



HAL
open science

PRSC: from PG to RDF and back, using schemas

Julian Bruyat, Pierre-Antoine Champin, Lionel Médini, Frédérique Laforest

► To cite this version:

Julian Bruyat, Pierre-Antoine Champin, Lionel Médini, Frédérique Laforest. PRSC: from PG to RDF and back, using schemas. *Semantic Web – Interoperability, Usability, Applicability*, 2024. hal-04772450

HAL Id: hal-04772450

<https://hal.science/hal-04772450v1>

Submitted on 8 Nov 2024

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

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

PRSC: from PG to RDF and back, using schemas

Julian Bruyat ^{a,*}, Pierre-Antoine Champin ^{a,b,c}, Lionel Médini ^a and Frédérique Laforest ^a

^a INSA Lyon, CNRS, Université Claude Bernard Lyon 1, LIRIS, UMR5205, 69621 Villeurbanne, France

E-mails: julian.bruyat@liris.cnrs.fr, lionel.medini@liris.cnrs.fr, frederique.laforest@liris.cnrs.fr

^b W3C, Sophia Antipolis, France

E-mail: pierre-antoine@w3.org

^c University Côte d'Azur, Inria, CNRS, I3S UMR 7271, France

Editor: Stefan Schlobach, Vrije Universiteit Amsterdam, Netherlands

Solicited reviews: Olaf Hartig, Linköping University, Sweden; anonymous reviewer; anonymous reviewer

Abstract.

Property graphs (PG) and RDF graphs are two popular database graph models, but they are not interoperable: data modeled in PG cannot be directly integrated with other data modeled in RDF. This lack of interoperability also impedes the use of the tools of one model when data are modeled in the other.

In this paper, we propose PRSC, a configurable conversion to transform a PG into an RDF graph. This conversion relies on PG schemas and user-defined mappings called PRSC contexts. We also formally prove that a subset of PRSC contexts, called well-behaved contexts, can be used to reverse back to the original PG, and provide the related algorithm. Algorithms for conversion and reversion are available as open-source implementations.

Keywords: Property Graph, RDF graph, Conversion, Schema

1. Introduction

Graphs are a popular data model in which knowledge is represented through objects and links between these objects. Today, there are two mainly used models of graphs: Property Graphs and RDF.

Property Graphs (PGs) are a family of implementations, in which data are represented with nodes and edges, and labels and properties (key-value pairs) can be attached to these nodes and edges. Property Graphs are not a uniform model: some implementations like Neo4j ¹ only allow exactly one label for each edge. Most PG engines offer easy-to-use graph query languages like Cypher [1] and Gremlin [2] that rely on graph traversal. While no uniform standard has been settled yet, the *Property Graph needs a Schema Working Group* ² is working towards defining a schema language for PG, and a unified formalization of the PG model. In the rest of this paper, following other authors [3], the variability of implementations is neglected and PGs are considered as a uniform model.

*Corresponding author. E-mail: julian.bruyat@liris.cnrs.fr.

¹<https://neo4j.com/>

²<https://www.w3.org/Data/events/data-ws-2019/assets/position/Juan%20Sequeda.txt>

Another popular graph model is the Resource Description Framework (RDF) model [4]. In this model, data are represented with triples that represent links between resources. The resources and the links between them are identified through Internationalized Resource Identifiers (IRI). This data model is a W3C standard, and has been studied through a large quantity of works, like RDFS [5] and OWL [6] for inference, or SHACL [7] for data validation. This model has been extended by RDF-star [8] that helps writing properties on triple terms in a more concise manner, but does not provide exactly the same modeling capabilities as PGs. For example, multiple links with the same predicate between two resources are still not supported and requires using extra constructs, like the occurrence pattern described in the “Triples and occurrences” section of the RDF-star community group final report [9].

While PG and RDF are both based on the idea of using graph data, the choice of one removes the ability to use the tools developed for the other. The maintainers of Amazon Neptune, a graph database service that can support both models independently, report that their users choose the solution that best suits their current use case, then struggle because they are stuck with the tools of this model even if the tools of the other model would better answer the new business problem they have [10]. Generally speaking, this diversity of graph models, and more precisely the lack of interoperability, hinders graph database adoption.

Numerous authors [11, 12] highlight the benefits of interoperable systems, and in particular of interoperability between PGs and RDF graphs [13–15]. We can distinguish two levels of interoperability: syntactic interoperability and semantic interoperability. The former consists in having common data formats, for example JSON or CSV. The latter consists in having data not only in a common data format, but also that use shared vocabularies. Tim Berners-Lee theorized the 5-star Open-data model³. In this model, datasets are given a number of stars depending on how interoperable they are. RDF graphs are by design the best candidate to elevate the ranking of data in the 5-star data-model, in particular through the use of shared vocabularies. So, while internally developers may choose RDF or PG databases based on their preferences and the convenience of associated tools, RDF is better suited as a pivot model for interoperable exchange of data.

Among the many questions risen by the interoperability between PGs and RDF graphs, the scenario we want to support is the conversion from PGs to RDF (syntactic interoperability) without information loss, and while promoting the use of existing shared vocabularies (semantic interoperability). The requirement of using of existing shared vocabularies requires a converter to be highly customizable, in particular in terms of how the produced triples are structured. On the other hand, the requirement of not losing information, *i.e.* producing reversible conversions, requires to formally study the possible conversions.

In a previous paper [16], we introduced the motivations behind PREC (PG to RDF Experimental Converter), a user-configured mapping from Property Graphs (PG) to RDF graphs and proposed a mapping language, specialized for the PG to RDF conversion problem, to let the user describe how to convert the node labels, the edges and the properties of the original PG to an RDF graph. By converting the data stored in PGs into RDF, users are then able to use all the tools available for RDF. In this paper, we introduce a new mapping language, named PRSC⁴, driven by a schema and a description of how to convert the elements of the types in the schema to RDF. This mapping language is formally defined, and conditions under which the conversion produced by PRSC is reversible are also defined. Compared to the previous mapping language, this language contains fewer terms and concepts, and should be easier to use. The PRSC engine is available under the MIT license⁵ and can connect both to a Cypher endpoint and a Gremlin endpoint.

The rest of this paper is organized as follows. Section 2 gives an overview of PRSC to understand its principles. Section 3 gives generic formal definitions of PGs and RDF graphs. Section 4 gives a formal definition of a PRSC conversion, which is essentially a formal definition of Section 2. Section 5 studies reversibility, and proves that some contexts can convert PG to RDF graph without information loss. Section 6 discusses the existing works to make easier interoperability between PGs and RDF graphs relatively to PRSC. Section 7 discusses the proposed solution and describes some future works.

³<https://5stardata.info/en/>

⁴PG to RDF: Schema-driven Converter, pronounced “presque”

⁵<https://github.com/BruJu/PREC>, <https://npmjs.com/package/prec>

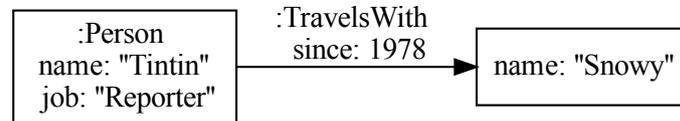


Fig. 1. A small PG about Tintin that serves as a running example in this paper

Listing 1 An example of an RDF-star graph in Turtle Format

```

1  _:n1 rdf:type ex:Person .
2  _:n1 foaf:name "Tintin" .
3  _:n1 ex:profession "Reporter" .
4  _:n1 ex:isTeammateOf _:n2 .
5  << _:n1 ex:isTeammateOf _:n2 >> ex:since 1978 .
6  _:n2 foaf:name "Snowy" .
  
```

2. PRSC in practice

The Property Graph exposed on Figure 1 describes the relationship between Tintin and Snowy. It is composed of two nodes. The first one holds the label `Person` and two properties: the first property has “name” as its key and “Tintin” as its value, the second has “job” as its key and “Reporter” as its value, or more simply its name is “Tintin” and its job is “Reporter”. The other node only has one property: the name “Snowy”. These two nodes are connected by an edge that holds one label, `TravelsWith`, and a property that tells that it is “since” “1978”.

A similar example represented in RDF-star is exposed on Listing 1. Most information that was in the PG is represented by the triples in lines 1-4 and 6. The information about since when Tintin travels with Snowy is represented through a nested RDF-star triple.

Using the user-defined mapping, PRSC is able to convert the PG in Figure 1 into the RDF-star graph in Listing 1, and more generally any Property Graph with the same schema into the corresponding RDF graph. In this paper, we consider that Property Graphs with the same schema as the one in Figure 1 are Property Graphs where all nodes have either (1) the “Person” label, a “name” property and a “job” property, (2) no label and a “name” property, (3) and where all edges have the “TravelWith” label and a “since” property. The mapping the user must provide to the PRSC engine is exposed on Listing 2 and is in Turtle-star format [17]. Rules are split in two parts:

- The target part that describes which elements of the Property Graph are targeted. The target is described depending on three criteria: (1) whether the element must be an edge or a node, (2) the labels and (3) the properties of the element.
- The production part that describes the triples to produce with a list of *template triples*. Values in the *pvar* namespace are mapped to the blank node in the resulting RDF graph. The literals that use *valueOf* as their datatype are converted to the property values in the RDF graph.

The mapping, named *PRSC context*, exposed on Listing 2 reads as follows:

- The first rule is named `_:PersonRule` (line 1)
 - * The rule is used for all PG nodes (line 3) that only have the node label “Person” (line 4) and have the properties “name” and “job” (line 5). In our example, the node corresponding to Tintin matches this description, but Snowy does not as it misses the Person label and the job property.
 - * It will produce three triples:

Listing 2 The PRSC context that maps the PG running example to the RDF graph running example

```

1
2
3 1  _:PersonRule
4 2  # Target: all nodes with label "Person" and two properties "name" and "job"
5 3  a prec:PRSCNodeRule ;
6 4  prec:label "Person" ;
7 5  prec:propertyKey "name", "job" ;
8 6  # Production part of the rule: a template graph
9 7  prec:produces
10 8  << pvar:self rdf:type      ex:Person >> ,
11 9  << pvar:self foaf:name     "name"^^prec:valueOf >> ,
12 10 << pvar:self ex:profession "job"^^prec:valueOf >> .
13
14 11
15 12  _:NamedRule
16 13  # Target: all nodes with no label and one property "name"
17 14  a prec:PRSCNodeRule ;
18 15  prec:propertyKey "name" ;
19 16  # Production part of the rule
20 17  prec:produces
21 18  << pvar:self foaf:name "name"^^prec:valueOf >> .
22
23 19
24 20  _:TravelsWithRule
25 21  # Target: all edges with the label "TravelsWith" and one property "since"
26 22  a prec:PRSCEdgeRule ;
27 23  prec:label "TravelsWith" ;
28 24  prec:propertyKey "since" ;
29 25  # Production part of the rule
30 26  prec:produces
31 27  << pvar:source ex:isTeammateOf pvar:destination >> ;
32 28  << << pvar:source ex:isTeammateOf pvar:destination >> ex:since "since"^^prec:valueOf >> .
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51

```

★ One triple with a blank node as its subject, `rdf:type` as its predicate and `ex:Person` as its object (line 8). Each node from the Property Graph is identified by a distinct blank node. In this example, `_:n1 rdf:type ex:Person` will be produced.

★ Another triple with the same blank node as its subject, `foaf:name` as its predicate and a literal that matches the value of the name property in the PG (line 9). The PRSC engine converts all literals whose datatype is `prec:valueOf` into the value of the corresponding property in the PG. In this example, `_:n1 rdf:type "Tintin"` will be produced.

★ One last triple is produced with the same blank node as its subject, `ex:profession` as its predicate and a literal corresponding to the value of the property job (line 10). In this example, `_:n1 ex:profession "Reporter"` will be produced.

– The second rule is named `_:NamedRule` (line 12).

- * It is applied to nodes (line 14) that have no labels and only one property: name (line 15). This is the case of the PG node used to describe Snowy but not the one that describes Tintin as it has an extra label and an extra property.
- * These PG nodes will be converted into one triple with a blank node that identifies the PG node as its subject, `foaf:name` as its predicate and the literal that correspond to the value of the name property as its object (line 18). In this example, the triple `_:n2 foaf:name "Snowy"` is produced.

– The third rule is named `_:TravelsWithRule` (line 20):

- * It is used to convert edges (line 22) whose only label is “TravelsWith” (line 23) and with one and only one property named “since” (line 24).
- * These edges are converted by producing a triple with the identifier of the source PG node as the subject, `ex:isTeammateOf` as the predicate and the identifier of the destination PG node as the object (line 27). In this example, the triple `_:n1 ex:isTeammateOf _:n2` is produced.
- * A triple with a quoted triple is created by the rule on line 28: the triple that was created by the line 27 is used on the subject position of the triples created by this triple, `ex:since` is used as the predicate and the value of the “since” property is used as the object. In our example, the triple `<< _:n1 ex:isTeammateOf _:n2 >> ex:since 1978` is produced.

* Note that in this example, `pvar : self` is not used in lines 27 and 28. If it was used, it would be mapped to a blank node that identifies the edge. The consequence of not using it is a smaller RDF graph, at the cost of if several PG edges with the “TravelsWith” label were present between the two same nodes, the RDF representation of these edges would have been merged.

Note that this mechanism of using quoted triples to describe templates in a pure Turtle-star file was already presented in our previous work [16]. Compared to R2RML [11], it lets the user describe the triples to produce with a syntax closer to the triples that will actually be produced, but it is currently impossible to express templated terms in any position. For example, if the user wants to generate a named node `<http://example.org/person/{name}>` in subject position in which the `{name}` part is substituted with the value of the name property, R2RML offers it as a native feature. As PRSC is currently unable to generate arbitrary IRIs, nodes and edges from the PG are always mapped to blank nodes and minting IRIs is left for future works.

3. General definitions

This section introduces some standard definitions, mostly inspired from previous works.

3.1. Notations and conventions

Let *Str* be the set of all strings. Strings are noted between quotes. For example, “*node*”, “*edge*”, “*since*” and “*Snowy*” are strings.

Definition 1 (Domain and image of a function). For all partial functions $f : D \rightarrow A$, *Dom* and *Img* are defined as follows:

- $Dom(f) = \{x \mid \exists y \in A, \text{ such that } f(x) = y\}$
- $Img(f) = \{y \mid \exists x \in D, \text{ such that } f(x) = y\}$.

Example 1. For the partial function $inverse : \mathbb{R} \rightarrow \mathbb{R}$, with $inverse(x) = 1/x$, $Dom(inverse) = Img(inverse) = \mathbb{R} - \{0\}$.

Let *S* be a set, we recall that 2^S denotes the set of all parts of *S*.

3.2. Compatible functions

For all functions *f*, we recall that they can be seen as sets: $f = \{(x, f(x)) \mid x \in Dom(f)\}$. For all sets *S* of 2-tuples, *S* can be seen as a function iff (if and only if) $\forall (x, y_1, y_2), (x, y_1) \in S \wedge (x, y_2) \in S \Rightarrow y_1 = y_2$.

Example 2. Consider the three functions f_1, f_2, f_3 exposed in Table 1.

Table 1
Some functions defined both with the usual function notation and with a set notation

Function notation	Set notation
$f_1(x) = \begin{cases} 0 & \text{if } x = 0 \\ 1 & \text{if } x = 1 \end{cases}$	$f_1 = \{(0, 0), (1, 1)\}$
$f_2(x) = \begin{cases} 66 & \text{if } x = -2 \\ 33 & \text{if } x = -1 \\ 0 & \text{if } x = 0 \end{cases}$	$f_2 = \{(-2, 66), (-1, 33), (0, 0)\}$
$f_3(x) = \begin{cases} 10 & \text{if } x = 0 \\ 1 & \text{if } x = 1 \end{cases}$	$f_3 = \{(0, 10), (1, 1)\}$

As f_1 , f_2 and f_3 can be defined with a set, it is possible to use the usual set operations.

The set $f_1 \cup f_2 = \{(-2, 66), (-1, 33), (0, 0), (1, 1)\}$ is a function: the first element of all tuples has a different value. Using a function notation, it may be written as:

$$(f_1 \cup f_2)(x) = \begin{cases} 66 & \text{if } x = -2 & [f_2(-2) = 66] \\ 33 & \text{if } x = -1 & [f_2(-1) = 33] \\ 0 & \text{if } x = 0 & [f_1(0) = f_2(0) = 0] \\ 1 & \text{if } x = 1 & [f_1(1) = 1] \end{cases}$$

On the opposite, $f_1 \cup f_3 = \{(0, 0), (0, 10), (1, 1)\}$ is not a function. Both $(0, 0)$ and $(0, 10)$ are members of the set $f_1 \cup f_3$, $(f_1 \cup f_3)(0)$ would be equal to both $f_1(0) = 0$ and $f_3(0) = 10$ which are different values.

Remark 1. Instead of using the notation $\{(x_0, f(x_0)), (x_1, f(x_1)), \dots\}$, the notation $\{x_0 \mapsto f(x_0), x_1 \mapsto f(x_1), \dots\}$ is sometimes used to clarify the fact that a set is a function. For example, f_3 may be noted as $f_3 = \{0 \mapsto 10, 1 \mapsto 1\}$.

Definition 2 (Functions compatibility). Two functions f and g are compatible iff $f \cup g$ is a function, i.e. $\forall (x, y_f, y_g), (x, y_f) \in f \wedge (x, y_g) \in g \Rightarrow y_f = y_g$.

In other words, two functions f and g are compatible iff for every common input, they share the same output i.e. $\forall x \in \text{Dom}(f) \cap \text{Dom}(g), f(x) = g(x)$.

3.3. Property Graph

Definition 3 (Property Graph). Following the definition of Angles [3], a property graph pg is defined as the tuple $(N_{pg}, E_{pg}, \text{src}_{pg}, \text{dest}_{pg}, \text{labels}_{pg}, \text{properties}_{pg})$, where:

- N_{pg} and E_{pg} are finite sets with $N_{pg} \cap E_{pg} = \emptyset$. N_{pg} and E_{pg} are respectively the set of nodes and the set of edges of the property graph pg .
- $\text{src}_{pg} : E_{pg} \rightarrow N_{pg}$ and $\text{dest}_{pg} : E_{pg} \rightarrow N_{pg}$ are two total functions. These two functions map each edge to its starting and destination nodes.
- $\text{labels}_{pg} : N_{pg} \cup E_{pg} \rightarrow 2^{\text{Str}}$ is a total function. This function maps the nodes and edges to their sets of labels.
- $\text{properties}_{pg} : (N_{pg} \cup E_{pg}) \times \text{Str} \rightarrow V$ is a partial function. This function describes the properties of the elements. V is the set of all possible property values. Considering that a property is a key-value pair, it expects two inputs: a node or an edge, and a property key. The output is the property value.

The set of all property graphs is denoted PGs .

In this paper, property graph nodes and edges are grouped under the term **element**.

In the rest of the paper, when any PG x is introduced, we allow ourselves to use the symbols $N_x, E_x, \text{src}_x, \text{dest}_x, \text{labels}_x$ and properties_x without explicitly defining x as the tuple $(N_x, E_x, \text{src}_x, \text{dest}_x, \text{labels}_x, \text{properties}_x)$.

