

# Algebraic Mapping Operators for Knowledge Graph Generation

Sitt Min Oo<sup>a,\*</sup>, Ben De Meester<sup>a</sup>, Ruben Taelman<sup>a</sup> and Pieter Colpaert<sup>a</sup>

<sup>a</sup> *Department of Engineering and Architecture, University of Ghent – imec, Ghent, Belgium*

*E-mails: x.sittminoo@ugent.be, ben.demeester@ugent.be, ruben.taelman@ugent.be, pieter.colpaert@ugent.be*

**Abstract.** Recent advancements in declarative knowledge graph generation have led to the development of multiple mapping languages, their various versions, and different mapping engines that can interpret these languages and execute the mapping process. The field has progressed to the extent that current studies are now more focused on optimizing the knowledge graph generation process. Although different mapping engines share the common functionality of generating knowledge graphs from heterogeneous data sources, sharing the various optimization techniques and features of these engines remains challenging due to the lack of formal operational semantics for the general mapping processes. A set of algebraic mapping operators can provide the necessary operational semantics for general mapping processes, establish a theoretical foundation for mapping languages, and facilitate the introduction and evaluation of a compliant implementation, that is capable of interpreting and executing multiple mapping languages. In this paper, we propose such an algebra based on the SPARQL algebra. This allows us to maximally reuse established definitions, and further bridge the world of knowledge graph generation with query engines. To evaluate that our work is not limited to a single specific mapping language, we translated mapping languages ShExML and RML to our mapping plan composed of algebraic mapping operators. The results of our completeness evaluation shows that our algebraic operators cover the operational semantics of RML and partially for ShExML. To fully cover ShExML, further analysis into ShExML's concise operational semantics is needed (e.g. for joining data from two input sources). For performance evaluation, our proof-of-concept algebraic mapping engine has a consistent memory usage of around 500 MB across the different workloads, and achieved second place in the Knowledge Graph Construction Workshop's performance challenge. Algebraic mapping operators decouple mapping engines from the mapping languages, enabling multilingual mapping engines. Furthermore, the mapping plan can incorporate optimization techniques as a separate process from the mapping itself, allowing us to benefit from state-of-the-art mapping process optimizations. The proposed set of algebraic mapping operators will lay the foundation for future studies on the theoretical analysis of complexity and expressiveness of mapping languages, and will provide consistency in the execution semantics of mapping engines. Furthermore, the alignment of our algebra with SPARQL will enable further research into advanced methods such as virtualization, enabling heterogeneous data querying.

**Keywords:** Mapping Algebra, Semantic Web, Mapping Language, Knowledge Graph Generation

## 1. Introduction

There exist several *use-case-agnostic and declarative* mapping languages [1–7] to generate a Knowledge Graph (KG) using the Resource Description Framework (RDF) [8]. *Mapping rules* – described using such mapping language – specify the *mapping process*: how to generate a KG from existing semi-structured data [8]. These mapping languages' increasing popularity and importance is signified by the establishment of the W3C Knowledge

---

\*Corresponding author. E-mail: x.sittminoo@ugent.be.

Graph Construction Community Group, and a diverse ecosystem of mapping engines based on these mapping languages [8]. For example, several mapping engines based on the RDF Mapping Language (RML) [7] have been implemented [3, 9–13], and there are plans to establish a RML W3C Working Group in the future [14].

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 [15] or restricted to a particular mapping language [16, 17]. 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 [18] 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 [18] to formulate the mapping algebra allows us to maximally reuse established definitions, and closes the gap between KG *materialization* and KG *virtualization* (i.e. translating SPARQL to a data source query language such as SQL to provide on-the-fly KG generation [19]).

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 knowledge graph 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, e.g. D2RQ [20], R2RML [1], TARQL [21], and SML [2]. For example, R2RML only allows mapping relational databases to RDF using Turtle syntax [22], while TARQL maps only CSV data to RDF using SPARQL syntax [23].

On the other hand, *heterogeneous mapping languages* support multiple input data source types, i.e. multiple data formats (CSV, JSON, XML, etc. . .) and multiple data access (websocket, files, databases, etc. . .). Heterogeneous mapping languages can be further categorized [8]: 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, e.g. XSPARQL [24] combines XQuery and SPARQL, SPARQL-Generate [5] extends SPARQL with generation-specific operators, and SPARQL-Anything [17] applies extended SPARQL (CONSTRUCT) queries over an input data meta-model called Façade-X [25].

*Dedicated* mapping languages extend existing mapping languages or use custom syntax. RML [3, 7], xR2RML [26], and D2RML [27] 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 Group<sup>1</sup> for ongoing standardization [7]. D-REPR [4] is the only dedicated mapping language with its own syntax.

Currently, ShExML [6] is the only *constraint-based* language based on the ShEx [28] syntax with extensive modifications and with a focus on making a user-friendly language.

## 2.2. Formalizations

Several formalization works have been conducted to formalize the aforementioned languages. For *homogeneous mapping languages*, Stadler et al. [2] 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. [19] 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. [29] and improving the works of SparqlMap [30]. 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 [24] 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 [18]. For SPARQL-Anything, it formally describes the heterogeneous input data with their RDF meta-model Façade-X [25], and uses SPARQL to query the RDF meta-model [17], leading to SPARQL-Anything having provided formal operational semantics of its mapping process. SPARQL-Generate [5] 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 [31] extended RML's syntax to also allow mapping of nested heterogeneous data, provided an informal operational semantics of how it works. Iglesias et al. [15], and Arenas et al. [9] 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 on the formalizations 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 to introduce general operational semantics of a mapping process.

<sup>1</sup><https://www.w3.org/community/kg-construct/>

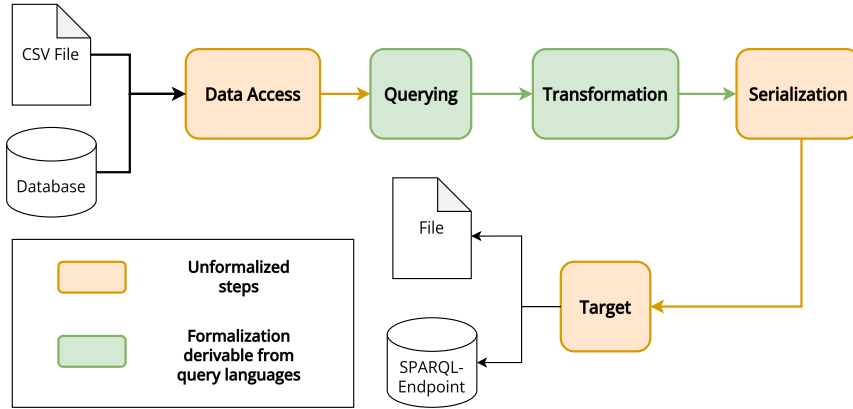


Figure 1. Breakdown of the mapping process: the *green* labelled steps can be derived from existing query algebra, and the *orange* labelled steps need to be formally defined. The exemplary Data Access step can read data from a relational database, and a CSV file while the exemplary Target step, writes the generated output to both a file and a SPARQL endpoint.

### 2.3. Discussion

We observe that existing formalizations are typically partial, in the case of query-based languages partially relying on SPARQL semantics, and exclusively applicable to a specific syntax: they do not provide a unified model for a generic mapping process independent of the mapping language. Currently, there is only a single work on providing a unified model for mapping languages through an ontology [14]. However, it is not formal and does not provide operational semantics. Hence, no KG generation operational semantics independent of the mapping language currently exists.

The lack of such operational semantics impedes using an optimization technique proposed in one mapping language in 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 [15] and Morph-KGC [9] employ their own concept of operational semantics for optimizing 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 followings steps: i) *data access*, ii) *data querying*, iii) *data transformations*, iv) *data serialization*, and v) *target output*. Figure 1 shows the breakdown of the mapping process.

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 [18, 23] 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 *newvalues* through operations such as string concatenation. An example for

a querying operator is the SPARQL's *projection* operator which we redefine 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*,  *$\theta$ -join*, *Left outer-join*, *Union*, *Serialize*, and *Target*. The algebra described in this section substantially extends our previous work [32], introduces seven more operators, and improves the mapping tuple definition. Since this work is inspired by SPARQL algebra, existing definitions and terms by Perez et al. [18] 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 [18]. 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*  $\mu$  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 noted as  $\Omega$  [18].

