RML-view-to-CSV: A Proof-of-Concept Implementation for RML Logical Views

Els de Vleeschauwer[1], Pano Maria[2], Ben De Meester[1] and Pieter Colpaert[1]

[1]*IDLab, Dept. Electronics & Information Systems, Ghent University – imec, Belgium*
[2]*Skemu, Schiedam, The Netherlands*

Abstract

Although the W3C Community Group on Knowledge Graph Construction (KGC)'s work on the modular RDF Mapping Language (RML) specification has taken great strides, open issues and respective solution proposals remain. Some of these issues are (i) inability to handle hierarchy in nested data, (ii) limited join functionality, and (iii) inability to handle mixed data formats. To combat these issues, the RML Logical Views module is proposed. However, proper but efficient validation of this module requires an implementation that allows short development cycles. In this workshop paper, we propose a proof-of-concept RML Logical Views implementation, independent of and complementary to existing RML mapping engines. Our proof-of-concept covers three important features of the new RML Logical Views module: (i) flattening of nested data, (ii) extended joining of data sources, and (iii) handling mixed data formats. Our implementation supports one nested source format (JSON) and one tabular source format (CSV), and can be used independently, as preprocessor, by any RML Engine. With this implementation, we successfully executed the available relevant test cases of the RML Logical Views module. Additionally, we measured the knowledge graph construction times on GTFS-Madrid-Bench. When we included our implementation in the knowledge graph construction pipeline and replaced referencing object maps by joins in RML Logical Views, we noticed considerable execution time reductions. We conclude that the RML Logical Views specification can be implemented, and can solve needs that were not yet solvable by RML. The current implementation can already be realized as a modular part of a knowledge graph construction process. Although boosting performance was not the aim of our work, our implementation reduces the execution time of GTFS-Madrid-Bench scale 100 by 35% when combined with Morph-KGC or Carml. RMLStreamer, when used alone, times out after two hours on this task, but, in conjunction with our implementation, completes it in 236 seconds. We hope this proof-of-concept inspires the developers of existing RML engines to integrate the RML Logical Views module and benefit from its features.

Keywords

RML Logical View, flattening, joining, mixed content, proof-of-concept

1. Introduction

The W3C Community Group on Knowledge Graph Construction (KGC)[1] works on a declarative approach to construct RDF graph data from existing, heterogeneous data sources. The group recently proposed a new modular specification, ontology, and accompanying SHACL shapes for the RDF Mapping Language (RML)[2], including novel features which increase its expressiveness and empowers practitioners to define mapping rules for constructing RDF graph data that were previously unattainable [1]. Nevertheless, challenges such as handling hierarchy of nested data, more flexible joining (also across data hierarchies), and handling data sources that mix source formats (e.g., a table that contains a column storing data as JSON) remain unsolved.

As of July 2023, a dedicated task force of the W3C Community Group on KGC works on an additional RML module: RML Logical Views[3]. This module aims to resolve the aforementioned challenges by allowing to specify a *logical view*: a flattened, source format-agnostic view over one or more existing data sources.

RML Logical views are new in RML, and still under development. The specification[4] is not finalized and still subject to change. Hence, there are no implementations available to validate the feasibility of the theoretical concepts. Proper but efficient validation of the RML Logical Views module requires an implementation that allows short development cycles.

A proof-of-concept implementation can validate if the proposed RML constructions are implementable, and reveal ambiguities in the specification. This allows for corrective iterations during the development of the specification, and can also support the creation of test cases.

In this workshop paper, we present RML-view-to-CSV: a proof-of-concept implementation made available under MIT license and designed following state-of-the-art best practices. RML-view-to-CSV materializes all RML Logical Views specified in a given set of RML mapping rules as CSV files, and rewrites these RML mapping rules to RML mapping rules without RML Logical Views[5]. Any RML engine that supports CSV files, can then use these resulting RML mapping rules and the generated CSV files to construct RDF graph data using existing RML constructs. The implementation supports two source file formats (CSV and JSON) and supports the following features: flattening of nested data, handling of mixed data formats, and more flexible joining of data sources (also across data hierarchies).

We will continue to share our code and the results and discussion of our evaluations as inspiration for the developers who want to implement RML Logical Views directly in RML engines.