Example 3 (Running example of a Property Graph). The PG exposed on Figure 1 can formally be defined as the PG denoted TT with

- $N_{TT} = \{n_1, n_2\}; E_{TT} = \{e_1\}$
- $\text{src}_{TT} = \{e_1 \mapsto n_1\}; \text{dest}_{TT} = \{e_1 \mapsto n_2\}$
- $\text{labels}_{TT} = \{n_1 \mapsto \{\text{"Person"}\}; n_2 \mapsto \emptyset; e_1 \mapsto \{\text{"TravelsWith"}\}$
- $\text{properties}_{TT} = \left\{ \begin{array}{l} (n_1, \text{"name"}) \mapsto \text{"Tintin"}; (n_1, \text{"job"}) \mapsto \text{"Reporter"} \\ (n_2, \text{"name"}) \mapsto \text{"Snowy"}; (e_1, \text{"since"}) \mapsto 1978 \end{array} \right\}$

Definition 4 (The empty PG). The empty PG, which is the PG that contains no nodes and no edges, is formalized as follows: pg_\emptyset with $N_{pg_\emptyset} = E_{pg_\emptyset} = \emptyset, \text{src}_{pg_\emptyset} = \text{dest}_{pg_\emptyset} = \text{labels}_{pg_\emptyset} = \emptyset \rightarrow \emptyset$ and $\text{properties}_{pg_\emptyset} : \emptyset \times \emptyset \rightarrow \emptyset$.

3.3.1. Renaming Property Graphs and isomorphism

The chosen formal definition of the running example is not the only one that is possible: for example an arbitrary element named a could have been used in place of n_1 as the first listed node identifier in Example 3.

Definition 5 (Renaming function). For all sets N_1, N_2, E_1, E_2 where $N_1 \cap E_1 = \emptyset$ and $N_2 \cap E_2 = \emptyset$, a renaming from N_1 and E_1 to N_2 and E_2 is a bijective function $\phi : N_1 \cup E_1 \rightarrow N_2 \cup E_2$ where $\forall n \in N_1, \phi(n) \in N_2 \wedge \forall e \in E_1, \phi(e) \in E_2$.

Example 4. Let TT' be another formal definition of the PG in Figure 1 that uses the sets $N_{TT'} = \{a, b\}$ and $E_{TT'} = \{c\}$.

An example of a renaming function ϕ_{TT} from $N_{TT} = \{n_1, n_2\} \cup E_{TT} = \{e_1\}$ to $N_{TT'} = \{a, b\} \cup E_{TT'} = \{c\}$ is $\phi_{TT} = \{n_1 \mapsto a; n_2 \mapsto b; e_1 \mapsto c\}$.

Definition 6 (Property Graph renaming). Let pg be a property graph, N' and E' be two sets and ϕ be a renaming function from N_{pg} and E_{pg} to N' and E' . We define the renamed PG $rename(\phi, pg) = pg'$ as follows:

- $N_{pg'} = N' = \{\phi(n) \mid n \in N_{pg}\}$
- $E_{pg'} = E' = \{\phi(e) \mid e \in E_{pg}\}$
- $src_{pg'} = \{e \mapsto \phi(src_{pg}(\phi^{-1}(e))) \mid e \in E_{pg'}\}$
- $dest_{pg'} = \{e \mapsto \phi(dest_{pg}(\phi^{-1}(e))) \mid e \in E_{pg'}\}$
- $labels_{pg'} = \{m \mapsto labels_{pg}(\phi^{-1}(m)) \mid m \in N_{pg'} \cup E_{pg'}\}$
- $properties_{pg'} = \{(m, key) \mapsto properties_{pg}(\phi^{-1}(m), key) \mid m \in N_{pg'} \cup E_{pg'} \wedge key \in Str\}$

Example 5. Let us consider back TT , the PG about Tintin defined in Example 3, TT' the other formalization of the same PG and ϕ_{TT} the renaming function introduced in the Example 4.

The PG produced by $rename(\phi_{TT}, TT) = TT'$ is

- $N_{TT'} = \{a, b\}; E_{TT'} = \{c\}$
- $src_{TT'} = \{c \mapsto a\}; dest_{TT'} = \{c \mapsto b\}$
- $labels_{TT'} = \{a \mapsto \{\text{"Person"}\}; b \mapsto \emptyset; c \mapsto \{\text{"TravelsWith"}\}\}$
- $properties_{TT'} = \left\{ \begin{array}{l} (a, \text{"name"}) \mapsto \text{"Tintin"}; (a, \text{"job"}) \mapsto \text{"Reporter"} \\ (b, \text{"name"}) \mapsto \text{"Snowy"}; (c, \text{"since"}) \mapsto 1978 \end{array} \right\}$

Definition 7 (Isomorphic property graph). Two PGs pg and pg' are isomorphic iff there exists a renaming function ϕ from N_{pg} and E_{pg} to $N_{pg'}$ and $E_{pg'}$ such that $rename(\phi, pg) = pg'$. Note that because renaming functions are bijectives, it implies that ϕ^{-1} exists and $rename(\phi^{-1}, pg') = pg$.

Note that both TT and TT' from Examples 3 and 4 match the graphical representation given in Figure 1. An informal way to define the isomorphism between two PGs is to check if they have the same graphical representation.

Existing works [1, 18] on PG query languages focus on extracting the properties of some nodes and edges, and never look for the exact identity of the elements. It is therefore possible to affirm that the exact identity is not important, and that if two PGs are isomorphic, they are the same PG for practical matter.

3.4. RDF-star definition

Definition 8 (Atomic RDF terms). Let I be the infinite set of IRIs, $L = Str \times I$ be the set of literals and B be the infinite set of blank nodes. The sets I, L and B are disjoint.

IRIs, literals and blank nodes are grouped under the name "Atomic RDF terms".

Notation: In the examples, the IRIs, the elements of I , will be either noted as full IRIs between brackets, e.g. `<http://example.org/Tintin>` or by using prefixes to shorten the IRI e.g. `ex:Tintin`. The list of prefixes used in this paper is described in Table 2.

Literals, the elements of L , can be noted either by using the usual tuple notation, e.g. `("1978", xsd:integer)` or with the compact notation `"1978"xsd:integer`.

Table 2
List of prefixes used in this paper

Prefix	IRI
rdf	http://www.w3.org/1999/02/22-rdf-syntax-ns#
xsd	http://www.w3.org/2001/XMLSchema#
ex	http://example.org/
prec	http://bruy.at/prec#
pvar	http://bruy.at/prec-var#

Finally, the blank nodes, the elements of B , are denoted by blank node labels prefixed with the two symbols “_:”
e.g. `_:edge`, `_:2021` or `_:node35`.

Definition 9 (RDF(-star) triples and graphs). The set of all RDF triples⁶ is denoted $RdfTriples$ and is defined as follows:

- $\forall subject \in I \cup B, \forall predicate \in I, \forall object \in I \cup B \cup L, (subject, predicate, object) \in RdfTriples$.
- $\forall tsubject \in RdfTriples, \forall tobject \in RdfTriples$, and for all $subject, predicate$ and $object$ defined as above, $(tsubject, predicate, object)$, $(subject, predicate, tobject)$ and $(tsubject, predicate, tobject)$ are members of $RdfTriples$.

A subset of $RdfTriples$ is an RDF graph.

Both the atomic RDF terms defined in Definition 8 and RDF triples are *terms*. A triple used in another triple, in subject or object position, is a *quoted triple*. An RDF triple can not contain itself and can not be nested infinitely.

Example 6.

1. The triple $(ex:tintin, rdf:type, ex:Person)$ is an element of $RdfTriples$. Its Turtle representation is `ex:tintin rdf:type ex:Person ..`
2. The RDF graph exposed in Listing 1 is composed of 5 triples written in Turtle format. In our formalism, the second triple, `_:tintin foaf:name "Tintin"`, is $(_:tintin, foaf:name, "Tintin"^{xsd:string})$.
3. $(ex:tintin, ex:travelsWith, ex:snowy)$ is an element of $RdfTriples$.
 $((ex:tintin, ex:travelsWith, ex:snowy), ex:since, "1978"^{xsd:integer})$ is an element of $RdfTriples$ that has a quoted triple in subject position.

Definition 10 (Term membership). The \in operator is extended to triples to check if a term is part of a triple.

$$\forall term \in I \cup B \cup L \cup RdfTriples, \forall (s, p, o) \in RdfTriples, term \in (s, p, o) \Leftrightarrow \left[\begin{array}{l} term = s \\ \vee (s \in RdfTriples \wedge term \in s) \\ \vee term = p \\ \vee term = o \\ \vee (o \in RdfTriples \wedge term \in o) \end{array} \right]$$

Example 7 (Term membership examples).

- $rdf:type \in (ex:tintin, rdf:type, ex:Person)$.
- $ex:snowy \notin (ex:tintin, rdf:type, ex:Person)$.
- $_:n \in (_:n, rdf:type, ex:Person)$
- $_:e \notin (_:n, rdf:type, ex:Person)$

⁶For the sake of readability, although RDF-star is not yet part of the official RDF recommendation [4], we conflate RDF-star and RDF in this paper. When we mention an RDF triple or an RDF graph, we allow them to contain quoted triples.

- 1 – $xsd:string \in (xsd:string, ex:p, ex:o)$
- 2 – $xsd:string \notin (ex:tintin, ex:name, "Tintin"^{xsd:string})$
- 3 – $ex:tintin \in ((ex:tintin, ex:travelsWith, ex:snowy), ex:since, "1978"^{xsd:integer})$

Definition 11 (List of blank nodes used in an RDF graph). For every RDF graph rdg , $BNodes(rdg)$ is the set of blank nodes in rdg i.e. $\forall rdg \subseteq RdfTriples, BNodes(rdg) = \{bn \in B \mid \exists t \in rdg, bn \in t\}$.

Example 8. Let GTT be the RDF graph exposed on Listing 1. $BNodes(GTT) = \{_:tintin, _:snowy\}$

4. PRSC: mapping PGs to RDF graphs

PRSC enables the user to convert any Property Graph to an RDF graph by using user-defined templates.

4.1. Property graphs with blank nodes

By default, every node and edge of the PG to convert is mapped by PRSC to an arbitrary fresh blank node. This mapping is arbitrary because blank nodes are essentially interchangeable, and because they have no global identifiers which would allow to map to a specific blank node anyway. This mapping is used throughout the transformation. For the sake of conciseness, we introduce the following modelling trick.

We note that in the respective definitions of PGs and RDF, the sets N and E of nodes and edges (in any PG) and the global set B of blank nodes (in RDF), are very loosely characterized. The only constraints are that N and E are disjoint and finite, and that B is disjoint from the sets of IRIs and literals. Theoretically, nothing prevents a property graph to take its nodes and edges in the set B , in other words, to have $N \subset B$ and $E \subset B$.

Our arbitrary mapping from $N \cup E$ to fresh blank nodes can then be used as a renaming function to build a new PG that is isomorphic to the original one. As noted in Section 3.3.1, the two PGs are indistinguishable for any practical purpose, including the transformation to an RDF graph. Without loss of generality, all the definitions and algorithms in this paper expect PGs whose nodes and edges are elements of B . This amounts to assume that the arbitrary mapping to blank nodes has been set in advance, and is carried by the PG to transform: every element of $N \cup E$ is identified with the blank nodes it maps to.

Conversely, when we study the reversibility of some contexts in Section 5, we prove that the produced BPG is exactly the original one. Producing a PG with arbitrary elements, and proving that it was isomorphic to the original, would be much more complex.

Definition 12 (Blank Node Property Graph). $BPGs$ is the set of property graphs with blank nodes only (BPG), i.e. $BPGs = \{pg \in PGs \mid (N_{pg} \cup E_{pg}) \subseteq B\}$.

Example 9. By defining the renaming function $\phi_{BTT} = \{n_1 \mapsto _:n1; n_2 \mapsto _:n2; e_1 \mapsto _:e1\}$, it is possible to build the BPG $BTT = rename(\phi_{BTT}, TT)$:

- 39 – $N_{BTT} = \{_:n1, _:n2\}; E_{BTT} = \{_:e1\}$
- 40 – $src_{BTT} = \{_:e1 \mapsto _:n1\}; dest_{BTT} = \{_:e1 \mapsto _:n2\}$
- 41 – $labels_{BTT} = \{_:n1 \mapsto \{\text{"Person"}\}; _:n2 \mapsto \emptyset; _:e1 \mapsto \{\text{"TravelsWith"}\}\}$
- 42 – $properties_{BTT} = \left\{ \begin{array}{l} (_:n1, \text{"name"}) \mapsto \text{"Tintin"}; (_:n1, \text{"job"}) \mapsto \text{"Reporter"} \\ (_:n2, \text{"name"}) \mapsto \text{"Snowy"}; (_:e1, \text{"since"}) \mapsto 1978 \end{array} \right\}$

By construction, BTT is isomorphic to TT , and as $N_{BTT} \cup E_{BTT} \subseteq B$, $BTT \in BPGs$.

This may raise the question of the meaning of two BPGs that share the same blank nodes, in particular if one blank node is used as a node in one BPG and as an edge in the other. However, this question could also be raised for all PGs: what would it mean for two PGs to have common elements in their respective sets of nodes and edges? In general, it would not hold any semantics, as PG elements are considered locally for a given PG. However, in Section 5.3.4, we will define operators for PGs, even the ones without blank nodes, that will consider PGs with shared PG elements and their possible relations.

4.2. Type of a PG element and PG schemas

We define the type of a PG element and PG schemas as follows. Let pg be a PG.

Definition 13 (Property keys of an element). We recall that in Section 2 and in Definition 3, properties are described as key-value pairs.

$keys_{pg}$ is the function that maps a PG element (node or edge) of the PG pg to the list of property keys for which it has a value, i.e. $keys_{pg} : N_{pg} \cup E_{pg} \rightarrow 2^{Str}$, with $\forall m \in (N_{pg} \cup E_{pg}), keys_{pg}(m) = \{key \mid properties_{pg}(m, key) \text{ is defined}\}$.

Definition 14 (Type of a PG element). A type is a triple composed of 1) the *kind* of the PG element, i.e. if it is a node or an edge, 2) a set of labels and 3) a set of property keys. The set of all types is denoted $Types$ and is defined as $Types = (\{“node”, “edge”\} \times 2^{Str} \times 2^{Str})$.

The type of an element $m \in N_{pg} \cup E_{pg}$ is

$$typeof_{pg}(m) = \left(\begin{cases} “node” & \text{if } m \in N_{pg} \\ “edge” & \text{if } m \in E_{pg} \end{cases}, labels_{pg}(m), keys_{pg}(m) \right)$$

A set of PG types is named a *schema*.

The functions *kind*, *labels* and *keys* are defined for types such that $\forall type = (u, l, key) \in Types, kind(type) = u, labels(type) = l, keys(type) = key$.

Example 10. Table 3 shows the types of the PG elements in the running example.

m	$typeof_{BTT}(m)$
$_:n1$	$(“node”, \{“Person”\}, \{“name”, “job”\})$
$_:n2$	$(“node”, \emptyset, \{“name”\})$
$_:e1$	$(“edge”, \{“TravelsWith”\}, \{“since”\})$

Remark 2. If two PGs pg and pg' are isomorphic, their elements share the same type.

Indeed, by definition, there exists a renaming function ϕ from the elements of pg to the elements of pg' , and $\forall m \in N_{pg} \cup E_{pg}, typeof_{pg}(m) = typeof_{pg'}(\phi(m))$.

4.3. Template triples

PRSC resorts to a mechanism of templating: to produce an RDF graph from a PG, we use tuples of three elements, named template triples, that will be mapped to proper RDF triples.

Definition 15 (Placeholders). There are four distinct elements, not included in either of the previously defined sets, named *valueOf*, *?self*, *?source* and *?destination*⁷.

Let $pvars = \{?self, ?source, ?destination\}$. $pvars$ elements serve as placeholders that will be replaced by the blank nodes that represent the nodes and edges in the PG.

Let $P = \{(l, valueOf) \mid l \in Str\}$. Elements of P can be noted with the same syntax as literals, for example “name”^{valueOf} is the pair (“name”, valueOf). Each element of P serves as a placeholder to be replaced with an RDF literal that represents the value of a property in the PG.

⁷In practice, our implementation of this paper maps *valueOf* to the IRI `prec:valueOf` and all terms prefixed with `?` to the `pvar` namespace. Examples given in Turtle reflect the implementation instead of fully fitting the theoretical definitions.

Definition 16 (Template triples). A *template triple* is a member of *Templates* and is defined as follows:

- $\forall \text{subject} \in I \cup \text{pvars}, \forall \text{predicate} \in I, \forall \text{object} \in I \cup \text{pvars} \cup L \cup P, (\text{subject}, \text{predicate}, \text{object}) \in \text{Templates}.$
- $\forall \text{subject} \in \text{Templates}, \forall \text{object} \in \text{Templates},$ and for all *subject*, *predicate* and *object* defined as above, $(\text{tsubject}, \text{predicate}, \text{object}), (\text{subject}, \text{predicate}, \text{tobject})$ and $(\text{tsubject}, \text{predicate}, \text{tobject})$ are members of *Templates*.

Note that unlike RDF triples, the elements of *Templates* can not contain blank nodes but can contain placeholders: *pvars* members can be used in subject (first) and/or object (third) position as they will be mapped later to blank nodes, and members of *P* are allowed in object position as they will be mapped later to literals. Similarly to RDF triples, template triples can not contain themselves and can not be nested infinitely.

Any subset of *Templates* is named a **template graph**. The PRSC engine will use template graphs to produce RDF graphs.

Example 11. The triple $(\text{ex:tintin}, \text{rdf:type}, \text{ex:Person})$ is both an element of *RdfTriples* and an element of *Templates*. Indeed, both *RdfTriples* and *Templates* allow IRIs in the subject, predicate and object positions.

The triple $(\text{?self}, \text{rdf:type}, \text{ex:Person})$ is a member of *Templates* but not of *RdfTriples*. The first member is an element of *pvars*, which is only allowed in template triples. The two other members are IRIs, and $\text{pvars} \times I \times I$ is a subset of *Templates*.

The triple $(\text{ex:tintin}, \text{ex:name}, \text{"name"}^{\text{valueOf}})$ is also a member of *Templates* but not of *RdfTriples*. While the first two members are IRIs, the third member is an element of *P*. Again, elements of *P* are only allowed by template triples.

Definition 17 (Placeholders and template triple membership). The \in operator, which we extended in Definition 10, is further extended to placeholders and template triples:

$$\forall \text{term} \in I \cup B \cup L \cup \text{RdfTriples} \cup \text{Templates}, \forall (s, p, o) \in \text{RdfTriples} \cup \text{Templates},$$

$$\text{term} \in (s, p, o) \Leftrightarrow \left[\begin{array}{l} \text{term} = s \\ \vee (s \in \text{RdfTriples} \cup \text{Templates} \wedge \text{term} \in s) \\ \vee \text{term} = p \\ \vee \text{term} = o \\ \vee (o \in \text{RdfTriples} \cup \text{Templates} \wedge \text{term} \in o) \end{array} \right]$$

4.4. PRSC context

In this paper, the notion of PRSC context is the keystone to let the user drive the conversion from a PG to an RDF graph. It maps PG types to template graphs. The *prsc* algorithm proceeds by looping on each node and edge of the PG, computing its type, finding the associated template graph in the context, and replacing the placeholders of this template graph with data extracted from the PG to produce an RDF graph.