Two solution mappings  $\mu_1$  and  $\mu_2$  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  $\mu_1$  and  $\mu_2$ ,  $\mu_1 \cup \mu_2$ , is also a solution mapping. Given two multisets of solution mappings  $\Omega_1$  and  $\Omega_2$ , SPARQL algebra [18] defined the *join* ( $\bowtie$ ), the *union* ( $\cup$ ), and the *difference* ( $\setminus$ ) between  $\Omega_1$  and  $\Omega_2$  as follows [18]:

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

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

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

When writing sequences of the join and the union operators, parentheses can be avoided since both operators are associative and commutative [18]:

$$\Omega_1 \bowtie \Omega_2 \cup \Omega_3 = \Omega_1 \bowtie (\Omega_2 \cup \Omega_3) \quad (4)$$

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 [33], RML engines can export the generated RDF output to different data sinks based on the description provided in the 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* [32].

Table 1

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

Multiset of Solution Mappings				
Fragment	Solution Mapping	?name	?age	?email
$f_{personal}$	$\mu_1$	John Doe	23	john.doe@example.com
$f_{friends}$	$\mu_2$	Susan Sue	25	susan.sue@example.com
$f_{friends}$	$\mu_3$	Alice Joe	26	alice.joe@example.com

**Definition 1.** Let  $\Omega$  be a multiset of solution mappings. A **fragment**,  $f \in F$ , is a grouping of a submultiset [34] of  $\Omega$ . The set of fragments,  $F$ , is infinite and pairwise disjoint with the other sets 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  $\Omega$  be a multiset of solution mappings, and  $P(\Omega)$  the powerset of the multiset  $\Omega$  [34]. A **mapping tuple**,  $\xi$ , is a partial function which maps *fragments* to multisets of solutions mappings:  $\xi : F \rightarrow P(\Omega)$ . A multiset of mapping tuples is  $\Xi$ . Note, from now on, we shall use the notation  $\omega$ , 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,  $\omega$  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.

#### 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 [6] 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* fields<sup>2</sup>. RML Fields [31] 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 Views<sup>3</sup>.

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 [31].

**Definition 3.** An **iterator**,  $I$ , consists of one or more **fields**,  $\phi^I$ , an *iterator path*, and a *reference formulation*. An iterator extracts a list of records from the data source according to the given *iterator path* [31], while the *reference formulation* determines the data format [3] of the data source. This is the entry point to further extract nested data structures with *fields*.

**Definition 4.** **Fields**,  $\phi^I$ , consist of a *reference path*, a *name*, an optional *reference formulation* and an optionally nested one or more *subfields*,  $\phi_{sub}^I$ . The *reference path* is used to extract part of the data record generated from the iterator. The *name* of the field is used to provide an alias to the extracted part of the data record. Aligning with the

<sup>2</sup>ShExML pushed and popped fields: <https://shexml.hermiogarcia.com/spec/#pushed-and-popped-fields>

<sup>3</sup>RML Logical Views: <https://github.com/kg-construct/rml-lv>.

Table 2  
Representation of the mapping tuples generated by the source operator using data from Listing 1

Multiset of Solution Mappings						
Fragment	Solution Mapping	?name	?age	?email	?pet.type	?pet.name
$f_{default}$	$\mu_1$	John Doe	23	john.doe@example.com	dog	Bax
$f_{default}$	$\mu_2$	Susan Sue	25	susan.sue@example.com		

definition of solution mappings, the *name* of the field determines the variables of the solution mappings,  $v \in \text{dom}(\mu)$ . The associated value,  $t$ , is then assigned to the  $\text{range}(\mu)$ . 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 [31]. *Fields* can be nested with *subfields* to 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  $\phi_I$ . The **source** operator generates a multiset of mapping tuples,  $\xi \in \Xi$ , where a default fragment  $f_0$  is mapped to a multiset of solution mappings  $\omega$ . Data access configuration  $C_{access}$  consists of metadata information about utilizing the data source (e.g. connection ports and credentials for Kafka<sup>4</sup>). Iterators enable querying of the data source to extract data records and generate our solution mappings  $\mu$ . 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 is to align the data records with the definition of solution mappings  $\mu : V \rightarrow T$ .

$f_0 =$  a default fragment

$\mu =$  flattened data record iterated according to  $I$  and fields  $\phi^I$

$\omega =$  a multiset containing  $\mu$

(5)

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

To clarify the workings of the source operator with root and subiterators, Example 1 shows how a source operator generates mapping tuples from a simple JSON file with nested data.

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" : 23,
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  $I$ :

<sup>4</sup>Kafka: <https://kafka.apache.org/>

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
$f_{default}$	$\mu_1$	John Doe	dog	Bax
$f_{default}$	$\mu_2$	Susan Sue		

- $C$ : path to the JSON file.
- $I$ :  $\$.people[*]$  with reference formulation "JSONPath"

The iterator  $I$  also contains the following fields:

- $\phi_1^I$ : with name "?name" and reference \$.name
- $\phi_2^I$ : with name "?age" and reference \$.age
- $\phi_3^I$ : with name "?email" and reference \$.email
- $\phi_4^I$ : with name "?pet.type" and reference \$.pet.type
- $\phi_5^I$ : with name "?pet.name" and reference \$.pet.name

The iterators are using the JSONPath reference formulation<sup>5</sup>. 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.

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.

#### 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  $\mu$ , associated with the mapping tuple  $\xi$ , according to  $P$ .

$$\text{Project}(\mu, P) = \mu \text{ restricted to variables in } P$$

$$\text{Project}(\omega, P) = \{\text{Project}(\mu, P) \mid \forall \mu \in \omega\}$$

$$\text{Project}(\xi, P) = \{(f, \text{Project}(\omega, P)) \mid \forall (f, \omega) \in \xi\}$$

$$\text{Project}(\Xi, P) = \{\text{Project}(\xi, P) \mid \xi \in \Xi\}$$

(6)

**Example 2.** A projection operator, configured with a set of variables  $\{"?name", "?pet.name", "?pet.type"\} \in P$ , 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 (BMI) 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.

<sup>5</sup>JSONPath: <https://goessner.net/articles/JsonPath/>



Table 4  
Extended mapping tuple as described in Example 3

Multiset of Solution Mappings					
Fragment	Solution Mapping	?name	?pet.type	?pet.name	?firstname_iri
$f_{default}$	$\mu_1$	John Doe	dog	Bax	<http://example.com/John>
$f_{default}$	$\mu_2$	Susan Sue			<http://example.com/Susan>

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

$$\begin{aligned}
\text{Extend}(\mu, E) &= \mu \cup \{(v_{new}, value) \mid (v_{new}, expr) \in E, v_{new} \notin dom(\mu) \text{ and } value = expr(\mu)\} \\
\text{Extend}(\omega, E) &= \{\text{Extend}(\mu, E) \mid \forall \mu \in \omega\} \\
\text{Extend}(\xi, E) &= \{(f, \text{Extend}(\omega, E)) \mid \forall (f, \omega) \in \xi\} \\
\text{Extend}(\Xi, E) &= \{\text{Extend}(\xi, E) \mid \xi \in \Xi\}
\end{aligned} \tag{7}$$

**Example 3.** Provided  $\{(?firstname\_iri, irify(\_))\} \in E$ . 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.

**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  $\Xi$ , renames the variables of the solution mappings associated with the mapping tuples as follows: if  $\mu \in range(\xi)$ , for each  $(v_a, v_b) \in R$  rename  $v_a \rightarrow v_b$  if  $v_a \in dom(\mu)$ . 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  $\forall v \in dom(\mu)$  as  $s_{alias}||v$ .

$$\begin{aligned}
\text{Rename}(\mu, R) &= \{(v_b, d) \mid \forall (v_a, v_b) \in R, \forall (v, d) \in \mu, v = v_a\} \cup \\
&\quad \{(v, d) \mid \forall (v_a, v_b) \in R, \forall (v, d) \in \mu, v \neq v_a\} \\
\text{Rename}(\omega, R) &= \{\text{Rename}(\mu, R) \mid \forall \mu \in \omega\} \\
\text{Rename}(\xi, R) &= \{(f, \text{Rename}(\omega, R)) \mid \forall (f, \omega) \in \xi\} \\
\text{Rename}(\Xi, R) &= \{\text{Rename}(\xi, R) \mid \xi \in \Xi\}
\end{aligned} \tag{8}$$

$$\begin{aligned}
\text{Rename}(\mu, s_{alias}) &= \{(v_{aliased}, d) \mid \forall (v, d) \in \mu, v_{aliased} = s_{alias}||v\} \\
\text{Rename}(\omega, s_{alias}) &= \{\text{Rename}(\mu, s_{alias}) \mid \forall \mu \in \omega\} \\
\text{Rename}(\xi, s_{alias}) &= \{(f, \text{Rename}(\omega, s_{alias})) \mid \forall (f, \omega) \in \xi\} \\
\text{Rename}(\Xi, s_{alias}) &= \{\text{Rename}(\xi, s_{alias}) \mid \xi \in \Xi\}
\end{aligned} \tag{9}$$

