

Algebraic Mapping Operators for Knowledge Graph Generation

Semantic Web
Vol. 16(5) 1–29
© The Author(s) 2025
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/22104968251361350
journals.sagepub.com/home/swj



Sitt Min Oo¹ , Ben De Meester¹, Ruben Taelman¹ and Pieter Colpaert¹

Abstract

Recent advancements in declarative knowledge graph generation have introduced multiple mapping languages and engines, causing a shift in studies towards optimizing the knowledge graph generation process. Although these engines commonly generate the knowledge graphs from heterogeneous data sources, sharing the optimization techniques and features remains challenging due to the lack of formal operational semantics. To address this, we propose a set of algebraic mapping operators that define operational semantics for general mapping processes. This algebra, based on the SPARQL algebra, enables reuse of established definitions and strengthens the link between knowledge graph generation and query engines. To evaluate language independence we translated mapping languages ShExML and the RDF Mapping Language (RML) into our algebraic mapping plan. Our completeness evaluation shows that our algebraic operators cover the operational semantics of RML and partially support ShExML. Additional analysis is required to cover additional features of ShExML such as joining data from two input sources. For performance evaluation, our proof-of-concept algebraic mapping engine exhibits consistent and low memory usage across workloads, getting second place in the Knowledge Graph Construction Workshop's performance challenge. Algebraic mapping operators decouple mapping engines from specific languages, enabling multilingual mapping engines and allowing optimization techniques to be applied independently of the mapping process. This work lays the foundation for theoretical analysis of complexity and expressiveness of mapping languages and enforces consistency in execution semantics of mapping engines. Furthermore, aligning our algebra with SPARQL opens the door to advanced methods such as virtualization for querying heterogeneous data sources.

Keywords

mapping algebra, semantic web, mapping language, knowledge graph generation

Received: August 14, 2024; accepted: May 23, 2025

Editor: Oscar Corcho, Ontology Engineering Group, Departamento de Inteligencia Artificial, Universidad Politécnica de Madrid, Spain, ocorcho@fi.upm.es

Solicited reviews: Mario Scrocca, Knowledge Engineer, Polytechnic University of Milan, Italy; Jose Emilio Labra-Gayo, University of Oviedo, Spain

1 Introduction

There exist several *use-case-agnostic and declarative* mapping languages (Dimou et al., 2014; García-González et al., 2020; Iglesias-Molina et al., 2023; Lefrançois et al., 2017; Stadler et al., 2015; Sundara et al., 2012; Vu et al., 2019) to generate a Knowledge graph (KG) using the resource description framework (RDF) (Van Assche et al., 2022). *Mapping rules* – described using such mapping language – specify the *mapping process*: How to generate a KG from existing semi-structured data (Van Assche et al., 2022). These mapping languages' increasing popularity and importance is signified by the establishment of the W3C Knowledge Graph Construction Community Group, and a diverse ecosystem of mapping

¹Department of Engineering and Architecture, University of Ghent – imec, Ghent, Belgium

Corresponding Author:

Sitt Min Oo, Department of Engineering and Architecture, University of Ghent – imec, Ghent, Belgium.

Email: x.sittminoo@ugent.be



engines based on these mapping languages (Van Assche et al., 2022). For example, several mapping engines based on the RDF Mapping Language (RML) (Iglesias-Molina et al., 2023) have been implemented (Arenas-Guerrero et al., 2022; Dimou et al., 2014; Haesendonck et al., 2019; Iglesias et al., 2020; Min Oo et al., 2022; Simsek et al., 2019), and there are plans to establish a RML W3C Working Group in the future (Iglesias-Molina et al., 2024).

However, these mapping languages typically provide no formal description of the *operational semantics* through *mapping plans*: The formally defined steps a mapping process should follow to generate a KG based on mapping rules. This prevents the static analysis of the mapping languages to prove the correctness of the mapping process, and analyse the expressiveness and the complexity of the mapping languages. Attempts have been made to formalize these mapping languages, however, these formalizations are mostly used in the context of proving the correctness of an optimization technique (Iglesias et al., 2023) or restricted to a particular mapping language (Asprino et al., 2023; Lopes et al., 2011). To the best of our knowledge, there is currently no research on the theoretical foundations, independent of specific mapping languages, that is applicable to multiple mapping languages.

Hence, different mapping engine implementations typically *individually* infer the operational semantics of the mapping plan based on the *syntax* of the mapping language. As a result, the operational semantics of the mapping engines are different even when using the same mapping rules. These individual inferences of the operational semantics lead to a slow-down in mapping engine development, with repetitive implementations of the same operational semantics in different flavours. Additionally, performance optimizations are scattered across different engine implementations; all incompatible with each other due to the aforementioned individual inferences of the operational semantics.

In this work, we provide such a mapping language-independent theoretical foundation by defining a set of algebraic operators for the mapping process called *mapping algebra*, adapting the algebra (Pérez et al., 2009) of SPARQL, an RDF data query language. These algebraic mapping operators can be composed together to form a mapping plan. We provide a proof-of-concept implementation on par – both in performance as in functionality – with existing mapping engine implementations. This shows how to translate mapping rules from multiple mapping languages into a mapping plan, and thereby providing operational semantics for mapping languages.

This approach can lead to more aligned mapping engines, and allows proving correctness across different mapping languages. Adapting the SPARQL algebra (Pérez et al., 2009) to formulate the mapping algebra allows us to maximally reuse established definitions.

The outline of this paper is as follows: In Section 2, we explore the state-of-the-art on mapping languages and their operational semantics. In Section 3, we discuss our methodology and its rationale. In Section 4, we provide the formal terms and definitions of the set of algebraic mapping operators used to construct the mapping plan. In Section 5, we introduce a reference implementation that consists of an algebraic mapping translator and a proof of concept engine. In Section 6, we provide an overview of our evaluation methodology and results, to show the viability of the algebraic mapping operators without it impeding the performance of a mapping engine. Finally, we conclude and discuss future work in Section 7.

2 Related Works

In this section, we provide an overview of existing mapping languages (Section 2.1) and existing formalizations (Section 2.2), and provide concluding discussions (Section 2.3).

2.1 Mapping Languages

Declarative mapping languages allow describing how to generate a KG by mapping existing data. These languages can be categorized based on the number of input data source types that they support: (i) Homogeneous mapping languages or (ii) heterogeneous mapping languages.

On the one hand, *homogeneous mapping languages* only support one input data source type. For example, relational databases to RDF with languages like D2RQ (Bizer & Seaborne, 2004), R2RML (Sundara et al., 2012) and SML (Stadler et al., 2015), or CSV data to RDF with TARQL (Cyganiak, 2012).

On the other hand, *heterogeneous mapping languages* support multiple input data source types, that is, multiple data formats (CSV, JSON, XML, etc. . .) and multiple data access (websocket, files, databases, etc. . .). Heterogeneous mapping languages can be further categorized (Van Assche et al., 2022): (i) *Query-based mapping languages*, (ii) *dedicated mapping languages*, and (iii) *constraint-based mapping languages*.

Query-based mapping languages extend or adapt the SPARQL syntax to map heterogeneous data to a KG, for example, XSPARQL (Bischof et al., 2012) combines XQuery and SPARQL, SPARQL-Generate (Lefrançois et al., 2017) extends

SPARQL with generation-specific operators, and SPARQL-Anything (Asprino et al., 2023) applies extended SPARQL (CONSTRUCT) queries over an input data meta-model called Façade-X (Daga et al., 2021).

Dedicated mapping languages extend existing mapping languages or use custom syntax. RML (Dimou et al., 2014; Iglesias-Molina et al., 2023), xR2RML (Michel et al., 2015), and D2RML (Chortaras & Stamou, 2018) are examples of dedicated mapping languages which extend R2RML to support more than just relational databases. Amongst them, RML is the most matured mapping language, taken up by the W3C KG-Construct Community Group¹ for ongoing standardization (Iglesias-Molina et al., 2023). D-REPR (Vu et al., 2019) is the only dedicated mapping language with its own syntax.

Currently, ShExML (García-González et al., 2020) is the only *constraint-based* mapping language based on the ShEx (Prud'hommeaux et al., 2018) syntax with extensive modifications and with a focus on making a user-friendly language. Although ShExML itself does not use constraints internally during the mapping of data to RDF, we follow the categorization of Van Assche et al. (2022).

2.2 Formalizations

Several formalization works have been conducted to formalize the aforementioned languages. For *homogeneous mapping languages*, Stadler et al. (2015) provide a unified model, focussed on relational databases as input data source type. However, there are 2 limitations to their approach. First, the formalizations are only applied to the authors' own mapping language, SML. Last, the translation of R2RML to SML – to demonstrate the completeness of their unified model – is informally described. Therefore, it is uncertain if the model is suitable for providing operational semantics for a generic mapping process. Priyatna et al. (2014) provide formalizations for the translation of SPARQL to SQL based on the mapping definitions in an R2RML document, extending the works of Chebotko et al. (2009) and improving the works of SparqlMap (Unbehauen et al., 2013). Since the formalization is not applied directly upon R2RML, it only provides partial operational semantics to R2RML.

For *heterogeneous mapping languages*, on the one hand, query-based mapping languages benefit from the usage of SPARQL semantics, which results in having partial operational semantics out-of-the-box. XSPARQL (Bischof et al., 2012) provides partial operational semantics on the combination of XQuery and SPARQL, using SPARQL's semantics. Similarly, SPARQL-Anything and SPARQL-Generate extend the existing SPARQL syntax to ensure that their operational semantics inherit SPARQL's well-defined semantics (Pérez et al., 2009). For SPARQL-Anything, it formally describes the heterogeneous input data with their RDF meta-model Façade-X (Daga et al., 2021), and uses SPARQL to query the RDF meta-model (Asprino et al., 2023), leading to SPARQL-Anything having provided formal operational semantics of its mapping process. SPARQL-Generate (Lefrançois et al., 2017) also provided the operational semantics for KG generation from heterogeneous data based on its extended SPARQL syntax.

On the other hand, dedicated mapping languages such as RML requires the authors to analyze the mapping language syntax and formulate their own operational semantics. RML Fields (Delva et al., 2021) extended RML's syntax to also allow mapping of nested heterogeneous data, provided an informal operational semantics of how it works. Iglesias et al. (2023), and Arenas-Guerrero et al. (2022) provide formalizations for RML, used to optimize the mapping process by grouping the execution order using the concept of *mapping assertions*, and *mapping partitions* respectively. *Mapping assertions* are formalized using Horn clauses, whereas *mapping partitions* are formalized using set theory. This mismatch in formalization techniques makes it difficult to determine the similarity between the two optimization approaches. The formalizations employed in both works are successfully used to prove the optimization techniques; however, they are tailored to RML and do not introduce general operational semantics of a mapping process.

In parallel, the first author of this paper has conducted a complementary study to formally capture the semantics of mapping languages (Min Oo & Hartig, 2025). This parallel work also investigates to capture the semantics of mapping heterogeneous types of data sources to KG based on RDF, however, it focuses on providing a theoretical foundation using a variation of relational algebra, and provides an algorithm for translating RML v1.1.2 documents² into their algebraic expressions, thus also capturing the semantics of RML. Although this theoretical foundation is formally proven, its scope is smaller, focussing on well-defining a small set of core operators instead of more exhaustively covering current KG generation features. For example, extensions such as RML Fields (Delva et al., 2021) and Logical Targets (Van Assche et al., 2021) are not taken into consideration. Thus, this theoretical foundation has been shown to capture the semantics of RML v1.1.2 and is not yet applied to languages other than RML.

2.3 Discussion

We observe that existing formalizations are: (i) Typically partial, in the case of query-based languages that partially rely on SPARQL semantics, (ii) exclusively applicable to a specific syntax, or (iii) applicable to general mapping semantics but only capture a subset of the semantics of existing mapping languages.

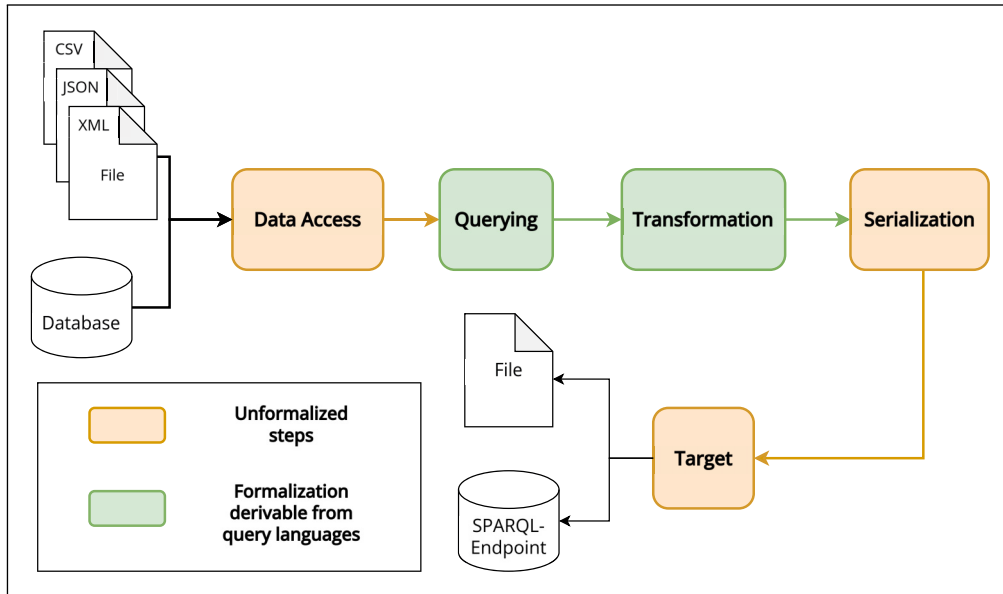


Figure 1. Breakdown of the Mapping Process: The *green* Labelled Steps can be Derived from Existing Query Algebra, and the *orange* Labelled Steps Need to be Formally Defined. The Exemplary Data Access Step can Read Data from a Relational Database, and a CSV File while the Exemplary Target Step, Writes the Generated Output to both a File and a SPARQL Endpoint.

The lack of operational semantics, applicable to multiple mapping languages, impedes using an optimization technique proposed in one mapping language on another mapping language. Furthermore, static analysis cannot be executed for verification of the mapping rules described using these languages. For example, mapping engines such as SDM-RDFizer (Iglesias et al., 2023) and Morph-KGC (Arenas-Guerrero et al., 2022) employ their own concept of operational semantics to optimize the mapping process. This makes it impossible to formally verify the combination of both techniques without a translation between the two operational semantics.

3 Methodology

In this section, we discuss our methodology for developing a mapping algebra and the rationale behind our design choices in selecting a specific set of algebraic mapping operators to represent our mapping algebra.

A mapping process can typically be decomposed into the following steps (Figure 1): (i) *Data access*, (ii) *data querying*, (iii) *data transformations*, (iv) *data serialization*, and (v) *target output*.

