e_1) \parallel eval_F(d_F, e_2) \tag{44}$$

Filter.

$$eval_F(d_F, \sigma) = eval_F(d_F, \sigma_F(\mathcal{E}, e)) = [m \in eval_F(d_F, e) \wedge eval_e(\mathcal{E}, m) \sim_{EBV} true] \quad (45)$$

Difference.

$$eval_F(d_F, \setminus) = eval_F(d_F, \setminus_F(e_1, e_2, e_3)) = \Omega \quad (46)$$

$$\mathcal{D}(\Omega) = [m \mid m \in \Omega_1, \forall m_2 \in \Omega_2, \Omega_1 = eval_F(d_F, e_1), \Omega_2 = eval_F(d_F, e_2), \begin{cases} \neg(m \sim_F m_2) \\ m \sim_F m_2 \wedge eval_e(\{e_3\}, \vee_F(m, m_2)) \sim_{EBV} false \end{cases}]$$

$$\Omega(m) = \setminus_F(\Omega_1, \Omega_2, e_3)(m) = \Omega_1(m)$$

Left outer join (optional).

$$eval_F(d_F, \bowtie) = eval_F(d_F, \bowtie_F(e_1, e_2, e_3)) = eval_F(d_F, \cup_F(\sigma_F(\{e_3\}, \bowtie_F(e_1, e_2)), \setminus_F(e_1, e_2, e_3))) \quad (47)$$

Group by.

$$eval_F(d_F, Gr) = eval_F(d_F, Gr_F(\mathcal{E}_n, \Omega)) = eval_F(d_F, Gr_F(\mathcal{E}_n, eval_F(d_F, e))) \quad (48)$$

Aggregation.

$$eval_F(d_F, Agg) = eval_F(d_F, Agg_F(\mathcal{E}_n, AggFun_F, \mathcal{S}, Gr_F)) = eval_F(d_F, Agg_F(\mathcal{E}_n, AggFun_F, \mathcal{S}, eval_F(d_F, Gr_F))) \quad (49)$$

Aggregate join.

$$\begin{aligned} eval_F(d_F, AggJ) &= eval_F(d_F, AggJ_F(\mathcal{A}\mathcal{E}_n)) = eval_F(d_F, AggJ_F(Agg_1, \dots, Agg_k)) \\ &= eval_F(d_F, AggJ_F(eval_F(d_F, Agg_1), \dots, eval_F(d_F, Agg_k))) \end{aligned} \quad (50)$$

B.3.2 Condensed algebra semantics.

Quad pattern.

$$eval_C(d_C, QP) = \begin{cases} [m \mid \exists t_g \exists V, m^{\checkmark} (?ag) = (t_g, V) \wedge d_C(m'(?s), m'(?p), m'(?o), t_g) = V \wedge V \neq \emptyset \\ \text{if } ?ag \neq ?s \vee ?ag \neq ?p \vee ?ag \neq ?o \\ [m \mid \exists t_g \exists v, m'(?ag) = VI(v, t_g) \wedge v \in d_C(m'(?s), m'(?p), m'(?o), t_g)] \\ \text{otherwise} \\ \text{if } ?ag \text{ is a variable} \\ eval_C(d_C, QP_C(?s, ?p, ?o, ag)) = [m \mid \exists t_g \exists v, VI(v, t_g) = ag \wedge (m'(?s), m'(?p), m'(?o), t_g, v) \in d_F] \text{ if } ag \text{ is an IRI} \end{cases} \quad (51)$$

Join.

$$eval_C(d_C, \bowtie) = eval_C(d_C, \bowtie_C(e_1, e_2)) = \Omega \quad (52)$$

$$\mathcal{D}(\Omega) = [m \mid \exists m_1, \exists m_2, m_1 \in \Omega_1, m_2 \in \Omega_2, m = \vee_C(m_1, m_2)]$$

$$\Omega(m) = \sum_{m_1 \in \mathcal{D}(\Omega_1), m_2 \in \mathcal{D}(\Omega_2), m = \vee_C(m_1, m_2)} \Omega_1(m_1) \times \Omega_2(m_2)$$

Union.

$$eval_C(d_C, \cup) = eval_C(d_C, \cup_C(e_1, e_2)) = eval_C(d_C, e_1) \parallel eval_C(d_C, e_2) \quad (53)$$

Filter.

$$eval_C(d_C, \sigma) = eval_C(d_C, \sigma_C(\mathcal{E}, e)) = [m \mid m \in eval_C(d_C, e) \wedge eval_e(\mathcal{E}, m) \sim_{EBV} true] \quad (54)$$

Difference.

$$eval_C(d_C, \setminus) = eval_C(d_C, \setminus_C(e_1, e_2, e_3)) = \Omega \quad (55)$$

$$\mathcal{D}(\Omega) = [m_1 \mid m_1 \in \Omega_1, \forall m_2 \in \Omega_2, \Omega_1 = eval_C(d_C, e_1), \Omega_2 = eval_C(d_C, e_2), \begin{cases} \neg(m_1 \sim_C m_2) \\ m_1 \sim_C m_2 \wedge eval_e(\{e_3\}, \vee_C(m_1, m_2)) \sim_{EBV} false \end{cases}]$$

$$\Omega(m) = \setminus_C(\Omega_1, \Omega_2, e_3)(m) = \Omega_1(m)$$

Left outer join (optional).