Table 5  
Output of the rename operator as described in Example 4

Multiset of Solution Mappings					
Fragment	Solution Mapping	?fullname	?pet.type	?pet.name	?firstname_iri
$f_{default}$	$\mu_1$	John Doe	dog	Bax	<http://example.com/John>
$f_{default}$	$\mu_2$	Susan Sue			<http://example.com/Susan>

Readers would also 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 chaining to represent the rename operation is formally defined as follows.

$$\begin{aligned}
 R &= \{(v_{a1}, v_{b1}) \dots (v_{an}, v_{bn}) \mid n \in \mathbb{N}\} \\
 P_{renamed} &= \{v \mid v \in \text{dom}(\mu)\} \cup \{v_b \mid \forall (v_a, v_b) \in R\} / \{v_a \mid \forall (v_a, v_b) \in R\} \\
 E_{rename} &= \{(v_b, \text{copyData}(v_a)) \mid \forall (v_a, v_b) \in R\} \\
 \text{Rename}(\Xi, R) &= \text{Project}(\text{Extend}(\Xi, E_{rename}), P_{renamed})
 \end{aligned} \tag{10}$$

**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.

**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  $\delta$  on the mapping tuple  $\xi = f \rightarrow \mu$ , and map it into a new mapping tuple  $\xi_{mapped} = f_{new} \rightarrow \mu$  if  $\text{dom}(\delta) \subseteq \text{dom}(\xi)$  and  $f_{new} \in \text{range}(\delta)$ . When  $\text{dom}(\delta) \not\subseteq \text{dom}(\xi)$ , fragmenter operator acts like an identity function. More formally, it is defined as follows.

$$\begin{aligned}
 \delta &= \{(f_{old}, f_{new}) \mid f_{old}, f_{new} \in F\} \\
 \delta(\xi) &= \{(f, \omega) \mid (f, \omega) \in \xi, f \neq f_{old}\} \cup \{(f_{new}, \omega) \mid (f_{old}, \omega) \in \xi, (f_{old}, f_{new}) \in \delta\} \\
 \text{Fragment}(\xi, \delta) &= \begin{cases} \delta(\xi) & \text{if } \text{dom}(\delta) \subseteq \text{dom}(\xi) \\ \xi & \text{otherwise} \end{cases} \\
 \text{Fragment}(\Xi, \delta) &= \{\text{Fragment}(\xi, \delta) \mid \xi \in \Xi\}
 \end{aligned} \tag{11}$$

**Example 5.** Continuing with the output of the *extend* operator as shown in Table 4. A fragmenter operator with  $\delta : \{(f_{default}, f_{contacts})\}$  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,  $f_{contacts}$ .

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
$f_{contacts}$	$\mu_1$	John Doe	dog	Bax	<http://example.com/John>
$f_{contacts}$	$\mu_2$	Susan Sue			<http://example.com/Susan>

#### 4.7. Binary operators

Previously defined operators are unary: they only work on a single multiset of mapping tuples  $\Xi$ . To combine two multisets of mapping tuples  $\Xi_1$ , and  $\Xi_2$ , 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  *$\theta$ -join*, and the *left-join* between  $\Xi_1$  and  $\Xi_2$  as follows.

**Definition 10.** Two mapping tuples,  $\xi_1 \in \Xi_1$  and  $\xi_2 \in \Xi_2$ , are **compatible**, if and only if,  $\forall f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2)$ , for the associated multisets of solution mappings  $\xi_1(f) = \omega_1$  and  $\xi_2(f) = \omega_2$ ,  $\exists \mu_1 \in \omega_1, \exists \mu_2 \in \omega_2$  where  $\mu_1$  and  $\mu_2$  are *compatible* so that  $\omega_1 \bowtie \omega_2 \neq \emptyset$ . See Section 4.1 for solution mappings compatibility. If  $\text{dom}(\xi_1)$  and  $\text{dom}(\xi_2)$  are disjoint,  $\xi_1$  and  $\xi_2$  are compatible.

##### 4.7.1. Natural join

**Definition 11.** **Natural join** is a binary operator that combines two multisets of mapping tuples  $\Xi_1$  and  $\Xi_2$  if they are *compatible* (Definition 10). It produces mapping tuples,  $\xi_1 \bowtie \xi_2$ , which is a combination of two mapping tuples,  $\xi_1 \in \Xi_1$  and  $\xi_2 \in \Xi_2$  that have common fragments,  $\forall f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2)$ , for which the associated multisets of solution mappings  $\omega_1 = \xi_1(f)$  and  $\omega_2 = \xi_2(f)$  are joined as  $\omega_1 \bowtie \omega_2$  according to Equation 2. It is formally defined as follows.

$$\begin{aligned} \text{NatJoin}(\xi_1, \xi_2) &= \{(f, \omega_1 \bowtie \omega_2) \mid \forall f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2), \xi_1(f) = \omega_1, \xi_2(f) = \omega_2\} \\ \text{NatJoin}(\Xi_1, \Xi_2) &= \{\text{NatJoin}(\xi_1, \xi_2) \mid \xi_1 \in \Xi_1, \xi_2 \in \Xi_2, \xi_1 \text{ and } \xi_2 \text{ are compatible}\} \end{aligned} \quad (12)$$

Natural join operator joins mapping tuples if they have common *fragments*, and they 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 for the *equality* condition on the common variables and fragments: it does not use other predicate functions such as  $\leq$  or  $\geq$ .

##### 4.7.2. $\theta$ -join

*$\theta$ -join* enables the use of a predicate function,  $\theta$ , 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  *$\theta$ -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,  $\mu_a$  and  $\mu_b$  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  $\mu_a$  and  $\mu_b$  are compatible, the mapping plan planner should apply the rename operator before the  *$\theta$ -join* operator to ensure that none of variables in  $\mu_a$  and  $\mu_b$  are equal to each other. It is defined as follows.

**Definition 12.** Given a binary predicate function,  $v_1 \theta_{v_2} : \Omega \times \Omega \rightarrow \text{boolean}$ , 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  $\Xi_1$  and  $\Xi_2$ . Provided that  $\forall v_a \in \text{dom}(\mu_1), \forall v_b \in \text{dom}(\mu_2), v_a \neq v_b$ . Theta-join then combines two mapping tuples  $\xi_1 \in \Xi_1$ , and  $\xi_2 \in \Xi_2$ , if for the same fragment  $f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2)$ , the following condition is satisfied:  $\mu_1 \in \omega_1, \mu_2 \in \omega_2, v_1 \theta_{v_2}(\mu_1, \mu_2) = \text{true}$  with  $\omega_1 = \xi_1(f)$  and  $\omega_2 = \xi_2(f)$ .

$${}_{v_1 \theta_{v_2}}(\mu_1, \mu_2) = \begin{cases} true, & \text{if } \mu_1(v_1) = \mu_2(v_2) \\ false, & \text{otherwise} \end{cases}$$

$$\text{ThetaJoin}(\omega_1, \omega_2, {}_{v_1 \theta_{v_2}}) = \{\mu_1 \cup \mu_2 \mid \mu_1 \in \omega_1, \mu_2 \in \omega_2, {}_{v_1 \theta_{v_2}}(\mu_1, \mu_2) \text{ evaluates to true}\} \quad (13)$$

$$\text{ThetaJoin}(\xi_1, \xi_2, {}_{v_1 \theta_{v_2}}) = \{(f, \text{ThetaJoin}(\omega_1, \omega_2, {}_{v_1 \theta_{v_2}})) \mid \forall f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2), \\ \xi_1(f) = \omega_1, \xi_2(f) = \omega_2\}$$

$$\text{ThetaJoin}(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) = \{\text{ThetaJoin}(\xi_1, \xi_2, {}_{v_1 \theta_{v_2}}) \mid \xi_1 \in \Xi_1, \xi_2 \in \Xi_2\}$$

Natural and  $\theta$ -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,  $\Xi_1$  and  $\Xi_2$ , with  $\xi_1 \in \Xi_1, \xi_2 \in \Xi_2, \omega_1 \in \xi_1, \omega_2 \in \xi_2, \mu_1 \in \omega_1, \mu_2 \in \omega_2$ , and a predicate function,  ${}_{v_1 \theta_{v_2}}$ . Similar to  $\theta$ -join, **left outer-join** requires the application of a rename operator beforehand to ensure compatibility between the solution mappings. If  $\Xi_1$  and  $\Xi_2$  are *incompatible*, **left outer-join** keeps the mapping tuples from  $\Xi_1$  but drops everything from  $\Xi_2$ . If  $\Xi_1$  and  $\Xi_2$  are *compatible*, **Left outer-join** combines two multisets of mapping tuples,  $\Xi_1$  and  $\Xi_2$  based on the *boolean condition* after evaluating  ${}_{v_1 \theta_{v_2}}(\mu_1, \mu_2)$  as follows. If it is true, it behaves the same as the  $\theta$ -join operator producing  $\mu_1 \cup \mu_2$  for the associated mapping tuple  $\xi$ . Otherwise, only  $\mu_1$  is added to the mapping tuple  $\xi$  while  $\mu_2$  is dropped. Thus, **left outer-join** operator can be broken down into a taking a union of two steps: the union of the  $\theta$ -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 [18]. More formally, it is defined as follows.