After discussing related work (Section 2) and explaining the main principles of logical views and how they are specified using RML Logical Views (Section 3), we describe our approach and the implemented features (Section 4). Finally, we evaluate (Section 5), and conclude (Section 6).

---

[1]https://www.w3.org/community/kg-construct/

[2]https://w3id.org/rml/portal/

[3]https://github.com/kg-construct/rml-lv

[4]https://kg-construct.github.io/rml-lv/dev.html

[5]https://github.com/RMLio/rml-view-to-csv

2. Related work

In this section, we first list solutions that were formulated in the past and influenced the development of RML Logical Views (Section 2.1). Then, we describe similar modular implementation approaches (Section 2.2).

2.1. Related proposals to extend RML

Partial solutions that also apply the concept of a (logical) view were formulated in the past to address (parts of) the challenges presented in this paper [2, 3]. RML Fields is a proposed solution to deal with nested data and mixed content in RML [2]. The proposal includes a language construct and algorithm, but no implementation. RML Logical Views is a revised extension of this proposal. Another solution is expanding RML's existing SQL views to general tabular sources instead of solely relational databases [3], to better support transformation functions, complex joins, and mixed content in RML. These views are formulated as SQL queries over tabular sources. The proposal is implemented as an extension of Morph-KGC[6]: a state-of-the-art RML mapping engine implemented using the pandas Python package [4][7].

Other proposed solutions include Facade-X [5], which directly maps the data source into RDF graph data (so not a flattened, but a graph-based source-agnostic view), allowing to, e.g. support joining over hierarchy via an iteration index[8]; and xR2RML, which (among others) supports mixed syntax paths for handling data sources that mix source formats [6].

2.2. RML engine module implementations

State-of-the-art RML engine module implementations typically provide (integrated or separate) preprocessing steps. SDM-RDFizer introduces a preprocessing step for grouping mapping rules, which leads to optimizations due to parallelization [7]. FunMap [8] is an interpreter of RML+FnO (the RML module that allows to describe data transformations in the mapping process), that converts a data integration system defined using RML+FnO into an equivalent one where RML mappings are function-free.

3. RML Logical Views Primer

Within RML, RDF triple generation is defined using *triples maps*. Within these triples maps, *expressions* are used to fill in data values in specified places in the triples. On which data these triples are generated, is currently specified using a *logical source*: a construct to describe how to extract *iterations* out of heterogeneous data sources, e.g. extracting individual rows from CSV files or extracting specific objects from JSON data from a Web API response.

For example, when mapping person data from a local CSV file that contains a person's (national) ID and name, the logical source describes how to extract rows as iterations from that CSV file. On these iterations, specific ID and name expressions are evaluated to create triples,

---

[6]https://github.com/morph-kgc/morph-kgc/releases/tag/2.3.0
[7]https://pandas.pydata.org/
[8]https://github.com/kg-construct/mapping-challenges/issues/43

using the ID expression as part of the subject identifier, and the name expression as a literal object.

With RML Logical Views, the W3C Community Group on KGC aims to describe a virtual flattened view on top of a logical source.

In this section, we give an example-based explanation of the main RML Logical Views terminology and functionalities. The examples[9] (Figures 1 to 3) include an RML mapping illustrating the declarative description of a logical view, the source data for the logical source used as source for the logical view, and an intermediate representation of the logical view. The headers of the columns of this intermediate representation can be used as reference expression in expression maps.

3.1. Flattening of nested data structures

Problem: References to nested data structures, like JSON or XML, may return multiple values. These values can be composite: they may again contain multiple values. RML defines mapping constructs that produce results by combining the results of other mapping constructs in a specific order. For example, a triples map combines the results of a subject map and a predicate-object map in that order. Another example is a template expression, which combines character strings and zero or more reference expressions in declared order. When mapping constructs produce multiple results, the combining mapping constructs will apply an n-ary Cartesian product[10] over the sets of results, maintaining the order of the mapping constructs. In the case of nested data structures, this may cause the generation of results that do not match the source hierarchy, since values are combined irrespective of it.

Furthermore, there is varying expressiveness in data source expression and query languages, and many languages have limited support for hierarchy traversal. For example, JSONPath has no operator to refer to an ancestor in the document hierarchy.

This limits the ability of RML to map nested data.

Solution: Fields can be nested (Figure 1). The value of the child field is derived from the value of its parent. The referenceable name of a child field is a concatenation of the referenceable name of its parent, a dot, and the field name of the child. There is no limit to the degree of nesting or the amount of children per parent.