Definition 18 (PRSC Context). A PRSC context $\text{ctx} : \text{Types} \rightarrow 2^{\text{Templates}}$ is a partial function that maps types to template graphs.

All template graphs must be *valid*, i.e. for all types, the placeholders used in the associated template graph must be consistent with the type: (1) for any given property key, its associated placeholder may only be used in template graphs associated with types that contain the property key, for example the placeholder $\text{"name"}^{\text{valueOf}}$ may only be used if the property key "name" is in the type associated to this template; (2) and templates associated to node types are not allowed to use the placeholders ?source and ?destination , as they are related to the source or the destination of an edge.

Formally, all template graphs used by a context *ctx* are valid iff $\forall \text{type} \in \text{Dom}(\text{ctx})$:

1. $\forall (\text{key}, \text{valueOf}) \in P, (\exists \text{tp} \in \text{ctx}(\text{type}) \mid (\text{key}, \text{valueOf}) \in \text{tp}) \Rightarrow \text{key} \in \text{keys}(\text{type}).$

2. $(\text{kind}(\text{type}) = \text{"node"}) \Rightarrow [\nexists tp \in \text{ctx}(\text{type}) \mid ?\text{source} \in tp \vee ?\text{destination} \in tp]$.

The set of all context functions is denoted Ctx .

Definition 19 (Complete PRSC contexts for a given PG). A PRSC context is said complete for a property graph $pg \in \text{BPGs}$ iff there is a template graph defined for each type used in pg . The set of all complete contexts for a PG pg is noted $\text{Ctx}_{pg} = \{\text{ctx} \in \text{Ctx} \mid \forall m \in N_{pg} \cup E_{pg}, \text{typeof}_{pg}(m) \in \text{Dom}(\text{ctx})\}$.

Remark 3. Note that the type system described in Definition 14 is trivial to resolve as the type of a PG element m , denoted by $\text{typeof}_{pg}(m)$, only depends on its kind (node or edge), the list of its labels and the list of its property keys. For this reason, checking if a PRSC context ctx is complete for a property graph pg is also trivial as it consists in computing the type of each PG element of pg and checking if all types are in the domain of ctx .

Example 12. Table 4 exposes an example of a complete ctx context function for our running example. First, the function is a context as all template graphs are valid: The placeholders "name" valueOf and "job" valueOf are only used in types with the associated property key. The fact that the property key "since" in the third type has no associated placeholder occurrence in the template graph does not invalidate the context. The placeholders $?source$ and $?destination$ are only used in the third type, which is an edge type. Then, all three types used by our running example have an associated template graph, so it is a complete context for the PG exposed in Example 1.

Table 4
An example of a complete context for the Tintin Property Graph.

<i>type</i>	<i>ctx(type)</i>
	$(?self, rdf:type, ex:Person)$
$(\text{"node"}, \{\text{"Person"}\}, \{\text{"name"}, \text{"job"}\})$	$(?self, foaf:name, \text{"name" valueOf})$ $(?self, ex:profession, \text{"job" valueOf})$
$(\text{"node"}, \emptyset, \{\text{"name"}\})$	$(?self, foaf:name, \text{"name" valueOf})$
$(\text{"edge"}, \{\text{"TravelsWith"}\}, \{\text{"since"}\})$	$(?source, ex:isTeammateOf, ?destination)$

Example 13. The function ctx exposed in Table 5 is not complete for the PG BTT as its domain lacks the type of $_:n2$ and the type of $_:e1$.

Table 5
An incomplete context for the Tintin PG

<i>type</i>	<i>ctx(type)</i>
	$(?self, rdf:type, ex:Person)$
$(\text{"node"}, \{\text{"Person"}\}, \{\text{"name"}, \text{"job"}\})$	$(?self, foaf:name, \text{"name" valueOf})$ $(?self, ex:profession, \text{"job" valueOf})$

Example 14. The function ctx exposed in Table 6 is not a context because "surname" , which is used in the template graph mapped to the first listed type $(\text{"node"}, \{\text{"Person"}\}, \{\text{"name"}, \text{"job"}\})$, is not a value in $\{\text{"name"}, \text{"job"}\}$.

Table 6
A function that is not a context

<i>type</i>	<i>ctx(type)</i>
$(\text{"node"}, \{\text{"Person"}\}, \{\text{"name"}, \text{"job"}\})$	$(?self, ex:familyName, \text{"surname" valueOf})$
$(\text{"node"}, \emptyset, \{\text{"name"}\})$	$(?self, foaf:name, \text{"name" valueOf})$
$(\text{"edge"}, \{\text{"TravelsWith"}\}, \{\text{"since"}\})$	$(?source, ex:isTeammateOf, ?destination)$

4.5. Application of a PRSC context on a PG

We now define formally the conversion operated by PRSC. A PRSC conversion of a PG depends on a chosen context $ctx \in Ctx$.

Definition 20 (Property value conversion). For the conversion of property values to literals, we consider that we have a fixed total injective function $toLiteral : V \rightarrow L$, common for all PGs and contexts. We suppose that $toLiteral$ is reversible, i.e. we are able to compute $toLiteral^{-1}$.

Definition 21 (The $prsc$ function). The operation that produces an RDF graph from the application of a PRSC context $ctx \in Ctx_{pg}$ on a property graph $pg \in BPGs$ is noted $prsc(pg, ctx)$. The result of the $prsc$ function is the union of the RDF graph built by converting all elements of the PG, into RDF. The conversion of a single element is materialized by the $build$ function.

$\forall tps \subseteq Templates, \forall pg \in BPGs, \forall m \in N_{pg} \cup E_{pg}, build(tps, pg, m) = \{\beta_{pg,m}(tp) \mid tp \in tps\}$ with $\beta_{pg,m}$ defined as follows:

$$\beta_{pg,m} : \begin{cases} Templates & \rightarrow RdfTriples \\ P \cup L & \rightarrow L \\ I & \rightarrow I \\ pvars & \rightarrow B \end{cases}$$

$$\beta_{pg,m}(x) = \begin{cases} (\beta_{pg,m}(x_s), \beta_{pg,m}(x_p), \beta_{pg,m}(x_o)) & \text{if } x = (x_s, x_p, x_o) \in Templates \\ x & \text{if } x \in L \cup I \\ m & \text{if } x = ?self \\ src_{pg}(m) & \text{if } x = ?source \wedge m \in E_{pg} \\ dest_{pg}(m) & \text{if } x = ?destination \wedge m \in E_{pg} \\ toLiteral(properties_{pg}(m, key)) & \text{if } x = (key, valueOf) \in P \\ undefined & \text{otherwise} \end{cases}$$

As said previously, the result of $prsc$ is the union of the graphs produced by $build$, i.e.

$$prsc(pg, ctx) = \bigcup_{m \in N_{pg} \cup E_{pg}} build(ctx(typeof_{pg}(m)), pg, m)$$

Example 15. Table 7 exposes the resolution of $prsc$ on the running example.

Table 7
Application of a PRSC context on the running example

m	$typeof_{BTT}(m)$	$ctx(typeof_{BTT}(m))$	$build(ctx(typeof_{BTT}(m)), BTT, m)$
$_ :n1$	("node", {"Person"}, {"name", "job"})	(?self, rdf:type, ex:Person) (?self, foaf:name, "name" ^{valueOf}) (?self, ex:profession, "job" ^{valueOf})	(_ :n1, rdf:type, ex:Person) (_ :n1, foaf:name, "Tintin") (_ :n1, ex:profession, "Reporter")
$_ :n2$	("node", \emptyset , {"name"})	(?self, foaf:name, "name" ^{valueOf})	(_ :n2, foaf:name, "Snowy")
$_ :e1$	("edge", {"TravelsWith"}, {"since"})	(?source, ex:isTeammateOf, ?destination)	(_ :n1, ex:isTeammateOf, _ :n2)

Algorithm 1: The *prsc* function**Input:** $pg \in BPGs, ctx \in Ctx_{pg}$ **Output:** An RDF graph**1 Main Function *prsc*(pg, ctx):****2** $rdf \leftarrow \{\}$ **3 forall PG element $m \in N_{pg} \cup E_{pg}$ do****4** $tps \leftarrow ctx(typeof_{pg}(m))$

/* build function */

5 $built \leftarrow \{\}$ **6 forall $tp \in tps$ do****7** $built \leftarrow built \cup \{\beta(tp, pg, m)\}$ /* $\beta_{tp,pg}(m)$ is necessarily defined, since $ctx \in Ctx_{pg}$ */**8** $rdf \leftarrow rdf \cup built$ **9** **return** rdf /* In the formal definition, pg and m are implicitly passed to β */**10 Function $\beta(tp, pg, m)$:****11** **if** $tp \in Templates$ **then****12** $(s, p, o) \leftarrow tp$ **13** **return** $(\beta(s, pg, m), \beta(p, pg, m), \beta(o, pg, m))$ **14** **else if** $tp \in L$ **then return** tp **15** **else if** $tp \in I$ **then return** tp **16** **else if** $tp \in P$ **then****17** $(key, valueOf) \leftarrow tp$ **18** **return** $toLiteral(properties_{pg}(m, key))$ **19** **else if** $tp \in pvars$ **then****20** **if** $tp = ?self$ **then return** m **21** **if** $m \in E_{pg}$ **then****22** **switch** tp **do****23** **case** $?source$ **do return** $src_{pg}(m)$ **24** **case** $?destination$ **do return** $dest_{pg}(m)$ /* $\beta_{tp,pg}(m)$ is undefined according to Definition 21 */**25** **raise** $Error(Undefined\ behavior)$ The resolution of $_ :n2$ is as follows: $build(ctx(typeof_{BTT}(_ :n2), BTT, _ :n2))$ $= build(ctx(("node", \emptyset, \{"name"\}), BTT, _ :n2))$

Resolution of the type

 $= build(\{(?self, foaf:name, "name"^{prec:valueOf})\}, BTT, _ :n2)$ Application of ctx $= \{(_ :n2, foaf:name, toLiteral(properties_{BTT}(_ :n2, "name")))\}$ Application of $build$ $= \{(_ :n2, foaf:name, toLiteral("Snowy"))\}$

Evaluation of the property "name"

 $= \{(_ :n2, foaf:name, "Snowy"^{xsd:string})\}$ Application of $toLiteral$ Algorithm 1 gives an algorithmic view of the *prsc* function presented by Definition 21.

4.6. Complexity analysis

In this section, we discuss the different metrics that can be used to evaluate the complexity and evaluate the complexity of the *prsc* function.

4.6.1. Functions considered constant

For a given PG *pg*, the complexity of the functions *src_{pg}*, *dest_{pg}*, *labels_{pg}* and *properties_{pg}* are considered constant.

The complexity of the functions *keys_{pg}*, the *toLiteral* and the *toLiteral*⁻¹ is also considered constant.

Evaluating if something is a member of a given set, for example if an entity is part of the set *N_{pg}*, is generally considered to be constant time thanks to hash maps⁸.

4.6.2. Considered metrics

For a given PG *pg* and a given context *ctx*, the following metrics are considered:

- The number of nodes and edges in *pg*, denoted *NbOfPGElements*.

$$NbOfPGElements = |N_{pg} \cup E_{pg}|$$

- The size of the biggest template graph, denoted *BiggestTemplateSize*.

$$BiggestTemplateSize = \max_{type \in Dom(ctx)} (|ctx(type)|)$$

- The complexity of the types in the context, denoted *TypeComplexity*, reflected by the number of labels and the number of properties of the type with the highest number of labels and properties. Note that since the context has to be valid, *i.e.* all elements of the PG must have their type in the context, *TypeComplexity* is also an upper bound of the type complexity of the types in the PG.

$$TypeComplexity = 1 + \max_{type \in Dom(ctx)} (|labels_{pg}(type)| + |keys_{pg}(type)|)$$

- The number of types supported by the context *NbTypes* = $|Dom(ctx)|$.

In RDF-star, quoted triples can be used as subject or object of other triples, without limit on how deeply triples can be nested. In practice, however, it is rare to have more than one level of nesting. Usually, users are expected to use atomic RDF triples like `_:tintin :travelsWith _:haddock` or to use RDF-star triples with a depth of one like `<< _:tintin :travelsWith _:haddock >> :since 1978`. We therefore consider the *depth* of any triple to be bound by a constant. As a consequence, in all functions processing terms and triples recursively (such as β in Algorithm 1), we can ignore the recursion depth in the complexity analysis.

In all complexity analyses, all metrics are considered non null. Indeed, if there are no elements or if the biggest template graph is empty, the produced RDF graph will be empty so this case is not interesting. As we add one to the number of labels and properties, the type complexity can never be zero, even if the context only supports nodes and edges with no labels and no properties.

4.6.3. Complexity of *ctx* calls in the *prsc* function

For each given PG element *m*, the complexity of a call to *ctx(typeof_{pg}(*m*))* is $\mathcal{O}(TypeComplexity * \ln(TypeComplexity))$:

- The type of *m* in the PG *pg* must be computed. *typeof_{pg}(*m*)* has a complexity of $\mathcal{O}(1)$:
 - * Evaluating if *m* is a node or an edge is constant as checking if an element is a member of a set is constant.
 - * Calls to *labels_{pg}(*m*)* and *keys_{pg}(*m*)* are considered constants in Section 4.6.1.

⁸Inserting and searching in a hash map is not strictly speaking a constant time operation but has an *amortized* constant complexity, and is linear in the worst case.

- In the complexity analysis, we consider that ctx is implemented as a hash map from the types to the template graphs. A ctx call would first need to compute the hash of the type. To do so, it has to look at all the labels and properties in the type in a deterministic order; for this, the labels and keys needs to be sorted, which has a complexity of $\mathcal{O}(TypeComplexity * \ln(TypeComplexity))$. After the hash has been computed, the cost of retrieving the template graph has an amortized constant complexity. The overall complexity of a ctx call is $\mathcal{O}(TypeComplexity * \ln(TypeComplexity))$.

4.6.4. Complexity of $prsc$

Given a PG $pg \in BPGs$ and a context $ctx \in Ctx_{pg}$,

- Calls to $ctx(type_{pg}(m))$ in line 4 have a complexity of $\mathcal{O}(TypeComplexity * \ln(TypeComplexity))$.
- Calls to the β function in line 7 are constant, as it has been assumed that the depth of the most nested triple is low enough to be ignored and the operations it performs are in constant time.
- There are two for loops, one iterating on all PG elements ($NbOfPGElements$) and one iterating on all template triples of a template graph ($BiggestTemplateSize$). All instructions in the $prsc$ but the one on line 4 are computed in constant time. Line 4 is inside the first loop but outside the second loop.

The $prsc$ function has an $\mathcal{O}(NbOfPGElements * (BiggestTemplateSize + TypeComplexity * \ln(TypeComplexity)))$ complexity.

5. PRSC reversibility

When PGs are converted into RDF graphs, an often desired property is to not have any information loss. To determine whenever or not a conversion induces information loss is to check if the conversion is reversible, *i.e.* if from the output, it is possible to compute back the input. The reversion is studied relatively to the used PRSC contexts: the PRSC context is used as both an input of both the PRSC algorithm and the reversion algorithm. In other words, we consider that the information stored in the PRSC context do not need to be stored in the produced RDF graph to produce a reversible conversion.

This section first shows that not all PRSC contexts are reversible. Then, properties are exhibited about PRSC contexts, leading to a description of a subset of reversible PRSC contexts, *i.e.* contexts that we prove do not induce information loss.

5.1. Reversibility in this paper

In this paper, we call a function f reversible if we can find back x in practice from $f(x)$. This implies that:

- The function f must be injective. Indeed, if two different values x and x' can produce the same value y , it is impossible to know if the value responsible for producing y was x or x' .
- The inverse function f^{-1} must be computable and tractable in reasonable time. By counter-example, a public-key encryption function is supposed to be injective. It is theoretically possible, although prohibitively costly, to decipher a given message by applying the encryption function on all possible inputs until the result is the original encrypted message. This is not the kind of “reversibility” we are interested in.

We say that a context ctx is reversible if for any PG $pg \in BPGs$ such that the context ctx is complete for the PG pg , it is possible to find back pg from the context ctx and the result of $prsc(pg, ctx)$.

More formally, when studying reversibility, we want to check if for a given $ctx \in Ctx$, we are able to define a tractable function $prsc_{ctx}^{-1}$ such that $\forall pg \in BPGs, [ctx \in Ctx_{pg} \Rightarrow prsc_{ctx}^{-1}(prsc(pg, ctx)) = pg]$.

Example 16 (A trivially non-reversible context). Consider the context ctx_{\emptyset} such that for all types, it returns the empty template graph, *i.e.* $\forall type \in Types, ctx_{\emptyset}(type) = \emptyset$. As it is complete for all property graphs, it is possible to use this context on any property graph. However, applying the context ctx_{\emptyset} produces the empty RDF graph. Therefore, the use of the context ctx_{\emptyset} makes the function $prsc$ not injective, and therefore not reversible.

Example 17 (A more realistic example of a non-reversible context). Another example of a non-reversible context is the context exposed in Table 4: while this context can be applied on PGs in which edges have the “since” property, the value of this property will never appear in the produced RDF graph.

As not all contexts are reversible, the next sections focus on characterizing some contexts that produce reversible conversions.

5.2. Well-behaved contexts

5.2.1. Characterization function

To be able to reverse back to the original PG, we need a way to distinguish the triples that may have been produced by a given member of *Templates* from the ones that cannot have been produced by it. For this purpose, this section introduces the κ function. This function must verify that, for every triple template tp and every triple t , $\kappa(t) = \kappa(tp)$ if and only if t can be produced from tp by the β function. It would then follow that two template triples that may produce the same triple have the same image through κ .

Definition 22 (Characterization function).

The κ function maps:

- All template triples to a super set of triples that it is able to generate.
- All RDF triples t to a super-set of the RDF triples that a template triple that may generate the triple t may also generate. For example, a literal may be generated by any element of P . An element of P may generate any literal. Therefore, the κ function maps all literals to the set of all literals.

$$\kappa : \begin{cases} Templates \cup RdfTriples & \rightarrow 2^{RdfTriples} \\ L \cup P & \rightarrow \{L\} \\ I & \rightarrow 2^I \\ B \cup pvars & \rightarrow \{B\} \end{cases}$$

$$\kappa(x) = \begin{cases} \kappa(s) \times \kappa(p) \times \kappa(o) & \text{if } x = (s, p, o) \in RdfTriples \cup Templates \\ L & \text{if } x \in L \cup P \\ \{x\} & \text{if } x \in I \\ B & \text{if } x \in B \cup pvars \end{cases}$$

The κ function is extended to all template graphs and RDF graphs xs as $\kappa(xs) = \bigcup_{t \in xs} \kappa(t)$.

Example 18 (κ applied to the running example from Figure 1).