One way to define the operational semantics of the mapping process is through the definitions of algebraic operators, similar to the query languages. It is beneficial to extend and adapt the SPARQL algebra with new algebraic operators as a foundation for defining algebraic mapping operators in a mapping algebra: (i) We already have query-based mapping languages successfully leveraging SPARQL’s semantics (Section 2.2) and (ii) this approach allows to align dedicated and query-based mapping languages, allowing for a generic mapping plan that can support multiple mapping languages.

To apply the SPARQL algebra approach, we must define the corresponding algebraic operators for the 5 steps in the mapping process (Figure 1). Formal operational semantics of the *querying* and *transformation* operators are well-defined in querying languages such as SPARQL (Pérez et al., 2009; Seaborne & Harris, 2013) and SQL. Querying operators are defined, in this work, as operators that do not change the *value* associated with an attribute of a data record, in contrast to the transformation operators, which derive new *values* through operations such as string concatenation. An example of a querying operator is SPARQL’s *projection* operator, which we redefine in Section 4.3 with minor modifications. The same is also done for other querying operators such as *join*, and *rename* in Section 4.7 and 4.5 respectively. Similarly, we adapt SPARQL’s *extend* operator (Section 4.4) to provide operational semantics for the transformation step. The remaining steps – *data access*, *data serialization*, and *target output* – need to be defined formally with an extended set of algebraic operators which we define in Section 4.2, 4.8, and 4.9, respectively. As shown in Figure 1, a target operator can write the generated RDF into multiple data sinks, such as a file or a SPARQL endpoint. In the original SPARQL algebra, there exists no algebraic operators which could do a *fan-out* operation and direct the output to multiple downstream operators. Thus,

we extended SPARQL algebra with the concept of *fragment* (Definition 1) and the *fragmenter operator* (Section 4.6) for handling this situation.

4 Mapping Algebra

In this section, we introduce our mapping algebra. We first describe the needed terms and then define the algebraic mapping operators: *Source*, *Projection*, *Extend*, *Rename*, *Fragmenter*, *Natural join*, *θ -join*, *Left outer-join*, *Union*, *Serialize*, and *Target*. The algebra described in this section substantially extends our previous work (Min Oo et al., 2023), introduces seven more operators, and improves the mapping tuple definition. Since this work is inspired by SPARQL algebra, existing definitions and terms by Pérez et al. (2009) will be reused where possible. We only briefly introduce the notations and concepts re-used from SPARQL algebra, readers are referred to the literature for more in-depth definitions (Pérez et al., 2009). Throughout this section, examples will be provided by applying these algebraic mapping operators sequentially on a small dataset.

4.1 Preliminaries

The following pairwise disjoint infinite sets are used in the definition of mapping algebra: V (variables), I (IRIs), B (RDF blank nodes), and L (RDF literals). A *solution mapping* μ is a mapping from variables V with associated data values of type $I \cup B \cup L$. More formally, it is defined as partial function $\mu : V \rightarrow T$ with $T = I \cup B \cup L$. A multiset of solution mapping is denoted as Ω (Pérez et al., 2009). Note that the term *mapping* in the solution mapping follows the mathematical notion of a function in the general sense: If f is a *mapping* function $f : X \rightarrow Y$, then f is a subset of $X \times Y$ consisting of all pairs $(x, f(x))$ for all $x \in X$ and that $f(x) \in Y$ (Halmos, 1998). This notion is different from the domain of mapping languages in KG construction community where *mapping* refers to the mapping of data in a specific format to RDF.

Two solution mappings μ_1 and μ_2 are *compatible*, if and only if, $\forall v \in \text{dom}(\mu_1) \cap \text{dom}(\mu_2), \mu_1(v) = \mu_2(v)$; extending this, the union of μ_1 and μ_2 , $\mu_1 \cup \mu_2$, is also a solution mapping. Given two multisets of solution mappings Ω_1 and Ω_2 , SPARQL algebra (Pérez et al., 2009) defined the *join* (\bowtie), the *union* (\cup), and the *difference* (\setminus) between Ω_1 and Ω_2 as follows (Pérez et al., 2009):

$$\Omega_1 \cup \Omega_2 = \{ \mu \mid \mu \in \Omega_1 \text{ or } \mu \in \Omega_2 \} \quad (1)$$

$$\Omega_1 \bowtie \Omega_2 = \{ \mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1, \mu_2 \in \Omega_2 \text{ and } \mu_1, \mu_2 \text{ are compatible.} \} \quad (2)$$

$$\Omega_1 \setminus \Omega_2 = \{ \mu_1 \in \Omega_1 \mid \forall \mu_2 \in \Omega_2, \mu_1 \text{ and } \mu_2 \text{ are not compatible.} \} \quad (3)$$

Now, we are ready to define the tuple type that our algebraic operators will operate upon. Unlike querying in SPARQL, mapping languages enable users to *fragment* the generated data into different data sinks (e.g. multiple files or web sockets). For example, using Logical Targets (Van Assche et al., 2021), RML engines can export the generated RDF output to different data sinks according to the description provided in Logical Targets. In contrast, query languages do not allow users to specify where to export the queried data. Thus, we introduce the concept of *fragments* (Min Oo et al., 2023), which uses the concepts of the *multiset*, and the *submultiset* as defined by Blizard (1988). Intuitively, the multiset can be seen as an unordered bag of elements where the bag can contain duplicates of elements. Then, the *submultiset* can be seen as a subset, S , of a multiset, M , where S itself is also a multiset. Finally, the *powerset*'s definition is taken from set theory (Bagaria, 2023).

Definition 1. Let Ω be a multiset of solution mappings. A **fragment**, $f \in F$, is a grouping of a submultiset of Ω . The set of fragments, F , is infinite and pairwise disjoint with the other sets, V, I, B , and L , defined in Section 4.1.

Using the definition of *fragments*, the core data model of the mapping algebra, the *mapping tuple* is defined as follows.

Definition 2. Let Ω be a multiset of solution mappings, and $P(\Omega)$ the powerset of the multiset Ω . A **mapping tuple**, t , is a partial function which maps fragments to multisets of solutions mappings $t : F \rightarrow P(\Omega)$. A multiset of mapping tuples is denoted as Γ . Note, from now on, we shall use the notation ω , as an element of the powerset $P(\Omega)$ with $\omega \in P(\Omega)$, to make the definitions of the algebraic operators more accessible to read. In this case, ω itself is also a multiset of solution mappings.

Utilizing fragments in the mapping tuple enables mapping processes to broadcast solution mappings across multiple downstream operators. Furthermore, it enables the partitioning of the solution mappings during construction based on either user defined conditions or some abstract concept such as personal or friend's information as shown in Table 1.

Table 1. Two Mapping Tuples Describing Information Related to John.

Fragment	Multiset of Solution Mappings			
	Solution Mapping	?Name	?Age	?Email
f_{personal}	μ_1	John Doe	23	john.doe@example.com
f_{friends}	μ_2	Susan Sue	25	susan.sue@example.com
f_{friends}	μ_3	Alice Joe	26	alice.joe@example.com

The tuples are fragmented according to *personal* information about John and information about John's *friends*.

4.2 Source

In an algebraic mapping plan, the *source* operators are the leaf nodes of the mapping plan: They generate the required mapping tuples for further processing by the downstream algebraic operators.

A way to extract data records from heterogeneous data formats needs to be defined to generate the mapping tuples. Especially, when the input data has a nested data structure. ShExML (García-González et al., 2020) enables data records extraction from nested data structures, using *iterators* and *fields*. Furthermore, it also enables referencing of data records on different hierarchical level through the use of the *pushed* and the *popped* fields³. RML Fields (Delva et al., 2021) expanded upon the concept of iterators and fields used in ShExML to also allow data records extraction from nested heterogeneous data formats. For example, RML Fields can extract a JSON data record in a CSV table cell using a nested *reference formulation*, unlike ShExML where only nested data structures of the same data format is supported. This work has been recently continued as RML Logical Views⁴.

To support the extraction of data records from a nested data structure, we define the *iterators* and *fields*, as part of the *source* operator, as follows. The definition is similar to the work of RML Fields.

Definition 3. An *iterator*, I , consists of one or more *fields*, $\phi^{r, v_{\text{alias}}}$, an *iterator path*, and a *reference formulation*. An *iterator* extracts a list of records from the data source according to the given *iterator path* (Delva et al., 2021), while the *reference formulation* determines the data format (Dimou et al., 2014) of the data source. This is the entry point to further extract nested data structures with *fields*.

Definition 4. Given an *iterator* I , a *reference path* r , an *alias name* v_{alias} , an *optional reference formulation* and zero or more *fields* $\emptyset \subseteq \Phi$. *Fields* applied on an *iterator* I and zero or more *fields* Φ as $\phi^{r, v_{\text{alias}}}(I, \Phi)$, uses the given *reference path* r to generate one or more records from each record generated by the *iterator* I . The *alias name* v_{alias} of the *field* is used as a *variable* to generate a *solution mapping* μ such that $\mu(v_{\text{alias}})$ is the data record extracted by $\phi^{r, v_{\text{alias}}}(I, \Phi)$. The *reference formulation* is optionally used to determine the data format of the part of the data record the associated *field* is extracting. This enables extraction of data records with mixed data formats in nested data structures as defined in RML Fields. The set of *fields* Φ , are the *subfields* that further extract nested data structures at a deeper hierarchical level.

Definition 5. Given a *data access configuration* C_{access} , and an *iterator* I which consists of one or more *fields* $\emptyset \subseteq \Phi$. The *source* operator generates a *multiset of mapping tuples*, $t \in \Gamma$. Since the generated mapping tuples do not belong to a fragment yet, we define a *default fragment*, f_0 , as the fragment to which **all** the mapping tuples initially belong to. *Data access configuration* C_{access} consists of *metadata information* about utilizing the data source (e.g. connection ports and credentials for Apache Kafka⁵). *Iterators* enable querying of the data source to extract data records and generate our *solution mappings* μ . The *field names* are used as the *variables*, and the associated extracted data value is generated as a *Literal*. The extracted data value can have a *datatype* in the *Literal* if the *datatype* can be inferred from the data source. This mapping between source data values and corresponding *RDF datatypes* are (sometimes implicitly) defined per mapping language⁶. For example, in RML, when extracting a data value from a JSON file typed as a JSON boolean, the extracted data value will have the corresponding `xsd:boolean` datatype. This is to align the data records with the definition of *solution mappings* $\mu : V \rightarrow T$.

f_0 = a default fragment

μ = flattened data record iterated according to I and fields Φ

ω = a multiset containing μ

$$\text{Source}(C, I) = \{t \mid t = f_0 \rightarrow \omega\} \quad (4)$$

Table 2. Representation of the Mapping Tuples Generated by the Source Operator Using Data from Listing 1.

Fragment	Multiset of Solution Mappings					
	Solution Mapping	?Name	?Age	?Email	?Pet.Type	?Pet.Name
f_{default}	μ_1	John Doe	23	john.doe@example.com	Dog	Bax
f_{default}	μ_2	Susan Sue	25	susan.sue@example.com		

To clarify the workings of the source operator, iterator and fields, Example 1 shows how mapping tuples are generated from a simple JSON file with nested data using a source operator with its iterator and fields.

Listing 1. Example input data in JSON format.

```

1 {
2   'people': [
3     {
4       'name' : 'John Doe',
5       'age'  : 23,
6       'email': 'john.doe@example.com',
7       'pet'  :
8         {
9           'type': 'dog',
10          'name': 'Bax'
11        }
12     },
13     {
14       'name' : 'Susan Sue',
15       'age'  : 25,
16       'email': 'susan.sue@example.com',
17     }
18   ]
19 }
```

Example 1. As an example, provided with an input data source such as the JSON file in Listing 1, the source operator could be configured with the following C and I :

- C : path to the JSON file.
- I : With reference formulation “JSONPath” and iterator path $\$.people[*]$

The iterator I also contains the following fields Φ applied as:

- $\phi_1^{r,v_{\text{alias}}}(I, \emptyset)$: with $r = \$.name$ and $v_{\text{alias}} = “?name”$
- $\phi_2^{r,v_{\text{alias}}}(I, \emptyset)$: with $r = \$.age$ and $v_{\text{alias}} = “?age”$
- $\phi_3^{r,v_{\text{alias}}}(I, \emptyset)$: with $r = \$.email$ and $v_{\text{alias}} = “?email”$
- $\phi_4^{r,v_{\text{alias}}}(I, \{\phi_5^{r,v_{\text{alias}}}, \phi_6^{r,v_{\text{alias}}}\})$: with $r = \$.pet$ and $v_{\text{alias}} = “?pet”$
- $\phi_5^{r,v_{\text{alias}}}(I, \emptyset)$: with $r = \$.type$ and $v_{\text{alias}} = “?type”$
- $\phi_6^{r,v_{\text{alias}}}(I, \emptyset)$: with $r = \$.name$ and $v_{\text{alias}} = “?name”$

The iterators are using the JSONPath reference formulation (Gössner et al., 2024). The source operator generates a mapping tuple as shown in Table 2 with a default fragment f_{default} . The fields are executed relative to the iterator, I , and assign the extracted value to the variable described by the name of the field. Execution of $\phi_4^{r,v_{\text{alias}}}(I, \{\phi_5^{r,v_{\text{alias}}}, \phi_6^{r,v_{\text{alias}}}\})$ provides the context for the execution of both $\phi_5^{r,v_{\text{alias}}}(I, \emptyset)$ and $\phi_6^{r,v_{\text{alias}}}(I, \emptyset)$ such that two data records are generated for the variables “?pet.type” and “?pet.name” respectively.

After the source operator generates the mapping tuples from heterogeneous data sources, these mapping tuples are further processed and transformed by the intermediate algebraic mapping operators. The intermediate algebraic mapping operators are defined as follows: *Projection*, *Rename*, *Extend*, *Fragmenter*, and the various *Join* operators.

Table 3. Projection Operator from Example 2 Applied on the Mapping Tuple Shown in Table 2.

Fragment	Multiset of Solution Mappings			
	Solution Mapping	?Name	?Pet.Type	?Pet.Name
f_{default}	μ_1	John Doe	Dog	Bax
f_{default}	μ_2	Susan Sue		

4.3 Projection

In SPARQL algebra, a *projection* restricts solution mappings to a set of variables. This is useful in reducing the amount of data that needs to be further processed by the downstream operators. The corresponding *projection* operator in the mapping algebra is defined as follows.

Definition 6. Given a set of variables $P \subseteq V$, the **projection** operator restricts the variables in the solution mappings μ , associated with the mapping tuple t , according to P .

$$\begin{aligned}
 \text{Project}(\mu, P) &= \mu \text{ restricted to variables in } P \\
 \text{Project}(\omega, P) &= \{\text{Project}(\mu, P) \mid \forall \mu \in \omega\} \\
 \text{Project}(t, P) &= \{(f, \text{Project}(\omega, P)) \mid \forall (f, \omega) \in t\} \\
 \text{Project}(\Gamma, P) &= \{\text{Project}(t, P) \mid t \in \Gamma\}
 \end{aligned} \tag{5}$$

Example 2. A projection operator, configured with a set of variables “?name”, “?pet.type”, “?pet.name” $\in P$, applied on the generated mapping tuple in Table 2 generates a mapping tuple in Table 3.