3.2. Handling of mixed data formats

Problem: Data in one format can contain multiple or composite values stored in another format [2][11]. RML can handle only one reference formulation per logical source.

Solution: For every field, the reference formulation can be adapted (Figure 2). The default reference formulation for a field is the reference formulation of its parent, or of the source of the logical view for the fields at the root of the iteration. Thus, every iteration level can iterate over data in a different format.

---

[9]Prefixes are omitted but can be found on https://prefix.cc

[10]https://w3id.org/rml/core/spec#dfn-n-ary-cartesian-product

[11]https://webusers.i3s.unice.fr/~fmichel/xr2rml_specification_v5.1.html_Toc466307454

```
1  :jsonSource a rml:LogicalSource ;
2    rml:source "json_data.json" ;
3    rml:referenceFormulation rml:JSONPath ;
4    rml:iterator "$.people[*]" .
5  :jsonView a rml:LogicalView ;
6    rml:source :jsonSource ;
7    rml:field [
8      rml:fieldName "name" ;
9      rml:reference "$.name" ; ] ;
10   rml:field [
11     rml:fieldName "item" ;
12     rml:reference "$.items[*]" ;
13     rml:field [
14       rml:fieldName "type" ;
15       rml:reference "$.type" ; ] ;
16     rml:field [
17       rml:fieldName "weight" ;
18       rml:reference "$.weight" ; ] ; ] .
```

```
1  { "people": [
2      { "name": "alice",
3        "items": [
4          { "type": "sword",
5            "weight": 1500 },
6          { "type": "shield",
7            "weight": 2500 }
8        ] },
9      { "name": "bob",
10       "items": [
11         { "type": "flower",
12           "weight": 15 }
13       ] }
14   ] }
```

(a) mapping_json_view.ttl (part 1)   (b) json_data.json

name	item	item.type	item.weight
alice	{"type": "sword", "weight": 1500}	sword	1500
alice	{"type": "shield", "weight": 2500}	shield	2500
bob	{"type": "flower", "weight": 15}	flower	15

(c) Intermediate representation of :jsonView

Figure 1: Example of a logical view on a logical source containing nested JSON data

3.3. Extended joining of data sources

Problem: RML restricts join operations to referencing object maps, hence it is not possible to combine data from two sources in one term, use data from a join on another position than the object, or generate a literal using data from a join [3]. Moreover, RML cannot join correctly across hierarchies [2].

Solution: A logical view can be extended with fields from one or more other logical views as a result of a join operation (Figure 3). The logical iteration will be adapted according to the type of join specified (*left join* or *inner join*). Any needed flattening of hierarchical data is done in the logical view before applying the join operation. Data that originally comes from a different data source is thus treated equally in a joined logical view, allowing more flexibility as to where to apply that joined data.

4. Approach and Implementation

We have designed our proof-of-concept implementation as a standalone application, independent of and complementary to existing RML mapping engines, that can be used as a preprocessing step in a knowledge graph construction pipeline (Figure 4). As such, we designed our proof-of-

```
1  :mixedSource a rml:LogicalSource ;
2      rml:source "./mixed_data.csv";
3      rml:referenceFormulation rml:CSV .
4  :mixedView a rml:LogicalView ;
5    rml:source :mixedSource;
6    rml:field [
7      rml:fieldName "name" ;
8      rml:reference "name" ; ] ;
9    rml:field [
10     rml:fieldName "item" ;
11     rml:reference "item" ;
12     rml:referenceFormulation rml:JSONPath ;
13     rml:field [
14       rml:fieldName "type" ;
15       rml:reference "$.type" ; ] ;
16     rml:field [
17       rml:fieldName "weight";
18       rml:reference "$.weight" ;] ; ] .
```

```
1  name, item
2  alice,"{""type"":""sword"",""weight"":2500}"
3  alice,"{""type"":""shield"",""weight"":1500}"
4  bob,"{""type"":""flower"",""weight"": 15 }"
```

(a) mapping_mixed_view.ttl   (b) mixed_data.json

name	item	item.type	item.weight
alice	{"type": "sword", "weight": 1500}	sword	1500
alice	{"type": "shield", "weight": 2500}	shield	2500
bob	{"type": "flower", "weight": 15}	flower	15

(c) Intermediate representation of :mixedView

Figure 2: Example of a logical view on a logical source containing nested JSON data

concept following the state-of-the-art best practices of, amongst others, FunMap [8].