$$eval_C(d_C, \bowtie) = eval_C(d_C, \bowtie_C(e_1, e_2, e_3)) = eval_C(d_C, \cup_C(\sigma_C(\{e_3\}, \bowtie_C(e_1, e_2)), \setminus_C(e_1, e_2, e_3))) \quad (56)$$

Group by.

$$eval_C(d_C, Gr) = eval_C(d_C, Gr_C(\mathcal{E}_n, \Omega)) = eval_C(d_C, Gr_C(\mathcal{E}_n, eval_C(d_C, e))) \quad (57)$$

Aggregation.

$$eval_C(d_C, Agg) = eval_C(d_C, Agg_C(\mathcal{E}_n, AggFunc_C, \mathcal{S}, Gr_C)) = eval_C(d_C, Agg_C(\mathcal{E}_n, AggFunc_C, \mathcal{S}, eval_C(d_C, Gr_C))) \quad (58)$$

Aggregate join.

$$\begin{aligned} eval_C(d_C, AggJ) &= eval_C(d_C, AggJ_C(\mathcal{A}\mathcal{E}_n)) = eval_C(d_C, AggJ_C(Agg_1, \dots, Agg_k)) \\ &= eval_C(d_C, AggJ_C(eval_C(d_C, Agg_1), \dots, eval_C(d_C, Agg_k))) \end{aligned} \quad (59)$$

B.3.3 Results equivalence.

CANONICAL REPRESENTATION OF AN EVALUATION. Let's use the condensed algebra $A_C = \tau(A_F)$ be the condensed algebra expression derived from a flat algebra expression A_F after its translation. The lemma states: If $seq_F \cong eval_C(d_C, A_C)$ then $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma, \Gamma')(A_C))$.

The proof proceeds by induction on the number of elementary variable representation transformations required to get from Γ to Γ' . An elementary transformation changes the representation of a single variable to the next representation in the defined order $\psi <_r \zeta <_r \iota <_r \nu$.

Base Case: $\Gamma = \Gamma'$. In this case, the number of transformations is 0. The environment transformation function $\downarrow_r(\Gamma, \Gamma')$ is the identity function when $\Gamma = \Gamma'$, because $\{var \mid \Gamma(var) \neq \Gamma'(var)\} = \emptyset$. Thus, $\downarrow_r(\Gamma, \Gamma')(e) = e$. The lemma statement becomes: If $seq_F \cong eval_C(d_C, e)$ then $seq_F \cong eval_C(d_C, e)$, which is trivially true.

Inductive Hypothesis (IH): Assume that for any pair of environments (Γ, Γ') such that $\Gamma <_r \Gamma'$ and the transformation $\downarrow_r(\Gamma, \Gamma')$ involves k elementary steps, the lemma holds. That is, if $seq_F \cong eval_C(d_C, e)$, then $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma, \Gamma')(e))$.

Inductive Step: Consider a transformation $\downarrow_r(\Gamma, \Gamma')$ that involves $k + 1$ elementary steps, where $\Gamma <_r \Gamma'$. We can decompose this transformation by identifying an intermediate environment Γ_1 such that $\Gamma <_r \Gamma_1 \leq_r \Gamma'$, where $\downarrow_r(\Gamma, \Gamma_1)$ represents a single elementary transformation (e.g., changing the representation of the variable var_0 from $\Gamma(var_0)$ to $\Gamma_1(var_0)$), and $\downarrow_r(\Gamma_1, \Gamma')$ involves k elementary steps. Thus, $\downarrow_r(\Gamma, \Gamma')(e) = \downarrow_r(\Gamma_1, \Gamma')(\downarrow_r(\Gamma, \Gamma_1)(e))$.

Let $e_1 = \downarrow_r(\Gamma, \Gamma_1)(e)$. The evaluation $eval_C(d_C, e_1)$ produces a solution sequence that is equivalent to applying the type downgrade transformation $\downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}$ to the solution sequence resulting from $eval_C(d_C, e)$. Let $seq_C = eval_C(d_C, e)$ and $seq_{C_1} = eval_C(d_C, e_1)$. By definition of the operator \downarrow_r on expressions, seq_{C_1} is equivalent to $\downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}(seq_C)$ and $\downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}(seq_C) = \bigcirc_{m \in seq_C} (\downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}(m))$.

We get $seq_F \cong seq_C$. By Definition 4.45 (Flattening) and Definition 4.46 (Permutation), this means seq_F is permutation-equivalent to $\bigcup_{m \in seq_C} \Xi(m)$.

We need to show that $seq_C \cong seq_{C_1}$ in terms of their flat representation equivalence to seq_F . Here we use the fact that flattening a match gives the same set of flattened matches as if we first applied type reduction to $m(var_0)$ to get a set of (potentially multiple) mappings $\{m'\}$ (one for each value in $\tau_{\Gamma(var_0) \rightarrow \Gamma_1(var_0)}(m(var_0))$), and then flattening each m' and taking their union.

That is: $\Xi(m) = \bigcup_{m' \in \downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}(m)} \Xi(m')$. Therefore, $\bigcup_{m \in seq_C} \Xi(m) = \bigcup_{m \in seq_C} \left(\bigcup_{m' \in \downarrow_{\Gamma \rightarrow \Gamma_1}^{var_0}(m)} \Xi(m') \right)$. This means $seq_F \cong seq_{C_1}$. So, $seq_F \cong eval_C(d_C, e_1)$.