4.4 Extend

A core operation of the mapping process is the derivation of new values using existing values in the data record. For example, given a data record containing the weight and height of a person, we can calculate the body mass index of the person using their weight and height. The following definition of the *extend* operator enables the mapping algebra to derive new values from existing values.

Definition 7. Given a set of pairs $(v_{\text{new}}, \text{expr}) \in E$, with $v_{\text{new}} \notin \text{dom}(\mu)$, $v_{\text{new}} \in V$ and $\text{expr} : \Omega \rightarrow T$ an extend expression. The **extend** operator derives a new value, value by evaluating the extend expression expr on the solution mapping ω such that $\text{expr}(\omega) = \text{value}$. It extends the solution mapping with the generated value, which is coupled to the new variable v_{new} such that $\omega(v_{\text{new}}) = \text{value}$. If evaluating expr causes an error and $v_{\text{new}} \notin \text{dom}(\mu)$, the extend operator behaves like an identity operator. It is undefined if the variable restriction is violated, which means $v_{\text{new}} \in \text{dom}(\mu)$. Formally, it is defined as follows.

$$\begin{aligned}
 \text{Extend}(\mu, E) &= \mu \cup \{(v_{\text{new}}, \text{value}) \mid (v_{\text{new}}, \text{expr}) \in E, v_{\text{new}} \notin \text{dom}(\mu) \text{ and } \text{value} = \text{expr}(\mu)\} \\
 \text{Extend}(\omega, E) &= \{\text{Extend}(\mu, E) \mid \forall \mu \in \omega\} \\
 \text{Extend}(t, E) &= \{(f, \text{Extend}(\omega, E)) \mid \forall (f, \omega) \in t\} \\
 \text{Extend}(\Gamma, E) &= \{\text{Extend}(t, E) \mid t \in \Gamma\}
 \end{aligned} \tag{6}$$

Different mapping languages can introduce different extend expressions.

Example 3. Provided $\{(?name_iri, \text{encodeIri}^{?name})\} \in E$ with $\text{encodeIri}^{?name}$ a (custom RML) expression that returns an IRI using the ?name attribute from a given solution mapping μ by IRI encoding the data value. The extend operator applied to the mapping tuple shown in Table 3 generates the mapping tuple shown in Table 4.

4.5 Rename

In order to avoid variable collision when processing mapping tuples, an algebraic operator must be able to rename the variables inside the solution mappings associated with the mapping tuples. The *rename* operator, which introduces aliasing of the existing variables in the solution mappings, is defined as follows.

Table 4. Extended Mapping Tuple as Described in Example 3.

Fragment	Multiset of Solution Mappings				
	Solution Mapping	?Name	?Pet.Type	?Pet.Name	?Name_iri
f_{default}	μ_1	John Doe	Dog	Bax	John%20Doe
f_{default}	μ_2	Susan Sue			Susan%20Sue

Table 5. Output of the Rename Operator as Described in Example 4.

Fragment	Multiset of Solution Mappings				
	Solution Mapping	?Fullname	?Pet.Type	?Pet.Name	?Firstname_iri
f_{default}	μ_1	John Doe	Dog	Bax	¡http://example.com/John¿
f_{default}	μ_2	Susan Sue			¡http://example.com/Susan¿

Definition 8. Given a set of pairs of variables $\{(v_{a1}, v_{b1})..(v_{an}, v_{bn})\} \in R$. The **rename** operator, applied on a multiset of mapping tuples Γ , renames the variables of the solution mappings associated with the mapping tuples as follows: If $\mu \in \text{range}(t)$, for each $(v_a, v_b) \in R$ rename $v_a \rightarrow v_b$ if $v_a \in \text{dom}(\mu)$. If the rename operator is configured with an alias string s_{alias} instead of R , the rename operator will concatenate s_{alias} as the suffix for all $v \in \text{dom}(\mu)$ as $s_{\text{alias}}\|v$.

$$\begin{aligned}
\text{Rename}(\mu, R) &= \{(v_b, d) \mid \forall (v_a, v_b) \in R, \forall (v, d) \in \mu, v = v_a\} \cup \\
&\quad \{(v, d) \mid \forall (v_a, v_b) \in R, \forall (v, d) \in \mu, v \neq v_a\} \\
\text{Rename}(\omega, R) &= \{\text{Rename}(\mu, R) \mid \forall \mu \in \omega\} \\
\text{Rename}(t, R) &= \{(f, \text{Rename}(\omega, R)) \mid \forall (f, \omega) \in t\} \\
\text{Rename}(\Gamma, R) &= \{\text{Rename}(t, R) \mid t \in \Gamma\} \tag{7}
\end{aligned}$$

$$\begin{aligned}
\text{Rename}(\mu, s_{\text{alias}}) &= \{(v_{\text{aliased}}, d) \mid \forall (v, d) \in \mu, v_{\text{aliased}} = s_{\text{alias}}\|v\} \\
\text{Rename}(\omega, s_{\text{alias}}) &= \{\text{Rename}(\mu, s_{\text{alias}}) \mid \forall \mu \in \omega\} \\
\text{Rename}(t, s_{\text{alias}}) &= \{(f, \text{Rename}(\omega, s_{\text{alias}})) \mid \forall (f, \omega) \in t\} \\
\text{Rename}(\Gamma, s_{\text{alias}}) &= \{\text{Rename}(t, s_{\text{alias}}) \mid t \in \Gamma\} \tag{8}
\end{aligned}$$

Readers may realize that the *rename* operator can also be derived by first extending the solution mapping with a new variable (Extension), copying the value associated with the old variable, and finally projecting away the old variable (Projection). We defined the *rename* operator to describe the execution of the rename operation in one operator instead of two operators. This reduces the complexity and redundancy of the generated mapping plan using the operators. The *extend* and *project* operator chain to represent the rename operator is formally defined as follows.

$$\begin{aligned}
R &= \{(v_{a1}, v_{b1}) \dots (v_{an}, v_{bn}) \mid n \in \mathbb{N}\} \\
P_{\text{renamed}} &= \{v \mid v \in \text{dom}(\mu)\} \cup \{v_b \mid \forall (v_a, v_b) \in R\} / \{v_a \mid \forall (v_a, v_b) \in R\} \\
\text{copyData}(v) &= \text{copies the value of } \mu(v) \text{ when evaluated by the extend operator on a solution mapping } \mu \\
E_{\text{rename}} &= \{(v_b, \text{copyData}(v_a)) \mid \forall (v_a, v_b) \in R\} \\
\text{Rename}(\Gamma, R) &= \text{Project}(\text{Extend}(\Gamma, E_{\text{rename}}), P_{\text{renamed}}) \tag{9}
\end{aligned}$$

Example 4. Provided with a set of variable pairs $\{(?name, ?fullname)\} \in R$. Applying the rename operator on the mapping tuples in Table 4, generates the mapping tuples in Table 5. The old variable, *?name*, in the solution mappings is renamed to *?fullname*.

4.6 Fragmenter

The aforementioned operators do not manipulate the *fragments* part of the mapping tuples, but only process the associated solution mappings. To manipulate the fragments of the mapping tuples, the *fragment* operator is defined as follows.

Table 6. Fragmentation of Mapping Tuple as Described in Example 3.

Fragment	Multiset of Solution Mappings				
	Solution Mapping	?Fullname	?Pet.Type	?Pet.Name	?Firstname_iri
f_{contacts}	μ_1	John Doe	Dog	Bax	$\text{\textbackslash}http://example.com/John\text{\textbackslash}$
f_{contacts}	μ_2	Susan Sue			$\text{\textbackslash}http://example.com/Susan\text{\textbackslash}$

Definition 9. The **fragmenter** operator fragments the mapping tuple into a new fragment f_{new} . Given a partial transformation function, $\delta : F \rightarrow F$. A fragmenter operator applies δ on the mapping tuple $t = f \rightarrow \mu$, and map it into a new mapping tuple $t_{\text{mapped}} = f_{\text{new}} \rightarrow \mu$ if $\text{dom}(\delta) \subseteq \text{dom}(t)$ and $f_{\text{new}} \in \text{range}(\delta)$. When $\text{dom}(\delta) \not\subseteq \text{dom}(t)$, fragmenter operator acts like an identity function. More formally, it is defined as follows.

$$\begin{aligned}
\delta &= \{(f_{\text{old}}, f_{\text{new}}) \mid f_{\text{old}}, f_{\text{new}} \in F\} \\
\delta(t) &= \{(f, \omega) \mid (f, \omega) \in t, f \neq f_{\text{old}}\} \cup \{(f_{\text{new}}, \omega) \mid (f_{\text{old}}, \omega) \in t, (f_{\text{old}}, f_{\text{new}}) \in \delta\} \\
\text{Fragment}(t, \delta) &= \begin{cases} \delta(t) & \text{if } \text{dom}(\delta) \subseteq \text{dom}(t) \\ t & \text{otherwise} \end{cases} \\
\text{Fragment}(\Gamma, \delta) &= \{\text{Fragment}(t, \delta) \mid t \in \Gamma\}
\end{aligned} \tag{10}$$

Example 5. Continuing with the output of the extend operator as shown in Table 4. A fragmenter operator with $\delta : \{(f_{\text{default}}, f_{\text{contacts}})\}$ applied on the mapping tuples generates new mapping tuples shown in Table 6, where the old fragment, f_{default} , is mapped to the new fragment, f_{contacts} .

4.7 Binary Operators

Previously defined operators are unary: They only work on a single multiset of mapping tuples Γ . To combine two multisets of mapping tuples Γ_1 , and Γ_2 , binary algebraic operators need to be defined, along with the definition of *compatibility* between the mapping tuples. Thus, to support the binary operations for combining mapping tuples, the mapping algebra defines the *compatibility of mapping tuples*, the *natural join*, the *θ -join*, and the *left-join* between Γ_1 and Γ_2 as follows.

Definition 10. Two mapping tuples, $t_1 \in \Gamma_1$ and $t_2 \in \Gamma_2$, are **compatible**, if and only if, $\forall f \in \text{dom}(t_1) \cap \text{dom}(t_2)$, for the associated multisets of solution mappings $t_1(f) = \Omega_1$ and $t_2(f) = \Omega_2$, $\exists \mu_1 \in \Omega_1, \exists \mu_2 \in \Omega_2$ where μ_1 and μ_2 are compatible so that $\Omega_1 \bowtie \Omega_2 \neq \emptyset$. See Section 4.1 for solution mappings compatibility. If $\text{dom}(t_1)$ and $\text{dom}(t_2)$ are disjoint, t_1 and t_2 are compatible.

4.7.1 Natural Join.

Definition 11. **Natural join** is a binary operator that combines two multisets of mapping tuples Γ_1 and Γ_2 if they are compatible (Definition 10). It produces mapping tuples, $t_1 \bowtie t_2$, which is a combination of two mapping tuples, $t_1 \in \Gamma_1$ and $t_2 \in \Gamma_2$ that have common fragments, $\forall f \in \text{dom}(t_1) \cap \text{dom}(t_2)$, for which the associated multisets of solution mappings $\Omega_1 = t_1(f)$ and $\Omega_2 = t_2(f)$ are joined as $\Omega_1 \bowtie \Omega_2$ according to Equation 2. It is formally defined as follows.

$$\begin{aligned}
\text{NatJoin}(t_1, t_2) &= \{(f, \Omega_1 \bowtie \Omega_2) \mid \forall f \in \text{dom}(t_1) \cap \text{dom}(t_2), t_1(f) = \Omega_1, t_2(f) = \Omega_2\} \\
\text{NatJoin}(\Gamma_1, \Gamma_2) &= \{\text{NatJoin}(t_1, t_2) \mid t_1 \in \Gamma_1, t_2 \in \Gamma_2, t_1 \text{ and } t_2 \text{ are compatible}\}
\end{aligned} \tag{11}$$

Natural join operator joins mapping tuples if they have common *fragments*, and also have equal values for all the common *variables* of the associated solution mapping of the fragment. It does not allow users to join mapping tuples based on the different variables of $\text{dom}(\mu_1)$ and $\text{dom}(\mu_2)$. Furthermore, natural join only checks the *equality* condition on the common variables and fragments: It does not use other predicate functions such as \leq or \geq .

4.7.2 θ -Join. θ -join enables the use of a predicate function, θ , on the specified variables to join two multisets of mapping tuples. Thus, it is a more general form of the natural join operator. In order to execute θ -join, the following conditions must be satisfied: $\forall v_a \in \mu_a \wedge \forall v_b \in \mu_b. v_a \neq v_b$. Otherwise, μ_a and μ_b can be incompatible and cannot be joined as

$\mu_a \cup \mu_b$ (see Section 4.1 for compatibility definition). Thus, to make sure that μ_a and μ_b are compatible, the mapping plan planner should apply the rename operator before the θ -join operator to ensure that none of variables in μ_a and μ_b are equal to each other. It is defined as follows.

Definition 12. Given a binary predicate function, $\theta_{v_1}^{v_2} : \Omega \times \Omega \rightarrow \text{boolean}$, with variables v_1 , and v_2 , where $v_1 \in \text{dom}(\mu_1)$, $v_2 \in \text{dom}(\mu_2)$, and two multisets of mapping tuples Γ_1 and Γ_2 . Provided that $\forall v_a \in \text{dom}(\mu_1), \forall v_b \in \text{dom}(\mu_2), v_a \neq v_b$. θ -join then combines two mapping tuples $t_1 \in \Gamma_1$, and $t_2 \in \Gamma_2$, if for the same fragment $f \in \text{dom}(t_1) \cap \text{dom}(t_2)$, the following condition is satisfied: $\mu_1 \in \Omega_1, \mu_2 \in \Omega_2, \theta_{v_1}^{v_2}(\mu_1, \mu_2) = \text{true}$ with $\Omega_1 = t_1(f)$ and $\Omega_2 = t_2(f)$.

$$\begin{aligned} \theta_{v_1}^{v_2}(\mu_1, \mu_2) &= \begin{cases} \text{true,} & \text{if } \mu_1(v_1) = \mu_2(v_2) \\ \text{false,} & \text{otherwise} \end{cases} \\ \text{ThetaJoin}(\Omega_1, \Omega_2, \theta_{v_1}^{v_2}) &= \{ \mu_1 \cup \mu_2 \mid \mu_1 \in \Omega_1, \mu_2 \in \Omega_2, \theta_{v_1}^{v_2}(\mu_1, \mu_2) \text{ evaluates to true } \} \\ \text{ThetaJoin}(t_1, t_2, \theta_{v_1}^{v_2}) &= \{ (f, \text{ThetaJoin}(\Omega_1, \Omega_2, \theta_{v_1}^{v_2})) \mid \forall f \in \text{dom}(t_1) \cap \text{dom}(t_2), \\ &\quad t_1(f) = \Omega_1, t_2(f) = \Omega_2 \} \\ \text{ThetaJoin}(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) &= \{ \text{ThetaJoin}(t_1, t_2, \theta_{v_1}^{v_2}) \mid t_1 \in \Gamma_1, t_2 \in \Gamma_2 \} \end{aligned} \quad (12)$$

Natural and θ -join filters out solution mappings which do not satisfy the predicate function. The filtering of solution mappings results in the loss of information which might require extra processing steps to retain them. Thus, a new binary algebraic operator, which retains the solution mappings that do not satisfy the given predicate function, needs to be defined. SPARQL and relational algebra have definitions for such a group of binary operators called *outer-joins*.