Our implementation, named RML-view-to-CSV, is available online[12] under the permissive MIT license. Our code is built on top of pandas [9], a Python library with data frames for manipulation of structured data sets.

At the moment of writing, our implementation supports one nested source format (JSON), one tabular source format (CSV), and the three important features of the new RML Logical Views module: flattening of nested data, more flexible joining of data sources (also across data hierarchies), and handling of mixed data formats. Its main functionality is the materialization of logical views (Section 4.1). We added two optional functionalities: the elimination of referencing object maps (Section 4.2) and the optimization of logical views based on the linked triples maps (Section 4.3).

4.1. Materialization of logical views

RML-view-to-CSV takes as input a given set of RML mapping rules and the source data used in these mapping rules. It produces CSV files with the intermediate representation of every

---

[12]https://github.com/RMLio/rml-view-to-csv/ (v0.0.0)

```
1  :csvSource a rml:LogicalSource ;
2      rml:source "./csv_data.csv";
3      rml:referenceFormulation rml:CSV .
4  :joinedView a rml:LogicalView ;
5    rml:source :csvSource;
6    rml:field [
7      rml:fieldName "name" ;
8      rml:reference "name" ;
9    ] ;
10   rml:field [
11     rml:fieldName "birthyear" ;
12     rml:reference "birthyear" ;
13   ] ;
14   rml:leftJoin [
15     rml:parentLogicalView :jsonView ;
16     rml:joinCondition [
17       rml:parent "name" ;
18       rml:child "name" ;
19     ] ;
20     rml:field [
21       rml:fieldName "item_type" ;
22       rml:reference "item.type" ;
23     ] ; ] .
```

(a) mapping.ttl (part 3)

```
1  name,birthyear
2  alice,1995
3  bob,1999
4  tobias,2005
```

(b) csv_data.csv

name	birthyear	item_type
alice	1995	sword
alice	1995	shield
bob	1999	flower
tobias	2005	

(c) Intermediate representation of :joinedView

Figure 3: Example of a logical view on a logical source containing CSV data and extended with data from the logical view from Figure 1


Figure 4: RML-view-to-CSV materializes all logical views specified in a given set of RML mapping rules as CSV files, and rewrites these RML mapping rules to RML mapping rules without logical views. The original source data, the materialized logical views and the new RML mapping without logical views are input for an RML engine to generate RDF graph data.

logical view specified in the RML mapping rules, and a new set of RML mapping rules in which all RML Logical Views are replaced by logical sources (Figure 5).

4.2. Elimination of referencing object maps

With the introduction of logical views, joins of data sources can be executed either within the triples map (via referencing object maps) or within the logical view. We added an option to

```
1  :jsonSource a rml:LogicalSource; rml:source "json_data.json";
2      rml:referenceFormulation rml:JSONPath; rml:iterator "$.people[*]".
3
4  :jsonView a rml:LogicalSource; rml:source "./view0.csv";
5      rml:referenceFormulation rml:CSV.
```

(a) MappingWithoutViews.ttl

```
1  name,item,item.type,item.weight
2  alice,"{""type"": ""sword"", ""weight"": 1500}",sword,1500
3  alice,"{""type"": ""shield"", ""weight"": 2500}",shield,2500
4  bob,"{""type"": ""flower"", ""weight"": 15}",flower,15
```

(b) view0.csv

Figure 5: Output of RML-view-to-CSV when using Figure 1a and Figure 1b as input

RML-view-to-CSV that eliminates all joins defined in triples maps. We applied the following algorithm: (i) rewrite logical sources that point to the same physical source as a single logical source with the same identifier; (ii) eliminate self-joins, i.e. rewrite referencing object map without join condition as normal object maps when all join condition demand equality between identical fields, and all references from the parent subject or all references from the child term maps occur in the join conditions; and (iii) rewrite all remaining referencing object maps as an equivalent combination of new logical views and a new triples map (Figure 6). This option allows us to test our proof-of-concept implementation with existing RML benchmarks even though these existing benchmarks do not include RML Logical Views, as logical views are new in RML.