$$\begin{aligned} \text{Difference}_\theta(\omega_1, \omega_2, {}_{v_1 \theta_{v_2}}) &= (\omega_1 \setminus \omega_2) \cup \{\mu_1 \mid \mu_1 \in \omega_1, \exists \mu_2 \in \omega_2, {}_{v_1 \theta_{v_2}}(\mu_1, \mu_2) \text{ evaluates to } false\} \\ \text{Difference}_\theta(\xi_1, \xi_2, {}_{v_1 \theta_{v_2}}) &= \{(f, \text{Difference}_\theta(\omega_1, \omega_2, {}_{v_1 \theta_{v_2}})) \mid \forall f \in \text{dom}(\xi_1) \cap \text{dom}(\xi_2), \\ &\quad \omega_1 = \xi_1(f), \omega_2 = \xi_2(f)\} \\ &\quad \cup \{(f, \omega) \mid \forall f \in \text{dom}(\xi_1) \setminus \text{dom}(\xi_2), \omega = \xi_1(f)\} \\ \text{Difference}_\theta(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) &= \{\text{Difference}_\theta(\xi_1, \xi_2, {}_{v_1 \theta_{v_2}}) \mid \xi_1 \in \Xi_1, \xi_2 \in \Xi_2, \xi_1 \text{ and } \xi_2 \text{ are compatible.}\} \\ \text{Difference}(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) &= \{\xi_1 \mid \xi_1 \in \Xi_1, \forall \xi_2 \in \Xi_2, \xi_1 \text{ and } \xi_2 \text{ are not compatible}\} \cup \\ &\quad \text{Difference}_\theta(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) \\ \text{LeftJoin}(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) &= \text{ThetaJoin}(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) \cup \text{Difference}(\Xi_1, \Xi_2, {}_{v_1 \theta_{v_2}}) \end{aligned} \quad (14)$$

We provide the following three examples, where the three aforementioned join operators are applied on the mapping tuples shown in Table 6 and Table 7 as  $\Xi_1$  and  $\Xi_2$  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

Table 7

Mapping tuples generated from another data source about pets

Multiset of Solution Mappings				
Fragment	Solution Mapping	?type	?name	?age
$\hat{f}_{contacts}$	$\mu_{a1}$	dog	Bax	10
$\hat{f}_{contacts}$	$\mu_{a2}$	cat	Coco	3
$\hat{f}_{contacts}$	$\mu_{a3}$	dog	Max	5

Table 8

Output of the natural join operator as described in Example 6. Only the common variables  $?pet.name$  and  $?pet.type$  are checked.

Multiset of Solution Mappings						
Fragment	Solution Mapping	?fullname	?pet.type	?pet.name	?firstname_iri	?pet.age
$\hat{f}_{contacts}$	$\mu_1 \cup \mu_{a1}$	John Doe	dog	Bax	<http://example.com/John>	10

Table 9

Output of  $\theta$ -join operator as described in Example 7. Only the variables  $?pet.type$  and  $?animal.type$  are checked for equality.

Multiset of Solution Mappings							
Fragment	Solution Mapping	?fullname	?pet.type	...	?animal_age	?animal_type	?animal_name
$\hat{f}_{contacts}$	$\mu_1 \cup \mu_{a1}$	John Doe	dog	...	10	dog	Bax
$\hat{f}_{contacts}$	$\mu_1 \cup \mu_{a3}$	John Doe	dog	...	5	dog	Max

join operator 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  $\mu_1$  and  $\mu_{a1}$  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  $\theta$ -join operator. Provided with the predicate function,  $?pet.type \theta ?animal.type$ , for equality check on the variable  $?pet.type$  and the variable  $?animal.type$ .  $\theta$ -join operator, applied on  $\Xi_1$  and  $\Xi_2$ , produces the output in Table 9. Unlike the natural join operator, the  $\theta$ -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,  $?pet.type \theta ?animal.type$ , for equality check on the variable  $?pet.type$  and the variable  $?animal.type$ . Left outer-join operator, applied on  $\Xi_1$  and  $\Xi_2$ , produces the output in Table 10. In this example, the left outer-join retains the solution mapping  $\mu_2$ , even though it does not satisfy the predicate function. For solution mapping  $\mu_1$ , 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,  $\Xi_1$  and  $\Xi_2$ . *Union* operator produces a new multiset containing mapping tuples from either  $\Xi_1$  or  $\Xi_2$ . More formally, it is defined as follows.

Table 10

Output of left-join operator as described in Example 8. Only the variables *?pet.type* and *?animal\_type* are checked for equality and all solution mappings from  $\Xi_1$  are retained.

Multiset of Solution Mappings							
Fragment	Solution Mapping	?fullname	?pet.type	...	?animal_age	?animal_type	?animal_name
<i>f_contacts</i>	$\mu_1 \cup \mu_{a1}$	John Doe	dog	...	10	dog	Bax
<i>f_contacts</i>	$\mu_1 \cup \mu_{a3}$	John Doe	dog	...	5	dog	Max
<i>f_contacts</i>	$\mu_2$	Susan Sue		...			

$$\begin{aligned}
 Union(\xi_1, \xi_2) = & \{(f, \omega_1 \cup \omega_2) \mid \forall f \in dom(\xi_1) \cap dom(\xi_2), \omega_1 = \xi_1(f), \omega_2 = \xi_2(f)\} \\
 & \cup \{(f_1, \omega) \in \xi_1 \mid f_1 \notin dom(\xi_2)\} \\
 & \cup \{(f_2, \omega) \in \xi_2 \mid f_2 \notin dom(\xi_1)\}
 \end{aligned} \tag{15}$$

$$Union(\Xi_1, \Xi_2) = \{Union(\xi_1, \xi_2) \mid \xi_1 \in \Xi_1, \xi_2 \in \Xi_2\}$$

#### 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. A serializer expression configured with  $C, \Psi_C : \Omega \rightarrow D$ , is provided to the serialize operator based on the desired output format.  $\Psi_C$  is the function which generates the data by using the input mapping tuples. Configuration on the data format specific serialization is given to the serializer operator in the form of  $C$ . Since we focus on the generation of knowledge graphs in this work, the configuration  $C$  for the serializer operator uses *quad patterns* (QPs) as defined in SPARQL<sup>6</sup>. 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,  $\Psi_C$  is, in this case, defined as a straightforward string templating function, which binds the variables in the *quad pattern* with values from the associated solution mappings in the input mapping tuples.

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

$$\begin{aligned}
 C &= \text{QPs representing the triples or quads output} \\
 \Psi_C &= \text{replace the variables in the QPs based on the input solution mapping } \mu \\
 E &= \{(v_{serialized}, \Psi_C)\} \\
 P &= \{v_{serialized}\}
 \end{aligned} \tag{16}$$

$$\text{Serialize}(\Xi, \Psi_C) = \text{Project}(\text{Extend}(\Xi, E), P)$$

<sup>6</sup>SPARQL Quad Pattern: <https://www.w3.org/TR/sparql11-query/#rQuadPattern>

Table 11  
Output of the serialize operator as described in Example 9.

Multiset of Solution Mappings		
Fragment	Solution Mapping	?serialized_output
$f_{contacts}$	$\mu_1$	<http://example.com/John> <http://example.com/name> "John Doe"; <http://example.com/petName> "Max".

Listing 2: Example QP configuration for the serialize operator

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

**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  $v_{serialized}$  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<sup>7</sup>. 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,  $v_{serialized}$  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,  $\xi \in \Xi$ , by writing all the values  $d = \mu(v_{target})$  to the data sink  $T$ ,  $\forall (f, \omega) \in \xi$  where  $f = f_{target}$ .

$$\begin{aligned}
 \text{Target}(v_{target}, \omega, T) &= \{\text{write } d \text{ to data sink } T \mid \forall \mu \in \omega, (v, d) \in \mu, v = v_{target}\} \\
 \text{Target}(f_{target}, v_{target}, \xi, T) &= \{\text{Target}(v_{target}, \omega, T) \mid (f, \omega) \in \xi, f = f_{target}\} \\
 \text{Target}(f_{target}, v_{target}, \Xi, T) &= \{\text{Target}(f_{target}, v_{target}, \xi, T) \mid \forall \xi \in \Xi\}
 \end{aligned} \tag{17}$$

**Example 10.** Given the input mapping tuples shown in Table 11 as  $\Xi$ , and a configuration  $T$  specifying a file path  $/target/output.nt$ . The **target**( $f_{contacts}, ?serialized\_output, \Xi$ ) will write the serialized triples in  $\mu_1(?serialized\_output)$  to the file  $/target/output.nt$ . If multiple solution mappings exist, the target operator will append the serialized data to the file.