4.7.3 Left Outer-Join. *Left outer-join* operator is a binary operator that retains the solution mappings from the *left* multisets of mapping tuples even if the given solution mappings do not satisfy the predicate function. Other outer-join operators, such as *right outer-join*, and *full outer-join* operators, can be derived from the left outer-join operator.

Definition 13. Given two multisets of mapping tuples, Γ_1 and Γ_2 , with $t_1 \in \Gamma_1, t_2 \in \Gamma_2, \Omega_1 \in t_1, \Omega_2 \in t_2, \mu_1 \in \Omega_1, \mu_2 \in \Omega_2$, and a predicate function, $\theta_{v_1}^{v_2}$. Similar to θ -join, **left outer-join** requires the application of a rename operator beforehand to ensure compatibility between the solution mappings. If Γ_1 and Γ_2 are incompatible, **left outer-join** keeps the mapping tuples from Γ_1 but drops everything from Γ_2 . If Γ_1 and Γ_2 are compatible, **Left outer-join** combines two multisets of mapping tuples, Γ_1 and Γ_2 based on the boolean condition after evaluating $\theta_{v_1}^{v_2}(\mu_1, \mu_2)$ as follows. If it is true, it behaves the same as the θ -join operator producing $\mu_1 \cup \mu_2$ for the associated mapping tuple t . Otherwise, only μ_1 is added to the mapping tuple t while μ_2 is dropped. Thus, **left outer-join** operator can be broken down into a taking a union of two steps: The union of the θ -join operator and the difference operator. The difference operator is used internally to define the left outer-join operator, similar to the definition of left-join in SPARQL algebra (Pérez et al., 2009). More formally, it is defined as follows.

$$\begin{aligned} \text{Difference}_\theta(\Omega_1, \Omega_2, \theta_{v_1}^{v_2}) &= (\Omega_1 \setminus \Omega_2) \cup \{ \mu_1 \mid \mu_1 \in \Omega_1, \exists \mu_2 \in \Omega_2, \theta_{v_1}^{v_2}(\mu_1, \mu_2) \text{ evaluates to false } \} \\ \text{Difference}_\theta(t_1, t_2, \theta_{v_1}^{v_2}) &= \{ (f, \text{Difference}_\theta(\Omega_1, \Omega_2, \theta_{v_1}^{v_2})) \mid \forall f \in \text{dom}(t_1) \cap \text{dom}(t_2), \\ &\quad \Omega_1 = t_1(f), \Omega_2 = t_2(f) \} \\ &\quad \cup \{ (f, \omega) \mid \forall f \in \text{dom}(t_1) \setminus \text{dom}(t_2), \omega = t_1(f) \} \\ \text{Difference}_\theta(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) &= \{ \text{Difference}_\theta(t_1, t_2, \theta_{v_1}^{v_2}) \mid t_1 \in \Gamma_1, t_2 \in \Gamma_2, t_1 \text{ and } t_2 \text{ are compatible. } \} \\ \text{Difference}(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) &= \{ t_1 \mid t_1 \in \Gamma_1, \forall t_2 \in \Gamma_2, t_1 \text{ and } t_2 \text{ are not compatible} \} \cup \\ &\quad \text{Difference}_\theta(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) \\ \text{LeftJoin}(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) &= \text{ThetaJoin}(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) \cup \text{Difference}(\Gamma_1, \Gamma_2, \theta_{v_1}^{v_2}) \end{aligned} \quad (13)$$

We provide the following three examples, where the three aforementioned join operators are applied on the mapping tuples shown in Tables 6 and 7 as Γ_1 and Γ_2 , respectively.

Example 6. Since natural join assumes solution mappings to have common variables, this example adjusts the solution mappings in Table 7 by renaming the variables of the solution mappings with the prefix “?pet.” The natural join operator

Table 7. Mapping Tuples Generated from Another Data Source About Pets.

Fragment	Multiset of Solution Mappings			
	Solution Mapping	?Type	?Name	?Age
f_{contacts}	μ_{a1}	Dog	Bax	10
f_{contacts}	μ_{a2}	Cat	Coco	3
f_{contacts}	μ_{a3}	Dog	Max	5

Table 8. Output of the Natural Join Operator as Described in Example 6.

Fragment	Multiset of Solution Mappings					
	Solution Mapping	?Fullname	?Pet.Type	?Pet.Name	?Firstname_iri	?Pet.Age
f_{contacts}	$\mu_1 \cup \mu_{a1}$	John Doe	Dog	Bax	!http://example.com/John?	10

Only the common variables $?pet.name$ and $?pet.type$ are checked.

Table 9. Output of θ -Join Operator as Described in Example 7. Only the Variables $?Pet.type$ and $?Animal_type$ are Checked for Equality.

Fragment	Multiset of Solution Mappings						
	Solution Mapping	?Fullname	?Pet.Type	...	?Animal_Age	?Animal_Type	?Animal_Name
f_{contacts}	$\mu_1 \cup \mu_{a1}$	John Doe	Dog	...	10	Dog	Bax
f_{contacts}	$\mu_1 \cup \mu_{a3}$	John Doe	Dog	...	5	Dog	Max

Table 10. Output of Left-Join Operator as Described in Example 8.

Fragment	Multiset of Solution Mappings						
	Solution Mapping	?Fullname	?Pet.Type	...	?Animal_Age	?Animal_Type	?Animal_Name
f_{contacts}	$\mu_1 \cup \mu_{a1}$	John Doe	Dog	...	10	Dog	Bax
f_{contacts}	$\mu_1 \cup \mu_{a3}$	John Doe	Dog	...	5	Dog	Max
f_{contacts}	μ_2	Susan Sue

Only the variables $?pet.type$ and $?animal_type$ are checked for equality and all solution mappings from Γ_1 are retained.

joins the mapping tuples from Table 6 and the adjusted Table 7 (renamed with the suffix “?pet.”), and it produces the output as shown in Table 8. The natural join merges solution mappings based on the value of the common variables. In the example, the common variables between the two different multisets of solution mappings are $?pet.name$ and $?pet.type$. Since only μ_1 and μ_{a1} have the same values for the variables $?pet.name$ and $?pet.type$, the output only contains $\mu_1 \cup \mu_{a1}$ and drops the other solution mappings.

Example 7. To satisfy the precondition that the variables in the solution mappings should not collide, we rename the variables in the solution mappings from Table 7 with a prefix string “animal_”. In practice, the renaming is done by using a rename operator right before the θ -join operator. Provided with the predicate function, $\theta_{?pet.type}^{?animal_type}$, for equality check on the variable $?pet.type$ and the variable $?animal_type$. θ -join operator, applied on Γ_1 and Γ_2 , produces the output in Table 9. Unlike the natural join operator, the θ -join operator joins the mapping solutions only if they satisfy the conditions of the provided predicate function: In this case, an equality check on the variables $?pet.type$ and $?animal_type$.

Example 8. Provided with the predicate function, $\theta_{?pet.type}^{?animal_type}$, for equality check on the variable $?pet.type$ and the variable $?animal_type$. Left outer-join operator, applied on Γ_1 and Γ_2 , produces the output in Table 10. In this example, the left outer-join retains the solution mapping μ_2 , even though it does not satisfy the predicate function. For solution mapping μ_1 , it produces the same result as Table 9.

4.7.4 Union. The aforementioned join operators have a limitation where they cannot merge mapping tuples without comparing the actual data values in the solution mappings. In order to collect mapping tuples from multiple operators, without data value comparisons, we need to define a **union** operator. The algebraic mapping **union** operator is based on SPARQL's *union* operation. It is defined as follows.

Definition 14. Given two multisets of mapping tuples, Γ_1 and Γ_2 . Union operator produces a new multiset containing mapping tuples from either Γ_1 or Γ_2 . More formally, it is defined as follows.

$$\begin{aligned} \text{Union}(t_1, t_2) &= \{(f, \Omega_1 \cup \Omega_2) \mid \forall f \in \text{dom}(t_1) \cap \text{dom}(t_2), \Omega_1 = t_1(f), \Omega_2 = t_2(f)\} \\ &\quad \cup \{(f_1, \omega) \in t_1 \mid f_1 \notin \text{dom}(t_2)\} \\ &\quad \cup \{(f_2, \omega) \in t_2 \mid f_2 \notin \text{dom}(t_1)\} \\ \text{Union}(\Gamma_1, \Gamma_2) &= \{\text{Union}(t_1, t_2) \mid t_1 \in \Gamma_1, t_2 \in \Gamma_2\} \end{aligned} \quad (14)$$

4.8 Serialize

The aforementioned operators process and transform the mapping tuples generated from heterogeneous data sources, they do not define how to process the mapping tuples to the target data format. **Serialize** operator enables the transformation of mapping tuples to the target data format. In order to keep the operator algebraic, the output of the serializer operator is a special mapping tuple containing a solution mapping with only a single variable *?serialized_output* containing the serialized data. The serialization of the mapping tuple, according to a serialization format, is achieved through the evaluation of a specific extend expression defined as *serializer expression*. One would also notice that due to the *multiset definitions* of the mapping tuples and solution mappings, duplicate outputs can be generated if the cardinality of a solution mapping is more than one. However, this can easily be rectified with the introduction of a deduplication operator as future work.

4.8.1 Serializer Expression Ψ_C . A serializer expression, $\Psi_C : \Omega \rightarrow T$, is an extend expression (Definition 7) that generates an RDF Literal containing the serialized data from a mapping tuple according to the serialization configuration C . In this work, we focus on the generation of KGs. Thus, the serialization configuration, C , for the serializer operator is the *quad patterns* (QPs) as defined in SPARQL⁷. Since the blank nodes labels may be created using the *extend* operator and bound to a variable, we do not require the functionality to handle blank nodes separately in the serializer operator. Thus, when the serializer operator evaluates Ψ_C on a mapping tuple t , it binds for each variable, $v \in V$, in the *quad pattern* of C , with (i.e., similar to the CONSTRUCT query from SPARQL⁸).

Definition 15. Provided with the serializer expression $\Psi_C : \Omega \rightarrow T$, C the serialization configuration, and a multiset of mapping tuples Γ . The **serialize** operator is defined using the extend and projection operator. It first applies the extend operator configured with $(v_{\text{serialized}}, \Psi_C) \in E$, $\text{Extend}(\Gamma, E)$. Afterwards, it projects the extended mapping tuples using the projection operator, $\text{Project}(\Gamma, P)$ configured with $P = \{v_{\text{serialized}}\}$. It is defined as follows.

$$\begin{aligned} C &= \text{QPs representing the triples or quads output} \\ \Psi_C &= \text{replace the variables in the QPs based on the input solution mapping } \mu \\ E &= \{(v_{\text{serialized}}, \Psi_C)\} \\ P &= \{v_{\text{serialized}}\} \\ \text{Serialize}(\Gamma, \Psi_C) &= \text{Project}(\text{Extend}(\Gamma, E), P) \end{aligned} \quad (15)$$

Note that, similar to the rename operator, the serialize operator is defined as a fixed chaining of a Projection operator (Section 4.3) after the Extend operator (Section 4.4). We defined the serialize operator to describe the execution of the serialization operation in one operator instead of two operators. This reduces the complexity and redundancy of the generated mapping plan.

Listing 2. Example QP configuration for the serialize operator.

```

1 ?firstname_iri <http://example.com/name> ?fullname ;
2 <http://example.com/petName> ?pet_name.

```

Table 11. Output of the Serialize Operator as Described in Example 9.

Fragment	Multiset of Solution Mappings	
	Solution Mapping	?Serialized_Output
f_{contacts}	μ_1	$\text{¡http://example.com/John¿ } \text{¡http://example.com/name¿ } \text{“John Doe”};$ $\text{¡http://example.com/petName¿ } \text{“Max”}.$

Example 9. *Provided with a QP configuration C as shown in Listing 2. Serialize operator evaluates the QP on each solution mappings from the mapping tuple and binds the variables $\{?firstname_iri, ?fullname, ?animal_name\}$ with the associated values. In this example, the variable $v_{\text{serialized}}$ is $?serialized_output$. The output of the serialize operator, applied on the input mapping tuples in Table 8, is shown in Table 11.*

4.9 Target

Finally, depending on the configuration of the data sink, the serialized data is written to heterogeneous data sinks such as files, websockets or Apache Kafka topics. In mapping algebra, the *fragments* of the mapping tuple determine where the associated solution mappings will be written to. **Target** operator writes the data value of the specified variable, $v_{\text{serialized}}$ from the solution mapping to a target data sink associated with a particular mapping tuple with the target fragment f_{target} . If the target sink does not exist or an error occurs during the process of writing the data to the sink, the default error handling procedure of the *target* operator is to stop the whole mapping process and shows the cause of the error. An optional configuration can be provided to the target operator to change the default error handling behaviour, such as silencing the errors to continue the mapping process as much as possible.

Definition 16. *Provided with a target fragment f_{target} , a target variable v_{target} , and a configuration of data sink T . **Target** operator process the mapping tuples, $t \in \Gamma$, by writing all the values $d = \mu(v_{\text{target}})$ to the data sink T , $\forall (f, \omega) \in t$ where $f = f_{\text{target}}$.*

$$\begin{aligned}
 \text{Target}(v_{\text{target}}, \omega, T) &= \{\text{write } d \text{ to data sink } T \mid \forall \mu \in \omega, (v, d) \in \mu, v = v_{\text{target}}\} \\
 \text{Target}(f_{\text{target}}, v_{\text{target}}, t, T) &= \{\text{Target}(v_{\text{target}}, \omega, T) \mid (f, \omega) \in t, f = f_{\text{target}}\} \\
 \text{Target}(f_{\text{target}}, v_{\text{target}}, \Gamma, T) &= \{\text{Target}(f_{\text{target}}, v_{\text{target}}, t, T) \mid \forall t \in \Gamma\}
 \end{aligned} \tag{16}$$

Example 10. *Given the input mapping tuples shown in Table 11 as Γ , and a configuration T specifying a file path $/\text{target}/\text{output}.\text{nt}$. Applying **Target** $(f_{\text{contacts}}, ?\text{serialized_output}, \Gamma, T)$ will write the serialized triples in $\mu_1(?serialized_output)$ to the file $/\text{target}/\text{output}.\text{nt}$ specified by target configuration T .*

5 Implementation

As a reference implementation utilizing the aforementioned algebraic mapping operators, we implemented an algebraic mapping *translator* and a proof-of-concept *engine*. The translator translates mapping rules in different mapping languages to a uniform mapping plan consisting of the algebraic mapping operators, while the proof of concept engine executes the mapping plan to generate RDF statements from heterogeneous data sources. As our mapping algebra extends the SPARQL algebra (and thus naturally aligns with query-based mapping languages), we choose RML and ShExML from the categories of *dedicated mapping languages* and *constraint-based languages*, respectively, to implement our translation algorithms described in Section 5.1 and Section 5.2. We specifically generate the mapping plan for the following versions: (i) RML v1.1.2 (Dimou et al., 2014)⁹, and (ii) ShExML v2020 (García-González et al., 2020).