4.3. Optimization of logical views

By default, RML-view-to-CSV does not take the content of triples maps into account. The materialized logical views represent all fields and logical iterations of the declared RML Logical Views. Thanks to this behaviour, we can verify if the processing of the logical views remains aligned with the specification. However, logical views can contain fields and logical iterations that are not used by any triples map. As the size of the source data impacts the knowledge graph construction process, we added an option in RML-view-to-CSV to eliminate unnecessary fields and logical iterations. With this option selected, RML-view-to-CSV first removes fields that are not used in any triples map linked to the logical view. Then, all duplicate logical iterations are removed, except when any linked triples map produces blank nodes that are not based on a field from the logical view (as this latter case results in always generate a new unique identifier, hence non-duplicate logical iterations).

5. Evaluation

We tested our proof-of-concept implementation against the test cases defined in the RML Logical Views modules, and added an evaluation on the GTFS-Madrid-Bench [10].

```
1  :map_services1_0 a rml:TriplesMap; rml:logicalSource :source_5.
2      rml:subjectMap [rr:template "http://gtfs/services/{service_id}"].
3
4  :map_trips_0 a rml:TriplesMap; rml:logicalSource :source_1;
5      rml:subjectMap [rml:template "http://gtfs/trips/{trip_id}"];
6      rml:predicateObjectMap [
7          rml:predicateMap [rml:constant gtfs:service];
8          rml:objectMap [
9              rml:parentTriplesMap :map_services1_0];
10             rml:joinCondition [rml:child "service_id"; rml:parent "service_id"]].
```

(a) MappingWithRefObjectMap.ttl

```
1  :new_map_0 a rml:TriplesMap; rml:logicalSource :new_child_view_0;
2      rml:subjectMap [rml:template "http://gtfs/trips/{trip_id}"];
3      rml:predicateObjectMap [
4          rml:predicateMap [rml:constant gtfs:service];
5          rml:objectMap [rml:template "http://gtfs/services/{service_id_from_parent}"]] .
6
7  :new_child_view_0 a rml:LogicalView; rml:source :source_1;
8      rml:field
9          [rml:fieldName "trip_id"; rml:reference "trip_id"],
10         [rml:fieldName "service_id"; rml:reference "service_id"];
11     rml:leftJoin [a rml:ViewJoin; rml:parentLogicalView :new_parent_view_0;
12         rml:joinCondition [rml:child "service_id"; rml:parent "service_id"];}
13         rml:field [rml:fieldName "service_id_from_parent"; rml:reference "service_id"]].
14
15 :new_parent_view_0 a rml:LogicalView; rml:source :source_5;
16     rml:field [rml:fieldName "service_id"; rml:reference "service_id"].
```

(b) MappingWithoutRefObjectMap.ttl

Figure 6: The referencing object map from MappingWithRefObjectMap.ttl is replaced by an equivalent combination of two new logical views and one new triples map in MappingWithoutRefObjectMap.ttl.

5.1. Test cases in the RML Logical Views module

With RML-view-to-CSV as preprocessor to RMLMapper, we generated the expected output for the relevant test cases listed in the RML Logical Views modules [13] (i.e. test cases RMLLV0001 to RMLLV0004). We excluded test case RMLLV0005 as it makes use of field indexes, which we have not yet included in our implementation due to an unclarity about the expected behaviour [14]. During the evaluation of the test cases we noticed and corrected human mistakes in the mappings and expected output [15]. This confirms the need and benefit of a proof-of-concept implementation during the development phase of a new RML module: a proof-of-concept implementation helps to spot ambiguities in the specification, ontology and shapes as well as errors in the test cases in an early development stage.