We have $seq_F \cong eval_C(d_C, e_1)$, and $e_1 = \downarrow_r(\Gamma, \Gamma_1)(e)$. The remaining part of the transformation is $\downarrow_r(\Gamma_1, \Gamma')$, which involves k elementary steps. By the IH (applied to e_1 with environments Γ_1 and Γ'): if $seq_F \cong eval_C(d_C, e_1)$, then $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma_1, \Gamma')(e_1))$.

Since we have established that $seq_F \cong eval_C(d_C, e_1)$, we can conclude: $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma_1, \Gamma')(e_1))$. Substituting $e_1 = \downarrow_r(\Gamma, \Gamma_1)(e)$, we get: $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma_1, \Gamma')(\downarrow_r(\Gamma, \Gamma_1)(e)))$. By the compositional nature of \downarrow_r , this is equivalent to: $seq_F \cong eval_C(d_C, \downarrow_r(\Gamma, \Gamma')(e))$. This completes the inductive step. Therefore, by the principle of induction, the lemma holds for any transformation $\downarrow_r(\Gamma, \Gamma')$ where $\Gamma <_r \Gamma'$. \square

EVALUATION EQUIVALENCE. We want to prove $eval_F(d_F, A_F) \cong eval_C(d_C, \tau(A_F))$. Let's proceed with the proof by induction on the structure of SPARQL algebra operators.

Base Case: we want to show that $eval_F(d_F, QP_F(s, p, o, g)) \cong eval_C(d_C, \tau(QP_F(s, p, o, g)))$. We rewrite this formula by applying the τ function: $eval_F(d_F, QP_F(s, p, o, g)) \cong eval_C(d_C, QP_C(s, p, o, g))$.

Let $m_F \in eval_F(d_F, QP_F(s, p, o, g))$, by definition, there exists $(t_s, t_p, t_o, t_g, v) \in d_F$ such that: $m_F(s) = t_s, m_F(p) = t_p, m_F(o) = t_o, m_F(g) = VI(v, t_g)$. We have $v \in d_C(t_s, t_p, t_o, t_g)$ then $(t_s, t_p, t_o, t_g) \in \mathcal{D}(d_C)$. This means that $m \in eval_C(d_C, QP_C(s, p, o, g))$ such that $m'(s) = t_s, m'(p) = t_p, m'(o) = t_o$ and $m^\zeta(g) = (t_g, V)$ with $V = d_C(t_s, t_p, t_o, t_g)$.

Let $var \in \mathcal{D}(m_F)$ then:

- if $var \in \{s, p, o\}$, $m_F(var) = m'(var) \wedge m^\zeta(var) \in v_m(var)$.
- if $var = g$, $m_F = VI(v, t_g), v \in V \wedge m^\zeta(var) = (t_g, V)$ then $m_F(var) \in v_m(var)$ in addition, $\mathcal{D}(m_F) = \mathcal{D}(m^v) \cup \mathcal{D}(m') \cup \mathcal{D}(m^\zeta) \cup \mathcal{D}(m^\psi)$. This means that $m_F \in \Xi(m)$.

IH: we now assume that the theorem holds for subexpressions and prove it for each SPARQL algebra operator.

Inductive Step: we want to show that $eval_F(d_F, e) \cong eval_C(d_C, \tau(e))$ for each operator e in the SPARQL algebra. We will show that the theorem holds for each operator in the SPARQL algebra.

Join operator: we need to show that $eval_F(d_F, \bowtie_F (e_1, e_2)) \cong eval_C(d_C, \tau(\bowtie_F (e_1, e_2)))$. The rewriting of this formula is: $eval_F(d_F, \bowtie_F (e_1, e_2)) \cong eval_C(d_C, \bowtie_C (\tau(e_1), \tau(e_2)))$.

Let $m_F \in eval_F(d_F, \bowtie_F (e_1, e_2))$, $\exists m_{F_1} \in eval_F(d_F, e_1)$, $\exists m_{F_2} \in eval_F(d_F, e_2)$, with $m_F = \vee_F(m_{F_1}, m_{F_2})$. By the IH: $eval_F(d_F, e_1) \cong eval_C(d_C, \tau(e_1))$. So, $\exists m_1 \in eval_C(d_C, \tau(e_1))$, $m_{F_1} \in \Xi(m_1)$.

It is given that $eval_F(d_F, e_2) \cong eval_C(d_C, \tau(e_2))$. So, $\exists m_2 \in eval_C(d_C, \tau(e_2))$, $m_{F_2} \in \Xi(m_2)$. Since $m_F = \vee_F(m_{F_1}, m_{F_2})$, the variables common to m_{F_1} and m_{F_2} map to the same value in m_F .

If $m_{F_1} \in \Xi(m_1)$ and $m_{F_2} \in \Xi(m_2)$ and $m_{F_1} \sim_F m_{F_2}$, then it must be that $m_1 \sim_C m_2$ (compatibility in the condensed model holds if it holds in the flat model for the corresponding expanded mappings). Then, $m = \vee_C(m_1, m_2)$ (the merge in the condensed model) is defined.

Moreover, $m_F \in \Xi(m)$, since the merge operations in both models are consistent with respect to the flattening function Ξ . Thus, $m \in eval_C(d_C, \bowtie_C (\tau(e_1), \tau(e_2)))$, and then $eval_F(d_F, \bowtie_F (e_1, e_2)) \cong eval_C(d_C, \bowtie_C (\tau(e_1), \tau(e_2)))$.