We choose RML v1.1.2 as RML is the prevalent declarative mapping language (Van Assche et al., 2022) – as evidenced by its ongoing support via the W3C Knowledge Graph Construction Community Group¹⁰ – and the v1.1.2 version is mature with large implementation coverage¹¹. The more recent version of RML (Iglesias-Molina et al., 2023) is backwards compatible with the previous version¹², hence, we do not expect breaking changes. We choose ShExML v2020 as it is independent of the other major mapping language families used in this work

This selection thus shows that our mapping algebra covers the semantics of at least one mapping language from the 3 categories of mapping languages mentioned in Section 2: RML-based languages are represented by RML v1.1.2 to compare against ShExML, and SPARQL-based languages are represented by the algebra definitions on which we based our algebraic operators on. This way, we ensure that our approach is not dependent upon a single mapping language.

Our proof-of-concept is used to demonstrate current coverage and practical feasibility of our proposed mapping algebra, but is not exhaustive in its current form, specifically concerning the ShExML translation. Syntactic sugar such as *Query* declarations are not supported since it is used to make the ShExML document more human-readable and do not add or change mapping steps. We currently only support one ShExML transformation operation (string concatenation), and no joins. We do not support the ShExML v2020 join – currently known as *substitution*¹³ – since it was defined with the usage of both the *UNION* and the *JOIN* keywords without detailed clarification on the operational semantics (García-González et al., 2020).

The algebraic mapping translator is implemented in Rust¹⁴ and utilizes Sophia (Champin, 2020) as a library for handling RDF types. The translator¹⁵ is called “Algebraic Mapping Loom: Weaving Mapping Languages” (“AlgeMapLoom”, v0.4.0), and contains different modules to implement a mapping language translator. The decision to use Rust as the implementation language is to leverage Rust’s cross compilation capabilities to enable the usage of the translator code on multiple operating systems. Furthermore, Rust has a rich ecosystem of libraries to support the generation of bindings for multiple programming languages, enabling developers to use the translator engine’s API from within their code. The output of the AlgeMapLoom translator is a graph data structure that is compliant with the Graphviz¹⁶ format in DOT language¹⁷. For ease of implementing the algebraic mapping engine, we also provide a JSON output of the translated mapping plan together with the JSON schema¹⁸.

The proof-of-concept algebraic mapping engine¹⁹ is implemented in JavaScript, called “RMLWeaver-JS” (v0.1.1), to show that the translated mapping plan is mapping and programming language agnostic. For this work, we only support processing CSV files²⁰ with RMLWeaver-JS. As the mapping plan is source-independent – except for the extensible source operator – only supporting CSV files is sufficient to prove the working of our proof-of-concept. We used the reactive programming paradigm²¹ when implementing RMLWeaver-JS to ensure that input data is processed in a streaming manner, resulting in lower memory usage. The following sections give an overview of the respective algorithms used to translate RML and ShExML into a mapping plan consisting of algebraic mapping operators.

5.1 RML Translation

RML v1.1.2 describes how triples maps (TPs) are used to generate RDF statements. A TP is linked to a source – called *logical source* in RML – which provides the necessary data to generate the RDF statements described by the TP. Multiple TPs can have the same RML *logical source*. For translating a logical source to a *Source* operator (Section 4.2), we extract the iterators and fields related to a logical source from the TP, as described in the work on RML Fields (Delva et al., 2021). Then, all *triples maps* are grouped according to their associated *Source* operators. The grouped TPs are then iterated over individually and the associated *term maps*²² are translated into corresponding *Projection*, *Join*, *Extend*, *Serialize*.

For each term map (subject, predicate, and object maps), the *references*²³ are extracted as *projection attributes* to create projection operators. The created *Projection* operators are applied directly after the previously created *Source* operators. This results in a *partially projected* mapping plan, used throughout the rest of the algorithm to build the final mapping plan representing the mapping process described by the RML document.

Once the *Projection* operators are created, the TPs are partitioned into groups with and without referencing object map²⁴ to be further translated into sub-mapping plans *with* and *without* joins.

For TP with referencing object maps, a *Fragmenter* operator and a *Join* operator is created to join the partially projected mapping plan involving the *child* and *parent* TP for each referencing object map associated with the TP. The *Fragmenter* operator is created with the fragment mapping, $(f_{\text{default}}, f_{\text{join}})$, to broadcast the output of the previous operator to go to both the *Join* operator and the other downstream operator not involved in the join operation. If multiple *Join* operators must be created, the previously created *Fragmenter* operator is updated with new fragment mappings to broadcast the output to more operators. The type of the *Join* operator is determined by the presence of *join conditions*²⁵. If the join condition exists, a θ -*join* operator is created using the attributes specified by the child and the parent references. Otherwise, a *natural join* operator is created. For TP without referencing object maps, the aforementioned step for the *Join* operator creation is skipped.

Then, information about the term maps are utilized to create *extend expressions*. For example, a constant-valued term map with the term type IRI is translated to generate a nested *extend expression* (Definition 7) where an IRI data typing extend expression is applied to the return value of the constant value generating extend expression (e.g. $\text{irify}(\text{constant}(\text{value}))$). The new variable, v_{new} , to which the corresponding extend expression is bound to, is generated uniquely for each term map. The *Extend* operator is generated from the aforementioned pairs of variables and extend expressions, and applied after the previously created operator (i.e. either a *Join* or a *Projection* operator, depending on the presence of the referencing object maps).

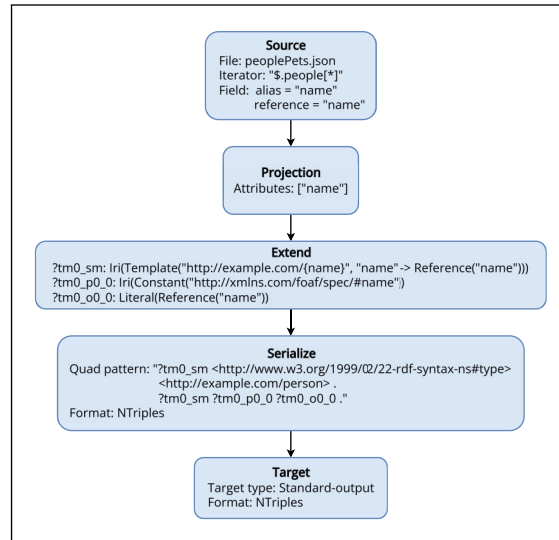


Figure 2. A Simple Mapping Plan Generated from the Mapping Process Described by the Resource Description Framework Mapping Language (RML) Document in Listing 3.

Finally, the *Serialize* operator is created based on the combination of term maps for the TP. Subject, predicate, object, and graph maps are used to generate *quad patterns*²⁶ with variables for each term.

The proposed RML *logical target* (Van Assche et al., 2021) is partially supported: Only the default logical target is interpreted by creating a default *Target* operator that pipes the generated RDF quads to the terminal's standard output, for all the mapping tuples having the *default fragment*.

5.2 ShExML Translation

Unlike RML, ShExML documents²⁷ have a structure split into two blocks: (i) *declarations* and (ii) *generators*. The *declarations* block contains individual lines defining *sources*, *iterators*, *prefixes*, and *expressions*. The declarations have the following structure: $jtype_i jvariable_i jstatement_i$. Each declaration is aliased with a *variable* which can be used within other declarations – introducing interdependency between different declarations. Thus, it is important to group related declarations together to generate our mapping plan. The *generators* block contains *shapes* and *graphs*, which in turn can contain nested shapes. The syntax for defining the graphs and shapes are the same with the ShEx specifications²⁸ with some modifications by ShExML.

First, unique combinations of the *source* and the *iterators* variables, used inside the ShExML *expressions*' statements, are used to generate the algebraic *Source* operators. One *Source* operator is generated for each unique combination of a source and an iterator variables. The generated *Source* operator is grouped with the variables of the expression definitions, which reference the same source as the *Source* operator. This results in pairs where every *Source* operator is paired with a set of expression variables. We shall annotate the set of expression variables as V_{expr} which will be referred to in the remaining algorithm steps.

For each pair of *Source* operator with V_{expr} , we (i) generate the RDF QPs that could be generated for the current *Source* operator and (ii) generate the relevant algebraic mapping operators.

RDF QPs are derived from the shapes and graphs in the *generator* block of ShExML document. These RDF QPs also contain metadata such as RDF data type or term type (IRI, Blank node, Literal) to aid in the generation of the value to be bound to the variables in the RDF QP. RDF QPs are generated, and added to the set of RDF QPs, if the subject node and the object node of a ShExML shape references one of the expression variables $v \in V_{\text{expr}}$. Furthermore, string templating extend expressions are generated for both the subject and the object nodes since the terms for these nodes need to be generated dynamically during the mapping process. The generated string templating extend expressions will be used later on for the configuration of the *Extend* operator. Predicate nodes in ShExML are predefined as a constant IRI. If there is a linking shape²⁹, and the expression variable used in the subject node of the nested shape is $v_{\text{subj}} \in V_{\text{expr}}$, an RDF QP is generated where the object term variable is the same as the subject term variable of the nested shape, and it is added to the set of RDF QPs. The generated set containing the RDF QPs are used later for the creation of the *Serialize* operator (Section 4.8).

Listing 3. Example RML document, with 1 subject and 1 predicate-object map, processing data from a JSON file. The input data is the JSON data shown in Listing 1.

```

1 @prefix rr: <http://www.w3.org/ns/r2rml#> .
2 @prefix rml: <http://semweb.mmlab.be/ns/rml{\xmlnum}> .
3 @prefix ql: <http://semweb.mmlab.be/ns/ql{\xmlnum}> .
4 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema{\xmlnum}> .
5 @base <http://example.com/base></http:> .
6 <TriplesMap1>
7   a rr:TriplesMap;
8   rml:logicalSource [
9     rml:source ‘‘peoplePets.json{\xmlquote}’’;
10    rml:referenceFormulation ql:JSONPath;
11    rml:iterator ‘‘XXXXXXXXXX.people[{\xmllast}]{\xmlquote}’’;
12   ];
13   rr:subjectMap [
14     rml:template ‘‘http://example.com/{name}{\xmlquote}’’;
15     rr:class <http://example.com/person>;
16   ];
17   rr:predicateObjectMap [
18     rr:predicate foaf:name;
19     rr:objectMap [ rml:reference ‘‘name{\xmlquote}’’ ];
20   ].

```

Once the RDF QPs are generated, we generate the relevant algebraic mapping operators based on the type of *expression declarations* with expression variable $v \in V_{\text{expr}}$. If the expression declarations are transformations, such as string operation³⁰, the relevant built-in extend expressions (Definition 7) are created. The extend expressions are paired with the ShExML expression variable to which the generated value will be bound. These pairs of variables and extend expressions are used to create an *Extend* operator which is applied on the *Source* operator currently being processed. This step is optional depending on the presence of the transformation expressions.

*Basic expressions*³¹ are translated after the transformation expressions. Basic expressions in ShExML behave just like a *Rename* operator (Section 4.5), where the values generated from the *statement_i* are aliased with the associated expression variable. Thus, a *Rename* operator is generated, with the rename pairs derived from the basic expressions whose $v \in V_{\text{expr}}$. The *Rename* operator is applied directly as the next step of the mapping plan. This step can also be optional depending on the presence of basic expressions.

Afterwards, the extend expressions (Definition 7) for type casting, derived from the metadata of the generated RDF QPs, are generated and paired with the corresponding variable in the RDF QP. For example, an IRI type casting statement is generated for a subject node and paired with the variable of the subject node in the RDF QP. The generated pairs of variables and extend expressions are used to create an *Extend* operator, with typecasting extend expressions. The created *Extend* operator is applied after the previous step.

Finally, the previously generated RDF QPs are used to generate the *Serialize* operator. Since ShExML can not specify targets, the default *Target* operator is created to pipe the generated RDF quads to the terminal’s standard output.

5.3 Mapping Plan

When applying the aforementioned algorithms to translate RML and ShExML documents, a mapping plan is generated. Figures 2 and 3 shows an example mapping plan generated from an RML document (Listing 3 and 4 respectively) using the algorithm described in Section 5.1, and Figure 4 shows the mapping plan generated from a ShExML document (Listing 5) using the algorithm described in Section 5.2.

5.3.1 RML Translation Result. When translating an RML document (Listing 3), the *Projection* operator projects the attributes referenced by the term maps in RML. The projected attribute is *name*, referenced in the subject map’s template and the object map’s reference. The *Extend* operator contains pairs of attributes and *extend expressions* for the generation of the terms required for the mapping process. In Figure 2, the *Extend* operator binds the variable *?tm0_sm* to the RDF generated by evaluating the extend expression *Iri(Template(“http://example.com/{name}”, “name” → Reference(“name”)))* resulting in an IRI with the string template, “http://example.com/name,” where the value for the variable “name” is retrieved from the current solution mapping being operated upon by the *Extend* operator. The QP, used by the *Serialize* operator for serializing the data into N-Triples, is derived from the usage of the subject and predicate object maps in

Listing 4. Example RML document where two triples maps are joined based on the attribute “name”. The input data is the JSON data shown in Listing 1.

```

1 @prefix rr: <http://www.w3.org/ns/r2rml#> .
2 @prefix rml: <http://semweb.mmlab.be/ns/rml#> .
3 @prefix ql: <http://semweb.mmlab.be/ns/ql#> .
4 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
5 @prefix foaf: <http://xmlns.com/foaf/spec/#> .
6 @prefix ex: <http://example.com></http:> .
7 @base <http://example.com/base></http:> .
8 <TriplesMap1>
9   a rr:TriplesMap;
10  rml:logicalSource [
11    rml:source “peoplePets.json”;
12    _rml:referenceFormulation ql:JSONPath;
13    _rml:iterator “$.people[*]”
14  ];
15  rr:subjectMap [
16    rr:template “http://example.com/{ name}”;
17    _rr:class <http://example.com/person>;
18  ];
19  _rr:predicateObjectMap [
20    _rr:predicate ex:hasPet;
21    _rr:objectMap [
22      _rr:parentTriplesMap
23      <TriplesMap2>;
24      _rr:joinCondition [
25        _rr:child “name”;
26        rr:parent “name”;
27      ];
28    ];
29  ];
30 <TriplesMap2>
31   a _rr:TriplesMap;
32   _rml:logicalSource [
33     rml:source “peoplePets.json”;
34     rml:referenceFormulation ql:JSONPath;
35     rml:iterator “$.people[*]”
36   ];
37   _rr:subjectMap [
38     rr:template “http://example.com/{ pet.name}”;
39     rr:class <http://example.com/pet>;
40   ];
41   rr:predicateObjectMap [
42     rr:predicate foaf:name;
43     rr:objectMap [ rml:reference “pet.name” ];
44   ].

```

the RML document. As there are no logical targets specified in the example RML document, the default *Target* operator is used: The generated N-Triples are piped to the standard output of the terminal.