---

[13] https://github.com/kg-construct/rml-lv/tree/main/test-cases
[14] https://github.com/kg-construct/rml-lv/issues/20
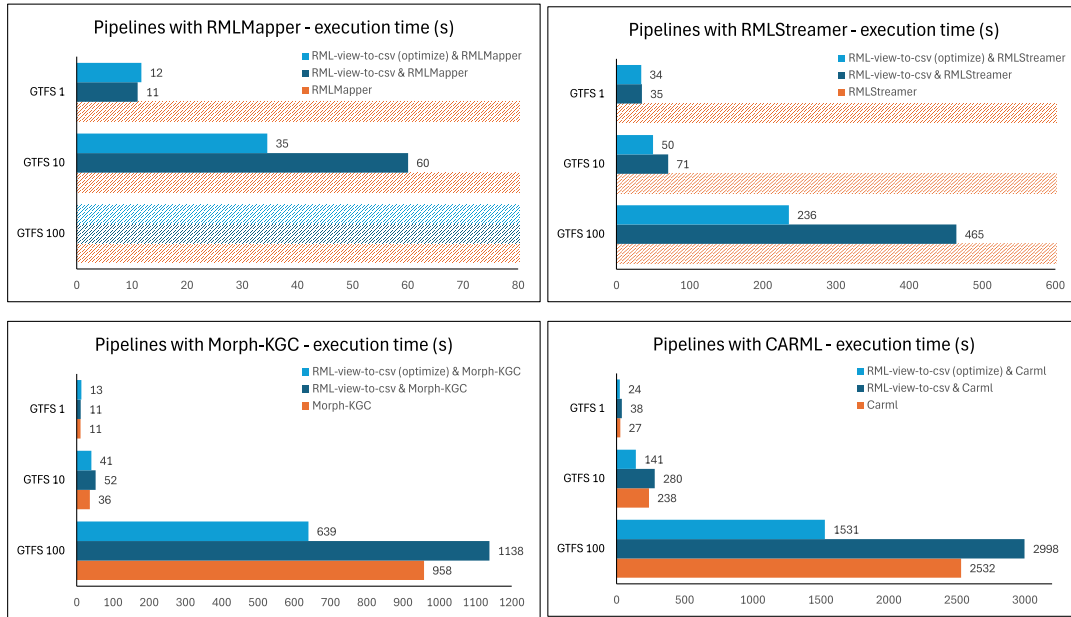[15] https://github.com/kg-construct/rml-lv/pull/22

Figure 7: GTFS-Madrid-Bench scale 1, 10 and 100 executed with three pipelines (only RML engine, RML-view-to-csv and RML engine, and RML-view-to-CSV with optimization and RML engine) and four RML engines (RMLMapper, RMLStreamer, Morph-KGC and Carml). For some pipelines, RMLMapper exceeds the available heap memory, while RMLStreamer times out after one hour. The pipelines with RML-view-to-CSV with optimization as preprocessor have the lowest execution time.

5.2. GTFS-Madrid-Bench

We tested our proof-of-concept implementation on the GTFS-Madrid-Bench, comparing the execution of joins by our implementation versus the execution joins by existing RML engines, i.e. RMLMapper, RMLStreamer, Morph-KGC, and Carml [16] (Figure 7). Our test setup included three pipelines: RML engine only, RML-view-to-CSV and RML engine, and RML-view-to-CSV with optimization (Section 4.3) and RML engine. In the latter two pipelines, we first eliminated all referencing object maps in the GTFS-Madrid-Bench mapping, moving them to logical views whenever applicable (Section 4.2). We measured the knowledge graph construction time per pipeline and per RML engine for scales 1, 10, and 100 of the GTFS-Madrid-Bench, with CSV as source format, using a device with following specifications: 2 x Hexacore Intel E5645 (2.4GHz) CPU, 24GB RAM, 1x 250GB harddisk.

RMLMapper and RMLStreamer cannot generate any output for the GTFS-Madrid-Bench within one hour when we count on these RML engines to execute the joins. However, when delegating the joins to RML-view-to-CSV, these mapping engines were able to generate correct output, with timings similar or better to using a state-of-the-art RML engine like Morph-KGC. We note that RMLMapper still cannot handle GTFS scale 100: the RMLMapper loads all data in memory during mapping, and the testing device ran out of memory during GTFS scale 100.

---

[16]https://github.com/RMLio/rmlmapper-java (v6.3.0), https://doi.org/10.5281/zenodo.3887065 (v2.5.0), and https://doi.org/10.5281/zenodo.5543552 (v2.6.4), https://github.com/carml/carml (V0.4.10), respectively.

Figure 8: With the optimization option RML-view-to-CSV eliminates unnecessary logical view fields and duplicate logical iterations. This reduces its executions time with a factor of six on the largest scale measured.

We also note that the elimination of unnecessary fields and duplicate logical iterations (RML-view-to CSV with optimization) reduces the execution time of RML-view-to-CSV by a factor of six for scale 100 (Figure 8), and leads to the fastest pipelines in combination with all tested engines.