- $\kappa(?source) = B$, $\kappa(_:nI) = B$.
- $\kappa(foaf:name) = \{foaf:name\}$.
- $\kappa("name"^{valueOf}) = L$, $\kappa("Tintin") = L$.
- $\kappa((?self, foaf:name, "name"^{valueOf})) = B \times \{foaf:name\} \times L$.
- $\kappa((_:nI, foaf:name, "Tintin")) = B \times \{foaf:name\} \times L$.
- Note that
 - * $\kappa((_:nI, foaf:name, "Tintin")) = \kappa((?self, foaf:name, "name"^{valueOf}))$
 - * $(_:nI, foaf:name, "Tintin") \in \kappa((?self, foaf:name, "name"^{valueOf}))$
- $\kappa((?source, ex:isTeammateOf, ?destination)) = B \times \{ex:isTeammateOf\} \times B$

$$\begin{aligned} & - \kappa(((?source, ex:isTeammateOf, ?destination), ex:since, "since"^{valueOf})) \\ & = (B \times \{ex:isTeammateOf\} \times B) \times \{ex:since\} \times L \end{aligned}$$

Table 8 provides an example of applying κ on the running example context of Table 4.

Remark 4 (κ on terms and triples is, as expected, a super-set of the possible generated values). When comparing the definition of the κ function with the β functions defined in Section 4.5, it appears that:

- For elements in B , $pvars$, L and P , the image of κ is equal to the corresponding image set of the β function.
- For elements in I , the image of κ is equal to a singleton containing that element; β maps any IRI to itself.
- If the given term is a triple, the image of κ is the cross product of the application of the κ function to the terms that compose the RDF triple. As β on triples recursively applies itself to the three terms in the triple, we can see that $\forall \beta, \forall triple, \beta(triple) \in \kappa(triple)$.

Therefore, **if x is a term or an RDF triple, for any β function, $\beta(x) \in \kappa(x)$.**

Remark 5 (The result of *build* is, as expected, a subset of the result of κ). The *build* function, from which *prsc* is defined, uses β on each template triple. After β is applied, the union of the singletons containing each triple is computed. This is similar to the definition of κ on a set of triples.

From Remark 4, it can be deduced that **if tps is a set of template triples, $\forall pg, \forall m, build(tps, pg, m) \subseteq \kappa(tps)$.**

Remark 6 (A template and its produced values share the same image through κ). When using the κ function, elements in B and $pvars$ both map to B , and elements in L and P both map to L . Elements in I are wrapped into a singleton and both *RdfTriples* and *Templates* apply the function recursively on their members.

When using the β function:

- Elements in $pvars$ map for all PGs $pg \in BPGs$ to elements of N_{pg} and E_{pg} , which are both subsets of B .
- Elements in P map to elements in $Img(toLiteral)$, which is a subset of L .
- Elements in L and I are mapped to themselves.
- Elements in *Templates* apply the β function recursively on their members.

Therefore, $\forall tp \in Templates, \kappa(\beta(tp)) = \kappa(tp)$

As mentioned previously, the role of κ is to allow us to determine whether two template triples with placeholders may produce the same triple. It maps all placeholders to a super-set⁹ of all elements they can generate with the *build* function. All RDF Triples are mapped by the κ function to a subset of *RdfTriples* they are a member of.

Lemma 1. If a triple is generated by a template graph, then there exists a template triple with the same image through κ .

$$\forall pg \in BPGs, \forall m \in (N_{pg} \cup E_{pg}), \forall tps \subseteq Templates, \forall td \in build(tps, pg, m), \exists tp \in tps \mid \kappa(td) = \kappa(tp)$$

Proof. Per the Definition 21 of *build*, a triple can only be generated by a template graph by the application of β to one of its template triples. Per Remark 6, the generated triple and the corresponding template triple have the same image through κ . \square

Definition 23 (*unique* template triple). A template triple tp is *unique* in a set of template triples if no other template triple in the set has the same image through κ equal to tp .

It is defined as follows with $tp \in tps \subset Templates$:

$$unique(tp, tps) = (\forall tp' \in tps, \kappa(tp) = \kappa(tp') \Leftrightarrow tp = tp')$$

⁹Note that as κ maps to a super set, it may catch false positives. For example, P can only generate elements in $Img(toLiteral)$, but the κ function considers that all elements of L can be generated from P . For the scope of this paper, κ catching false positives is considered acceptable, as we are only trying to prove the reversibility of a given class of contexts, rather than to characterize the whole class of reversible contexts.

Combined with Remark 6, what $unique(tp, tps)$ tells us is that any triple, with the same image through κ as tp , can not have been generated by any other element of tps than tp itself. This leads us to Theorem 1 below.

Theorem 1 (Triples produced by a *unique* template triple). In the result of the *build* function, if a data triple and a *unique* template triple have the same value through κ , then the data triple must have been produced by this template triple.

$\forall pg \in BPGs, \forall ctx \in Ctx_{pg}, \forall m \in (N_{pg} \cup E_{pg}),$ let $tps = ctx(typeof_{pg}(m)), \forall td \in build(tps, pg, m), \forall tp \in tps:$

$$unique(tp, tps) \wedge \kappa(td) = \kappa(tp) \Rightarrow td \in build(\{tp\}, pg, m)$$

Proof. We prove the theorem by contradiction.

Let us suppose that:

- (A) $td \in build(tps, pg, b)$
- (B1) $unique(tp, tps)$, i.e. $(\forall tp' \in tps, \kappa(tp) = \kappa(tp') \Rightarrow tp = tp')$
- (B2) $\kappa(td) = \kappa(tp)$
- (C) $td \notin build(\{tp\}, pg, b)$

$$td \in build(tps - \{tp\}, pg, b) \quad [(A) \text{ and } (C)]$$

$$\Rightarrow \exists tdp \in tps - \{tp\}, \kappa(tdp) = \kappa(td) \quad [\text{Lemma 1}]$$

$$\Rightarrow \exists tdp \in tps - \{tp\}, \kappa(tdp) = \kappa(tp) \quad [(B2)]$$

$$\Rightarrow \exists tdp \in tps - \{tp\}, tdp = tp \quad [(B1)]$$

$$\Rightarrow tp \in tps - \{tp\}$$

tp can not be a member of the set $tps - \{tp\}$, as it explicitly exclude it. As we reached a contradiction, it means that $td \in build(\{tp\}, pg, b)$. □

Theorem 1 allows us to link an RDF triple to the unique template triple that produced it. Then by comparing the terms of the RDF triple to the corresponding placeholders in the template triple, we will be able to reconstruct the original PG.

5.2.2. Well-behaved PRSC context

In this section, we define a subset of contexts that we call *well-behaved PRSC contexts*. In the next section, we will prove that these contexts are reversible.

Definition 24 (Well-behaved contexts). A PRSC context ctx is well-behaved if it conforms to those 3 criteria:

$\forall type \in Dom(ctx),$ let $tps = ctx(type)$

1. *Element provenance*: all generated triples must contain the blank node that identifies the node or the edge it comes from. This is achieved by using the *?self* placeholder in all template triples:

$$- \forall tp \in tps, ?self \in tp$$

2. *Signature template triple*: tps contains at least one template triple, called its *signature* and noted $sign_{ctx}(type)$, that will produce triples that no other template in ctx can produce. This will allow, for each blank node in the produced RDF graph, to identify its type in the original PG.

$$- \exists sign_{ctx}(type) \in tps, \forall x \in Dom(ctx), \kappa(sign_{ctx}(type)) \subseteq \kappa(ctx(x)) \Rightarrow x = type$$

3. *No value loss*: for all elements in the PG, we do not want to lose information stored in properties, nor for edges, the source and destination node. Each of these pieces of information must be present in an unambiguously recognizable triple pattern.

- $\forall \text{key} \in \text{keys}(\text{type}), \exists \text{tp} \in \text{tps} \mid \text{unique}(\text{tp}, \text{tps}) \wedge (\text{key}, \text{valueOf}) \in \text{tp}$
- $\text{kind}(\text{type}) = \text{"edge"} \Leftrightarrow \exists \text{tp} \in \text{tps} \mid \text{unique}(\text{tp}, \text{tps}) \wedge ?\text{source} \in \text{tp}$
- $\text{kind}(\text{type}) = \text{"edge"} \Leftrightarrow \exists \text{tp} \in \text{tps} \mid \text{unique}(\text{tp}, \text{tps}) \wedge ?\text{destination} \in \text{tp}$

The set of all well-behaved contexts is Ctx^+ , and the set of all well-behaved contexts for a PG pg is Ctx_{pg}^+ .
 $\text{Ctx}^+ \subset \text{Ctx}$ and $\text{Ctx}_{\text{pg}}^+ = \text{Ctx}^+ \cap \text{Ctx}_{\text{pg}}$.

Remark 7 (Handling multiple sign_{ctx} candidates). In the case where there are multiple template triples candidates to become the signature template triple, the choice of the signature template triple among the candidates is generally not important.

To make the choice deterministic, it could be considered that the chosen signature template triple is the first in lexicographic order. In the case of the presented algorithms, the choice of the signature template triple is not important, and will lead to the same output.

Remark 8 (The template graphs used in well-behaved contexts are not empty). A well-behaved context cannot map a type to an empty template graph: the *signature template triple* criterion ensures that every template graph contains at least one template triple: $\forall \text{tps} \in \text{Img}(\text{ctx}), \exists \text{tp} \in \text{tps} \Leftrightarrow |\text{tps}| \geq 1$.

Remark 9 (Inside a well-behaved context, all template graphs are different from all others). For any well-behaved context ctx , two types cannot share the same template graph. Indeed, if two types share the same template graph, *i.e.* there are two types type1 and type2 with $\text{type1} \neq \text{type2}$ such that $\text{ctx}(\text{type1}) = \text{ctx}(\text{type2})$, it would contradict the *signature template triple* criterion as it would lead to $\text{type1} = \text{type2}$.

Example 19. Table 8 studies the context used in our running example, exposed in Example 12.

Table 8
The running example context with the corresponding values through κ

type	$\text{ctx}(\text{type})$	$\kappa(\text{ctx}(\text{type}))$
$\text{tn1} = (\text{"node"}, \{\text{"Person"}\}, \{\text{"name"}, \text{"job"}\})$	$(?self, \text{rdf:type}, \text{ex:Person})$ $(?self, \text{foaf:name}, \text{"name"}^{\text{valueOf}})$ $(?self, \text{ex:profession}, \text{"job"}^{\text{valueOf}})$	$(B \times \{\text{rdf:type}\} \times \{\text{ex:Person}\})$ $\cup (B \times \{\text{foaf:name}\} \times L)$ $\cup (B \times \{\text{ex:profession}\} \times L)$
$\text{tn2} = (\text{"node"}, \emptyset, \{\text{"name"}\})$	$(?self, \text{foaf:name}, \text{"name"}^{\text{valueOf}})$	$(B \times \{\text{foaf:name}\} \times L)$
$\text{te1} = (\text{"edge"}, \{\text{"TravelsWith"}\}, \{\text{"since"}\})$	$(?source, \text{:isTeammateOf}, ?destination)$	$(B \times \{\text{:isTeammateOf}\} \times B)$

- The type tn1 matches all criteria of a well-behaved PRSC context:
 - * All triples contain $?self$.
 - * At least one template triple is a signature: the image through κ of $(?self, \text{rdf:type}, \text{ex:Person})$ is not contained in the image through κ of other types. It is also the case of $(?self, \text{ex:profession}, \text{"job"}^{\text{valueOf}})$.
 - * The properties "name" and "job" have a *unique* template triple inside $\kappa(\text{ctx}(\text{tn1}))$.
- The type tn2 violates the *signature template triple* criterion as $(?self, \text{foaf:name}, \text{"name"}^{\text{valueOf}})$, its only template triple, is shared with the type tn1 ,
- The type te1 violates the *element provenance* criterion as $?self$ is missing. It also violates the *no value loss* criterion as the term $\text{"since"}^{\text{valueOf}}$ is missing from any template triple.

For all these reasons, this context is not well-behaved.

Example 20 (A well-behaved context for the running example). Let ctx_{TTWB} be the function described in Table 9. In this new context, an arbitrary ex:NamedEntity IRI is used to sign the PG nodes that have no labels and only a name, and a classic RDF reification is used to model the PG edges.

This context is well-behaved:

Table 9

An example of a complete and well-behaved context for the Tintin Property Graph.

<i>type</i>	<i>ctx(type)</i>
	(<i>?self</i> , <i>rdf:type</i> , <i>ex:Person</i>)★
("node", {"Person"}, {"name", "job"})	(<i>?self</i> , <i>foaf:name</i> , "name" <i>valueOf</i>)
	(<i>?self</i> , <i>ex:profession</i> , "job" <i>valueOf</i>)★
("node", ∅, {"name"})	(<i>?self</i> , <i>foaf:name</i> , "name" <i>valueOf</i>)
	(<i>?self</i> , <i>rdf:type</i> , <i>ex:NamedEntity</i>)★
	(<i>?self</i> , <i>rdf:subject</i> , <i>?source</i>)★
("edge", {"TravelsWith"}, {"since"})	(<i>?self</i> , <i>rdf:object</i> , <i>?destination</i>)★
	(<i>?self</i> , <i>rdf:predicate</i> , <i>ex:TravelsWith</i>)★
	(<i>?self</i> , <i>ex:since</i> , "since" <i>valueOf</i>)★

- *?self* appears in all triples,
- Template triples that are signature are marked with a ★. At least one signature triple appears for each type,
- All property keys have a *unique* template triple.

Listing 3 is the RDF graph produced by the application of the context ctx_{TTWB} on the PG *BTT*. Each part that starts with a # denotes the application of a *build* application to the PG element described in the comment. The elements are ordered in the same order as their type in Table 9, and the RDF triples and the template triples that produced them are also in the same order.

Listing 3 The RDF graph produced by the application of the well-behaved context ctx_{TTWB} on the running example PG *BTT*.

```

1 # From _:n1
2 _:n1 rdf:type ex:Person .
3 _:n1 foaf:name "Tintin" .
4 _:n1 ex:profession "Reporter" .
5 # From _:n2
6 _:n2 foaf:name "Snowy" .
7 _:n2 rdf:type ex:NamedEntity .
8 # From _:e1
9 _:e1 rdf:subject _:n1 .
10 _:e1 rdf:object _:n2 .
11 _:e1 rdf:predicate ex:TravelsWith .
12 _:e1 ex:since 1978 .

```

Remark 10 (Relationship between the empty PG and the empty RDF graph with well-behaved PRSC context). For all well-behaved PRSC contexts, the only PG that can produce the empty RDF graph is the empty PG:

$$\forall pg \in BPGs, ctx \in Ctx_{pg}^+, |prsc(pg, ctx)| = 0 \Leftrightarrow pg = pg_{\emptyset}$$

Indeed, Remark 8 ensures that the template graphs are non-empty. So any application of the *build* function with any well-behaved context produces at least one RDF triple. As the produced RDF graph is the union of the graphs produced by the use of *build* on each node and edge, the only way to have an empty result is to have no node nor edge in the property graph.

5.2.3. Complexity analysis and implementation

PRSC well-behaved contexts will be proved to be reversible in Section 5.3, meaning that producing an RDF graph from them will preserve all information stored in the PG. It is therefore important to be able to determine if in practice, it is possible to compute if a PRSC context is well-behaved.

Lemma 2. The values through κ of two given terms are either disjoint or equal:

Proof. Consider any atomic RDF term t :

- If $t \in I$, its value through κ is the singleton composed of the element t . Other terms can not map κ to a super-set of the singleton $\{t\}$, in particular no term can be mapped to I .
- If $t \in L \cup P$, the value through κ is L . No other term can be mapped to a super-set or a subset of L .
- If $t \in B \cup pvars$, the value through κ is B . No other term can be mapped to a super-set or a subset of B .

As L , I and B are pairwise disjoint, for two given atomic RDF terms, the value through κ is either disjoint or equal.

For two given RDF triples composed of atomic terms, their value through κ are equals to the Cartesian product of the value through κ of the components. As the values through κ of their components are either equals or disjoint, the values through κ of the triples are also either equals or disjoint. By induction, this is true for any two RDF triples, even if their subject or object are also triples. □

Remark 11 (Implementing κ and complexity analysis). The function κ is defined to return sets, some of them being infinite sets. While this definition is useful to prove different theorems in this paper, it is not practical from an implementation perspective.

Let λ and δ be two distinct values that are not members of the set I . We propose below an alternative function κ_{impl} to be used instead of κ in algorithms:

$$\kappa_{impl}(x) = \begin{cases} \{\kappa(triple) \mid triple \in x\} & \text{if } x \in RdfTriples \cup Templates \\ (\kappa_{impl}(s), \kappa_{impl}(p), \kappa_{impl}(o)) & \text{if } x = (s, p, o) \in RdfTriples \cup Templates \\ \lambda & \text{if } x \in L \cup P \\ x & \text{if } x \in I \\ \delta & \text{if } x \in B \cup pvars \end{cases}$$

Compared to Definition 22, we replaced:

- the singleton $\{x\}$ with x , in the case where $x \in I$,
- the sets L and B with two constants λ and δ that are not elements of I ,
- the cross product with a simple triple of the values returned for each element of x when x is a triple.

The complexity of the $\kappa_{impl}(x)$ is:

- For any x that is not a triple nor a graph, calls to this function can be done in constant time, by simply checking the type of x .
- When x is a triple, calls to this function involves recursive calls up to the depth of x , which we consider to be bounded by a constant (see Section 4.6.2). So it is also done in constant time.
- When x is a graph, calls to this function involves calling κ_{impl} on each triple of the graph. As the call on a triple is constant, the call on the graph x has a linear complexity depending on the size of the graph.

Note that:

- For two triples, checking if their value through κ_{impl} are equals can be done in constant time.
- Thanks to Lemma 2, checking if the value through κ_{impl} of a triple t is included in the value through κ_{impl} of a graph tps can be done in linear time by iterating on each triple tp of the graph tps and comparing the values through κ_{impl} of the triples t and tp .

Remark 12 (Complexity of checking if a PRSC context is a well-behaved). The first task to check if a context ctx is well-behaved consists in computing the value through κ of all triples used in it. As the depth of a template triple is considered to be negligible, the complexity is the number of template triples, bounded to the number of types multiplied by the size of the biggest template graph: $\mathcal{O}(NbTypes * BiggestTemplateSize)$.

After the value through κ of all template triples have been computed, for each type, we need to check if the type complies with the three criterion exposed in the Definition 24.

- 1 – The *element provenance criterion* consists in checking if *?self* is in all templates triples of all type. This task
2 has an $\mathcal{O}(1)$ complexity for each template triple and an overall $\mathcal{O}(NbTypes * BiggestTemplateSize)$ complexity
3 for the whole context.
- 4 – The *signature template triple* consists in checking if there is at least one signature template triple in the
5 template graph of all types, *i.e.* checking if the value through κ of one of the template triples of each type
6 is not contained in the set of the value through κ of the other types template graph. As hash sets make the
7 membership check constant, this task has an $\mathcal{O}(NbTypes)$ complexity for a single template triple candidate,
8 and an overall $\mathcal{O}(NbTypes * BiggestTemplateSize * NbTypes)$ for the whole context.
- 9 – For a given type, checking the *no value loss* criterion consists in checking if a *unique* template triple can be
10 found in the template graph for each placeholder, *i.e.* a placeholder corresponding to each property keys in the
11 type; and if the PG element is an edge, the *?source* and *?destination* placeholders must also be found. Thanks
12 to hash sets, checking if a template triple is *unique* inside its template graph is constant. Implementing the test
13 by following the definition leads to an $\mathcal{O}(TypeComplexity * BiggestTemplateSize)$ complexity for each type
14 and an $\mathcal{O}(NbTypes * TypeComplexity * BiggestTemplateSize)$ complexity for the whole context.