Union operator: we want to show that $eval_F(d_F, \cup_F(e_1, e_2)) \cong eval_C(d_C, \tau(\cup_F(e_1, e_2)))$. We rewrite this formula: $eval_F(d_F, \cup_F(e_1, e_2)) \cong eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$.

Let $m_F \in eval_F(d_F, \cup_F(e_1, e_2))$. This implies that either $m_F \in eval_F(d_F, e_1)$ or $m_F \in eval_F(d_F, e_2)$.

- Case $m_F \in eval_F(d_F, e_1)$: By the IH, $eval_F(d_F, e_1) \cong eval_C(d_C, \tau(e_1))$. So, $\exists m_1 \in eval_C(d_C, \tau(e_1))$ such that $m_F \in \Xi(m_1)$. Since $m_1 \in eval_C(d_C, \tau(e_1))$, it is also true that: $m_1 \in eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$ by the definition of the union operator in the condensed algebra. Thus, $m_F \in \Xi(m_1) \Rightarrow m_F \in eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$.
- Case $m_F \in eval_F(d_F, e_2)$: By the IH, $eval_F(d_F, e_2) \cong eval_C(d_C, \tau(e_2))$. So, $\exists m_2 \in eval_C(d_C, \tau(e_2))$ such that $m_F \in \Xi(m_2)$. Since $m_2 \in eval_C(d_C, \tau(e_2))$, it is also true that: $m_2 \in eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$ by the definition of the union operator in the condensed algebra. Thus, $m_F \in \Xi(m_2) \Rightarrow m_F \in eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$.

We suppose $m \in eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$. Then, either $m \in eval_C(d_C, \tau(e_1))$ or $m \in eval_C(d_C, \tau(e_2))$.

If $m \in eval_C(d_C, \tau(e_1))$, then by IH, for any $m_F \in \Xi(m)$, we have $m_F \in eval_F(d_F, e_1)$, and $m_F \in eval_F(d_F, \cup_F(e_1, e_2))$. If $m \in eval_C(d_C, \tau(e_2))$, then by IH, for any $m_F \in \Xi(m)$, we have $m_F \in eval_F(d_F, e_2)$, and $m_F \in eval_F(d_F, \cup_F(e_1, e_2))$.

Then, $eval_F(d_F, \cup_F(e_1, e_2)) \cong eval_C(d_C, \cup_C(\tau(e_1), \tau(e_2)))$.

Filter operator: $eval_F(d_F, \sigma_F(\mathcal{E}, e)) \cong eval_C(d_C, \tau(\sigma_F(\mathcal{E}, e)))$ must be demonstrated. We rewrite this formula with the result of the τ function: $eval_F(d_F, \sigma_F(\mathcal{E}, e)) \cong eval_C(d_C, \sigma_C(\mathcal{E}, \tau(e)))$.

Let $m_F \in eval_F(d_F, \sigma_F(\mathcal{E}, e))$. By definition, this means $m_F \in eval_F(d_F, e)$ and $eval_e(\mathcal{E}, m_F) \sim_{EBV} true$. By the IH for the sub-expression e , $eval_F(d_F, e) \cong eval_C(d_C, \tau(e))$. So, there exists $m \in eval_C(d_C, \tau(e))$ such that $m_F \in \Xi(m)$.

Because the filter condition \mathcal{E} depends only on the values of variables, and the flattening function Ξ preserves these values (i.e., for any m' represented by m , $eval_e(\mathcal{E}, m_F)$ is equal to $eval_e(\mathcal{E}, m')$ for the variables in \mathcal{E}), it follows that if $eval_e(\mathcal{E}, m_F) \sim_{EBV} true$, then $eval_e(\mathcal{E}, m) \sim_{EBV} true$ also holds.

Thus, m will also satisfy the filter in the condensed algebra: $m \in eval_C(d_C, \sigma_C(\mathcal{E}, \tau(e)))$. This means that for every m_F satisfying the filter in the flat model, there's a corresponding m satisfying the filter in the condensed model, and $m_F \in \Xi(m)$.

Conversely, let $m \in eval_C(d_C, \sigma_C(\mathcal{E}, \tau(e)))$. This means $m \in eval_C(d_C, \tau(e))$ and $eval_e(\mathcal{E}, m) \sim_{EBV} true$. By IH, for any $m_F \in \Xi(m)$, we have $m_F \in eval_F(d_F, e)$. Since $eval_e(\mathcal{E}, m) \sim_{EBV} true$, and the values relevant to \mathcal{E} are preserved by Ξ , it holds that $eval_e(\mathcal{E}, m_F) \sim_{EBV} true$. Then, $m_F \in eval_F(d_F, \sigma_F(\mathcal{E}, e))$. This establishes that $eval_F(d_F, \sigma_F(\mathcal{E}, e)) \cong eval_C(d_C, \sigma_C(\mathcal{E}, \tau(e)))$.

Diff operator: $eval_F(d_F, \setminus_F(e_1, e_2, e_3)) \cong eval_C(d_C, \tau(\setminus_F(e_1, e_2, e_3)))$ must be demonstrated. We rewrite this formula with the result of the τ function: $eval_F(d_F, \setminus_F(e_1, e_2, e_3)) \cong eval_C(d_C, \setminus_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Let $m_F \in eval_F(d_F, \setminus_F(e_1, e_2, e_3))$. By definition, $m_F \in eval_F(d_F, e_1)$ and $\forall m'_F \in eval_F(d_F, e_2)$, either $m_F \sim_F m'_F$ does not hold, or if it does, then it must give $eval_e(\{e_3\}, \vee_F(m_F, m'_F)) \sim_{EBV} false$.

