



HAL
open science

Transforming Property Graphs

Angela Bonifati, Filip Murlak, Yann Ramusat

► **To cite this version:**

Angela Bonifati, Filip Murlak, Yann Ramusat. Transforming Property Graphs. 50th International Conference on Very Large Databases, Aug 2024, Guangzhou, China. pp.2906-2918, 10.14778/3681954.3681972 . hal-04777556

HAL Id: hal-04777556

<https://hal.science/hal-04777556v1>

Submitted on 12 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.

Transforming Property Graphs (Extended Version)*

Angela Bonifati
Lyon 1 Univ., Liris CNRS & IUF
angela.bonifati@univ-lyon1.fr

Filip Murlak
Univ. of Warsaw
fmurlak@mimuw.edu.pl

Yann Ramusat
Lyon 1 Univ., Liris CNRS
yann.ramusat@liris.cnrs.fr

ABSTRACT

In this paper, we study a declarative framework for specifying transformations of property graphs. In order to express such transformations, we leverage queries formulated in the Graph Pattern Calculus (GPC), which is an abstraction of the common core of recent standard graph query languages, GQL and SQL/PGQ. In contrast to previous frameworks targeting graph topology only, we focus on the impact of data values on the transformations—which is crucial in addressing users’ needs. In particular, we study the complexity of checking if the transformation rules do not specify conflicting values for properties, and we show this is closely related to the satisfiability problem for GPC. We prove that both problems are PSPACE-complete.

We have implemented our framework in openCypher. We show the flexibility and usability of our framework by leveraging an existing data integration benchmark, adapting it to our needs. We also evaluate the incurred overhead of detecting potential inconsistencies at run-time, and the impact of several optimization tools in a Cypher-based graph database, by providing a comprehensive comparison of different implementation variants. The results of our experimental study show that our framework exhibits large practical benefits for transforming property graphs compared to ad-hoc transformation scripts.

PVLDB Reference Format: Angela Bonifati, Filip Murlak, and Yann Ramusat. Transforming Property Graphs (Extended Version)¹. PVLDB, 14(1): XXX-XXX, 2020.
doi:XX.XX/XXX.XX

PVLDB Artifact Availability:


The source code, data, and/or other artifacts have been made available at <https://github.com/yannramusat/TPG>.

1 INTRODUCTION

Query languages for property graphs—those supported by existing systems, such as Neo4j’s openCypher [23] or Oracle’s PGQL [32], and those codified in international standards, such as GQL and SQL/PGQ [22]—define their semantics in terms of sets of tuples. This is inadequate for data interoperability tasks such as data migration or data integration, where outputs of some queries are to be fed directly to other queries. To support this kind of *composability*, queries should be able to output property graphs rather than sets of tuples. Such queries can be seen as *transformations*, turning an input property graph into an output property graph.

Interoperability of graph data has received little attention so far, compared to the relational and XML data models [3]. Notable research in the area [6, 11] relies on the simplified graph data model that had been devised to provide the foundations for querying the

topology of graphs with formalisms such as conjunctive regular path queries (CRPQs) [8] or regular queries [33]. As the simplified graph data model ignores the presence of properties (key-value pairs stored in nodes and edges), it is too far from the property graph models used in graph databases such as Neo4j or Tigergraph, and cannot be a foundation for practical solutions. These currently rely on opaque external libraries, such as Neo4j’s APOC [29], or involve complex handcrafted queries [30], as illustrated below.



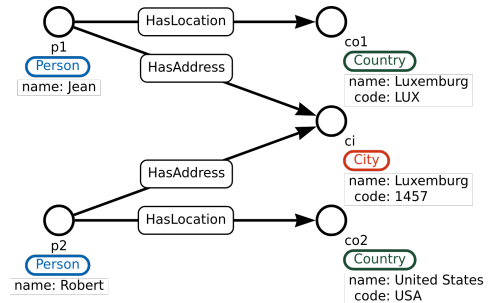
(i) Input property graph G containing ingested relational data.

```

1 MATCH (u:User)
2 MATCH (a:Address) WHERE a.aid = u.address
3 MATCH (l:Location) WHERE l.aid = u.address
4 WITH u, collect(a) AS Addresses, collect(l) AS Locations
5 CREATE (p:Person)
6 SET p.name = u.name
7 WITH p, Addresses, Locations
8 UNWIND Addresses AS a
9 MERGE (ci:City {name: a.cityName})
10 SET ci.code = a.cityCode
11 MERGE (p)-[:HasAddress]->(ci)
12 WITH p, Locations
13 UNWIND Locations AS l
14 MERGE (co:Country {name: l.countryName})
15 SET co.code = l.countryCode
16 MERGE (p)-[:HasLocation]->(co)

```

(ii) Ad-hoc transformation script in openCypher.



(iii) Resulting output property graph H .

Figure 1: Ad-hoc transformation of raw ingested data.