15 The final complexity of checking if a context is a well-behaved PRSC context is:

$$16 \quad \mathcal{O}(NbTypes * BiggestTemplateSize * (NbTypes + TypeComplexity))$$

17 5.3. Reversion algorithm

20 Algorithm 2 aims to convert an RDF graph, that was produced from a PG and a known well-behaved context,
21 into the original PG.

22 It is a four steps algorithm: 1) it finds the elements of the PG, by assuming they are the same as the blank node
23 in the RDF graph, 2) it gives a type to all PG elements with the *FindTypeOfElements* function in Algorithm 3¹⁰,
24 3) it assigns each triple to a single PG element, corresponding to the production of the *build* function, with the
25 *AssociateTriplesWithElements* function in Algorithm 4, and 4) it looks for the source, destination and properties of
26 all elements with the *buildpg* function in Algorithm 5.

27 Further subsections prove that for all $ctx \in Ctx^+$, for all PGs pg , applying these algorithms to $rdf = prsc(pg, ctx)$
28 actually produces pg , meaning that the reversion algorithm is a sound and complete implementation of $prsc^{-1}$ for
29 well-behaved contexts. Applying this algorithm to an arbitrary RDF graph and/or an arbitrary context is out of the
30 scope of this paper.

31 5.3.1. Finding the elements of the PG

32 The first step of the algorithm relies on the assumption that the blank nodes of the RDF graph and the elements
33 of the original PG are the same.

34 **Algorithm 2:** The main algorithm to convert back an RDF graph into a PG by using a context

35 **Input:** $rdf \subset RDFTriples, ctx \in Ctx^+$

36 **Output:** A BPG or error

37 **1 Main Function** *RDFToPG*(rdf, ctx):

38 **2** $Elements \leftarrow BNodes(rdf)$

39 **3** $typeof \leftarrow FindTypeOfElements(rdf, ctx, Elements)$

40 **4** $builtfrom \leftarrow AssociateTriplesWithElements(rdf, Elements, typeof)$

41 **5** **return** $buildpg(ctx, Elements, typeof, builtfrom)$

42 ¹⁰To help the comprehension of Algorithm 3, we recall that for a given set A , the mathematical notation $\exists! a \in A, somepredicate(a)$ means
43 that in the set A , there is one and only one element, denoted by a , that matches *somepredicate*. By extension, $\exists! a \in A$ means that there is one
44 and only one element in the set A that is denoted by a .

Algorithm 3: Associate the elements of the future PG with their types**Input:** $rdf \subset RDFTriples, ctx \in Ctx^+, Elements = BNodes(rdf)$ **Output:** A mapping between $Elements$ and $Dom(ctx)$ or error**1 Function FindTypeOfElements($rdf, ctx, Elements$):**

```

2   typeof  $\leftarrow \{\}$ 
3   forall element  $m \in Elements$  do
4       /* Find possible types */
5       candtypesnodes  $\leftarrow \{\}$ 
6       candtypesedges  $\leftarrow \{\}$ 
7       forall triple  $t \in rdf \mid m \in t$  do
8           forall type  $\in Dom(ctx)$  do
9               if  $\kappa(sign_{ctx}(type)) = \kappa(t)$  then
10                  if kind(type) = "node" then
11                      candtypesnodes  $\leftarrow candtypes_{nodes} \cup \{type\}$ 
12                  else
13                      candtypesedges  $\leftarrow candtypes_{edges} \cup \{type\}$ 
14
15           /* Choose a type */
16           if  $(\exists! type \in candtypes_{nodes})$  or  $(\exists! type \in candtypes_{edges} \text{ and } candtypes_{nodes} = \emptyset)$  then
17               | typeof( $m$ )  $\leftarrow type$ 
18           else
19               | raise Error(No type found)
20
21   return typeof

```

Algorithm 4: Associate each triple to the element that has produced it**Input:** $rdf \subset RDFTriples, Elements = BNodes(rdf), typeof : Elements \mapsto Type$ **Output:** A mapping $Elements \rightarrow 2^{RdfTriples}$ or error**1 Function AssociateTriplesWithElements($rdf, Elements, typeof$):**

```

2   builtfrom  $\leftarrow \{\}$ 
3   forall  $b \in Elements$  do builtfrom( $b$ )  $\leftarrow \{\}$ 
4   forall  $td \in rdf$  do
5       bns  $\leftarrow \{term \in td \mid term \in B\}$ 
6       if  $(\exists! b \in bns)$  or  $(\exists! b \in bns \mid kind(typeof(b)) = "edge")$  then
7           | builtfrom( $b$ )  $\leftarrow builtfrom(b) \cup \{td\}$ 
8       else
9           /* No blank node in bns, or multiple PG nodes but no PG edges, or multiple PG
10              edges */
11           | raise Error(No element provenance)
12
13   return builtfrom

```

Algorithm 5: Produce a PG from the previous analysis of the elements and triples.

Input: $ctx \in Ctx^+, Elements \subset B, typeof : Elements \rightarrow Type, builtfrom : Elements \rightarrow 2^{RdfTriples}$

Output: A BPG or error

```

1 Function buildpg(ctx, Elements, typeof, builtfrom):
2   g is initialized to the empty PG
3   forall  $b \in Elements$  do
4      $labels_g(b) \leftarrow labels(typeof(b))$ 
5     if  $kind(typeof(b)) = \text{"edge"}$  then
6        $src_g(b) \leftarrow extract(?source, builtfrom(b), ctx(typeof(b)))$ 
7        $dest_g(b) \leftarrow extract(?destination, builtfrom(b), ctx(typeof(b)))$ 
8        $N_g \leftarrow N_g \cup \{src_g(b), dest_g(b)\}$ 
9        $E_g \leftarrow E_g \cup \{b\}$ 
10    else
11       $N_g \leftarrow N_g \cup \{b\}$ 
12    forall  $key \in keys(typeof(b))$  do
13       $properties_g(b, key) \leftarrow extract(key, builtfrom(b), ctx(typeof(b)))$ 
14  return g

15 Function extract(placeholder, tds, tps):
16   $values \leftarrow \{\}$ 
17  forall  $tp \in tps \mid unique(tp, tps) \wedge placeholder \in tp$  do
18     $samekappa \leftarrow \{td \in tds \mid \kappa(td) = \kappa(tp)\}$ 
19    if  $\|samekappa\| \neq 1$  then raise Error(Unique data triple is not unique)
20     $td \leftarrow$  the only element in samekappa
21     $answer \leftarrow$  The term from td that is at the same place as placeholder in tp
22     $values \leftarrow values \cup \{answer\}$ 
23  if  $\|values\| \neq 1$  then raise Error(Not exactly one value for a placeholder)
24   $answer \leftarrow$  The only member of values
25  if  $placeholder \in P$  then
26    return  $toLiteral^{-1}(answer)$ 
27  else
28    return answer

```

Theorem 2 (Equality between the elements of a PG and the blank nodes of the RDF graph). If the RDF graph *rdf* has been produced from a PG *pg* and a PRSC well-behaved context *ctx*, then the set of all blank nodes of *rdf* is the set of PG elements of *pg*.

$$\forall pg \in BPGs, ctx \in Ctx_{pg}^+, rdf = prsc(pg, ctx), N_{pg} \cup E_{pg} = BNodes(rdf)$$

Proof.

- The *build* function, described in Section 4.5, produces specific triples depending on the given template. The template graphs cannot contain blank nodes: the blank node produced by *prsc* are forced to be the elements of the converted BPG. So $BNodes(rdf) \subseteq N_{pg} \cup E_{pg}$.

- From Remark 8, we know that $ctx(typeof_{pg}(m))$ contains at least one triple pattern tp . Combined with the *element provenance* criterion from Definition 24-1, we know that $?self \in tp$. When β is applied to tp , a triple that contains m is forced to appear, meaning that $N_{pg} \cup E_{pg} \subseteq BNodes(rdf)$.

□

Theorem 2 proves the correctness of the $Elements \leftarrow BNodes(rdf)$ step in Algorithm 2.

5.3.2. Finding the type related to each element

In this part of the proof, we show that the *FindTypeOfElements* function from Algorithm 3 is correct, *i.e.* it is able to find back the right type of all elements m in the original pg graph.

Lemma 3. If a data triple shares the same value through κ as one of the signature triples of a type (Definition 24-2), then the element from which the data triple was produced must be of this type:

$$\forall td \in rdf, \forall type \in Dom(ctx), \forall m \in N_{pg} \cup E_{pg},$$

$$[\kappa(td) = \kappa(sign_{ctx}(type)) \wedge td \in build(ctx(typeof_{pg}(m)), pg, m)] \Rightarrow typeof_{pg}(m) = type$$

Proof. $\forall td \in rdf, \forall type \in Dom(ctx), \forall m \in N_{pg} \cup E_{pg}$

Assuming (A) $\kappa(td) = \kappa(sign_{ctx}(type))$

$td \in build(ctx(typeof_{pg}(m)), pg, m)$

$\Rightarrow \exists tp \in ctx(typeof_{pg}(m)) \mid \kappa(td) = \kappa(tp)$ [Lemma 1]

$\Rightarrow \exists tp \in ctx(typeof_{pg}(m)) \mid \kappa(sign_{ctx}(type)) = \kappa(tp)$ [A]

$\Rightarrow \exists tp \in ctx(typeof_{pg}(m)) \mid \kappa(sign_{ctx}(type)) = \kappa(tp) \subseteq \kappa(ctx(typeof_{pg}(m)))$ $\left[\begin{array}{l} tp \in ctx(typeof_{pg}(m)) \\ \text{and by construction of } \kappa \end{array} \right]$

$\Rightarrow typeof_{pg}(m) = type$ $\left[\begin{array}{l} \text{Signature template triple} \\ \text{in Definition 24-2} \end{array} \right]$

□

Definition 25 (Formalizing *candtypes*). For a given blank node/PG element b , $candtypes_{nodes}$ and $candtypes_{edges}$, introduced in Algorithm 3, can be formally defined as:

$candtypes_{nodes}(b) = \{type \in Dom(ctx) \mid kind(type) = \text{“node”}$

$\wedge \exists td \in rdf \mid b \in td \wedge \kappa(sign_{ctx}(type)) = \kappa(td)\}$

$candtypes_{edges}(b) = \{type \in Dom(ctx) \mid kind(type) = \text{“edge”}$

$\wedge \exists td \in rdf \mid b \in td \wedge \kappa(sign_{ctx}(type)) = \kappa(td)\}$

They give the set of all node types and edge types, respectively, for which one of their signature triple could have produced a triple with b .

Theorem 3 (*candtypes* correctness). Even though the *candtypes* functions are defined by only used the used context and the produced RDF graph, they can be used to compute the type of any blank node in the original PG:

- $\forall b \in N_{pg}, candtypes_{nodes}(b) = \{typeof_{pg}(b)\}$
- $\forall b \in E_{pg}, candtypes_{nodes}(b) = \emptyset$ and $candtypes_{edges}(b) = \{typeof_{pg}(b)\}$.

Table 10
A simple view of Theorem 3

	$ candtypes_{nodes}(b) $	$ candtypes_{edges}(b) $
$b \in N_{pg}$	1	any
$b \in E_{pg}$	0	1

Table 10 provides an overview of the cardinality of the different *candtypes* sets.

Proof.

$\forall b \in BNodes(rdf), \forall type \in candtypes_{nodes}(b)$

Per Definition 25, $kind(type) = \text{“node”} \wedge \exists td \in rdf \mid b \in td \wedge \kappa(sign_{ctx}(type)) = \kappa(td)$.

We are going to restrict the portion of the graph *rdf* where such triples *td* may be located:

$$\begin{aligned}
 & td \in rdf \\
 \Leftrightarrow & td \in \bigcup_{m \in N_{pg} \cup E_{pg}} build(ctx(typeof(m)), pg, m) && \text{[Definition of } rdf / prsc\text{]} \\
 \Rightarrow & td \in \bigcup_{m \in N_{pg} \cup E_{pg} \mid typeof(m)=type} build(ctx(type), pg, m) && \left[\begin{array}{l} \kappa(td) = \kappa(sign_{ctx}(type)) \\ \text{and Lemma 3} \end{array} \right] \\
 \Rightarrow & td \in \bigcup_{m \in N_{pg} \mid typeof(m)=type} build(ctx(type), pg, m) && [kind(type) = \text{“node”}]
 \end{aligned}$$

- We see that all triples *td* contributing to $candtype_{nodes}(b)$ must have been produced by the signature triple template applied to a node from the PG. Also remember that *td* must contain *b*.
- If $b \in N_{pg}$, then the signature triple of $ctx(typeof_{pg}(b))$ must have generated a *td* containing *b* (since it must contain *?self*, according to Definition 24-1), so $typeof_{pg}(b) \in candtype_{nodes}(b)$. Furthermore, no other node can produce a *td* containing *b* (*?self* is the only blank node placeholder in node type templates), so $candtype_{nodes}(b)$ can not contain any other type. Therefore $candtype_{nodes}(b) = \{typeof_{pg}(b)\}$.
- If $b \in E_{pg}$, it is impossible to produce the blank node *b* from any node $m \in N_{pg}$ (again, *?self* is the only blank node placeholder in node type templates). No *td* containing *b* can be found, so $candtype_{nodes}(b)$ is empty.

The reasoning for $candtypes_{edges}(b)$ when *b* is an edge is similar to the one for $candtypes_{nodes}(b)$ when *b* is a node: only *b* can produce triples containing itself, and it will, because having at least one signature triple with *?self* is imposed by Definition 24. So $candtype_{edges} = \{typeof_{pg}(b)\}$.

Finally, a blank node $b \in N_{pg}$ can appear in any number of triples that share the same value through κ with an edge *signature template triple*: an edge *signature template triple* can contain *?source* or *?destination*, that can be mapped to any node depending on the PG. So $candtype_{edges}(b)$ can contain an arbitrary number of types in that case. \square

Remark 13. Theorem 3 not only shows that the *FindTypeOfElements* function in Algorithm 3 will always find the right $typeof_{pg}$ function by using *candtypes*, *i.e.* that it is computable from *rdf* and *ctx*, but Table 10 also explicitly shows that the *Error(No type found)* scenario can not appear if the RDF graph was produced from a PG, making the *FindTypeOfElements* function both sound and complete.

Remark 14 (Using different signatures to determine the types). As mentioned previously, the choice of the signature template triple $sign_{ctx}(type)$ of a given type is not important for the reversion algorithm. Choosing one or another signature template triple only leads to other data triples being used to determine the type of the elements. However, the end result does not change.

5.3.3. Finding the generated triples for each PG element

For each PG element, given the produced RDF graph and the type of this PG element, we are able to compute the list of RDF triples produced from this PG element. In other words, Algorithm 4 correctly partitions the RDF graph into sub-graphs describing each element of the original PG.

Theorem 4. In Algorithm 4, assuming that the passed value of *typeof* is equal to $typeof_{pg}$, $\forall m \in N_{pg} \cup E_{pg}$, $build(ctx(typeof_{pg}(m)), pg, m) = builtfrom(m)$.

Proof. As $rdf = \bigcup_{m \in N_{pg} \cup E_{pg}} build(ctx(typeof(m)), pg, m)$, each triple $td \in rdf$ is a member of at least one $build(ctx(typeof(m)), pg, m)$. For all triples $td \in build(ctx(typeof(m)), pg, m)$, the *element provenance* criterion ensures that $m \in td$. So the first step that consists in listing in the set *bns* the blank nodes in td , and consider that m is part of the set *bns* is correct: the actual element m is in the set.

The algorithm associates each triple td with a single $builtfrom(m)$:

- Let's first recall that blank nodes in *rdf* can only be produced via the placeholders from *pvars*: *?self*, *?source*, and *?destination*. Let's also recall that every triple pattern in the Well-behaved context *ctx* must contain *?self* (per the *element provenance* criterion).
- If *bns* contains only one blank node m , then m must come from *?self*, and the corresponding triple pattern must then belong to $ctx(typeof(m))$. td must then have been produced by $build(ctx(typeof(m)), pg, m)$ so putting it in $builtfrom(m)$ is correct.
- If *bns* contains multiple blank nodes, td must have been produced by a template triples with several placeholders from *pvars*.
 - * Node template graphs can contain only one placeholder from *pvars*: *?self*. No $m \in N_{pg}$ could then have produced td . It follows that td must have been produced by the template graph of an edge.
 - * Edge template graphs can contain several placeholders from *pvars*. But by definition of β , only *?self* can be mapped to an edge (when $m \in E_{pg}$); *?source* and *?destination* are always mapped to nodes. Of the multiple blank nodes in *bns*, exactly one of them, m , must therefore be an edge, and come from *?self*. Following the same reasoning as above, when *bns* contained only one blank node, we conclude that td must then have been produced by $build(ctx(typeof(m)), pg, m)$ and that putting it in $builtfrom(m)$ is correct.
- *Error(No element provenance)* will never be raised if *rdf* was produced by *prsc*: each triple will contain at least one blank node generated from *?self* (per the *element provenance* criterion), and if there are multiple blank nodes we showed that there must be only one edge blank node.

As each triple in *rdf* is attributed in $builtfrom(m)$ to the right element m that produced it from $build(ctx(typeof_{pg}(m)), pg, m)$, $\forall m \in N_{pg} \cup E_{pg}$, $builtfrom(m) = build(ctx(typeof_{pg}(m)), pg, m)$. □

5.3.4. Building the PG element

Projecting Property Graphs As an RDF graph is defined as a set of RDF triples, any subset of that set, as well as the union of two RDF graphs, are formally defined and are also RDF graphs. Algorithm 5 constructs back the original PG in an iterative manner. To prove its correctness, we need operators similar to \subset and \cup for RDF graphs, but for our formalization of PGs.

In this section, the projection of a Property Graph is defined by focusing only on a single **PG element**, node or edge. The concept of merging PGs, which is the inverse of the projection, is also defined.

Let pg be a PG.

Definition 26 (π projection of a Property Graph on an element). The π projection of a PG on a node is equal to the PG with only the node itself. The π projection of a PG on an edge is the edge, and its source and destination nodes without the labels and properties of these nodes.