By the IH for e_1 , $eval_F(d_F, e_1) \cong eval_C(d_C, \tau(e_1))$. So, there exists $m_1 \in eval_C(d_C, \tau(e_1))$ such that $m_F \in \Xi(m_1)$.

Also, by IH for e_2 , $eval_F(d_F, e_2) \cong eval_C(d_C, \tau(e_2))$. So for every $m'_F \in eval_F(d_F, e_2)$, $\exists m_2 \in eval_C(d_C, \tau(e_2))$ such that $m'_F \in \Xi(m_2)$.

It remains to show that for $m_1, \forall m_2 \in eval_C(d_C, \tau(e_2))$, either $m_1 \sim_C m_2$ does not hold, or, if it does, then $eval_e(\{e_3\}, \vee_C(m_1, m_2)) \sim_{EBV} false$ holds. Consider any $m_2 \in eval_C(d_C, \tau(e_2))$. For any $m'_F \in \Xi(m_2)$: if $\neg(m_F \sim_F m'_F)$ for all $m'_F \in \Xi(m_2)$, then it implies $\neg(m_1 \sim_C m_2)$. If there exists a m'_F such that $m_F \sim_F m'_F$, then it must also hold that $eval_e(\{e_3\}, \vee_F(m_F, m'_F)) \sim_{EBV} false$. Because the compatibility and evaluation of e_3 are preserved between the flat and condensed models (as established in the previous proofs for Join and Filter), this entails that $eval_e(\{e_3\}, \vee_C(m_1, m_2)) \sim_{EBV} false$. It follows that m_1 belongs to $eval_C(d_C, \setminus_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Conversely, let $m \in eval_C(d_C, \setminus_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

This means $m \in eval_C(d_C, \tau(e_1))$ and $\forall m' \in eval_C(d_C, \tau(e_2))$, either $\neg(m \sim_C m')$ or $(m \sim_C m' \wedge eval_e(\{e_3\}, \vee_C(m, m')) \sim_{EBV} false)$.

Let $m_F \in \Xi(m)$. By IH, $m_F \in eval_F(d_F, e_1)$. Consider any $m'_F \in eval_F(d_F, e_2)$. By IH, there is a $m'' \in eval_C(d_C, \tau(e_2))$ such that $m'_F \in \Xi(m'')$.

If $\neg(m \sim_C m'')$, this implies $\forall m'_F \in \Xi(m)$ and $m''_F \in \Xi(m'')$, that $\neg(m'_F \sim_F m''_F)$. Thus, $\neg(m_F \sim_F m'_F)$.

If $m \sim_C m''$ and $eval_e(\{e_3\}, \vee_C(m, m'')) \sim_{EBV} false$, then for each corresponding pair of flat mappings $m_F \in \Xi(m)$ and $m'_F \in \Xi(m'')$, it follows that $m_F \sim_F m'_F$ and also $eval_e(\{e_3\}, \vee_F(m_F, m'_F)) \sim_{EBV} false$.

Therefore, $m_F \in eval_F(d_F, \setminus_F(e_1, e_2, e_3))$. Then, we proved that $eval_F(d_F, \setminus_F(e_1, e_2, e_3)) \cong eval_C(d_C, \setminus_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Optional operator: $eval_F(d_F, \bowtie_F(e_1, e_2, e_3)) \cong eval_C(d_C, \tau(\bowtie_F(e_1, e_2, e_3)))$ must be demonstrated. We rewrite this formula with the result of the τ function: $eval_F(d_F, \bowtie_F(e_1, e_2, e_3)) \cong eval_C(d_C, \bowtie_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Let $m_F \in eval_F(d_F, \bowtie_F(e_1, e_2, e_3))$. This implies, by definition, that $m_F \in eval_F(d_F, \cup_F(e, e'))$, where $e = \sigma_F(\{e_3\}, \bowtie_F(e_1, e_2))$ and $e' = \setminus_F(e_1, e_2, e_3)$. As a result, either $m_F \in eval_F(d_F, e)$ or $m_F \in eval_F(d_F, e')$.

- Case $m_F \in eval_F(d_F, e)$: by the IH applied to the Filter and Join operators, we have $eval_F(d_F, e) \cong eval_C(d_C, \tau(e))$. Therefore, there exists $m \in eval_C(d_C, \tau(e))$ such that $m_F \in \Xi(m)$. Since $\tau(e) = \sigma_C(\{\tau(e_3)\}, \bowtie_C(\tau(e_1), \tau(e_2)))$, it follows that $m \in eval_C(d_C, \tau(e))$. Thus, m_F is included in $eval_C(d_C, \bowtie_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.
- Case $m_F \in eval_F(d_F, e')$: by IH applied to the Diff operator, $eval_F(d_F, e') \cong eval_C(d_C, \tau(e'))$. Therefore, there exists $m \in eval_C(d_C, \tau(e'))$ such that $m_F \in \Xi(m)$. Since $\tau(e') = \setminus_F(\tau(e_1), \tau(e_2), \tau(e_3))$, so $m \in eval_C(d_C, \cup_C(\tau(e), \tau(e')))$. Thus, m_F is covered by $eval_C(d_C, \bowtie_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Conversely, let $m_C \in eval_C(d_C, \bowtie_C(\tau(e_1), \tau(e_2), \tau(e_3)))$. This implies $m_C \in eval_C(d_C, \cup_C(\tau(A_F), \tau(B_F)))$. Thus, either $m_C \in eval_C(d_C, \tau(A_F))$ or $m_C \in eval_C(d_C, \tau(B_F))$.