## 5. Implementation

As a reference implementation utilizing the aforementioned algebraic mapping operators, we implemented an algebraic mapping *interpreter*, and a proof-of-concept *engine*. The interpreter translates mapping rules in different

<sup>7</sup>Apache Kafka: <https://kafka.apache.org/>

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. Since 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 v1.1.1 [3]<sup>8</sup>, and ii) ShExML [6] from 2020 (ShExML v2020).

We choose RML v1.1.1 as RML is the prevalent declarative mapping language [8] – as evidenced by its ongoing support via the W3C Knowledge Graph Construction Community Group<sup>9</sup> – and the v1.1.1 version is mature with large implementation coverage<sup>10</sup>. The more recent version of RML [7] is backwards compatible with the previous version<sup>11</sup>, 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 v1.1.1 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 language agnostic.

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 interpretation. 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*<sup>12</sup> – since it was defined with the usage of both the *UNION* and the *JOIN* keywords without detailed clarification on the operational semantics [6].

The algebraic mapping interpreter is implemented in Rust<sup>13</sup> and utilizes Sophia [35] as a library for handling RDF types. The interpreter<sup>14</sup> is called “Algebraic Mapping Loom: Weaving Mapping Languages” (“AlgeMapLoom”, v0.4.0), and contains different modules to implement a mapping language interpreter. The proof-of-concept algebraic mapping engine<sup>15</sup> 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 files<sup>16</sup> 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 employed the reactive programming paradigm<sup>17</sup> 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 v1.1.1 describes how triples maps are used to generate RDF statements. A triples map is linked to a source – called *logical source* in RML – which provides the necessary data to generate the RDF statements described by the triples map. Multiple triples maps 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 triples map, as described in the work on RML Fields [31].

<sup>8</sup><https://rml.io/specs/rml/v/1.1.1/>

<sup>9</sup><https://www.w3.org/community/kg-construct/>

<sup>10</sup><https://rml.io/implementation-report/>

<sup>11</sup><https://kg-construct.github.io/rml-resources/portal/backwards-compatibility.html>

<sup>12</sup><https://shexml.herminioarcar.com/spec/#substitution>

<sup>13</sup><https://www.rust-lang.org/>

<sup>14</sup><https://github.com/RMLio/algeMaploom-rs/releases/tag/v0.4.0>

<sup>15</sup><https://github.com/RMLio/rmlweaver-js/releases/tag/v0.1.1>

<sup>16</sup>With some limitations in error handling for malformed CSV records and handling deduplication.

<sup>17</sup><https://rxjs.dev/>



Then, all *triples maps* are grouped according to their associated *Source* operators. The grouped triples maps are then iterated over individually and the associated *term maps*<sup>18</sup> are translated into corresponding *Projection*, *Join*, *Extend*, *Serialize* operators as follows. For each term maps (subject, predicate, and object maps), the *references*<sup>19</sup> 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 triples maps are partitioned into groups with and without referencing object map<sup>20</sup> to be further translated into sub-mapping plans *with* and *without* joins.

For triples map 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* triples map for each referencing object map associated with the triples map. The *Fragmenter* operator is created with the fragment mapping,  $(f_{default}, f_{join})$ , 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*<sup>21</sup>. If the join condition exists, a  *$\theta$ -join* operator is created using the attributes specified by the child and the parent references. Otherwise, a *natural join* operator is created. For triples map without referencing object maps, the aforementioned step for the *Join* operator creation is skipped.

Then, information about the term maps are utilized to create *expression statements*. For example, a constant-valued term map with the term type IRI is translated to generate a nested *expression statement* (Definition 7) where an IRI data typing expression statement is applied to the return value of the constant value generating expression statement (e.g. *irify(constant(value))*). The new variable,  $v_{new}$ , to which the corresponding expression statement is bound to, is generated uniquely for each term map. The *Extend* operator is generated from the aforementioned pairs of variables and expression statements, 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).

Finally, the *Serialize* operator is created based on the combination of term maps for the triples map. Subject, predicate, object, and graph maps are used to generate *quad patterns*<sup>22</sup> with variables for each term.

The proposed RML *logical target* [33] 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 documents<sup>23</sup> 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: `<type> <variable> <statement>`. 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 specifications<sup>24</sup> with some modifications by ShExML.

First, unique combinations of the *source* and the *iterators* variables, used inside the *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

<sup>18</sup>RML term maps: <https://rml.io/specs/rml/#term-map>

<sup>19</sup>RML reference: <https://rml.io/specs/rml/#reference>

<sup>20</sup>Reference object map: <https://rml.io/specs/rml/#logical-join>

<sup>21</sup>RML join conditions: <https://rml.io/specs/rml/#join-condition>

<sup>22</sup>Quad Pattern: <https://www.w3.org/TR/sparql11-query/#rQuadPattern>

<sup>23</sup><https://github.com/herminiogg/ShExML>

<sup>24</sup><https://shex.io/shex-antics/>




Figure 2. A simple mapping plan generated from the straight forward mapping process described by the RML document in Listing 3

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 perform a two-step processing: i) generation of the RDF quad patterns that could be generated for the current *Source* operator, and ii) generation of the relevant algebraic mapping operators.

RDF quad patterns are derived from the shapes and graphs in the *generator* block of ShExML document. These RDF quad patterns 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 quad pattern. RDF quad patterns are generated, and added to the set of RDF quad patterns, if the subject node and the object node of a ShExML shape references one of the expression variables  $v \in V_{expr}$ . Predicate nodes in ShExML are predefined as a constant IRI. If there is a linking shape<sup>25</sup>, and the expression variable used in the subject node of the nested shape is  $v_{subj} \in V_{expr}$ , an RDF quad pattern 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 quad patterns. The generated set containing the RDF quad patterns are used later for the creation of the *Serialize* operator (Section 4.8).

Once the RDF quad patterns are generated, we generate the relevant algebraic mapping operators based on the type of *expression declarations* with expression variable  $v \in V_{expr}$ . If the expression declarations are transformations, such as string operation<sup>26</sup>, the relevant built-in expression statements (Definition 7) are created. The expression statements are paired with the ShExML expression variable to which the generated value will be bound. These pairs of variables and expression statements 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*<sup>27</sup> 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>` are aliased with the associated expression variable. Thus, a *Rename* operator is generated, with the rename pairs derived from the basic expressions whose  $v \in V_{expr}$ . 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 expression statements (Definition 7) for type casting, derived from the metadata of the generated RDF quad patterns, are generated and paired with the corresponding variable in the RDF quad pattern. 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 quad pattern. The generated pairs of variables and expression statements are used to create an *Extend* operator, with typecasting expression statements. The created *Extend* operator is applied after the previous step.

Finally, the previously generated RDF quad patterns 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

Once we apply the aforementioned algorithms to translate RML or ShExML documents, a mapping plan is generated. Figure 2 shows an example mapping plan that could be generated from RML document, Listing 3, using the algorithm described in Section 5.1. The *Projection* operator, in the figure, projects the attributes referenced by the term maps in RML. The projected attributes are *ID* and *Name*, which are referenced in the subject map's template

<sup>25</sup>ShExML linking shapes: <https://shexml.herminiojarcia.com/spec/#linking-shapes>

<sup>26</sup><https://shexml.herminiojarcia.com/spec/#string-operation>

<sup>27</sup><https://shexml.herminiojarcia.com/spec/#basic-expression>

Listing 3: Example RML document, with 1 subject and 1 predicate-object map, processing data from a CSV file.

```

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  @base <http://example.com/base/> .
6
7  <TriplesMap1>
8    a rr:TriplesMap;
9
10   rml:logicalSource [
11     rml:source "sport.csv";
12     rml:referenceFormulation ql:CSV
13   ];
14
15   rr:subjectMap [
16     rr:template "http://example.com/resource/sport_{ID}";
17     rr:class <http://example.com/ontology/Sport>;
18   ];
19
20   rr:predicateObjectMap [
21     rr:predicate rdfs:label ;
22     rr:objectMap [ rml:reference "Name" ];
23   ].

```

and the object map respectively. The *Extend* operator contains *expression statements* for type casting (IRIs, Literals) and string templating as specified by *rml:template* in the given example RML document. The QP (Listing 4), 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 the RML document. As there are no logical targets specified in the example RML document, the default *Target* operator is generated: the generated N-Triples are piped to the standard output of the terminal.

Listing 4: The QP derived from the term maps of RML document in Listing 3

```

1  ?sm      a      ?sm_class ;
2  rdfs:label ?om.