$\forall m \in N_{pg} \cup E_{pg}$, $\pi_m(pg)$ is a PG such as:

- 1 - If $m \in N_{pg}$, $N_{\pi_m(pg)} = \{m\}$, $E_{\pi_m(pg)} = \emptyset$, $src_{\pi_m(pg)} = dest_{\pi_m(pg)} = \emptyset \rightarrow \emptyset$ 1
- 2 - If $m \in E_{pg}$, $N_{\pi_m(pg)} = \{src_{pg}(m), dest_{pg}(m)\}$, $E_{\pi_m(pg)} = \{m\}$, $src_{\pi_m(pg)} = \{m \mapsto src_{pg}(m)\}$, $dest_{\pi_m(pg)} =$ 2
- 3 $\{m \mapsto dest_{pg}(m)\}$ 3
- 4 - $\forall x \in N_{\pi_m(pg)} \cup E_{\pi_m(pg)}$, $labels_{\pi_m(pg)}(x) = \begin{cases} labels_{pg}(x) & \text{if } x = m \\ \emptyset & \text{otherwise} \end{cases}$ 4
- 5 5
- 6 - $\forall key \in keys_{pg}(m)$, $properties_{\pi_m(pg)}(m, key) = properties_{pg}(m, key)$, all other values are undefined. 6
- 7 7

8 **Definition 27** (Property Graph merge operator \oplus). The merge operator \oplus is the inverse of the projection operator 8
 π . It can only be used on two PGs that are compatible, *i.e.* (1) a PG element defined as a node is not defined as an 9
edge in the other, (2) an edge defined in both PGs have the same source and destination in both, and (3) if in both 10
PGs, the value of a property key on the same PG element is defined, the values should be the same. The \oplus operator 11
builds a PG with the PG elements, labels and properties of both PGs. 12

13 We now define the \oplus merge operator on property graphs. $\forall (pg', pg'') \in PGs^2$, $\oplus(pg', pg'')$ (or $pg' \oplus pg''$) is 13
defined only if: 14

- 15 15
- 16 - $E_{pg'} \cap N_{pg''} = \emptyset \wedge N_{pg'} \cap E_{pg''} = \emptyset$ 16
- 17 - $src_{pg'}$ is compatible with $src_{pg''}$ and $dest_{pg'}$ is compatible with $dest_{pg''}$ (see compatibility definition in Sec- 17
tion 3.2). 18
- 19 - $properties_{pg'}$ is compatible with $properties_{pg''}$. 19
- 20 20

21 Its value is $\oplus(pg', pg'') = pg$ with: 21

- 22 22
- 23 - $N_{pg} = N_{pg'} \cup N_{pg''}$ 23
- 24 - $E_{pg} = E_{pg'} \cup E_{pg''}$ 24
- 25 - $src_{pg} : E_{pg} \rightarrow N_{pg}$, $src_{pg} = src_{pg'} \cup src_{pg''}$. 25
- 26 - $dest_{pg} : E_{pg} \rightarrow N_{pg}$, $dest_{pg} = dest_{pg'} \cup dest_{pg''}$. 26
- 27 - $\forall m \in N_{pg} \cup E_{pg}$, $labels_{pg}(m) = \begin{cases} labels_{pg'}(m) \cup labels_{pg''}(m) & \text{if both are defined} \\ labels_{pg'}(m) & \text{if only } labels_{pg'}(m) \text{ is defined} \\ labels_{pg''}(m) & \text{if only } labels_{pg''}(m) \text{ is defined} \end{cases}$ 27
- 28 28
- 29 - $properties_{pg} : (N_{pg} \cup E_{pg}) \times Str \rightarrow V$, $properties_{pg} = properties_{pg'} \cup properties_{pg''}$. 29
- 30 30
- 31 31

32 **Lemma 4.** \oplus is commutative, associative, and the neutral element is the empty PG pg_\emptyset 32

33 *Proof.* (Sketch) \oplus is defined by using the \cup operator, which is commutative, associative and whose neutral element 33
is \emptyset . The equivalent of \emptyset for PGs is pg_\emptyset . 34

35 **Theorem 5.** The \oplus merge of the π projection of a PG on all its PG elements is equal to the PG itself: 35

$$36 \forall pg \in PGs, pg = \bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg) \quad 36$$

37 *Proof.* The proof is provided in Appendix A. 37

38 While the \oplus operator has been primarily designed as the inverse operation of the projection operator π , it can 38
only merge two PG that are consistent between themselves. If a PG defines an entity as a node, the other PG can not 39
define it as an edge. If a PG has a given source and destination for a given edge, the other one can not define another 40
source or destination. Properties must also be consistent in both PGs. The presented version of the merge operator 41
can only add information, and in the case merging two PGs would lead to an inconsistent PG, the merge operator is 42
undefined. 43

1 *Relationship between the prsc function and projections* We are now going to redefine the *prsc* function using the
2 π operator.

3 The RDF graph built by *prsc* from a PG *pg* with a context *ctx* is equal to:

$$4 \quad rdf = \bigcup_{m \in N_{pg} \cup E_{pg}} build(ctx(typeof_{pg}(m)), pg, m)$$

5
6
7
8 The *build* function is defined in such a way that the RDF triples it produces from an element *m* are only influenced
9 by:

- 10 – *m* itself.
- 11 – Its labels, i.e. $labels_{pg}(m)$.
- 12 – Its property values, i.e. $\forall key, properties_{pg}(m, key)$.
- 13 – If *m* is an edge, its source and destination nodes, i.e. $src_{pg}(m)$ and $dest_{pg}(m)$.
- 14 – The template graph $ctx(typeof_{pg}(m))$.

15 Therefore the following equality can be asserted, $\forall pg \in BPGs, \forall m \in N_{pg} \cup E_{pg}, \forall ctx \in Ctx_{pg}$:

$$16 \quad build(ctx(typeof_{pg}(m)), pg, m) = build(ctx(typeof_{pg}(m)), \pi_m(pg), m)$$

17
18
19 $\pi_m(pg)$ can be considered as the minimal required Property Graph to produce the RDF triples related to the
20 element *m* in the PG *pg*. If we can prove that the reversion algorithm constructs all $\pi_m(pg)$ graphs and merges them
21 with the \oplus operator, then it means that we have properly reconstructed the *pg* PG.

22
23 *Completing the proof of the reversion algorithm* We are now back to proving that for all well-behaved contexts
24 *ctx*, the *RDFToPG* function presented in Algorithm 2 is an implementation of the $prsc^{-1}$ function. *pg* is a PG for
25 which we know the value of $prsc(pg, ctx) = rdf$. We are focusing on the last line of Algorithm 2, where the *buildpg*
26 function is invoked.

27 To prove the correctness of the *buildpg* function in Algorithm 5, starting from an empty PG *g*, we are going to
28 show that at each iteration, we are adding to the PG *g* the π projection of *pg* on an element *m*. After iterating on all
29 elements, as we merged the π projection of all PG elements, the PG *g* ends up being equal to the PG *pg* itself.

30
31 **Lemma 5** (Merging the projection of one PG element to the reconstructed PG). In Algorithm 5, assuming that
32 the *typeof* parameter is equal to $typeof_{pg}$ and *buildfrom* is a total function that maps all PG elements *b* to
33 $build(ctx(typeof_{pg}(b)), pg, b)$, at the end of an iteration of an element $b \in N_{pg} \cup E_{pg}$ after line 13, the computed PG
34 *g_{after}* is equal to $g_{before} \oplus \pi_b(pg)$, where *g_{before}* is the PG *g* at the beginning of the iteration between lines 3 and 4.

35
36 *Proof.* The PG $\pi_b(pg)$ is described in Table 11. Bold values are the ones for which we need to prove that we compute
37 the correct value: $src_g(b)$, $dest_g(b)$ and $properties_g(b, key)$. Other values are trivially correct by construction.

38
39 Table 11
40 Description of the PG projection that is built in Algorithm 5

	$b \in N_{pg}$	$b \in E_{pg}$
$N_{\pi_b(pg)}$	$\{b\}$	$Img(src_{\pi_b(pg)}) \cup Img(dest_{\pi_b(pg)})$
$E_{\pi_b(pg)}$	\emptyset	$\{b\}$
$src_{\pi_b(pg)}$	$\emptyset \rightarrow \emptyset$	$b \mapsto \mathbf{src_g(b)}$
$dest_{\pi_b(pg)}$	$\emptyset \rightarrow \emptyset$	$b \mapsto \mathbf{dest_g(b)}$
	$b \in N_{pg} \cup E_{pg}$	
$labels_{\pi_b(pg)}$	$\{b \mapsto labels(type)\}$	
$properties_{\pi_b(pg)}$	$\bigcup_{key \in keys(type)} \{(b, key) \mapsto \mathbf{properties_g(b, key)}\}$	

41
42
43
44
45
46
47
48
49
50 In the following, we want to check that $extract(?source, build(\pi_b(pg), pg, b), ctx(typeof(b)))$ properly returns
51 $src_{\pi_b(pg)}$. Proofs for $?destination / dest_{\pi_b(pg)}$ and $key \in keys_{typeof_{pg}(b)} / properties_{\pi_b(pg)}(b, key)$ are identical.

The *values* set is filled by iterating on all tp such that $unique(tp, tps) \wedge ?source \in tp$. The *no value loss* criterion ensures that at least one such template triple exists, so the loop in *extract* is iterated at least once.

Theorem 1 ensures that the built set *samekappa* in the loop of the *extract* function will always have 1 element, that we name td . *Error(Unique data triple is not unique)* may never be raised if *rdf* was produced by PRSC. By definition of the *build* function, $?source$ in tp and $src_{\pi_b(pg)}$ in td are at the same position.

After the loop, because only $src_{\pi_b(pg)}$ is added to *values* in the loop, *Error(Not exactly one value for a placeholder)* may never be raised.

The last instructions differ for $?source / ?destination$ and P . In the case of $?source$ and $?destination$, the obtained value is directly the value of the PG node; in the case of P , the obtained RDF literal needs to be converted into the proper PG property value, which is possible because $toLiteral^{-1}$ is assumed to be computable in Section 4.5.

extract properly computes the values that are missing in $\pi_b(pg)$. When these values are extracted, they are directly merged with the \cup operator into the g property graph. Values that were already known or can be computed from the values that were just extracted, *i.e.* $labels_{\pi_m(pg)}$, $N_{\pi_m(pg)}$ and $E_{\pi_m(pg)}$, are also merged into g .

As all values of $\pi_b(pg)$ are merged into g_{before} , $g_{after} = g_{before} \oplus \pi_m(pg)$

□

Remark 15 (Completeness of *buildpg*). In the case where *rdf* is built from a PG pg , the value that a placeholder is mapped to is the same everywhere, so we never run at the risk of encountering multiples values, *i.e.* *Error(Not exactly one value for a placeholder)* is never raised. Furthermore, the proof of Lemma 5 shows that *Error(Unique data triple is not unique)* may not be raised, because we know that each unique template triple has produced one data triple.

Theorem 6 (Merging the projection of all PG elements to the reconstructed PG). Under the same assumptions as Lemma 5, the PG returned by Algorithm 5 is the original pg , the PG that was used to produce the RDF graph $rdf = prsc(pg, ctx)$.

Proof. The PG g in the algorithm is initialized to pg_{\emptyset} . Lemma 5 shows that after each iteration in the loop with an element b , the PG g is \oplus -merged with the PG $\pi_b(pg)$. The loop iterates on all elements in the PG pg , so after all the iterations, the PG g is equal to:

$$\begin{aligned}
 g &= pg_{\emptyset} \oplus \bigoplus_{b \in N_{pg} \cup E_{pg}} \pi_b(pg) \\
 &= \bigoplus_{b \in N_{pg} \cup E_{pg}} \pi_b(pg) && [pg_{\emptyset} \text{ is the neutral element of } \oplus] \\
 &= pg && [\text{Theorem 5}]
 \end{aligned}$$

□

As *buildpg* in Algorithm 5 correctly reconstructs pg , and as its value is directly returned by the *RDFToPG* function in Algorithm 2, we have finally proven that the latter is a sound and complete implementation of the $prsc^{-1}$ function for any well-behaved PRSC context ctx .

5.3.5. Complexity analysis

Let us now discuss the complexity of the *RDFToPG* function described in Algorithm 2. In this discussion, a new metric is considered: the number of triples in the RDF graph: $NbTriples = |rdf|$. It is assumed that we have first checked if the context ctx is a well-behaved PRSC context, and computed the $sign_{ctx}$ function so it can now be called in constant time.

Extracting the list of blank nodes of an RDF graph on line 2 has a linear complexity of $\mathcal{O}(NbTriples)$.

The *FindTypeOfElements* function in Algorithm 3 uses three nested loops and only uses constant time operations: its complexity is $\mathcal{O}(NbOfPGElements * NbTriples * NbTypes)$.

Trivially, Algorithm 4 has a complexity of $\mathcal{O}(NbOfPGElements + NbTriples)$.

In Algorithm 5:

- Calls to $extract(placeholder, tds, tps)$ have a complexity of $\mathcal{O}(|tps|^2 * |tds|)$:
 - * It loops all triples in the template graph tps such that they are *unique*. Evaluating $unique(tp, tps)$ itself has an $\mathcal{O}(|tps|)$ complexity so the overall complexity of evaluating all $tp \in tps \mid unique(tp, tps)$ is $\mathcal{O}(|tps|^2)$.
 - * Inside the loop, building the *samekappa* set forces to loop on all tds , multiplying the complexity by a $|tds|$ factor.

In the context of the *buildpg* function, the *extract* function is always called with a template graph tps from the PRSC context and a sub-graph of the RDF graph rdf as tds . the complexity of the calls of the *extract* function in the *buildpg* function is $\mathcal{O}(BiggestTemplateSize^2 + NbTriples)$ ¹¹.

- The *buildpg* function loops on all PG elements.
 - * Notice that while $ctx(typeof(b))$ is called multiple times, it can be called once and then its value can be cached. Its cost is $\mathcal{O}(TypeComplexity * \ln(TypeComplexity))$ as mentioned in Section 4.6.3.
 - * There are at most $2 + TypeComplexity$ calls of the *extract* function. All of them have a $\mathcal{O}(BiggestTemplateSize * NbTriples)$ complexity.
 - * The complexity of each iteration is $\mathcal{O}(TypeComplexity * \ln(TypeComplexity) + TypeComplexity * BiggestTemplateSize^2 * NbTriples)$
- The overall complexity of the *buildpg* function is $\mathcal{O}(NbOfPGElements * TypeComplexity * (\ln(TypeComplexity) + BiggestTemplateSize^2 * NbTriples))$

The overall complexity of the *RDFToPG* function presented in Algorithm 2 presented in this section for an RDF graph produced from a PG and a well-behaved PRSC context is:

$$\begin{aligned}
 & \mathcal{O}(NbTriples \\
 & + NbTypes * NbTriples * BiggestTemplateSize \\
 & + NbOfPGElements + NbTriples \\
 & + NbOfPGElements * TypeComplexity * (\ln(TypeComplexity) + BiggestTemplateSize^2 * NbTriples)) \\
 = & \mathcal{O}(NbTypes * NbTriples * BiggestTemplateSize \\
 & + NbOfPGElements * TypeComplexity * (\ln(TypeComplexity) + BiggestTemplateSize^2 * NbTriples))
 \end{aligned}$$

The *RDFToPG* function is computable in polynomial time w.r.t. all the considered metrics, and is therefore considered as tractable. Furthermore, in our implementation, the algorithm has been optimized to significantly lower the complexity. However, for the sake of concision, we do not describe these optimizations in this paper.

5.4. Edge-unique extension

In many cases, there is only one edge of certain types between two nodes, like the “TravelWith” edge in our running example or for relationships like knowing someone, a parental relationship. . . For this type of edges, it is more intuitive to represent them with a simple RDF triple, and get rid of the blank node corresponding to the edge. However, Well-Behaved PRSC contexts require *?self* in edge templates. In this section, we propose an extension to allow *?self* to be missing in edge templates and still produce reversible conversions.

¹¹Note that if rdf has been generated by the *prsc* function, the number of triples in the graph tds is inferior or equal to the number of triples in the template graph tps as tds has been generated from tps . In this case, the complexity of calling *extract* in the context of the *buildpg* function is $\mathcal{O}(BiggestTemplateSize^3)$.

Table 12
A context for the Tintin PG with the “since” property

<i>type</i>	<i>ctx(type)</i>
	(<i>?self</i> , <i>rdf:type</i> , <i>ex:Person</i>)
(“node”, {“Person”}, {“name”, “job”})	(<i>?self</i> , <i>foaf:name</i> , “name” <i>valueOf</i>)
	(<i>?self</i> , <i>ex:profession</i> , “job” <i>valueOf</i>)
(“node”, ∅, {“name”})	(<i>?self</i> , <i>foaf:name</i> , “name” <i>valueOf</i>)
	(<i>?source</i> , <i>ex:isTeammateOf</i> , <i>?destination</i>)
(“edge”, {“TravelsWith”}, {“since”})	((<i>?source</i> , <i>ex:isTeammateOf</i> , <i>?destination</i>), <i>ex:since</i> , “since” <i>valueOf</i>)

Consider the Tintin PG exposed in Figure 1 and the context exposed in Table 12, which uses RDF-star to convert the “since” property. The output of PRSC from those two inputs is exposed in Listing 4. By looking at the produced RDF graph, it appears that the RDF graph captures all the information of the PG. More generally, RDF graphs produced by this context would always be reversible as long as the source PG does not contain multiple “TravelsWith” edges between two given nodes.

Listing 4 The output of PRSC for the Tintin PG and the context exposed in Table 12

```

19 1  % Tintin node
20 2  _:n1 rdf:type ex:Person .
21 3  _:n1 foaf:name "Tintin" .
22 4  _:n1 ex:profession "Reporter" .
23 5  % Snowy node
24 6  _:n2 foaf:name "Snowy" .
25 7  % TravelsWith edge
26 8  _:n1 ex:isTeammateOf _:n2 .
27 9  << _:n1 ex:isTeammateOf _:n2 >> ex:since 1978 .

```

Definition 28 (Edge-unique extension).

a) In a context *ctx*, an **edge-unique type** *edgeuniq* is an edge type such that:

- *ctx(edgeuniq)* complies with the *no value loss* criterion (as per Definition 24-3) and is not empty.
- For all template triples *tp* ∈ *ctx(edgeuniq)*:
 - * *?source* ∈ *tp* and *?destination* ∈ *tp*
 - * *tp* is a *signature template triple* (as per Definition 24-2), *i.e.* no other type has a template triple that shares its value through κ .
 - * *tp* is a *unique* template triple, *i.e.* no other template triple in *ctx(edgeuniq)* shares its value through κ .

b) A PG *pg* is said edge-unique valid for a context *ctx* if for all edge-unique types in the context, there is at most one edge of this type between two given nodes:

$$\forall e \in E_{pg}, \text{typeof}_{pg}(e) \text{ is an edge-unique type} \Rightarrow (\forall e' \in E_{pg}, \left[\begin{array}{l} \text{typeof}_{pg}(e) = \text{typeof}_{pg}(e') \\ \wedge \text{src}_{pg}(e) = \text{src}_{pg}(e') \\ \wedge \text{dest}_{pg}(e) = \text{dest}_{pg}(e') \end{array} \right] \Rightarrow e = e')$$

c) The *prscEdgeUnique* function is introduced to serve as a proxy to the *prsc* function to be applied only if the given PG is edge-unique valid relatively to the given context:

$$\text{prscEdgeUnique}(pg, ctx) = \begin{cases} \text{prsc}(pg, ctx) & \text{if } pg \text{ is edge-unique valid for } ctx \\ \text{undefined} & \text{otherwise} \end{cases}$$

Theorem 7 shows that $prscEdgeUnique$ is reversible up to an isomorphism.

Theorem 7. Let ctx be a context such that each type either a) matches the constraints of a type in a well-behaved PRSC context in Definition 24 or b) is an edge-unique type.