Translating the RML document in Listing 3 does not produce the *Fragmenter* nor the binary operators such as θ -join operator. These operators are only present when the RML document describes a join operation between two TP such as the example RML document in Listing 4. The mapping plan generated from the RML document in Listing 4 contains a *Fragmenter* operator to fragment the mapping tuples to two downstream operator; the *Projection* operator, and the θ -join operator. The *Projection* operator after the *Fragmenter* operator, in the downstream path from Source A, ensures that the mapping tuples only contain attributes required by the rest of the downstream operators.

5.3.2 ShExML Translation Result. There are two major differences between the mapping plan generated by RML translation and the ShExML translation. First, ShExML translation generates a *Rename* operator as part of the mapping plan due to the basic expression on line 12 in Listing 5. Lastly, the *Extend* operator has an extra variable binding for the generation of a constant IRI term namely $http://example.com/person_i$. This arises from the ShExML translation algorithm (Section 5.2) which generates extend expressions for both the subject and the object nodes in ShExML’s shapes and graphs block. The other operators are semantically similar to those generated from the RML document.

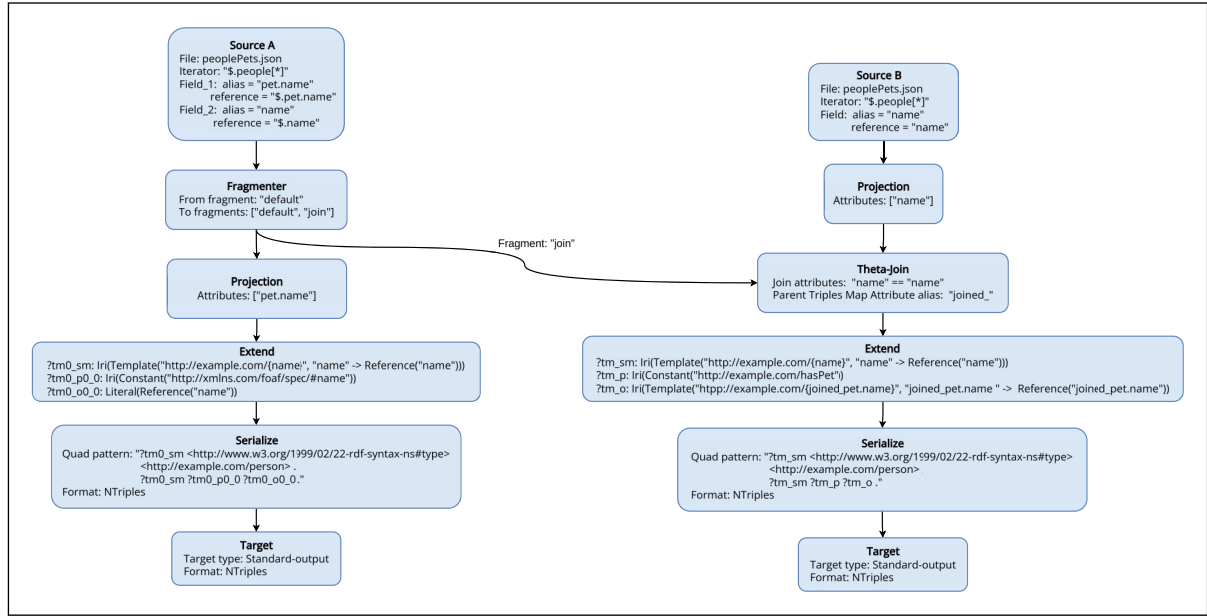


Figure 3. A Mapping Plan with Joins Generated from the Resource Description Framework Mapping Language (RML) Document in Listing 4.

6 Evaluation

This work introduces the definitions of algebraic mapping operators. We conduct an *empirical evaluation* of the algebraic mapping operators, using the aforementioned RML and ShExML mapping languages (Section 5). We implemented a reference algebraic mapping engine for the evaluation. Two types of empirical evaluation are carried out for this work: (i) Completeness of the defined algebraic mapping operators and (ii) the impact on performance of a mapping engine utilizing algebraic mapping operators. The first evaluation shows that this work is sufficient to create complete mapping engines. The second evaluation shows that utilizing algebraic mapping operators results in performance of real-world mapping engines comparable to the state of the art. Figure 5 shows an example execution pipeline consisting of AlgeMapLoom-rs and RMLWeaver-JS used to translate and execute RML document. The same pipeline setup is also used for ShExML evaluation.

6.1 Completeness of Algebraic Mapping Operators

Evaluating the completeness of the algebraic mapping operators is done by translating test cases, provided by the RML and ShExML reference implementations, into a mapping plan using the defined algebraic mapping operators. We then execute the generated mapping plan with our reference implementation, RMLWeaver-JS, and check the output of our implementation against the output of the reference implementations of RML and ShExML.

For RML, we use the RML v1.1.1 specification conformance test cases. Since RMLWeaver-JS only supports processing CSV files, we only chose the test cases using CSV files as input data source. The test cases for RML are available at the RMLWeaver-JS repository: <https://github.com/RMLio/rmlweaver-js/tree/v0.1.1/test/rml-mapper-test-cases-csv>.

For ShExML, the test cases provided with the reference implementation utilize heterogeneous data formats and sources such as a mix of JSON and XML or SPARQL endpoints. Since RMLWeaver-JS only supports input data files in CSV format, we adapted the test cases to utilize only CSV files as input. ShExML reference implementation's test cases evaluate multiple features of the ShExML language per test case. Therefore, we also split up the existing test cases into multiple smaller test cases to conduct a more granular evaluation of RMLWeaver-JS's execution of ShExML documents. For example, ShExML test case called *MultipleElementTest* tests for the usage of both multiple *Iterators*, and multiple *Basic Expressions* statements. We split it into two smaller test cases which evaluate the usage of multiple iterators and multiple basic expressions statements separately. This resulting set of test cases for ShExML are available at the RMLWeaver-JS repository: <https://github.com/RMLio/rmlweaver-js/tree/v0.1.1/test/shexml>.

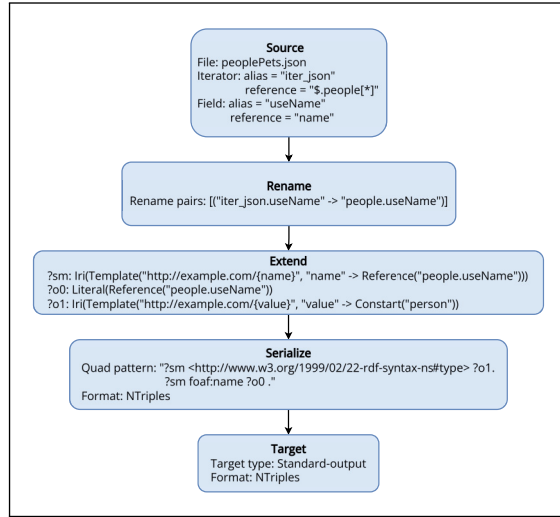


Figure 4. A Mapping Plan Generated from the ShExML Document in Listing 5.

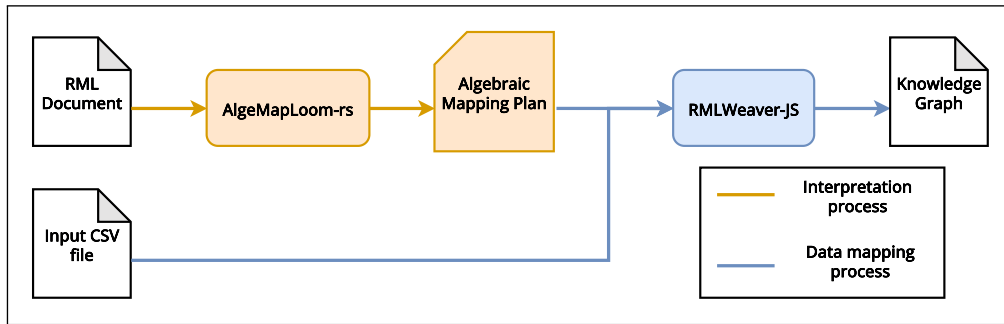


Figure 5. Algebraic Mapping Engine Pipeline where an RML Document is First Translated into an Algebraic Mapping Plan which is used to Generate the KG. RML: Resource Description Framework Mapping Language; KG: Knowledge Graph.

6.1.1 Results and Discussion. RMLWeaver-JS produces the same output as the reference RML implementation³² for all 39 out of the 39 RML CSV test cases (100%), covering 100% of the operational semantics of the RML CSV test cases (Table 12).

For the ShExML test cases, RMLWeaver-JS generates the same output as the reference v0.5.1 ShExML implementation³³.

Table 12 shows the number of features supported by RMLWeaver-JS in execution for both RML v1.1.1 and ShExML v2020. RMLWeaver-JS supports 16 out of 23 (65%) of ShExML features, and supports 11 out of 11 (100%) of RML v1.1.1 features. The exhaustive alignment of all ShExML features (such as Joins) to the introduced mapping algebra is future work. All testing code and results are published on GitHub³⁴.

6.2 Performance of An Algebraic Mapping Engine

To show that the implementation of an algebraic mapping engine does not have a large negative impact on the performance, we participated (Min Oo et al., 2024) in the first part of the performance track of the 2024 KG Construction Challenge³⁵. We highlight results of that participation in this paper, put it in context with the other participating mapping engines, and align the results with learnings concerning our proposed mapping algebra.

The performance track’s first part evaluates the mapping engines’ performance when handling diverse KG construction parameters with synthetic datasets, using RML mapping rules. Mapping engines are evaluated by changing the following properties of the data: The number of data records, data properties, duplicates, empty values, and input files. It also changes the following properties of the mapping rules: The number of subjects, predicates and objects, and finally, the number and type of joins used. For the measurements, the following metrics of the participating mapping engines are measured: (i) maximum RAM usage (GB), (ii) CPU usage (s), and (iii) execution time (s). The interpretation of CPU usage is as

Listing 5. Example ShExML document equivalent to the RML document in Listing 3.

```

1 PREFIX rr : <http://www.w3.org/ns/r2rml#>
2 PREFIX rml : <http://semweb.mmlab.be/ns/rml#>
3 PREFIX ql : <http://semweb.mmlab.be/ns/ql#>
4 PREFIX rdfs : <http://www.w3.org/2000/01/rdf-schema#>
5 PREFIX foaf : <http://xmlns.com/foaf/spec/#>
6 PREFIX : <http://example.com></http:>
7 SOURCE peoplePetsJSONFile <peoplePets.json>
8 ITERATOR iter_json <jsonpath: _XXXXXXXXX.people[*]> {
9   FIELD useName
10  <name>
11 }
12 EXPRESSION people <peoplePetsJSONFile.iter_json>
13 :Films :[people.useName] {
14   a :person ;
15   foaf:name [people.useName] ;
16 }

```

Table 12. Our Solution Supports 65% of ShExML v2020 Features and 100% of RML v1.1.1 Features.

ShExML v2020 Feature	Is Supported	RML v1.1.1 Feature	Is Supported
<i>Declarations</i>	7/8		2/2
Prefix	✓		
Source	✓	Logical source	✓
Query			
Iterator	✓	Logical iterator	✓
Nested iterator	✓		
Fields	✓		
Push fields	✓		
Pop fields	✓		
<i>Expressions</i>	3/7		2/2
Basic	✓		
Union	✓		
String operation	✓	Template-valued term maps	✓
Join		Referencing object map + join condition	✓
Matcher			
Autoincrement			
Dynamic function			
<i>Shapes & graphs</i>	6/8		7/7
Basic	✓	Reference-valued term maps	✓
Basic (constant)	✓	Constant-valued term maps	✓
Link shapes	✓	Referencing object map	✓
Matcher			
Datatypes static + dynamic	✓	Object map + datatype	✓
Langtype static + dynamic	✓	Object map + language tag	✓
Conditional + dynamic functions			
		Predicate map	✓
Graphs	✓	Graph map	✓

The features are aligned across the two languages in terms of their functionality in mapping heterogeneous data to RDF. For example, *string operations* in ShExML is equivalent to the usage of *template-valued term maps* in RML. RML: resource description framework Mapping Language; RDF: resource description framework.

follows: 100% CPU usage is achieved when the CPU usage time (in seconds) equals the product of the execution time and the number of CPU cores available on the machine. In our evaluation setup, 100% CPU usage is four times the execution time as our machine has 4 cores available.

There are 5 mapping engines, including RMLWeaver-JS, participating in the performance track of the challenge. The engines can be classified into two major groups: Those based on data processing frameworks (e.g. Apache Spark and Flink), and those without using data processing frameworks. RPT-Sansa (Stadler & Bin, 2024), Mapping-template (Scrocca et al., 2024), and RMLStreamer (Min Oo et al., 2022) are based on Apache Spark³⁶, Apache Velocity³⁷, and

Table 13. FlexRML is the Most Performant Engine When Constructing Knowledge Graph from Varying Triples Maps and Predicate-Object Maps.

Test Cases	Engines	Execution Time (s)	CPU Usage (s)	Peak RAM (GB)
1TM 15POM	Mapping-template	–	–	–
	FlexRML	6.54	6.81	0.47
	RMLWeaver-JS	11.26	13.21	0.54
	RPT-Sansa	43.25	133.92	4.50
	RMLStreamer	44.76	113.01	6.10
3TM 5POM	Mapping-template	–	–	–
	FlexRML	3.79	9.59	0.51
	RMLWeaver-JS	15.66	17.73	0.55
	RPT-Sansa	44.18	122.37	4.40
	RMLStreamer	43.52	116.28	6.06
15TM 1POM	Mapping-template	–	–	–
	FlexRML	6.34	18.02	0.46
	RMLWeaver-JS	42.57	46.65	0.55
	RPT-Sansa	48.68	99.80	3.99
	RMLStreamer	40.74	108.58	6.09

Note: Best performance highlighted in bold.

It performs best when the number of TMs is closer to the number of CPU cores on the machine. TM: triples map; RML: resource description framework Mapping Language; POM: predicate-object map.

Apache Flink³⁸ data processing frameworks respectively. The other 2 engines are FlexRML (Freund et al., 2024a, 2024b), implemented in C++, and our JavaScript implementation RMLWeaver-JS.

All the engines are evaluated on a virtual machine provided by the organizers, which has a standardized specification to ensure a fair evaluation. Each engine is provided with its own separate virtual machine for evaluation. The virtual machine has a 64 bit architecture, and it is configured with an Intel(R) Xeon(R) Gold 6161 CPU at 2.20GHz with 4 cores, 16765 MB of RAM memory, and 150 GB of storage space. The operating system of the machine is Ubuntu 22.04.03 LTS. The execution of the experiment is done using the tools provided by the challenge organizers (Van Assche et al., 2024), isolated via Docker container³⁹.

Since the results are significantly more verbose than the completeness evaluation in Section 6.1, we present our results and discuss their causes separately.

6.2.1 Results. The full results of the challenge for KG parameters can be found on Zenodo⁴⁰ by downloading the file *System-Results-Challenge-2024.zip*. These results are based on the submissions by the authors of their respective engines. Thus, we can not conclude whether the engine failed at executing the test case or the results are omitted from the list. We skip the presentation of the results for duplicates and empty values test cases, since RMLWeaver-JS does not deduplicate the generated triples nor handle empty values in the columns by ignoring them. We discuss the results of the other test cases and compare our performance against the other participants in the following paragraphs.