```

## 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 3 shows an example execution

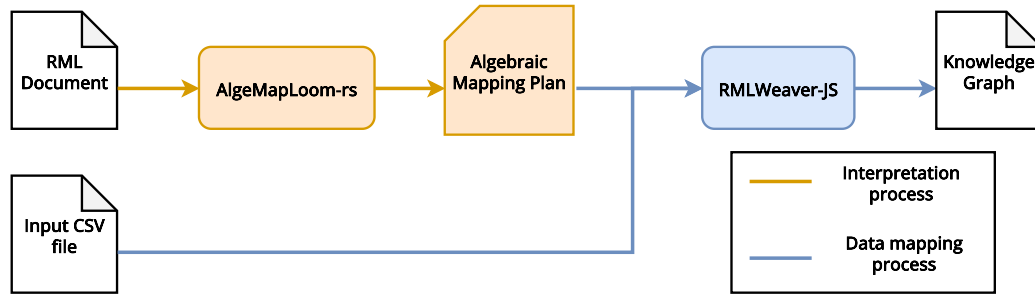


Figure 3. Algebraic mapping engine pipeline where an RML document is first translated into an algebraic mapping plan which is used to generate the KG.

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 evaluates 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 evaluates 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>.

#### 6.1.1. Results and discussion

RMLWeaver-JS produces the same output as the reference RML implementation<sup>28</sup> 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 adapted chosen ShExML test cases, RMLWeaver-JS generates the same output as the reference v0.5.1 ShExML implementation<sup>29</sup>. 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 13 out of 20 (65%) of ShExML features, and supports 11 out of 11 (100%) of RML v1.1.1 features. All testing code and results are published on GitHub<sup>30</sup>.

### 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 [36] in the first part of the performance track of the 2024 Knowledge Graph Construc-

<sup>28</sup><https://github.com/RMLio/rmlmapper-java/releases/tag/v7.0.0>, <https://doi.org/10.5281/zenodo.11518178>

<sup>29</sup><https://github.com/herminiogg/ShExML/releases/tag/v0.5.1>

<sup>30</sup>[https://github.com/RMLio/rmlweaver-js/blob/v0.1.1/test/rml\\_tests.js](https://github.com/RMLio/rmlweaver-js/blob/v0.1.1/test/rml_tests.js)

Table 12

Our solution supports 65% of ShExML v2020 features and 100% of RML v1.1.1 features. 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.

ShExML v2020 feature	Is supported	RML v1.1.1 feature	Is supported
<b>Declarations</b>	4/5		2/2
Prefix	✓		
Source	✓	Logical source	✓
Query			
Iterator	✓	Logical iterator	✓
Nested Iterator	✓		
<b>Expressions</b>	3/7		2/2
Basic	✓		
Union	✓		
String Operation	✓	Template-valued term maps	✓
Join		Referencing object map + join condition	✓
Matcher			
Autoincrement			
Dynamic Function			
<b>Shapes &amp; Graphs</b>	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	✓

tion Challenge<sup>31</sup>. We highlight results of that participation in this paper, extend on its descriptions, 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 knowledge graph 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 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 [37], Mapping-template [38],

<sup>31</sup><https://kg-construct.github.io/workshop/2024>

Table 13

FlexRML is the most performant engine when constructing knowledge graph from varying triples maps and predicate-object maps. It performs best when the number of triples maps (TM) is closer to the number of CPU cores on the machine.

Test cases	Engines	Execution time (s)	CPU usage (s)	Peak RAM (GB)
1TM 15POM	Mapping-template	-	-	-
	FlexRML	<b>6.54</b>	<b>6.81</b>	<b>0.47</b>
	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	<b>3.79</b>	<b>9.59</b>	<b>0.51</b>
	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	<b>6.34</b>	<b>18.02</b>	<b>0.46</b>
	RMLWeaver-JS	42.57	46.65	0.55
	RPT-Sansa	48.68	99.80	3.99
	RMLStreamer	40.74	108.58	6.09

and RMLStreamer [13] are based on Apache Spark<sup>32</sup>, Apache Velocity<sup>33</sup>, and Apache Flink<sup>34</sup> data processing frameworks respectively. The other 2 engines are FlexRML [39, 40], 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 an 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 [41], isolated via Docker container<sup>35</sup>.

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 knowledge graph parameters can be found on Zenodo<sup>36</sup> 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 triples map (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

<sup>32</sup><https://spark.apache.org/>

<sup>33</sup><https://velocity.apache.org/>

<sup>34</sup><https://flink.apache.org/>

<sup>35</sup>Docker: <https://www.docker.com/>

<sup>36</sup><https://zenodo.org/doi/10.5281/zenodo.10721874>

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	<b>1.73</b>	2.10
	FlexRML	<b>5.63</b>	5.80	<b>0.42</b>
	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	<b>15.26</b>	<b>3.73</b>	3.22
	FlexRML	43.66	43.82	0.79
	RMLWeaver-JS	65.63	69.78	<b>0.52</b>
	RPT-Sansa	100.24	351.04	11.14
	RMLStreamer	168.16	425.24	6.11
30 columns	Mapping-template	<b>39.34</b>	<b>7.80</b>	5.22
	FlexRML	137.30	140.21	1.72
	RMLWeaver-JS	172.80	194.09	<b>0.50</b>
	RPT-Sansa	319.38	1203.27	5.37
	RMLStreamer	462.99	1311.78	6.15

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	<b>0.55</b>	1.04
	FlexRML	<b>1.23</b>	1.33	<b>0.40</b>
	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	<b>4.95</b>	<b>1.16</b>	1.48
	FlexRML	8.28	8.40	<b>0.46</b>
	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	<b>943.34</b>	<b>963.50</b>	11.82
	RMLWeaver-JS	1116.40	1249.97	<b>0.54</b>
	RPT-Sansa	1569.04	5909.89	6.75
	RMLStreamer	1768.58	6918.19	6.17

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 16

Performance of the mapping engines, on a test case Join N-M, depends on  $\min(N, M)$ . 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	<b>15.35</b>	<b>19.57</b>	<b>0.56</b>
	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	<b>15.03</b>	<b>19.14</b>	<b>0.50</b>
	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	<b>23.41</b>	<b>32.29</b>	<b>0.59</b>
	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	<b>22.51</b>	<b>31.44</b>	<b>0.59</b>
	RMLWeaver-JS	81.64	90.94	0.79
	RPT-Sansa	47.73	149.37	4.83
	RMLStreamer	120.97	421.70	6.35

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 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 1Pom 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



1 a distinct source operator for each logical source definition identified. In the specific test case of 15 TMs and 1  
 2 POM, a total of 15 source operators are generated, with each TM associated with one of these source operators. 2  
 3 RMLWeaver-JS is implemented in JavaScript without worker threads, thus inherently single-threaded and can only 3  
 4 execute one task at a time. Thus, processing the CSV data file 15 times instead of just once is done sequentially. 4  
 5 This results in a noticeable increase in both execution time and CPU usage. 5

6 For the test cases regarding increasing number of properties and records, RMLWeaver-JS and RMLStreamer 6  
 7 managed to maintain constant memory even if the input data columns or rows increase. This consistent memory 7  
 8 usage is due to the way both engines process data. On one hand, RMLStreamer – built on the Apache Flink stream 8  
 9 processing framework – processes input CSV rows one at a time. Consequently, RMLStreamer maintains a con- 9  
 10 stant memory usage of around 6.1 GB, even for inputs with a higher number of columns and records. On the 10  
 11 other hand, RMLWeaver-JS, implemented based on reactive programming paradigm, also processes the input CSV 11  
 12 records one at a time, leading to the same constant memory usage. The substantial difference in memory usage, with 12  
 13 RMLWeaver-JS using around 0.5 GB and RMLStreamer using approximately 6 GB, is due to the overhead of the 13  
 14 underlying Apache Flink framework. Apache Flink allocates a fixed amount of heap memory for the Java Virtual 14  
 15 Machine on which RMLStreamer is executed. Thus, RMLWeaver-JS has a lower memory usage than RMLStreamer 15  
 16 due to the implementation not relying on data processing frameworks. 16