- If $m_C \in eval_C(d_C, \tau(A_F))$: then for all $m_F \in \Xi(m_C)$, by IH, $m_F \in eval_F(d_F, A_F)$. Thus, $m_F \in eval_F(d_F, \cup_F(A_F, B_F))$, which means $m_F \in eval_F(d_F, \bowtie_F(e_1, e_2, e_3))$.
- If $m_C \in eval_C(d_C, \tau(B_F))$: then for all $m_F \in \Xi(m_C)$, by IH, $m_F \in eval_F(d_F, B_F)$. Thus, $m_F \in eval_F(d_F, \cup_F(A_F, B_F))$, which means $m_F \in eval_F(d_F, \bowtie_F(e_1, e_2, e_3))$.

Therefore, $eval_F(d_F, \bowtie_F(e_1, e_2, e_3)) \cong eval_C(d_C, \bowtie_C(\tau(e_1), \tau(e_2), \tau(e_3)))$.

Group operator: $eval_F(d_F, Gr_F(\mathcal{E}_n, e)) \cong eval_C(d_C, \tau(Gr_C(\mathcal{E}_n, e)))$ must be demonstrated. We rewrite this formula with the result of the τ function: $eval_F(d_F, Gr_F(\mathcal{E}_n, e)) \cong eval_C(d_C, Gr_C([\tau(e_1), \dots, \tau(e_n)], \tau(e)))$.

Let's consider a solution sequence $\Omega_F = eval_F(d_F, e)$ and its condensed equivalent $\Omega_C = eval_C(d_C, \tau(e))$. By the inductive hypothesis, we have $\Omega_F \cong \Omega_C$. This means that for every flat mapping $m_F \in \Omega_F$, there is a corresponding condensed mapping $m \in \Omega_C$ such that $m_F \in \Xi(m)$, and vice versa.

The Group operator partitions the solution sequence based on the values of the grouping expressions \mathcal{E}_n . In the flat model, the grouping is performed on Ω_F , while in the condensed model, it is performed on Ω_C . We need to show that the resulting grouped partitions are equivalent.

Let $G_F = Gr_F(\mathcal{E}_n, \Omega_F)$ and $G_C = Gr_C([\tau(e_1), \dots, \tau(e_n)], \Omega_C)$.

For any group key k in the flat model, the corresponding group is $\{m_F \mid m_F \in \Omega_F \wedge eval_{\mathcal{E}}(\mathcal{E}_n, m_F) = k\}$. In the condensed model, the group key is the same, and the group is $\{m \mid m \in \Omega_C \wedge eval_{\mathcal{E}}([\tau(e_1), \dots, \tau(e_n)], m) = k\}$.

Since the evaluation of expressions is preserved between the flat and condensed mappings (as shown in Lemma 4.30), a flat mapping m_F and its corresponding condensed mapping m will produce the same group key k . Therefore, the set of flat mappings in a group in G_F corresponds exactly to the set of flat mappings that can be generated from the condensed mappings in the corresponding group in G_C .

Thus, the grouped structures are equivalent, and the theorem holds for the Group operator.

Aggregation operator: We must demonstrate that: $eval_F(d_F, Agg_F(\mathcal{E}_n, AggFun_F, \mathcal{S}, Gr_F)) \cong eval_C(d_C, \tau(Agg_C(\mathcal{E}_n, AggFunc, \mathcal{S}, Gr_C)))$. This can be rewritten using the translation function τ as: $eval_F(d_F, Agg_F(\mathcal{E}_n, AggFun_F, \mathcal{S}, Gr_F)) \cong eval_C(d_C, Agg_C(\mathcal{E}_n, AggFunc, \mathcal{S}, Gr_C))$.

From the IH for the Group operator, we know that the grouped structures are equivalent: $eval_F(d_F, Gr_F) \cong eval_C(d_C, Gr_C)$.

Let G_F and G_C be the grouped solution sequences for the flat and condensed models, respectively. By the inductive hypothesis, for each group $g_F \in G_F$ defined by a key k , there is a corresponding group $g_C \in G_C$ for the same key k , such that $\bigcup_{m \in g_C} \Xi(m)$ is a permutation of g_F .

The flat aggregation function $AggFun_F$ is applied to each group g_F . The condensed aggregation function $AggFunc$ is applied to each group g_C . We need to show that for each group, the result of the aggregation is the same.

For the COUNT aggregation, the flat model computes the number of mappings in g_F . In the condensed model, for each mapping $m \in g_C$, we compute $eval_e(card, m)$, which represents the number of flat mappings encoded by m . Summing $eval_e(card, m)$ over all $m \in g_C$ yields the same total as counting the mappings in g_F , ensuring equivalence between the two models.

For other aggregations like SUM, AVG, MIN, MAX, the values are extracted from the variables. Since the Ξ function preserves the values of the variables, the set of values passed to the aggregation function is the same for both the flat and condensed models. For example, for each mapping $m_F \in g_F$, there's a corresponding $m \in g_C$ such that $m_F \in \Xi(m)$ and the values of the relevant variables are the same.

This equivalence is formally supported by the principles of rewriting queries with aggregation functions using views, as established in the paper by Sarah Cohen [7], which shows that aggregations over base data can be equivalently computed over pre-aggregated views.