- For every two BPGs, pg' and pg'' , such that $prscEdgeUnique(pg', ctx) = prscEdgeUnique(pg'', ctx)$, pg' and pg'' are isomorphic.
- There is an algorithm such that for all BPGs pg , from the RDF graph $prscEdgeUnique(pg, ctx)$ and the context ctx , the algorithm computes a PG pg' such that $prscEdgeUnique(pg, ctx) = prscEdgeUnique(pg', ctx)$, i.e. from the produced RDF graph and the context, it is possible to compute a PG that is isomorphic to the original one.

Proof. (Sketch) The context ctx is composed of two parts: a) the well-behaved part and b) the edge-unique part. The well-behaved part has been proved to be reversible. As template triples used for edge-unique types are signatures, their value through κ is different from the triples produced from the value through κ of the triples of the well-behaved part: triples produced from edge-unique types are distinguishable from the rest of the RDF graph.

Denote W the set of all types in the well-behaved part and U the types in the edge-unique part. Let pg be a PG such that $rdf = prscEdgeUnique(pg, ctx)$ exists. It is possible to split pg using W and U :

$$pg = \underbrace{\bigoplus_{m \in N_{pg} \cup E_{pg} \mid typeof_{pg}(m) \in W} \pi_m(pg)}_{pg_W} \oplus \underbrace{\bigoplus_{u \in E_{pg} \mid typeof_{pg}(u) \in U} \pi_u(pg)}_{pg_U}$$

It is also possible to split rdf by defining an *isWellBehaved* predicate that uses κ to filter triples that come from types in the well-behaved part:

$$\forall td \in RdfTriples, isWellBehaved(td) \Leftrightarrow \exists type \in W, \exists tp \in ctx(type), \kappa(td) = \kappa(tp).$$

$$rdf = \underbrace{\{td \in rdf \mid isWellBehaved(td)\}}_{rdf_W} \cup \underbrace{\{td \in rdf \mid \neg isWellBehaved(td)\}}_{rdf_U}$$

From all the theorems on well-behaved contexts, there is a bijection between pg_W and rdf_W .

All template triples used in the template graph of edge-unique types are both signature and unique: from any triple in rdf_U , it is possible to find which template triple produced it. Consider an arbitrary edge u , whose type is an edge-unique type, i.e. $typeof_{pg}(u) \in U$. As edge-unique template graphs must also comply with the *no value loss* criterion, all properties, the source node and the destination node of u can be found in a non-ambiguous manner in rdf_U . The only missing information is the edge identity, i.e. the blank node u itself.

By using a fresh blank node for u , it is possible to build a PG isomorphic to $\pi_u(pg)$ from rdf_U , by extension, a PG isomorphic to pg_U from rdf_U , and by extension a PG isomorphic to pg from rdf . □

5.5. Discussion about the constraints on well-behaved PRSC contexts

In this section, we discuss the acceptability of the different constraints posed by PRSC well-behaved contexts in terms of usability. In other words, to what extent do they limit what can be achieved with PRSC?

The *no value loss* criterion on well-behaved contexts ensures that the data are still present and can be found unambiguously: as its name implies, this constraint is obviously required to avoid information loss. Therefore, it should not be perceived as overly constraining when building PRSC contexts.

The *signature template triple* is a method to force the user to type the resources, which is usually considered to be good practice. The type can either be explicit, through a triple with $rdf:type$ as the predicate, or implicit through

1 a property that is only used by this type. For example, the template graph for a type *Person* could contain a template 1
2 triple for the form $(?self, :personId, "pid" valueOf)$. The constraint of a signature composed of only one triple can be 2
3 considered too strong: one may want to write a context that works for all PGs. For example, many authors [8, 14] 3
4 propose to map each label to an RDF type or a literal used as the object of a specific predicate like `pgo:label`. 4
5 More generally, users may want to use a composite key to sign their types. For these kinds of mappings, our ap- 5
6 proach of identifying the type by finding a single signature template is not sufficient. It requires finding all the 6
7 signature template triples and deciding to which type they are associated, for example through a Formal Concept 7
8 Analysis process. This could be studied as a future extension of the PRSC reversion algorithm. 8
9

10 The *element provenance* constraint may hinder the integration of RDF data coming from a PG with regular 10
11 RDF data: it forces the user to keep the structure exposed in the PG, with blank nodes representing the underlying 11
12 structure of the PG. The edge-unique extension enables to leverage this constraint, by avoiding representing PG 12
13 edges as RDF nodes. 13
14

16 6. Related works 16

17 Many works already exist to address the interoperability between PGs and RDF. 17
18

19 *A common pivot for PGs and RDF* To achieve interoperability, some authors propose to store the data into an- 19
20 other data model, and then expose the data through usual PG and RDF APIs. Angles et al. propose multilayered 20
21 graphs [19], for which the OneGraph vision from Lassila et al. [10] is a more concrete version. These works propose 21
22 to describe the data with a list of edges, with the source of the edge, a label and the destination of the edge. All edges 22
23 are associated with an identifier, that can be used as the source or the destination of other edges. However, authors 23
24 note that several challenges are raised about the way to implement the interoperability between the OneGraph model 24
25 and the existing PG and RDF APIs. 25
26

27 In a Unified Relational Storage Scheme [13], Zhang et al. propose to store the data in relational databases. While 27
28 they specify how to store both models in a similar relational database structure, they do not mention how they align 28
29 the data that come from one model with the data that come from another, for example to match the PG label "Person" 29
30 with the RDF type `foaf:Person`. 30

31 The Singleton Property Graph model proposed by Nguyen et al. [20] is an abstract graph model that uses the RDF 31
32 Singleton Property pattern that can be implemented both with a PG and an RDF graph. They also describe how to 32
33 convert a regular RDF graph or a regular PG into a Singleton Property Graph. But the use of the Singleton Property 33
34 pattern induces the creation of many different predicates, which hinders the performance of many RDF database 34
35 systems as shown by Orlandi et al. [21]. 35
36

37 *From PGs to RDF* In terms of PG to RDF conversion, the most impactful work is probably RDF-star [8, 9, 22, 23], 37
38 an extension of the RDF model originally proposed by Olaf Hartig and Bryan Thompson to bridge the gap between 38
39 PGs and RDF by allowing the use of triples in the composition of other triples. Indeed, the most blatant difficulty 39
40 when converting PG to RDF is converting the edge properties. However, most PG engines support multi-edges, *i.e.* 40
41 two edges of the same type between the two same nodes. On the other hand, the naive approach consisting in using 41
42 the source node, the type of the edge and the destination node as respectively the subject, the predicate and the 42
43 object of an RDF triple would collapse the multi-edges. Converting each edge property to an RDF-star triple that 43
44 uses the former triple as its subject would lead to the properties of each multi-edge to be merged. Khayatbashi et 44
45 al. [24] study on a larger scale the different mappings described by Hartig and benchmark them. By allowing triples 45
46 to be used as the subject and the object of other triples, it is possible to emulate the edge properties of PGs. While 46
47 these mappings allow some kind of user customization, by letting them choosing the used IRIs, they never consider 47
48 using different model structures for different PG types during the same conversion. For example, consider a PG 48
49 with two types of edges: one edge type with the "knows" label and one with the "marriedTo" label. The mappings 49
50 described in this paper do not enable the user to model edges with the "knows" label as a predicate and edges with 50
51 the "marriedTo" label as a proper RDF resource. To tackle the edge property problem, Das et al. study how to use 51

1 already existing reification techniques to represent properties [25]: the modelings that do not rely on quads can be 1
2 used when writing a PRSC context. 2

3 Tomaszuk et al. propose the Property Graph Ontology (PGO) [14], an ontology to describe PGs in RDF. As this 3
4 solution only describes the structure of the PG in RDF, the produced data is forced to use this ontology, with the 4
5 exception of other already existing RDF ontologies without further transformations. Thanks to the Neosemantics¹² 5
6 plugin developed by Barrasa, Neo4j is able to benefit from RDF related tools like ontologies, and performs a 2-way 6
7 conversion from and to RDF-star data. However, the PG to RDF conversion performed by Barrasa tends to affirm 7
8 all triples it can, even for PG edges that may describe facts with a probability or that are time restricted: if the 8
9 marriage between Alice and Bob has ended in 2017, the triple `:Alice :marriedto :Bob` should probably 9
10 not be produced. 10

11 Gremlinator [26] allows users to query a PG and an RDF database by using the SPARQL language. This is a first 11
12 step towards federated queries. However, it supposes that data stored in the PG and data stored in the RDF graph 12
13 have a similar modeling, and it does not support RDF-star. 13

14 Instead of having a fixed mapping, our work on PREC [16] propose a mapping language named PREC to drive 14
15 the conversion from PG to RDF. Delva propose RML-star [27], an extension of RML [11] and R2RML [28] that 15
16 introduces new RML directives to generate RDF-star triples. As discussed in Section 2, the format in which the 16
17 template triples are described in this work is closer to the produced triples, at the cost of reducing the ability to 17
18 produce templated IRIs or terms. 18

19 To leverage the requirement for users to manually write the mapping, Fathy et al. proposed ProGoMap [15], an 19
20 engine that first generates a putative ontology for the terms in a PG, aligns this ontology with a real world ontology, 20
21 and finally converts the PG to an RDF graph with an RML mapping generated from the alignment. The authors 21
22 motivate the choice of automating the ontology alignment process by mentioning that writing mapping manually 22
23 can be time-consuming. While this may be true, we do not think that this hinders PRSC usefulness as 1) the schema 23
24 part of a PRSC context may be generated automatically, and nothing prevents a tool from generating the template 24
25 triples automatically, 2) a PG schema for which writing template triples is time-consuming may be an indicator that 25
26 the schema is too complex, and therefore that the structure of the data stored in the PG should be cleaned, not only 26
27 to make the conversion to RDF easier, but also as a benefit for the PG itself, and 3) while the process of aligning 27
28 terms of the putative ontologies to terms of real ontologies indeed increase data interoperability, it still relies on 28
29 the idea that the structure of the PG is somewhat close to the structure of the desired RDF graph. However, in this 29
30 paper, we advocate that it may not be the case, and that the desired method to model each edge type depends on the 30
31 semantics of held by each edge type. 31

32 In general, to the best of our knowledge, most existing works only tackle the syntactic aspect of converting PG to 32
33 RDF data. While this level of interoperability is appreciated, it is not enough to be able to properly use the converted 33
34 data in any existing RDF application without further modification of the RDF graph. Other existing works provide a 34
35 level of semantic interoperability. However, they tend to choose one of the many syntactic representations of edges 35
36 that exist, despite RDF ontologies having a great variety of patterns to model knowledge. PRSC gives full control 36
37 to the user, by letting them choosing both the syntactic representation and use shared vocabularies in the produced 37
38 RDF graph. The use of PG schemas in PRSC contexts guides users in the process of writing the context. The study 38
39 on PRSC well-behaved contexts, and the related discussion on their constraints in Section 5.5, provide information 39
40 on if the PRSC context written by the user leads to a reversible conversion or not, and by consequence, if any 40
41 information is lost in the process of converting the data stored in the PG into RDF. 41

42 *From RDF to PGs* Abuoda et al. [29] study the different RDF-star to PG approaches and identified two classes: the 42
43 RDF-topology preserving transformation which converts each term into a PG node, and the PG transformation that 43
44 converts literals into property values. They also evaluate the performance of these different approaches. The PRSC 44
45 reversion algorithm, and the general philosophy of this work clearly falls under the latter category. The former can 45
46 be considered as using a PG database to store RDF data. 46

47 Angles et al. [30] discuss different methods to transform an RDF graph into a PG. They propose different map- 47
48 pings, including an RDF-topology preserving one and a PG transformation. Atemezing and Hyunh [31] propose to 48
49 49

51 ¹²<https://github.com/neo4j-labs/neosemantics> 51

1 use a mapping similar to the former to publish and explore RDF data with PG tools, namely Neo4j. However, these 1
2 works offer little customization for the user. 2

3 With G2GML [32], Chiba et al. propose to convert RDF data by using queries: the output of the query is trans- 3
4 formed into a PG by describing a template PG, similar to a Cypher insert query. This approach can be considered to 4
5 be a counterpart of PRSC, but to convert RDF into PG. 5

6 *PG schemas* Finally, the “Property Graph needs a Schema” Working Group propose a formal definition of PG 6
7 schemas [33]. Some PG engines, like TigerGraph, are based on the use of schemas. For PG engines that do not 7
8 enforce a schema at creation, like Neo4j or Amazon Neptune, the schema may be extracted from the data, as 8
9 proposed by Bonifati et al. [34] or Beereen [35]. As PRSC uses schemas for mapping between PGs and RDF 9
10 graphs, these approaches may be used to automatically list the types existing in the PG to convert *i.e.* the target of 10
11 the rule part in Listing 2. Then the user would only have to provide the way to convert these types into RDF, *i.e.* the 11
12 template graph part. 12
13 13
14 14

15 7. Conclusion 15

16 16
17 This work improves interoperability between the two worlds of Property Graphs and RDF graphs. We have 17
18 presented PRSC, a mapping language to convert PGs into RDF graphs. A mapping, named PRSC context, is written 18
19 by the user and is driven by a schema: PG elements are converted according to their type. By letting the user 19
20 configure the conversion, we aim to better integrate PG data into already existing RDF graphs: the produced RDF 20
21 graphs can be made to use a specific vocabulary, or comply with specific shapes. 21

22 We have also proved that some PRSC contexts, named well-behaved PRSC contexts, are reversible: they do not 22
23 induce any information loss, and therefore it is possible to reverse back to the original PG from the produced RDF 23
24 graph. Finally, we broaden the realm of reversible contexts with the edge-unique extension. Other existing works 24
25 focus either on describing a syntactic construction, or providing a semantic construction that relies on a specific 25
26 syntactic pattern. PRSC lets the user specify the semantics, and proved the reversibility for the two most popular 26
27 methods to model edges: as an RDF resource in PRSC well-behaved contexts, and as a predicate through the edge- 27
28 unique extension. 28
29 29

30 For big PGs, fully converting them into RDF may not scale. For this reason, future works include studying how 30
31 to use PRSC context not only for PG conversion but also to convert SPARQL queries into the usual PG query 31
32 languages Cypher and Gremlin. This would not only address the scalability issues, but also avoid data duplication 32
33 and help for federated queries. 33

34 The expressiveness of PRSC contexts could also be extended. As it is currently presented, PRSC contexts are 34
35 unable to reproduce RDF graphs complying with some ontologies, for example the PG ontology [14]. To solve this 35
36 issue, PRSC contexts should be able to introduce new blank nodes, and not be limited to the ones in the original 36
37 PG. This would lead to new challenges, as the presented reversion algorithm relies on the fact that all blank nodes 37
38 are PG elements. 38

39 Other extensions on expressiveness may also be interesting. For examples, types in PRSC contexts are *closed*, in 39
40 the sense that a complying element must have exactly the properties of the type, barring any other. Allowing extra 40
41 properties in elements of the PGs would be useful, but raises the challenge of converting properties that are not 41
42 known in advance. 42
43 43

44 To let PG data further benefit from the tools that have been developed around RDF, PRSC could also be explored 44
45 in two directions. The use of blank nodes for the PG elements may not be suitable in all cases, especially in a linked 45
46 data context. PRSC could be extended to mint IRIs for nodes and edges of the PG, but this would require to extend 46
47 the mapping language to be able to express these template IRIs. It would also require an adaptation of the reversion 47
48 algorithm to be able to differentiate the minted IRIs from the “static” IRIs of the template, which would require 48
49 additional precautions on well-behaved contexts. 49

50 The provided reversion algorithm does not only work for RDF graphs that were produced by PRSC, but can 50
51 work on any compatible RDF graph. One way to use it would be to use PRSC to convert a PG to an RDF graph, 51

modify the produced RDF graph with RDF-specific tools, *e.g.* by using a reasoner designed for RDF graphs, and then transform back the RDF graph into a PG, which could be considered as equivalent as using a reasoner on a PG. However, this requires to formally characterize the modifications that can be performed on the RDF triples while maintaining the ability to convert it back to a PG.

Appendix A. Proof of projecting and merging back a Property Graph

In this section, we expose the proof for Theorem 5

A.1. Extra mathematical elements

Definition 29 (Restriction). For all functions f , for all sets X , $f|_X = \{(x, f(x)) \mid x \in X \cap \text{Dom}(f)\}$. $f|_X$ is called the restriction of the function f to the set X . In other words, the restriction of a function by a set X is equal to a function in which the domain is restricted to the elements of the set X .

Remark 16. The restriction of a function to its domain is equal to the function itself:

$$f|_{\text{Dom}(f)} = \{(x, f(x)) \mid x \in \text{Dom}(f)\} = f.$$

Remark 17. A functional definition of Definition 29 would be, for all functions f , $f|_X : x \mapsto f(x)$ if $x \in X \cap \text{Dom}(f)$, undefined otherwise.

Lemma 6. For all functions f , for all sets X_1 and X_2 , the union of the function restricted by the two sets is equal to the function restricted by the union of the two sets: $f|_{X_1} \cup f|_{X_2} = f|_{X_1 \cup X_2}$.

Proof.

$$\begin{aligned} & f|_{X_1} \cup f|_{X_2} \\ &= \{(x, f(x)) \mid x \in X_1 \cap \text{Dom}(f)\} \cup \{(x, f(x)) \mid x \in X_2 \cap \text{Dom}(f)\} \\ &= \{(x, f(x)) \mid x \in (X_1 \cap \text{Dom}(f)) \cup (X_2 \cap \text{Dom}(f))\} \\ &= \{(x, f(x)) \mid x \in (X_1 \cup X_2) \cap \text{Dom}(f)\} \\ &= f|_{X_1 \cup X_2} \end{aligned}$$

□

Remark 18. For all function f , for all sets X_1 and X_2 , $f|_{X_1}$ and $f|_{X_2}$ are always compatible.

Theorem 8. If the union of the sets X_i is a super-set of the domain of a function f , then the union of the function f restricted by each set X_i is equal to the function f itself: $(\text{Dom}(f) \subseteq \bigcup_{i=1}^n X_i) \Rightarrow (\bigcup_{i=1}^n f|_{X_i} = f)$.

Proof.

$$\begin{aligned} \bigcup_{i=1}^n f|_{X_i} &= \bigcup_{i=1}^n \{(x, f(x)) \mid x \in X_i \cap \text{Dom}(f)\} \\ &= \left\{ (x, f(x)) \mid x \in \bigcup_{i=1}^n (X_i \cap \text{Dom}(f)) \right\} \end{aligned}$$

$$\begin{aligned}
&= \left\{ (x, f(x)) \mid x \in \left(\bigcup_{i=1}^n X_i \right) \cap \text{Dom}(f) \right\} \\
&= \{(x, f(x)) \mid x \in \text{Dom}(f)\} = f
\end{aligned}$$

□

A.2. Redefinition of the projection

Remark 19. $\text{src}_{\pi_m(pg)}$, $\text{dest}_{\pi_m(pg)}$ and $\text{properties}_{\pi_m(pg)}$ can be redefined by using the restriction:

- $\text{src}_{\pi_m(pg)} = \text{src}_{pg}|_{\{m\}}$
- $\text{dest}_{\pi_m(pg)} = \text{dest}_{pg}|_{\{m\}}$
- $\text{properties}_{\pi_m(pg)} = \text{properties}_{pg}|_{\{(m, \text{str}) \mid \text{str} \in \text{Str}\}}$

Proof. For nodes, $m \in N_{pg}$ cannot be in the domain of src_{pg} , as their domain is a subset of E_{pg} . Therefore, $\text{src}_{pg}|_{\{m\}} = \emptyset \rightarrow \emptyset = \text{src}_{\pi_m(pg)}$.