17 As observed in Table 16, all engines demonstrate the same performance across two different test cases in Join N- 17  
 18 M, depending on the  $\min(N, M)$ . We attempt to provide an explanation for such behaviour for RMLWeaver-JS [36], 18  
 19 however, the same conclusion cannot be extended to other engines since we lack the details of their implementations. 19  
 20 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 20  
 21 from  $S_n$  eligible to be joined with  $M$  records from  $S_m$ . RMLWeaver-JS employs a simple hash-join algorithm to 21  
 22 join data from two different sources. It creates two hash maps – one for each source – for bookkeeping when joining 22  
 23 the CSV records. Assuming the data are going to be joined on an attribute  $A$ , and provided  $M < N$  with  $M$  records 23  
 24 coming from  $S_m$  and  $N$  records coming from  $S_n$ . In order for RMLWeaver-JS to achieve the same performance as 24  
 25 presented in this work for the joins,  $M$  records from  $S_m$  needs to arrive first at the join operator and be stored in 25  
 26 the hash map  $HashMap_{S_m}(A)$ . This ensures that the amortized cost of joining the  $N$  records from  $S_n$  is lower since 26  
 27 it only requires  $HashMap_{S_m}(A)$  to be looped through  $M$  times, where  $M < N$ . Otherwise, the amortized cost will 27  
 28 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 28  
 29  $M < N$ . The aforementioned explanation applies to both Join 5-5 and Join 10-5, where RMLWeaver-JS exhibits 29  
 30 similar performance in terms of execution time, CPU usage, and memory usage. 30

31 Summarizing the results, RMLWeaver-JS achieved the second place for the Track 2 performance challenge in 31  
 32 Knowledge Graph Construction Workshop, handing first place to FlexRML [41]. This achievement shows that an 32  
 33 algebraic mapping engine implemented in a garbage-collected programming language like JavaScript is efficient in 33  
 34 terms of CPU usage, memory usage, and execution time. A performant algebraic mapping engine in JavaScript for 34  
 35 web browsers could potentially empower web clients and servers with knowledge graphs generated from heteroge- 35  
 36 neous data sources. 36

## 37 7. Conclusion 37

38 38  
 39 39  
 40 40  
 41 In this paper, we presented a mapping language-independent mapping algebra consisting of algebraic mapping 41  
 42 operators *Source*, *Projection*, *Extend*, *Rename*, *Fragmenter*, *Natural join*,  $\theta$ -*join*, *Left outer-join*, *Union*, *Serialize*, 42  
 43 and *Target*. We empirically showed how this mapping algebra provides (partial) operational semantics by translating 43  
 44 mapping rules of two existing but very different mapping languages (RML and ShExML, 100% and 63% of feature 44  
 45 coverage, respectively), to a mapping plan consisting of our introduced algebraic mapping operators. We showed 45  
 46 practical feasibility of our approach via a proof-of-concept algebraic mapping engine, RMLWeaver-JS, achieving 46  
 47 second place in the Knowledge Graph Construction Workshop’s performance challenge. 47

48 For future work, we will exploit this mapping algebra for theoretical research in mapping processes: translating 48  
 49 mapping rules – in multiple mapping languages – into mapping plans conforming to our mapping algebra opens 49  
 50 the door to static analysis of mapping rules, e.g. for verification and optimization. Specifically, we will investigate a 50  
 51 *mapping plan optimizer* to improve the mapping process independent of the used mapping language. We will start 51

by considering existing optimizations, such as mapping partitions [9] and mapping assertions [15], and present and benchmark them 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 knowledge graph 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.

## References

- [1] S. Sundara, S. Das and R. Cyganiak, R2RML: RDB to RDF Mapping Language, W3C Recommendation, W3C, 2012, <https://www.w3.org/TR/2012/REC-r2rml-20120927/>.
- [2] C. Stadler, J. Unbehauen, P. Westphal, M.A. Sherif and J. Lehmann, Simplified RDB2RDF Mapping, in: *LDOW@WWW*, 2015. <https://api.semanticscholar.org/CorpusID:18692672>.
- [3] A. Dimou, M. Vander Sande, P. Colpaert, R. Verborgh, E. Mannens and R. Van de Walle, RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data, in: *Proceedings of the 7<sup>th</sup> Workshop on Linked Data on the Web*, C. Bizer, T. Heath, S. Auer and T. Berners-Lee, eds, CEUR Workshop Proceedings, Vol. 1184, CEUR, 2014. ISSN 16130073. [http://ceur-ws.org/Vol-1184/ldow2014\\_paper\\_01.pdf](http://ceur-ws.org/Vol-1184/ldow2014_paper_01.pdf).
- [4] B. Vu, J. Pujara and C.A. Knoblock, D-REPR: A Language for Describing and Mapping Diversely-Structured Data Sources to RDF, in: *Proceedings of the 10<sup>th</sup> International Conference on Knowledge Capture, K-CAP '19*, Association for Computing Machinery, New York, NY, USA, 2019, pp. 189–196. ISBN 9781450370080. doi:10.1145/3360901.3364449.
- [5] M. Lefrançois, A. Zimmermann and N. Bakerally, A SPARQL Extension for Generating RDF from Heterogeneous Formats, in: *The Semantic Web 14<sup>th</sup> International Conference, ESWC 2017, Portoroz, Slovenia, May 28 – June 1, 2017, Proceedings*, E. Blomqvist, D. Maynard, A. Gangemi, R. Hoekstra, P. Hitzler and O. Hartig, eds, Springer International Publishing, Portoroz, Slovenia, 2017, pp. 35–50. ISBN 978-3-319-58068-5. doi:10.1007/978-3-319-58068-5\_3. <http://www.maxime-lefrancois.info/docs/LefrancoisZimmermannBakerally-ESWC2017-Generate.pdf>.
- [6] H. García-González, I. Boneva, S. Staworko, J.E. Labra-Gayo and J.M.C. Lovelle, ShExML: improving the usability of heterogeneous data mapping languages for first-time users, *PeerJ Computer Science* **6** (2020), e318.
- [7] A. Iglesias-Molina, D. Van Assche, J. Arenas-Guerrero, B. De Meester, C. Debruyne, S. Jozashoori, P. Maria, F. Michel, D. Chaves-Fraga and A. Dimou, 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, 2023, pp. 152–175. ISSN 1611-3349. ISBN 9783031472435. doi:10.1007/978-3-031-47243-5\_9.
- [8] D. Van Assche, T. Delva, G. Haesendonck, P. Heyvaert, B. De Meester and A. Dimou, Declarative RDF graph generation from heterogeneous (semi-)structured data: A systematic literature review, *Journal of Web Semantics* (2022). doi:10.1016/j.websem.2022.100753.
- [9] J. Arenas-Guerrero, D. Chaves-Fraga, J. Toledo, M.S. Pérez and O. Corcho, Morph-KGC: Scalable knowledge graph materialization with mapping partitions, *Semantic Web* (2022), 1–20. doi:10.3233/sw-223135.
- [10] E. Iglesias, S. Jozashoori, D. Chaves-Fraga, D. Collarana and M.-E. Vidal, SDM-RDFizer: An RML Interpreter for the Efficient Creation of RDF Knowledge Graphs, in: *Proceedings of the 29<sup>th</sup> ACM International Conference on Information & Knowledge Management*, ACM, 2020. doi:10.1145/3340531.3412881.
- [11] U. Simsek, E. Kärle and D.A. Fensel, RocketRML - A NodeJS Implementation of a Use Case Specific RML Mapper, *ArXiv abs/1903.04969* (2019). doi:10.48550/ARXIV.1903.04969.
- [12] G. Haesendonck, W. Maroy, P. Heyvaert, R. Verborgh and A. Dimou, Parallel RDF generation from heterogeneous big data, in: *Proceedings of the International Workshop on Semantic Big Data - SBD '19*, S. Groppe and L. Gruenwald, eds, SBD '19, ACM Press, Amsterdam, Netherlands, 2019. ISBN 978-1-4503-6766-0. doi:10.1145/3323878.3325802. <https://biblio.ugent.be/publication/8619808/file/8659668.pdf>.
- [13] Sitt Min Oo, G. Haesendonck, B. De Meester and A. Dimou, RMLStreamer-SISO: An RDF Stream Generator from Streaming Heterogeneous Data, in: *The Semantic Web – ISWC 2022*, U. Sattler, A. Hogan, M. Keet, V. Presutti, J.P.A. Almeida, H. Takeda, P. Monnin, G. Pirrò and C. d'Amato, eds, Springer International Publishing, Cham, 2022, pp. 697–713, Springer. ISBN 978-3-031-19433-7. doi:10.1007/978-3-031-19433-7\_40.
- [14] A. Iglesias-Molina, A. Cimmino, E. Ruckhaus, D. Chaves-Fraga, R. García-Castro and O. Corcho, An ontological approach for representing declarative mapping languages, *Semantic Web* **15**(1) (2024), 191–221, Publisher: IOS Press. doi:10.3233/SW-223224. <https://content.iiospress.com/articles/semantic-web/sw223224>.
- [15] E. Iglesias, S. Jozashoori and M.-E. Vidal, Scaling up knowledge graph creation to large and heterogeneous data sources, *Journal of Web Semantics* **75** (2023). doi:10.1016/j.websem.2022.100755. <http://arxiv.org/abs/2201.09694>.
- [16] N. Lopes, S. Bischof, S. Decker and A. Polleres, 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, Citeseer, 2011, pp. 10–13.
- [17] L. Asprino, E. Daga, A. Gangemi and P. Mulholland, Knowledge graph construction with a façade: a unified method to access heterogeneous data sources on the web, *ACM Transactions on Internet Technology* **23**(1) (2023), 1–31.