Therefore, for any given group key, the aggregation results will be identical. Since this holds for all groups, the entire result of the aggregation operation is equivalent.

Aggregate join operator: We must demonstrate that $eval_F(d_F, AggJ_F(\mathcal{A}\mathcal{E}_n)) \cong eval_C(d_C, \tau(AggJ_C(\mathcal{A}\mathcal{E}_n)))$. This can be rewritten as: $eval_F(d_F, AggJ_F(AggF_1, \dots, AggF_k)) \cong eval_C(d_C, AggJ_C(\tau(AggF_1), \dots, \tau(AggF_k)))$.

The AggregateJoin operator, as per Definitions 4.35 and 4.41, combines the results from multiple aggregation operations based on their common group keys. From the proof for the Aggregation operator above, we have established that the result of each flat aggregation $AggF_i$ is equivalent to the result of its translated condensed counterpart $\tau(AggF_i) = AggC_i$. This means that for each aggregation, the set of $key \rightarrow val$ pairs produced by the flat evaluation is identical to the set produced by the condensed evaluation. Since both $AggJ_F(\mathcal{A}\mathcal{E}_n)$ and $AggJ_C(\mathcal{A}\mathcal{E}_n)$ operate on an identical set of input pairs and perform the same function of joining these pairs by their keys, their final output must also be equivalent. Therefore, the theorem holds for the Aggregate Join operator. \square

B.4 Model - Space usage

B.4.1 Space usage. The tables below compare the storage requirements for Blazegraph, Jena TDB, QuaQue-flat and QuaQue-condensed implementations.

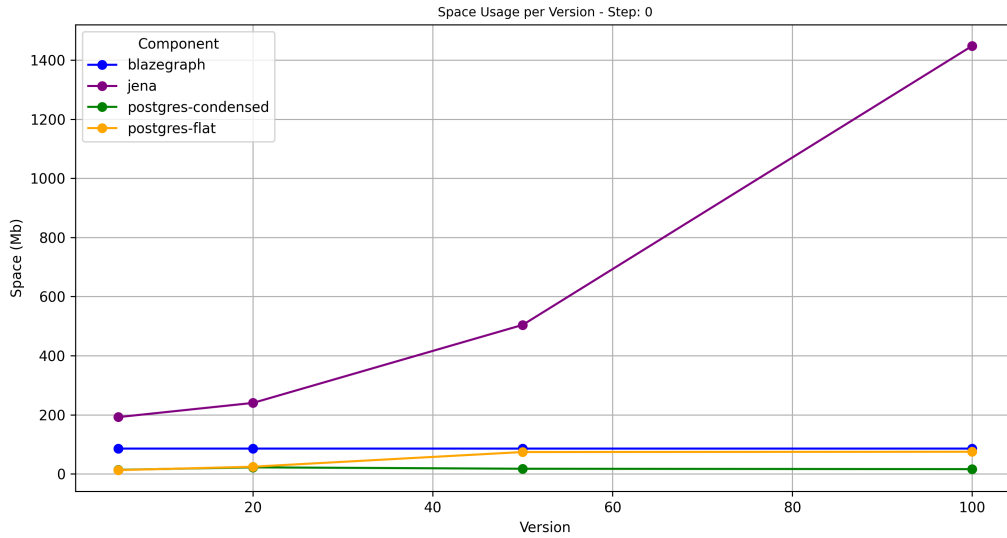


Figure 6: Comparison of space usage between condensed and flat representations for a non-evlutive dataset.

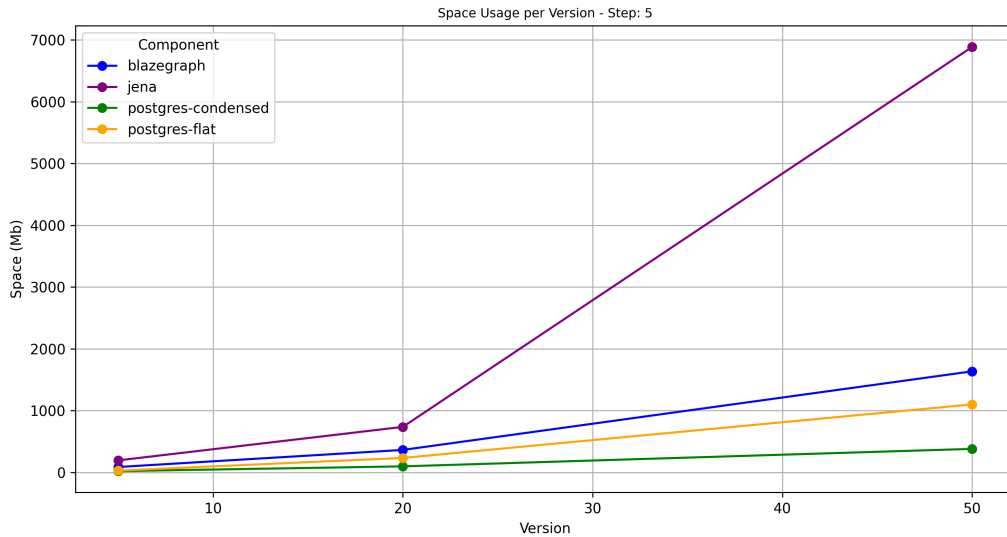


Figure 7: Comparison of space usage between condensed and flat representations for a evlutive dataset.