Table 13 shows the engines' performance when the number of TM and predicate-object maps (POM) in the RML document changes, while the input dataset size stays the same. For the other engines, there is no difference in CPU usage and execution time increase when compared to RMLWeaver-JS. Where for FlexRML execution time shortens and CPU usage maximizes for the test case with 3 TM and 5 POM, for RMLWeaver-JS execution time and CPU usage increases across the test cases. It even has a 4-fold increase for execution time and CPU usage for the test case with 15 TM and 1 POM.

Table 14 shows the results of the test cases to evaluate the performance impact by increasing the number of columns in the input CSV dataset. RMLStreamer and RMLWeaver-JS maintain constant memory usage, while for the other engines memory usage increases with the number of columns. RMLWeaver-JS maintains the lowest memory usage amongst engines for the 10 and 30 columns CSV data records test cases. For all engines, execution time and CPU usage increases with the number of columns. Mapping-template is the fastest in terms of execution time and lowest CPU usage. FlexRML is only slightly faster, by 0.08 seconds, than Mapping-template for the first test case with 1 column CSV data records.

Table 15 contains the results of the test case with increasing number of CSV data records. For the test case with 10K rows, FlexRML is the fastest engine (1.23 seconds) using the lowest memory (0.4 GB) while Mapping-template's CPU usage is the lowest at 0.55 seconds. Once the number of rows reaches 100k, Mapping-template becomes faster than FlexRML and uses noticeably less CPU time (8 times less than FlexRML), but memory usage increases with 0.44 GB

Table 14. Number of Properties in CSV Data Records has Little to no Impact on the Memory Usage of RMLWeaver-JS.

Test Cases	Engines	Execution Time (s)	CPU Usage (s)	Peak RAM (GB)
1 column	Mapping-template	5.71	1.73	2.10
	FlexRML	5.63	5.80	0.42
	RMLWeaver-JS	17.78	19.52	0.47
	RPT-Sansa	40.12	86.10	2.82
	RMLStreamer	37.11	93.08	6.09
10 columns	Mapping-template	15.26	3.73	3.22
	FlexRML	43.66	43.82	0.79
	RMLWeaver-JS	65.63	69.78	0.52
	RPT-Sansa	100.24	351.04	11.14
	RMLStreamer	168.16	425.24	6.11
30 columns	Mapping-template	39.34	7.80	5.22
	FlexRML	137.30	140.21	1.72
	RMLWeaver-JS	172.80	194.09	0.50
	RPT-Sansa	319.38	1203.27	5.37
	RMLStreamer	462.99	1311.78	6.15

Note: Best performance highlighted in bold.

RML: resource description framework Mapping Language.

Table 15. RMLWeaver-JS Manages to Keep Memory Usage Constant Around 0.5 GB with Increasing Input Data Size, While flexRML is the Fastest.

Test Cases	Engines	Execution Time (s)	CPU Usage (s)	Peak RAM (GB)
10K rows	Mapping-template	2.41	0.55	1.04
	FlexRML	1.23	1.33	0.40
	RMLWeaver-JS	2.56	3.34	0.48
	RPT-Sansa	33.06	97.43	1.68
	RMLStreamer	23.49	51.06	1.75
100K rows	Mapping-template	4.95	1.16	1.48
	FlexRML	8.28	8.40	0.46
	RMLWeaver-JS	13.51	14.76	0.49
	RPT-Sansa	48.35	152.83	5.12
	RMLStreamer	49.84	129.47	6.16
10M rows	Mapping-template	—	—	—
	FlexRML	943.34	963.50	11.82
	RMLWeaver-JS	1116.40	1249.97	0.54
	RPT-Sansa	1569.04	5909.89	6.75
	RMLStreamer	1768.58	6918.19	6.17

Note: Best performance highlighted in bold.

RML: resource description framework Mapping Language.

compared to the 10k rows test case. However, Mapping-template fails to produce any results for 10M rows of CSV data. Throughout the experiment, RMLWeaver-JS maintains a constant memory usage of approximately 0.5 GB even for the 10M rows test case. RMLStreamer uses the same amount of memory of around 6.1 GB for both 100k and 10M rows of CSV data. This is similar to the amount of memory RMLStreamer uses for the test cases in Table 14.

Table 16 provides the measurements of evaluating on the join related test cases where 100% of the data records are eligible to be joined. All engines maintain similar performance across test cases Join N-M based on $\min(N, M)$. For example, FlexRML has similar performance for test cases Join 5-5 and Join 5-10 across all metrics measured. Furthermore, FlexRML is the fastest engine across all join test cases with the lowest memory and CPU usage.

6.2.2 Discussion. Analysing the results presented in Section 6.2.1 reveals several potential improvements for implementing a more efficient algebraic mapping engine.

Compared to the other engines, RMLWeaver-JS exhibits an abnormal behaviour for the test cases where the number of TMs and POMs changes inversely in the RML document (Table 13). For example, RMLWeaver-JS is the only engine with a significant spike in execution time from *11.26 seconds* to *42.57 seconds* for the test cases 1TM 15POM, and 15TM

Table 16. Performance of the Mapping Engines, on a Test Case Join N-M, Depends on $\text{Min}(N, M)$. 100% of the Data Records are Eligible to be Joined in the Test Cases Presented.

Test Cases	Engines	Execution Time (s)	CPU Usage (s)	Peak RAM (GB)
Join 10-1	Mapping-template	–	–	–
	FlexRML	15.35	19.57	0.56
	RMLWeaver-JS	30.03	33.75	0.64
	RPT-Sansa	40.41	121.75	5.11
	RMLStreamer	66.98	222.97	6.36
Join 1-10	Mapping-template	–	–	–
	FlexRML	15.03	19.14	0.50
	RMLWeaver-JS	30.04	33.60	0.64
	RPT-Sansa	38.97	119.19	4.13
	RMLStreamer	60.14	195.16	6.36
Join 5-5	Mapping-template	–	–	–
	FlexRML	23.41	32.29	0.59
	RMLWeaver-JS	82.04	90.99	0.81
	RPT-Sansa	50.70	149.82	5.32
	RMLStreamer	109.65	403.17	6.36
Join 5-10	Mapping-template	–	–	–
	FlexRML	22.51	31.44	0.59
	RMLWeaver-JS	81.64	90.94	0.79
	RPT-Sansa	47.73	149.37	4.83
	RMLStreamer	120.97	421.70	6.35

Note: Best performance highlighted in bold.

RML: resource description framework Mapping Language.

IPOm respectively. This can be explained due to the manner in which AlgeMapLoom (Section 5) translates the RML document into the algebraic mapping plan used by RMLWeaver-JS. The test cases include multiple TMs, each with its own definition of a logical source, all referring to the same CSV data file and using the same iterator. Due to the lack of detection for semantically similar data sources, the interpreter generates a distinct source operator for each logical source definition identified. In the specific test case of 15 TMs and 1 POM, a total of 15 source operators are generated, with each TM associated with one of these source operators. RMLWeaver-JS is implemented in JavaScript without worker threads, thus inherently single-threaded and can only execute one task at a time. Thus, processing the CSV data file 15 times instead of just once is done sequentially. This results in a noticeable increase in both execution time and CPU usage.

For the test cases regarding increasing number of properties and records, RMLWeaver-JS and RMLStreamer managed to maintain constant memory even if the input data columns or rows increase. This consistent memory usage is due to the way both engines process data. On one hand, RMLStreamer – built on the Apache Flink stream processing framework – processes input CSV rows one at a time. Consequently, RMLStreamer maintains a constant memory usage of around 6.1 GB, even for inputs with a higher number of columns and records. On the other hand, RMLWeaver-JS, implemented based on reactive programming paradigm, also processes the input CSV records one at a time, leading to the same constant memory usage. The substantial difference in memory usage, with RMLWeaver-JS using around 0.5 GB and RMLStreamer using approximately 6 GB, is due to the overhead of the underlying Apache Flink framework. Apache Flink allocates a fixed amount of heap memory for the Java Virtual Machine on which RMLStreamer is executed. Thus, RMLWeaver-JS has a lower memory usage than RMLStreamer due to the implementation not relying on data processing frameworks.

As observed in Table 16, all engines demonstrate the same performance across two different test cases in Join N-M, depending on the $\text{min}(N, M)$. We attempt to provide an explanation for such behaviour for RMLWeaver-JS (Min Oo et al., 2024), however, the same conclusion cannot be extended to other engines since we lack the details of their implementations. The explanation is as follows. Join $N - M$ is a test case containing two sources S_n and S_m , where there are N records from S_n eligible to be joined with M records from S_m . RMLWeaver-JS employs a simple hash-join algorithm to join data from two different sources. It creates two hash maps – one for each source – for bookkeeping when joining the CSV records. Assuming the data are going to be joined on an attribute A , and provided $M < N$ with M records coming from S_m and N records coming from S_n . In order for RMLWeaver-JS to achieve the same performance as presented in this work for the joins, M records from S_m needs to arrive first at the join operator and be stored in the hash map $\text{HashMap}_{S_m}(A)$. This ensures that the amortized cost of joining the N records from S_n is lower since it only requires $\text{HashMap}_{S_m}(A)$ to be

looped through M times, where $M < N$. Otherwise, the amortized cost will be higher if N records from S_n arrives first, causing the $HashMap_{S_n}(A)$ to be looped through at least N times with $M < N$. The aforementioned explanation applies to both Join 5-5 and Join 10-5, where RMLWeaver-JS exhibits similar performance in terms of execution time, CPU usage, and memory usage.

Summarizing the results, RMLWeaver-JS achieved the second place for the Track 2 performance challenge in KG Construction Workshop, handing first place to FlexRML (Van Assche et al., 2024). This achievement shows that it is possible to implement a performant algebraic mapping engine in JavaScript for web browsers, potentially empowering web clients and servers with KGs generated from heterogeneous data sources.

7 Conclusion

In this paper, we presented a mapping language-independent mapping algebra consisting of algebraic mapping operators *Source*, *Projection*, *Extend*, *Rename*, *Fragmenter*, *Natural join*, θ -*join*, *Left outer-join*, *Union*, *Serialize*, and *Target*. We empirically showed how this mapping algebra provides (partial) operational semantics by translating mapping rules of two existing but very different mapping languages (RML and ShExML, 100% and 63% of feature coverage, respectively), to a mapping plan consisting of our introduced algebraic mapping operators. We showed practical feasibility of our approach via a proof-of-concept algebraic mapping engine, RMLWeaver-JS, achieving second place in the Knowledge Graph Construction Workshop’s performance challenge.

In our parallel work, we focussed on a provable theoretical foundation for core operators applied to RML v1.1.2, allowing us to prove equivalence classes, and thus, theoretical optimizations for RML v1.1.2 engines (Min Oo & Hartig, 2025). In this journal article, we focus on a more exhaustive set of language-agnostic operators, allowing us to build end-to-end language-agnostic engines.

For future work, we will investigate to which degree we can apply the theoretically founded operators and data model definitions of our parallel work into the more exhaustive data model and operators of this work, given the difference in scope and theoretical grounding (i.e., relational algebra in our parallel work versus SPARQL algebra in this work). The resulting mapping algebra could further be used to develop a formal mapping language with a formal grammar. Such a language will benefit from the operational semantics already defined by the algebra. To formalize the operators not investigated in our parallel work, the Algemaploom-rs output adhering to a specific JSON schema can serve as an inspiration.

We will exploit this mapping algebra for mapping process research: Translating mapping rules – in multiple mapping languages – into mapping plans conforming to our mapping algebra opens the door to static analysis of mapping rules, for example, for verification and optimization. After further completion of our proof-of-concept implementations, for example, researching the alignment of the join semantics in ShExML with our mapping algebra, we will investigate a *mapping plan optimizer* to improve the mapping process regardless of the mapping language used. We will start by considering existing optimizations, such as mapping partitions (Arenas-Guerrero et al., 2022) and mapping assertions (Iglesias et al., 2023), and we will exploit the theoretical equivalence classes found in our parallel work (Min Oo & Hartig, 2025) (i.e., projection pushing for intermediate result size reduction): These can be expanded towards our presented operators to have a more comprehensive set of optimization rules in our language-agnostic engine. We will present and benchmark these optimizations in a unified mapping algebra model.

With the advent of this mapping algebra and algebraic mapping engines, users are no longer locked into using a specific mapping language for KG generation. The performance across languages will become consistent by having an algebraic mapping engine that is multilingual, hence mapping language design can focus on functionality, decoupled from performance.


Funding


The authors received no financial support for the research, authorship and/or publication of this article.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

ORCID iDs

Sitt Min Oo  <https://orcid.org/0000-0001-9157-7507>

Ruben Taelman  <https://orcid.org/0000-0001-5118-256X>

Notes

1. <https://www.w3.org/community/kg-construct/>
2. RML v1.1.2 <https://rml.io/specs/rml/v/1.1.2/>
3. ShExML pushed and popped fields: <https://shexml.herminiojgarcia.com/spec/#pushed-and-popped-fields>
4. RML Logical Views: <https://github.com/kg-construct/rml-lv>.
5. Apache Kafka: <https://kafka.apache.org/>
6. For an explicit example, see <https://kg-construct.github.io/rml-io-registry/json-path/#natural-rdf-mapping>
7. SPARQL Quad Pattern: <https://www.w3.org/TR/sparql11-query/#rQuadPattern>
8. SPARQL Construct: <https://www.w3.org/TR/sparql11-query/#construct>
9. <https://rml.io/specs/rml/v/1.1.2/>
10. <https://www.w3.org/community/kg-construct/>
11. <https://rml.io/implementation-report/>
12. <https://kg-construct.github.io/rml-resources/portal/backwards-compatibility.html>
13. <https://shexml.herminiojgarcia.com/spec/#substitution>
14. <https://www.rust-lang.org/>
15. <https://github.com/RMLio/algemaploom-rs/releases/tag/v0.4.0>
16. <https://graphviz.org/>
17. <https://graphviz.org/doc/info/lang.html>
18. https://rmlio.github.io/algemaploom-rs/schema/schema_doc.html
19. <https://github.com/RMLio/rmlweaver-js/releases/tag/v0.1.1>
20. With some limitations in error handling for malformed CSV records and handling deduplication.
21. <https://rxjs.dev/>
22. RML term maps: <https://rml.io/specs/rml/#term-map>
23. RML reference: <https://rml.io/specs/rml/#reference>
24. Reference object map: <https://rml.io/specs/rml/#logical-join>
25. RML join conditions: <https://rml.io/specs/rml/#join-condition>
26. Quad Pattern: <https://www.w3.org/TR/sparql11-query/#rQuadPattern>
27. <https://github.com/herminiojg/ShExML>
28. <https://shex.io/shex-antics/>
29. ShExML linking shapes: <https://shexml.herminiojgarcia.com/spec/#linking-shapes>
30. <https://shexml.herminiojgarcia.com/spec/#string-operation>
31. <https://shexml.herminiojgarcia.com/spec/#basic-expression>
32. <https://github.com/RMLio/rmlmapper-java/releases/tag/v7.0.0>, <https://doi.org/10.5281/zenodo.11518178>
33. <https://github.com/herminiojg/ShExML/releases/tag/v0.5.1>
34. https://github.com/RMLio/rmlweaver-js/blob/v0.1.1/test/rml_tests.js
35. <https://kg-construct.github.io/workshop/2024>
36. <https://spark.apache.org/>
37. <https://velocity.apache.org/>
38. <https://flink.apache.org/>
39. Docker: <https://www.docker.com/>
40. <https://zenodo.org/doi/10.5281/zenodo.10721874>