Example 1.1. Figure 1 illustrates a graph transformation scenario, in which a user has imported relational data into the popular

¹A shorter version of this paper has been accepted for publication in VLDB 2024.

Neo4j graph database and would like to reshape it into a semantically meaningful property graph instance, to facilitate navigational querying. The relational data consists of three tables,

User(name, address), Address(aid, cityName, cityCode),
 Location(aid, countryName, countryCode),

with primary keys consisting of the underlined attributes, and having two foreign keys: *aid* references *address* in User from both Address and Location.

Figure 1 (i) shows a rudimentary property graph obtained after importing the relational data, using a generic ingestion method, such as Cypher’s LOAD CSV clause. In the resulting property graph, each node represents a single tuple of the relational instance, with the relation’s name represented as the label, and the attributes stored in the node’s properties. Note that there are no edges in this property graph: relationships between places, locations, and users are represented by way of foreign keys, just like in the original relational instance. Needless to say, this is not the best way to represent a relational instance as a property graph.

The user now wants to transform the instance in Figure 1 (i) into one that makes better use of the property graph data model by facilitating navigational operations in queries like “Which people live in the same city as Jean?”. The user intends to create a node for each person, city, and country, and replace foreign key references with explicit relationships. Figure 1 (ii) shows an implementation of this transformation in openCypher that closely follows a graph refactoring solution described in Neo4j’s GraphAcademy [30]. The reader will notice how difficult it is to relate the constructs of this query to the informal specification above. Even just making sense of the MERGE clauses interleaved with implicit grouping and list manipulations (UNWIND and collect) is a daunting task for an unacquainted user. But the query leverages other advanced idioms too. For instance, in Line 5, the script creates as many nodes of type Person as there are rows output by the previous WITH clause: one for each *u*, due to implicit grouping. In line 9, the script generates one City node for each *distinct* value found in property cityName across all *a*’s; this is because the property name is specified as a cityName in the MERGE clause. Similarly, in line 14, a single Country node is created for each *distinct* value found in property countryName.

Figure 1 (iii) shows the output property graph obtained by running the script on the input property graph from Figure 1 (i). It reveals that the ad-hoc transformation fails to account for the fact that cities are weak entities that cannot be identified by their name alone, and conflates Luxembourg in Europe with Luxembourg in the US. Detecting such errors is hard because openCypher lacks a transparent mechanism for specifying identities of created elements. ◀

As we have seen, ad-hoc transformation scripts are error-prone and hard to interpret and analyze. Moving away from handcrafted implementations to declarative specifications has long been recognized as pivotal for solving data programmability problems [10]. The aim of this work is to lay the theoretical foundations for the declarative specification of property graph transformations, and facilitate practical solutions for turning such specifications into executable scripts in modern property graph query languages. Constraint-based, fully declarative formalisms, such as schema mappings for relational [9, 17, 18] and graph [6, 24] data, allow multiple

target solutions, leading to ambiguous transformations [12, 15]. For property graphs, this makes the schema mapping problem undecidable even under strong restrictions [24]. We avoid this problem by focusing on transformations that return a unique, well-defined output instance for each input instance, thus facilitating direct execution.

We propose a rule-based formalism that allows the user to describe the output property graph based on the input property graph, by specifying not only labels, properties, and relationships between output elements, but also their identities. The formalism builds upon the *Graph Pattern Calculus* (GPC) [21], which is an abstraction of the common graph pattern matching features of GQL and SQL/PGQ [16]. GPC is adequate in terms of expressive power: it has ample facilities for querying properties and even on property-less graphs it goes well beyond classical formalisms such as RPQs and CRPQs. It is suitable for theoretical investigation owing to its concise syntax and rigorous semantics. It should also keep our proposal future-proof by ensuring out-of-the-box compatibility with the expected implementations of these standards. Until then, we can rely on the already implemented graph query languages, such as Neo4j’s openCypher [23] or Oracle’s PGQL [32], which were a strong inspiration for GQL. Indeed, the actual query language used in the rules can be seen as a parameter of the framework.

In contrast to the purely topological formalism of [11], specifications of property-aware transformations can easily become *inconsistent*, when they attempt to specify two different values for the same property of a given element. Detecting such conflicts naturally comes to the foreground of static analysis. As we show, this problem is tightly connected to the *satisfiability* problem for GPC+ (GPC extended with union and projection, also introduced in [21]), which is to decide if there is a property graph satisfying a given GPC+ query. Exploiting this connection, we establish tight complexity bounds for both these problems, showing that they are PSPACE-complete. To the best of our knowledge, this is the first static-analysis result on GPC. Given that query satisfiability is the work horse of static analysis throughout database theory, we believe that with the adoption of the GQL standard our result will find other uses. An immediate consequence for property graph transformations is that consistency cannot be checked statically due to the prohibitive cost, and conflicts must be handled dynamically, during the execution of the transformation.