For edges, $m \in E_{pg}$ is forced to be in the domain of src_{pg} , and its value is $\text{src}_{pg}(m)$. Therefore, $\text{src}_{pg}|_{\{m\}} = (m \mapsto \text{src}_{pg}(m)) = \text{src}_{\pi_m(pg)}$

The same reasoning applies for dest_{pg} .

The new definition of $\text{properties}_{\pi_m(pg)}$ that uses restrictions is immediate from the definition of the restriction. □

A.3. Proof of Theorem 5

Remark 20. The property graphs used in the following proof are described with formula. To help readability, for a given PG x , we allow ourselves to use the notation $N(x)$ instead of N_x . Similar notation will be used for $E(x)$, $\text{src}(x)$, $\text{dest}(x)$, $\text{labels}(x)$ and $\text{properties}(x)$. For example, $N_{\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)}$ will instead of noted $N(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg))$.

Proof. We first need to check if we can apply the \bigoplus operator, i.e. if the three conditions of Definition 27 are met:

- When the π function is applied, nodes remain nodes and edges remain edges. The \bigoplus operator also conserves this property. As $\forall m, N_{\pi_m(pg)} \subseteq N_{pg}$ and $E_{\pi_m(pg)} \subseteq E_{pg}$, the first condition is met.
- The definition of π (restriction of the original function), the definition of \bigoplus (union of the functions) and the Lemma 6 (the union of two restriction is a restriction) imply that the src , dest and properties are compatible.

As \bigoplus is commutative and associative, we can write the following decomposition: $\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg) = (\bigoplus_{m \in N_{pg}} \pi_m(pg)) \bigoplus (\bigoplus_{m \in E_{pg}} \pi_m(pg))$

To prove the theorem, we are going to check if it is true for all functions related to pg .

Edges (E_{pg}):

$$\begin{aligned}
E\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) &= \bigcup_{m \in N_{pg} \cup E_{pg}} E(\pi_m(pg)) && \text{[Definition of } \bigoplus \text{ on E]} \\
&= \left(\bigcup_{m \in N_{pg}} E(\pi_m(pg))\right) \cup \left(\bigcup_{m \in E_{pg}} E(\pi_m(pg))\right) \\
&= \left(\bigcup_{m \in N_{pg}} \emptyset\right) \cup \left(\bigcup_{m \in E_{pg}} \{m\}\right) = \bigcup_{m \in E_{pg}} \{m\} && \text{[Definition of } \pi \text{ on E]} \\
&= E_{pg}
\end{aligned}$$

Nodes (N_{pg}):

$$\begin{aligned}
& N\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) \\
&= \left(\bigcup_{m \in N_{pg}} N(\pi_m(pg))\right) \cup \left(\bigcup_{m \in E_{pg}} N(\pi_m(pg))\right) \\
&= N_{pg} \cup \left(\bigcup_{m \in E_{pg}} N(\pi_m(pg))\right)
\end{aligned}$$

To prove that the last expression above is equal to N_{pg} , we need to prove that $(\bigcup_{m \in E_{pg}} N(\pi_m(pg))) \subseteq N_{pg}$:

$$\forall m \in E_{pg}, N(\pi_m(pg)) = \{src_{pg}(m), dest_{pg}(m)\} \subseteq N_{pg} \Rightarrow \bigcup_{m \in E_{pg}} N(\pi_m(pg)) \subseteq \bigcup_{m \in E_{pg}} N_{pg} = N_{pg}$$

Source of the edges (src_{pg}):

$$\begin{aligned}
& src\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) \\
&= \bigcup_{m \in N_{pg} \cup E_{pg}} src_{pg}|_{\{m\}} \\
&= src_{pg} \quad \left[\text{per Theorem 8, since } \bigcup_{m \in N_{pg} \cup E_{pg}} \{m\} \supseteq E_{pg} = Dom(src_{pg})\right]
\end{aligned}$$

Destination of the edges ($dest_{pg}$): The proof for $dest_{pg}$ follows the same steps as the proof for src_{pg} .

Properties ($properties_{pg}$) The proof is very similar to src_{pg} .

Noticing that:

$$\begin{aligned}
& - \forall m \in N_{pg} \cup E_{pg}, properties(\pi_m(pg)) = properties_{pg}|_{\{(m, str) | str \in Str\}} \\
& - \bigcup_{m \in N_{pg} \cup E_{pg}} \{(m, s) | s \in Str\} = \{(m, str) | m \in N_{pg} \cup E_{pg} \wedge str \in Str\} = (N_{pg} \cup E_{pg}) \times Str \\
& \quad \supseteq Dom(properties_{pg})
\end{aligned}$$

we can reapply the same reasoning as for src_{pg} to find

$$properties\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) = properties_{pg}$$

Labels ($labels_{pg}$) The domain of definition of $labels(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg))$ is:

$$N\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) \cup E\left(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right) = N_{pg} \cup E_{pg}$$

The value of this function is $\forall x \in N_{pg} \cup E_{pg}$,

$$\begin{aligned}
labels\left(\bigcup_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)\right)(x) &= \bigcup_{m \in N_{pg} \cup E_{pg}} labels(\pi_m(pg))(x) \\
&\text{if } labels(\pi_m(pg))(x) \text{ is defined}
\end{aligned}$$

From the definition of π applied on $labels$, two outcomes are possible for $labels(\pi_m(pg))(x)$:

- For $m = x$, $labels(\pi_m(pg))(x) = labels_{pg}(x)$.
- For all other $m \neq x$, $labels(\pi_m(pg))(x)$ is either the empty set or undefined. In both cases, no extra value is contributed to $labels(\bigcup_{m \in N_{pg} \cup E_{pg}} \pi_m(pg))(x)$.

It can be concluded that $labels(\bigcup_{m \in N_{pg} \cup E_{pg}} \pi_m(pg))(x) = labels_{pg}(x)$, so $labels(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)) = labels_{pg}$.

Conclusion : We have demonstrated that

$\forall pg \in PGs, \forall f \in \{N_{pg}, E_{pg}, src_{pg}, dest_{pg}, labels_{pg}, properties_{pg}\}, f(pg) = f(\bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg))$ therefore
 $\forall pg \in PGs, pg = \bigoplus_{m \in N_{pg} \cup E_{pg}} \pi_m(pg)$

□

References

- [1] N. Francis, A. Green, P. Guagliardo, L. Libkin, T. Lindaaker, V. Marsault, S. Plantikow, M. Rydberg, P. Selmer and A. Taylor, Cypher: An Evolving Query Language for Property Graphs, in: *Proceedings of the 2018 International Conference on Management of Data*, SIGMOD '18, Association for Computing Machinery, New York, NY, USA, 2018, pp. 1433–1445. ISBN 9781450347037. doi:10.1145/3183713.3190657.
- [2] M.A. Rodriguez, The Gremlin graph traversal machine and language (invited talk), in: *Proceedings of the 15th Symposium on Database Programming Languages*, DBPL 2015, Association for Computing Machinery, New York, NY, USA, 2015, pp. 1–10. ISBN 9781450339025. doi:10.1145/2815072.2815073.
- [3] R. Angles, The Property Graph Database Model, in: *Proceedings of the 12th Alberto Mendelzon International Workshop on Foundations of Data Management, Cali, Colombia, May 21-25, 2018*, D. Olteanu and B. Poblete, eds, CEUR Workshop Proceedings, Vol. 2100, CEUR-WS.org, 2018. <https://ceur-ws.org/Vol-2100/paper26.pdf>.
- [4] D. Wood, M. Lanthaler and R. Cyganiak, RDF 1.1 Concepts and Abstract Syntax, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/2014/REC-rdf11-concepts-20140225/>.
- [5] D. Brickley and R. Guha, RDF Schema 1.1, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/2014/REC-rdf-schema-20140225/>.
- [6] D.L. McGuinness, F. Van Harmelen et al., OWL web ontology language overview, *W3C recommendation* 10(10) (2004), 2004.
- [7] D. Kontokostas and H. Knublauch, Shapes Constraint Language (SHACL), W3C Recommendation, W3C, 2017, <https://www.w3.org/TR/2017/REC-shacl-20170720/>.
- [8] O. Hartig and B. Thompson, Foundations of an alternative approach to reification in RDF, *arXiv preprint arXiv:1406.3399* (2014).
- [9] O. Hartig, P.-A. Champin, G. Kellogg and A. Seaborne, RDF-star and SPARQL-star, *W3C Community Group Report*, Online at <https://www.w3.org/2021/12/rdf-star.html> (2021).
- [10] O. Lassila, M. Schmidt, O. Hartig, B. Bebee, D. Bechberger, W. Broekema, A. Khandelwal, K. Lawrence, C. López-Enríquez, R. Sharda and B.B. Thompson, The OneGraph Vision: Challenges of Breaking the Graph Model Lock-In, *Semantic Web* 14(1) (2023), 125–134. doi:10.3233/SW-223273.
- [11] A. Dimou, M.V. Sande, P. Colpaert, R. Verborgh, E. Mannens and R.V. de Walle, RML: A Generic Language for Integrated RDF Mappings of Heterogeneous Data, in: *Proceedings of the Workshop on Linked Data on the Web co-located with the 23rd International World Wide Web Conference (WWW 2014)*, Seoul, Korea, April 8, 2014, C. Bizer, T. Heath, S. Auer and T. Berners-Lee, eds, CEUR Workshop Proceedings, Vol. 1184, CEUR-WS.org, 2014. https://ceur-ws.org/Vol-1184/ldow2014_paper_01.pdf.
- [12] M. Sporny, G. Kellogg and M. Lanthaler, JSON-LD 1.0, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/2014/REC-json-ld-20140116/>.
- [13] R. Zhang, P. Liu, X. Guo, S. Li and X. Wang, A Unified Relational Storage Scheme for RDF and Property Graphs, in: *Web Information Systems and Applications*, W. Ni, X. Wang, W. Song and Y. Li, eds, Springer International Publishing, Cham, 2019, pp. 418–429. ISBN 978-3-030-30952-7.
- [14] D. Tomaszuk, R. Angles and H. Thakkar, PGO: Describing property graphs in RDF, *IEEE Access* 8 (2020), 118355–118369. doi:10.1109/ACCESS.2020.3002018.
- [15] N. Fathy, W. Gad, N. Badr and M. Hashem, ProGOMap: Automatic Generation of Mappings From Property Graphs to Ontologies, *IEEE Access* 9 (2021), 113100–113116. doi:10.1109/ACCESS.2021.3104293.
- [16] J. Bruyat, P.-A. Champin, L. Médini and F. Laforest, PREC: semantic translation of property graphs, in: *1st workshop on Squaring the Circles on Graphs, SEMANTiCS*, Amsterdam, Netherlands, 2021. <https://hal.archives-ouvertes.fr/hal-03407785>.
- [17] E. Prud'hommeaux and G. Carothers, RDF 1.1 Turtle, W3C Recommendation, W3C, 2014, <https://www.w3.org/TR/2014/REC-turtle-20140225/>.
- [18] O. van Rest, S. Hong, J. Kim, X. Meng and H. Chafi, PGQL: a property graph query language, in: *Proceedings of the Fourth International Workshop on Graph Data Management Experiences and Systems*, GRADES '16, Association for Computing Machinery, New York, NY, USA, 2016. ISBN 9781450347808. doi:10.1145/2960414.2960421.

- [19] R. Angles, A. Hogan, O. Lassila, C. Rojas, D. Schwabe, P. Szekely and D. Vrgoč, Multilayer graphs: a unified data model for graph databases, in: *Proceedings of the 5th ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, GRADES-NDA '22, Association for Computing Machinery, New York, NY, USA, 2022. ISBN 9781450393843. doi:10.1145/3534540.3534696.
- [20] V. Nguyen, H.Y. Yip, H. Thakkar, Q. Li, E. Bolton and O. Bodenreider, Singleton Property Graph: Adding A Semantic Web Abstraction Layer to Graph Databases, in: *Proceedings of the Blockchain enabled Semantic Web Workshop (BlockSW) and Contextualized Knowledge Graphs (CKG) Workshop co-located with the 18th International Semantic Web Conference, BlockSW/CKG@ISWC 2019, Auckland, New Zealand, October 27, 2019*, R. Samavi, M.P. Consens, S. Khatchadourian, V. Nguyen, A.P. Sheth, J.M. Giménez-García and H. Thakkar, eds, CEUR Workshop Proceedings, Vol. 2599, CEUR-WS.org, 2019. https://ceur-ws.org/Vol-2599/CKG2019_paper_4.pdf.
- [21] F. Orlandi, D. Graux and D. O'Sullivan, Benchmarking RDF Metadata Representations: Reification, Singleton Property and RDF, in: *15th IEEE International Conference on Semantic Computing, ICSC 2021, Laguna Hills, CA, USA, January 27-29, 2021*, IEEE, 2021, pp. 233–240. doi:10.1109/ICSC50631.2021.00049.
- [22] O. Hartig, Foundations of RDF★ and SPARQL★ (An Alternative Approach to Statement-Level Metadata in RDF), in: *Proceedings of the 11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web, Montevideo, Uruguay, June 7-9, 2017*, J.L. Reutter and D. Srivastava, eds, CEUR Workshop Proceedings, Vol. 1912, CEUR-WS.org, 2017. <https://ceur-ws.org/Vol-1912/paper12.pdf>.
- [23] O. Hartig, Foundations to Query Labeled Property Graphs using SPARQL, in: *Joint Proceedings of the 1st International Workshop On Semantics For Transport and the 1st International Workshop on Approaches for Making Data Interoperable co-located with 15th Semantics Conference (SEMANTiCS 2019), Karlsruhe, Germany, September 9, 2019*, L. Kaffee, K.M. Endris, M. Vidal, M. Comerio, M. Sadeghi, D. Chaves-Fraga and P. Colpaert, eds, CEUR Workshop Proceedings, Vol. 2447, CEUR-WS.org, 2019. <https://ceur-ws.org/Vol-2447/paper3.pdf>.
- [24] S. Khayatbashi, S. Ferrada and O. Hartig, Converting property graphs to RDF: a preliminary study of the practical impact of different mappings, in: *GRADES-NDA '22: Proceedings of the 5th ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, Philadelphia, Pennsylvania, USA, 12 June 2022, V. Kalavri and S. Salihoglu, eds, ACM, 2022, pp. 10:1–10:9. doi:10.1145/3534540.3534695.
- [25] S. Das, J. Srinivasan, M. Perry, E.I. Chong and J. Banerjee, A Tale of Two Graphs: Property Graphs as RDF in Oracle, in: *Proceedings of the 17th International Conference on Extending Database Technology, EDBT 2014, Athens, Greece, March 24-28, 2014*, S. Amer-Yahia, V. Christophides, A. Kementsietsidis, M.N. Garofalakis, S. Idreos and V. Leroy, eds, OpenProceedings.org, 2014, pp. 762–773. doi:10.5441/002/EDBT.2014.82. <https://doi.org/10.5441/002/edbt.2014.82>.
- [26] H. Thakkar, D. Punjani, J. Lehmann and S. Auer, Two for one: querying property graph databases using SPARQL via gremlinator, in: *Proceedings of the 1st ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, Houston, TX, USA, June 10, 2018, A. Arora, A. Bhattacharya, G.H.L. Fletcher, J.L. Larriba-Pey, S. Roy and R. West, eds, ACM, 2018, pp. 12:1–12:5. doi:10.1145/3210259.3210271.
- [27] T. Delva, J. Arenas-Guerrero, A. Iglesias-Molina, Ó. Corcho, D. Chaves-Fraga and A. Dimou, RML-star: A Declarative Mapping Language for RDF-star Generation, in: *Proceedings of the ISWC 2021 Posters, Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 20th International Semantic Web Conference (ISWC 2021), Virtual Conference, October 24-28, 2021*, O. Seneviratne, C. Pesquita, J. Sequeda and L. Etcheverry, eds, CEUR Workshop Proceedings, Vol. 2980, CEUR-WS.org, 2021. <https://ceur-ws.org/Vol-2980/paper374.pdf>.
- [28] 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/>.
- [29] G. Abuoda, D. Dell'Aglio, A. Keen and K. Hose, Transforming RDF-star to Property Graphs: A Preliminary Analysis of Transformation Approaches, in: *Proceedings of the QuWeDa 2022: 6th Workshop on Storing, Querying and Benchmarking Knowledge Graphs co-located with 21st International Semantic Web Conference (ISWC 2022), Hangzhou, China, 23-27 October 2022*, M. Saleem and A.N. Ngomo, eds, CEUR Workshop Proceedings, Vol. 3279, CEUR-WS.org, 2022, pp. 17–32. <https://ceur-ws.org/Vol-3279/paper2.pdf>.
- [30] R. Angles, H. Thakkar and D. Tomaszuk, Mapping RDF databases to property graph databases, *IEEE Access* **8** (2020), 86091–86110. doi:10.1109/ACCESS.2020.2993117.
- [31] G.A. Atemez and A. Huynh, Knowledge Graph publication and browsing using Neo4J, in: *1st workshop on Squaring the Circles on Graphs, SEMANTiCS*, Amsterdam, Netherlands, 2021.
- [32] H. Chiba, R. Yamanaka and S. Matsumoto, G2GML: Graph to Graph Mapping Language for Bridging RDF and Property Graphs, in: *Proceedings of the ISWC 2020 Demos and Industry Tracks: From Novel Ideas to Industrial Practice co-located with 19th International Semantic Web Conference (ISWC 2020), Globally online, November 1-6, 2020 (UTC)*, K.L. Taylor, R.S. Gonçalves, F. Lécué and J. Yan, eds, CEUR Workshop Proceedings, Vol. 2721, CEUR-WS.org, 2020, pp. 363–368. <https://ceur-ws.org/Vol-2721/paper591.pdf>.
- [33] R. Angles, A. Bonifati, S. Dumbrava, G. Fletcher, A. Green, J. Hidders, B. Li, L. Libkin, V. Marsault, W. Martens, F. Murlak, S. Plantikow, O. Savkovic, M. Schmidt, J. Sequeda, S. Staworko, D. Tomaszuk, H. Voigt, D. Vrgoc, M. Wu and D. Zivkovic, PG-Schema: Schemas for Property Graphs, *Proc. ACM Manag. Data* **1**(2) (2023), 198:1–198:25. doi:10.1145/3589778.
- [34] A. Bonifati, S. Dumbrava, E. Martinez, F. Ghasemi, M. Jaffré, P. Luton and T. Pickles, DiscoPG: Property Graph Schema Discovery and Exploration, *Proc. VLDB Endow.* **15**(12) (2022), 3654–3657. doi:10.14778/3554821.3554867. <https://www.vldb.org/pvldb/vol15/p3654-bonifati.pdf>.
- [35] N. Beeren, Designing a Visual Tool for Property Graph Schema Extraction and Refinement: An Expert Study, *CoRR* **abs/2201.03643** (2022). <https://arxiv.org/abs/2201.03643>.