References

- Arenas-Guerrero, J., Chaves-Fraga, D., Toledo, J., Pérez, M. S., & Corcho, O. (2022). Morph-KGC: Scalable knowledge graph materialization with mapping partitions. *Semantic Web*, 15, 1–20. <https://doi.org/10.3233/sw-223135>
- Asprino, L., Daga, E., Gangemi, A., & Mulholland, P. (2023). Knowledge graph construction with a façade: A unified method to access heterogeneous data sources on the web. *ACM Transactions on Internet Technology*, 23(1), 1–31.
- Bagaria, J. (2023). Set theory. In E. N. Zalta & U. Nodelman (Eds.), *The stanford encyclopedia of philosophy*, Spring 2023 edn. Metaphysics Research Lab, Stanford University.
- Bischof, S., Decker, S., Krennwallner, T., Lopes, N., & Polleres, A. (2012). Mapping between RDF and XML with XSPARQL. *Journal on Data Semantics*, 1(3), 147–185.
- Bizer, C., & Seaborne, A. (2004). D2RQ-treating non-RDF databases as virtual RDF graphs. In *Proceedings of the 3rd international semantic web conference (ISWC2004)*, Vol. 2004. Springer Hiroshima.
- Blizard, W. D. (1988). Multiset theory. *Notre Dame Journal of Formal Logic*, 30(1), 36–66. <https://doi.org/10.1305/ndjfl/1093634995>

- Champin, P.-A. (2020). Sophia: A linked data and semantic web toolkit for rust, Taipei, TW. <https://www2020devtrack.github.io/site/schedule>
- Chebotko, A., Lu, S., & Fotouhi, F. (2009). Semantics preserving SPARQL-to-SQL translation. *Data & Knowledge Engineering*, 68(10), 973–1000. <https://doi.org/https://doi.org/10.1016/j.datak.2009.04.001>. <https://www.sciencedirect.com/science/article/pii/S0169023X09000469>
- Chortaras, A., & Stamou, G. (2018). D2RML: Integrating heterogeneous data and web services into custom RDF graphs. In *LDOw@WWW*. <https://api.semanticscholar.org/CorpusID:51950275>
- Cygniak, R. (2012). Tarql: SPARQL for Tables, GitHub. <https://github.com/tarql/tarql>
- Daga, E., Asprino, L., Mulholland, P., & Gangemi, A. (2021). Facade-X: An opinionated approach to SPARQL anything. In *Further with knowledge graphs – proceedings of the 17th international conference on semantic systems, 6–9 September 2021, Amsterdam, The Netherlands, Studies on the Semantic Web*, Vol. 53 (pp. 58–73). IOS Press. ISSN 18681158, 22150870. <https://doi.org/10.3233/SSW210035>
- Delva, T., Assche, D. V., Heyvaert, P., Meester, B., & Dimou, A. (2021). Integrating nested data into knowledge graphs with RML fields. <https://www.semanticscholar.org/paper/Integrating-Nested-Data-into-Knowledge-Graphs-with-Delva-Assche/cfd3929eb7eb98209acea307838be4c9ddc4d33c>
- Dimou, A., Van der Sande, M., Colpaert, P., Verborgh, R., Mannens, E., & Van de Walle, R. (2014). RML: A generic language for integrated RDF mappings of heterogeneous data. In C. Bizer, T. Heath, S. Auer & T. Berners-Lee (Eds.), *Proceedings of the 7th workshop on linked data on the web. CEUR workshop proceedings*, Vol. 1184. CEUR. ISSN 16130073. http://ceur-ws.org/Vol-1184/ldow2014_paper_01.pdf
- Freund, M., Schmid, S., Dorsch, R., & Harth, A. (2024a). FlexRML: A flexible and memory efficient knowledge graph materializer. In A. Meroño Peñuela, A. Dimou, R. Troncy, O. Hartig, M. Acosta, M. Alam, H. Paulheim & P. Lisena (Eds.), *The semantic web, Springer Nature Switzerland, Cham* (pp. 40–56). ISBN 978-3-031-60635-9.
- Freund, M., Schmid, S., Dorsch, R., & Harth, A. (2024b). Performance results of FlexRML in the KGCW challenge 2024. In D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles & D. V. Assche (Eds.), *Proceedings of the 5th international workshop on knowledge graph construction co-located with 21th extended semantic web conference (ESWC 2024), Hersonissos, Greece, May 27, 2024, CEUR Workshop Proceedings* Vol. 3718. CEUR-WS.org. <https://ceur-ws.org/Vol-3718/paper9.pdf>
- García-González, H., Boneva, I., Staworko, S., Labra-Gayo, J. E., & Lovelle, J. M. C. (2020). ShExML: Improving the usability of heterogeneous data mapping languages for first-time users. *PeerJ Computer Science*, 6, e318.
- Gössner, S., Normington, G., & Bormann, C. (2024). JSONPath: Query expressions for JSON, *Request for Comments*, RFC Editor. <https://doi.org/10.17487/RFC9535>. <https://www.rfc-editor.org/info/rfc9535>
- Haesendonck, G., Maroy, W., Heyvaert, P., Verborgh, R., & Dimou, A. (2019). Parallel RDF generation from heterogeneous big data. In S. Groppe & L. Gruenwald (Eds.), *Proceedings of the international workshop on semantic big data - SBD '19, SBD '19, ACM Press, Amsterdam, Netherlands*. ISBN 978-1-4503-6766-0. <https://doi.org/10.1145/3323878.3325802>. <https://biblio.ugent.be/publication/8619808/file/8659668.pdf>
- Halmos, P. R. (1998). Naive set theory, undergraduate texts in mathematics. Springer New York. ISBN 9780387900926. <https://books.google.be/books?id=x6cZBQ9qtgoC>
- Iglesias-Molina, A., Cimmino, A., Ruckhaus, E., Chaves-Fraga, D., García-Castro, R., & Corcho, O. (2024). An ontological approach for representing declarative mapping languages. *Semantic Web*, 15(1), 191–221. <https://doi.org/10.3233/SW-223224>. <https://content.iospress.com/articles/semantic-web/sw223224>
- Iglesias-Molina, A., Van Assche, D., Arenas-Guerrero, J., De Meester, B., Debruyne, C., Jozashoori, S., Maria, P., Michel, F., Chaves-Fraga, D., & Dimou, A. (2023). The RML ontology: A community-driven modular redesign after a decade of experience in mapping heterogeneous data to RDF. In *Proceedings of the international semantic web conference (ISWC), Lecture Notes in Computer Science*. Springer, Cham (pp. 152–175). ISSN 1611-3349. ISBN 9783031472435. https://doi.org/10.1007/978-3-031-47243-5_9
- Iglesias, E., Jozashoori, S., Chaves-Fraga, D., Collarana, D., & Vidal, M.-E. (2020). SDM-RDFizer: An RML interpreter for the efficient creation of rdf knowledge graphs. In *Proceedings of the 29th ACM international conference on information & knowledge management*. ACM. <https://doi.org/10.1145/3340531.3412881>
- Iglesias, E., Jozashoori, S., & Vidal, M.-E. (2023). Scaling up knowledge graph creation to large and heterogeneous data sources. *Journal of Web Semantics*, 75, 100755. <https://doi.org/10.1016/j.websem.2022.100755>. <http://arxiv.org/abs/2201.09694>
- Lefrançois, M., Zimmermann, A., & Bakerally, N. (2017). A SPARQL extension for generating RDF from heterogeneous formats. In E. Blomqvist, D. Maynard, A. Gangemi, R. Hoekstra, P. Hitzler & O. Hartig (Eds.), *The semantic web 14th international conference, ESWC 2017, Portorož, Slovenia, May 28 – June 1, 2017, Proceedings*. Springer International Publishing, Portoroz, Slovenia (pp. 35–50). ISBN 978-3-319-58068-5. https://doi.org/10.1007/978-3-319-58068-5_3. <http://www.maxime-lefrancois.info/docs/LefrancoisZimmermannBakerally-ESWC2017-Generate.pdf>

- Lopes, N., Bischof, S., Decker, S., & Polleres, A. (2011). On the semantics of heterogeneous querying of relational, XML and RDF data with XSPARQL. In *Proceedings of the 15th Portuguese conference on artificial intelligence (EPIA 2011), Lisbon, Portugal* (pp. 10–13). Citeseer.
- Michel, F., Djimenou, L., Faron-Zucker, C., & Montagnat, J. (2015). Translation of heterogeneous databases into RDF, and application to the construction of a SKOS taxonomical reference. In *International conference on web information systems and technologies* (pp. 275–296). Springer. https://doi.org/10.1007/978-3-319-30996-5_14
- Min Oo, S., De Meester, B., Taelman, R., & Colpaert, P. (2023). Towards algebraic mapping operators for knowledge graph construction (p. 5). ISBN 978-3-031-47239-8.
- Min Oo, S., Haesendonck, G., De Meester, B., & Dimou, A. (2022). RMLStreamer-SISO: An RDF stream generator from streaming heterogeneous data. In U. Sattler, A. Hogan, M. Keet, V. Presutti, J. P. A. Almeida, H. Takeda, P. Monnin, G. Pirò & C. d'Amato (Eds.), *The semantic web – ISWC 2022, Springer International Publishing, Cham* (pp. 697–713). Springer. ISBN 978-3-031-19433-7. https://doi.org/10.1007/978-3-031-19433-7_40
- Min Oo, S., & Hartig, O. (2025). An algebraic foundation for knowledge graph construction. In *Proceedings of the 22nd extended semantic web conference (ESWC), Springer Nature Switzerland*, extend version available at <https://arxiv.org/abs/2503.10385>
- Min Oo, S., Verbeken, T., & De Meester, B. (2024). RMLWeaver-JS: An algebraic mapping engine in the KGCW challenge 2024. In D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles & D. V. Assche (Eds.), *Proceedings of the 5th international workshop on knowledge graph construction co-located with 21th extended semantic web conference (ESWC 2024), Hersonissos, Greece, May 27, 2024, CEUR workshop proceedings* Vol. 3718. CEUR-WS.org. <https://ceur-ws.org/Vol-3718/paper8.pdf>
- Pérez, J., Arenas, M., & Gutierrez, C. (2009). Semantics and complexity of SPARQL. *ACM Transactions on Database Systems*, 34(3), 1–45. <https://doi.org/10.1145/1567274.1567278>
- Priyatna, F., Corcho, O., & Sequeda, J. (2014). Formalisation and experiences of R2RML-based SPARQL to SQL query translation using morph. In *Proceedings of the 23rd international conference on world wide web, WWW '14, Association for Computing Machinery, New York, NY, USA* (pp. 479–490). ISBN 9781450327442. <https://doi.org/10.1145/2566486.2567981>
- Prud'hommeaux, E., Boneva, I., Labra Gayo, J. E., & Kellogg, G. (2018). Shape expressions language 2.1, draft community group report, world wide web consortium (W3C). <http://shex.io/shex-semantic/>
- Scrocca, M., Carenini, A., Grassi, M., Comerio, M., & Celino, I. (2024). Not everybody speaks RDF: Knowledge conversion between different data representations. In D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles & D. V. Assche (Eds.), *Proceedings of the 5th international workshop on knowledge graph construction co-located with 21th extended semantic web conference (ESWC 2024), Hersonissos, Greece, May 27, 2024, CEUR workshop proceedings* Vol. 3718. CEUR-WS.org. <https://ceur-ws.org/Vol-3718/paper3.pdf>
- Seaborne, A., & Harris, S. (2013). SPARQL 1.1 query language, W3C recommendation, W3C. <https://www.w3.org/TR/2013/REC-sparql11-query-20130321/>
- Simsek, U., Kärle, E., & Fensel, D. A. (2019). RocketRML - A NodeJS implementation of a use case specific RML mapper. *ArXiv abs/1903.04969*. <https://doi.org/10.48550/ARXIV.1903.04969>
- Stadler, C., & Bin, S. (2024). KGCW2024 challenge report: RDFProcessingToolkit. In D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles & D. V. Assche (Eds.), *Proceedings of the 5th international workshop on knowledge graph construction co-located with 21th extended semantic web conference (ESWC 2024), Hersonissos, Greece, May 27, 2024, CEUR Workshop Proceedings* Vol. 3718. CEUR-WS.org. <https://ceur-ws.org/Vol-3718/paper13.pdf>
- Stadler, C., Unbehauen, J., Westphal, P., Sherif, M. A., & Lehmann, J. (2015). Simplified RDB2RDF mapping. In *LDOW@WWW*. <https://api.semanticscholar.org/CorpusID:18692672>
- Sundara, S., Das, S., & Cyganiak, R. (2012). R2RML: RDB to RDF mapping language, W3C recommendation, W3C. <https://www.w3.org/TR/2012/REC-r2rml-20120927/>
- Unbehauen, J., Stadler, C., & Auer, S. (2013). Optimizing SPARQL-to-SQL rewriting. In *Proceedings of international conference on information integration and web-based applications & services, IIWAS '13, Association for Computing Machinery, New York, NY, USA* (pp. 324–330). ISBN 9781450321136. <https://doi.org/10.1145/2539150.2539247>
- Van Assche, D., Chaves-Fraga, D., Dimou, A., Serles, U., & Iglesias, A. (2024). KGCW 2024 Challenge @ ESWC 2024, Zenodo. <https://doi.org/10.5281/zenodo.11577087>
- Van Assche, D., Delva, T., Haesendonck, G., Heyvaert, P., De Meester, B., & Dimou, A. (2022). Declarative RDF graph generation from heterogeneous (semi-)structured data: A systematic literature review. *Journal of Web Semantics*, 75, 100753. <https://doi.org/10.1016/j.websem.2022.100753>

- Van Assche, D., Haesendonck, G., De Mulder, G., Delva, T., Heyvaert, P., De Meester, B., & Dimou, A. (2021). Leveraging web of things W3C recommendations for knowledge graphs generation. In *Web engineering, 21st international conference, ICWE 2021, Proceedings* (pp. 337–352). https://doi.org/10.1007/978-3-030-74296-6_26. <https://dylanvanassche.be/assets/pdf/icwe2021-wot-logical-target.pdf>
- Vu, B., Pujara, J., & Knoblock, C. A. (2019). D-REPR: A language for describing and mapping diversely-structured data sources to RDF. In *Proceedings of the 10th international conference on knowledge capture, K-CAP '19*. Association for Computing Machinery, New York, NY, USA (pp. 189–196). ISBN 9781450370080. <https://doi.org/10.1145/3360901.3364449>