B.5 Dataset query

B.5.1 Query time. This section presents the results of dataset query performance across different systems and configurations. The tables below compare query execution times for both non-aggregative and aggregative queries, measured on Blazegraph, Jena TDB, QuaQue-flat,

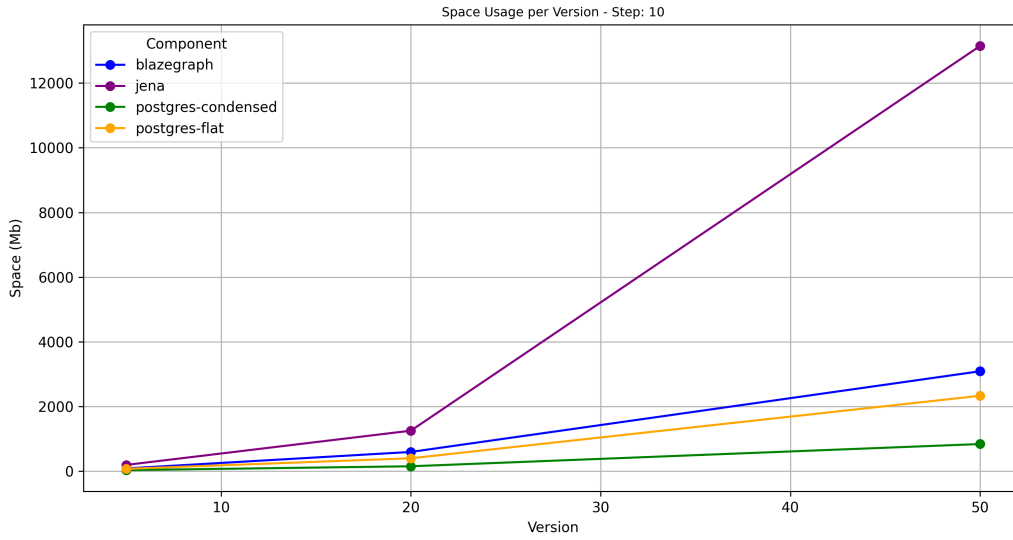


Figure 8: Comparison of space usage between condensed and flat representations for a highly evolutive dataset.

and QuaQue-condensed implementations. Metrics include the 75th percentile, 95th percentile, and median query times for varying dataset sizes and version counts, providing insight into the scalability and efficiency of each approach.

Table 11: Comparison of average times for non-aggregative queries across all components.

Configuration		Blazegraph			Jena TDB			QuaQue-condensed			QuaQue-flat		
Step	Version	75th P.	95th P.	Median	75th P.	95th P.	Median	75th P.	95th P.	Median	75th P.	95th P.	Median
0	5	86.30	121.40	69.40	83.33	104.81	71.40	75.30	98.47	59.85	64.20	101.57	46.50
0	20	79.75	106.26	66.65	101.70	130.23	88.35	116.55	141.84	100.35	70.40	90.86	50.65
0	50	114.28	143.27	91.40	188.93	222.07	170.65	197.40	224.91	180.30	56.88	76.23	42.30
0	100	174.43	215.38	153.00	233.60	277.34	218.35	330.38	357.05	307.80	74.50	115.66	56.00
5	5	94.00	139.58	69.60	86.55	121.19	72.20	131.90	163.56	109.55	57.27	82.99	48.85
5	20	241.60	283.02	220.25	326.52	369.75	289.05	457.15	496.12	429.40	79.03	124.72	59.05
5	50	933.20	1000.39	885.45	1341.42	1460.86	1283.75	1724.75	1844.65	1648.70	59.48	82.65	44.10
10	5	110.20	138.01	96.20	97.28	129.05	85.20	162.43	208.91	138.75	57.45	82.13	43.15
10	20	426.07	530.04	377.35	535.70	637.97	486.65	640.15	721.86	599.40	86.28	135.91	65.20
10	50	1889.95	1975.68	1826.15	2358.30	2497.15	2236.75	3056.75	3200.38	2893.05	65.20	91.95	47.60

Table 12: Comparison of average times for aggregative queries across all components.

Configuration		Blazegraph			Jena TDB			QuaQue-condensed			QuaQue-flat		
Step	Version	75th P.	95th P.	Median	75th P.	95th P.	Median	75th P.	95th P.	Median	75th P.	95th P.	Median
0	5	71.75	113.89	56.93	72.00	103.51	53.86	60.79	103.02	45.36	67.21	104.49	48.86
0	20	45.86	64.11	41.50	57.18	95.02	42.29	74.68	122.38	50.57	77.36	117.25	53.86
0	50	71.54	136.28	53.57	69.00	98.30	53.21	48.04	81.31	37.00	51.71	79.16	37.21
0	100	48.43	70.59	41.71	83.71	137.97	58.43	74.96	126.25	53.50	70.75	103.44	46.93
5	5	77.36	152.74	52.71	59.54	96.30	42.07	70.61	100.35	49.14	60.68	92.64	46.29
5	20	101.39	128.54	88.71	97.64	154.01	73.07	101.79	139.36	71.64	81.86	137.03	56.57
5	50	404.00	459.72	368.50	222.89	268.16	203.07	232.71	280.91	218.57	58.86	95.07	44.93
10	5	81.29	131.74	65.36	59.39	87.23	42.93	79.68	152.66	51.29	46.82	75.23	34.43
10	20	235.43	303.52	170.43	147.86	207.18	109.07	123.39	163.47	103.50	80.61	138.06	53.36
10	50	730.89	803.31	704.14	353.11	379.76	332.86	560.25	617.20	536.29	74.25	106.71	63.79