- [18] J. Pérez, M. Arenas and C. Gutierrez, Semantics and Complexity of SPARQL, *ACM Trans. Database Syst.* **34**(3) (2009). doi:10.1145/1567274.1567278.
- [19] F. Priyatna, O. Corcho and J. Sequeda, 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, 2014, pp. 479–490. ISBN 9781450327442. doi:10.1145/2566486.2567981.
- [20] C. Bizer and A. Seaborne, D2RQ-treating non-RDF databases as virtual RDF graphs, in: *Proceedings of the 3rd international semantic web conference (ISWC2004)*, Vol. 2004, Springer Hiroshima, 2004.
- [21] R. Cyganiak, Tarql: SPARQL for Tables, GitHub, 2012. <https://github.com/tarql/tarql>.
- [22] E. Prud'hommeaux and G. Carothers, RDF 1.1 Turtle, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/2014/REC-turtle-20140225/>.
- [23] A. Seaborne and S. Harris, SPARQL 1.1 Query Language, W3C Recommendation, W3C, 2013, <https://www.w3.org/TR/2013/REC-sparql11-query-20130321/>.
- [24] S. Bischof, S. Decker, T. Krennwallner, N. Lopes and A. Polleres, Mapping between RDF and XML with XSPARQL, *Journal on Data Semantics* **1**(3) (2012), 147–185.
- [25] E. Daga, L. Asprino, P. Mulholland and A. Gangemi, 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, IOS Press, 2021, pp. 58–73. ISSN 18681158, 22150870. doi:10.3233/SSW210035.
- [26] F. Michel, L. Djiméno, C. Faron-Zucker and J. Montagnat, 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*, Springer, 2015, pp. 275–296. doi:10.1007/978-3-319-30996-5\_14.
- [27] A. Chortaras and G. Stamou, D2RML: Integrating Heterogeneous Data and Web Services into Custom RDF Graphs, in: *LDOW@WWW*, 2018. <https://api.semanticscholar.org/CorpusID:51950275>.
- [28] E. Prud'hommeaux, I. Boneva, J.E. Labra Gayo and G. Kellogg, Shape Expressions Language 2.1, Draft Community Group Report, World Wide Web Consortium (W3C), 2018. <http://shex.io/shex-semantic/>.
- [29] A. Chebotko, S. Lu and F. Fotouhi, Semantics preserving SPARQL-to-SQL translation, *Data Knowledge Engineering* **68**(10) (2009), 973–1000. doi:<https://doi.org/10.1016/j.datak.2009.04.001>. <https://www.sciencedirect.com/science/article/pii/S0169023X09000469>.
- [30] J. Unbehauen, C. Stadler and S. Auer, 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, 2013, pp. 324–330. ISBN 9781450321136. doi:10.1145/2539150.2539247.
- [31] T. Delva, D.V. Assche, P. Heyvaert, B. Meester and A. Dimou, Integrating Nested Data into Knowledge Graphs with RML Fields, 2021. <https://www.semanticscholar.org/paper/Integrating-Nested-Data-into-Knowledge-Graphs-with-Delva-Assche/cfd3929eb7eb98209acea307838be4c9ddc4d33c>.
- [32] Min Oo, Sitt and De Meester, Ben and Taelman, Ruben and Colpaert, Pieter, Towards algebraic mapping operators for knowledge graph construction, 2023, p. 5. ISBN 978-3-031-47239-8.
- [33] D. Van Assche, G. Haesendonck, G. De Mulder, T. Delva, P. Heyvaert, B. De Meester and A. Dimou, Leveraging Web of Things W3C Recommendations for Knowledge Graphs Generation, in: *Web Engineering, 21st International Conference, ICWE 2021, Proceedings*, 2021, pp. 337–352. doi:10.1007/978-3-030-74296-6\_26. <https://dylanvanassche.be/assets/pdf/icwe2021-wot-logical-target.pdf>.
- [34] W.D. Blizard, Multiset Theory, *Notre Dame Journal of Formal Logic* **30**(1) (1988), 36–66. doi:10.1305/ndjfl/1093634995.
- [35] P.-A. Champin, Sophia: A Linked Data and Semantic Web toolkit for Rust, Taipei, TW, 2020. <https://www2020devtrack.github.io/site/schedule>.
- [36] S.M. Oo, T. Verbeke and B.D. Meester, RMLWeaver-JS: An algebraic mapping engine in the KGCW Challenge 2024, in: *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*, D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles and D.V. Assche, eds, CEUR Workshop Proceedings, Vol. 3718, CEUR-WS.org, 2024. <https://ceur-ws.org/Vol-3718/paper8.pdf>.
- [37] C. Stadler and S. Bin, KGCW2024 Challenge Report: RDFProcessingToolkit, in: *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*, D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles and D.V. Assche, eds, CEUR Workshop Proceedings, Vol. 3718, CEUR-WS.org, 2024. <https://ceur-ws.org/Vol-3718/paper13.pdf>.
- [38] M. Scrocca, A. Carenini, M. Grassi, M. Comerio and I. Celino, Not Everybody Speaks RDF: Knowledge Conversion between Different Data Representations, in: *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*, D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles and D.V. Assche, eds, CEUR Workshop Proceedings, Vol. 3718, CEUR-WS.org, 2024. <https://ceur-ws.org/Vol-3718/paper3.pdf>.
- [39] M. Freund, S. Schmid, R. Dorsch and A. Harth, FlexRML: A Flexible and Memory Efficient Knowledge Graph Materializer, in: *The Semantic Web, A. Meroño Peñuela, A. Dimou, R. Troncy, O. Hartig, M. Acosta, M. Alam, H. Paulheim and P. Lisena, eds, Springer Nature Switzerland, Cham, 2024*, pp. 40–56. ISBN 978-3-031-60635-9.
- [40] M. Freund, S. Schmid, R. Dorsch and A. Harth, Performance Results of FlexRML in the KGCW Challenge 2024, in: *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*, D. Chaves-Fraga, A. Dimou, A. Iglesias-Molina, U. Serles and D.V. Assche, eds, CEUR Workshop Proceedings, Vol. 3718, CEUR-WS.org, 2024. <https://ceur-ws.org/Vol-3718/paper9.pdf>.
- [41] D. Van Assche, D. Chaves-Fraga, A. Dimou, U. Serles and A. Iglesias, KGCW 2024 Challenge @ ESWC 2024, Zenodo, 2024. doi:10.5281/zenodo.11577087.

- 1 [42] R. Cyganiak, A relational algebra for SPARQL, *Digital Media Systems Laboratory HP Laboratories Bristol. HPL-2005-170* **35**(9) (2005). 1
- 2 [43] E. Prud'hommeaux and C.B. Aranda, SPARQL 1.1 Federated Query, W3C Recommendation, W3C, 2013, 2  
3 <https://www.w3.org/TR/2013/REC-sparql11-federated-query-20130321/>. 3
- 4 [44] D. Dell'Aglío, A. Polleres, N. Lopes and S. Bischof, Querying the Web of data with XSPARQL 1.1, *CEUR Workshop Proceedings* **1268** 4  
(2014), 113–118.
- 5 [45] D. Chaves-Fraga, F. Priyatna, A. Cimmino, J. Toledo, E. Ruckhaus and O. Corcho, GTFS-Madrid-Bench: A bench- 5  
6 mark for virtual knowledge graph access in the transport domain, *Journal of Web Semantics* **65** (2020), 100596. 6  
7 doi:<https://doi.org/10.1016/j.websem.2020.100596>. <https://www.sciencedirect.com/science/article/pii/S1570826820300354>. 7
- 8 [46] M. Scrocca, M. Comerio, A. Carenini and I. Celino, Turning Transport Data to Comply with EU Standards While Enabling a Multimodal 8  
9 Transport Knowledge Graph, *The Semantic Web – ISWC 2020* (2020), 411–429. ISBN 9783030624668. doi:10.1007/978-3-030-62466- 9  
8\_26.
- 10 [47] B. De Meester, T. Seymoens, A. Dimou and R. Verborgh, Implementation-independent function reuse, *Future Generation Com- 10  
11 puter Systems* **110** (2020), 946–959. doi:<https://doi.org/10.1016/j.future.2019.10.006>. [https://www.sciencedirect.com/science/article/pii/ 11  
12 S0167739X19303723](https://www.sciencedirect.com/science/article/pii/S0167739X19303723). 12

13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
511  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51