The combination of RML-view-to-CSV with optimization and RMLStreamer emerges as the most efficient approach[17]. This showcases the potential of modular mapping engines, delegating each task to the most suitable framework, i.e. the dataframes from pandas (used in RML-view-to-CSV) are optimized for data transformations and joins, while streaming and parallelization of Flink enables RMLStreamer to create RDF graph data with a linear scaling of execution time and CPU usage, proportional to the size of the input data, while maintaining a constant memory usage [11].

6. Conclusion

In this paper, we show how the RML Logical Views specification can be implemented and can solve needs that were not solvable yet by RML. The implementation can be realized as a modular part of a knowledge graph construction process.

Our proof-of-concept, RML-views-to-CSV, as a preprocessor to any RML engine (that supports CSV input) did not only help to validate and improve the RML Logical Views module, but benchmarks also show performance gains for handling joins between CSV sources. The modular approach showcases the potential of modular mapping engines, allowing to use specialized data structures and delegating each task to the most suitable framework, offering best-of-breed performance enhancements.

For future work, we will gradually add more features (e.g. indexes per field, data transformation functions, groups and aggregations) as they are discussed in the W3C Community Group on KGC and described in the RML Logical Views specification. We will also further investigate whether the detected performance improvements hold as well for sources other than CSV.

We will continue to share our code as inspiration for the developers who want to implement RML Logical Views directly in RML Engines.

---

[17]The current set-up of using a preprocessing materialization step prevents the RMLStreamer of currently using this optimization for streaming data.

Acknowledgments

References

[1] A. Iglesias-Molina, D. Van Assche, J. Arenas-Guerrero, B. De Meester, C. Debruyne, S. Joza-shoori, P. Maria, F. Michel, D. Chaves-Fraga, 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.

[2] T. Delva, D. Van Assche, P. Heyvaert, B. De Meester, A. Dimou, Integrating nested data into knowledge graphs with RML fields, in: D. Chaves-Fraga, A. Dimou, P. Heyvaert, F. Priyatna, J. Sequeda (Eds.), Proceedings of the 2nd International Workshop on Knowledge Graph Construction co-located with 18th Extended Semantic Web Conference (ESWC 2021), volume 2873, CEUR, 2021. URL: http://ceur-ws.org/Vol-2873/paper9.pdf.

[3] J. Arenas-Guerrero, A. Alobaid, M. Navas-Loro, M. S. Pérez, O. Corcho, Boosting knowledge graph generation from tabular data with RML views, in: The Semantic Web, Springer Nature Switzerland, 2023, pp. 484–501. doi:10.1007/978-3-031-33455-9_29.

[4] J. Arenas-Guerrero, D. Chaves-Fraga, J. Toledo, M. S. Pérez, O. Corcho, Morph-KGC: Scalable knowledge graph materialization with mapping partitions, Semantic Web (2022) 1–20. doi:10.3233/sw-223135.

[5] E. Daga, L. Asprino, P. Mulholland, 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, volume 53 of *Studies on the Semantic Web*, IOS Press, 2021, pp. 58–73. doi:10.3233/SSW210035.

[6] F. Michel, L. Djimenou, C. Faron-Zucker, 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.

[7] E. Iglesias, S. Jozashoori, M.-E. Vidal, Scaling up knowledge graph creation to large and heterogeneous data sources, Journal of Web Semantics 75 (2023). URL: http://arxiv.org/abs/2201.09694. doi:10.1016/j.websem.2022.100755.

[8] S. Jozashoori, D. Chaves-Fraga, E. Iglesias, M.-E. Vidal, O. Corcho, Funmap: Efficient execution of functional mappings for knowledge graph creation, in: International Semantic Web Conference, Springer, 2020, pp. 276–293.

[9] W. McKinney, pandas: a foundational python library for dataanalysis and statistics (2011).

[10] D. Chaves-Fraga, F. Priyatna, A. Cimmino, J. Toledo, E. Ruckhaus, O. Corcho, Gtfs-madrid-

bench: A benchmark for virtual knowledge graph access in the transport domain, Journal of Web Semantics 65 (2020) 100596. doi:10.1016/j.websem.2020.100596.

[11] E. de Vleeschauwer, G. Haesendonck, B. De Meester, P. Colpaert, The Knowledge Graph Construction 2023 Challenge - RMLStreamer submission, in: Knowledge Graph Construction Workshop, co-located with ESWC, 2023.