In order to prove that our formalism can serve as a foundation for practical data interoperability solutions, we provide a proof-of-concept implementation. As no existing query engine supports GQL yet, we rely on the Neo4j’s open-source implementation of openCypher [23, 26], which offers most of the functionalities of GQL described in [22]. We study the case when the rules are provided by the users and describe a generic, easily automated method of translating these rules into executable openCypher scripts, and apply it manually to selected realistic property graph transformations derived from real-world data integration scenarios of the iBench benchmark suite [4]. We perform a comprehensive experimental study gauging the efficiency of conflict detection and the effect of rule order and various optimizations on several implementation variants. We confirm that our implementation performs well in all scenarios and scales to large input data. We also demonstrate that our framework can be successfully applied in a concrete data

integration scenario on real-world data [20], and report the results of a small-scale user study confirming that our framework enhances readability and usability of transformations.

In summary, our main contributions are the following.

- We propose a comprehensive declarative formalism for specifying transformations of property graphs, compatible with SQL/PGQ and the upcoming GQL standard.
- We identify consistency as a key static-analysis problem, and show that it is interreducible with satisfiability of GPC+ queries and that both are PSPACE-complete.
- We provide a proof-of-concept implementation of our formalism in openCypher, and apply it to realistic scenarios of graph-shaped data transformations.
- We show experimentally that our solution scales to large input data, handles on-the-fly conflict detection with low overhead, and enhances readability and usability, without sacrificing performance.

The rest of paper is organized as follows. In Section 2, we recall the property graph data model along with GPC. In Section 3, we give syntax and semantics of our graph transformation formalism. In Section 4, we discuss the consistency in relation with satisfiability of GPC+ queries, and establish the complexity bounds. In Section 5, we describe our proof-of-concept implementation. In Section 6, we present both the experiments and the user study. In Section 7 and in Section 8, we discuss the related work and conclude the paper.

2 PRELIMINARIES

We briefly introduce the basic concepts of the property graph data model and the *Graph Pattern Calculus* (GPC) that we use in this paper. We mostly follow the definitions from [21].

2.1 Data model

Conforming to the formal specification originating from [21], a *property graph* G is a tuple $\langle N, E, \lambda, \text{src}, \text{tgt}, \delta \rangle$ where \mathcal{O} , \mathcal{L} , \mathcal{K} and Const are disjoint countable sets of object identifiers (ids), labels, keys (also called property names) and constants (data values), and

- $N \subset \mathcal{O}$ is the finite set of node ids in G ;
- $E \subset \mathcal{O}$ is the finite set of edge ids;
- N and E are disjoint;
- $\lambda : N \cup E \rightarrow 2^{\mathcal{L}}$ is a labeling function that associates to every id a (possibly empty) finite set of labels;
- $\text{src}, \text{tgt} : E \rightarrow N$ define the source and target of each edge;
- $\delta : (N \cup E) \times \mathcal{K} \rightarrow \text{Const}$ is a finite-domain partial function that associates a constant with an id and a key from \mathcal{K} .

The node ids and edge ids will be respectively called the *nodes* and *edges* of the property graph.

That is to say, a property graph is a *multigraph* in the sense that two vertices may be connected by more than one edge, even with these edges having the same label(s), and that loops are permitted. All the *elements* of the database (the nodes and the edges) store a finite set of property-value pairs, represented by δ .

A property graph is presented in Figure 1 (iii). It contains information about peoples' location such as the City and the Country they live in. We see that it contains one node with label City and two nodes with label Country; two edges with label HasLocation and

two edges with label HasAddress; all nodes have property *name*; and all nodes with label City or Country have an additional property *code*. Annotations $p1, p2, \dots, co2$ are node identifiers; edge identifiers are not shown.

2.2 Graph Pattern Calculus

In the following, we introduce GPC by means of examples. The reader can refer to [21] for more details on GPC, and to [22] for insight on how it will actually be used in GQL.

In Example 1.1, the user can retrieve from the property graph in Figure 1 (iii) “all people who live in the same city as a person named $\$name$ ” using the following GPC query:

$$\langle (x : \text{Person}) \xrightarrow{:\text{HasAddress}} (y : \text{City}) \xleftarrow{:\text{HasAddress}} (z : \text{Person}) \mid \langle name = \$name \rangle \rangle$$

This is an example of a *path pattern*, which is essentially a regular path query [8] augmented with *conditioning*: the filter $\langle name = \$name \rangle$ checks that the value of the property *name* is indeed the one sought. Given a property graph, this pattern returns the nodes that can be matched to y .

One can also use *graph patterns* in GPC (also called patterns or queries in this paper), which are conjunctions of path patterns. For example, the following query retrieves pairs of people living in the same city, such that one person knows, possibly indirectly, the other one:

$$\langle (x : \text{Person}) \xrightarrow{:\text{HasAddress}} (y : \text{City}) \xleftarrow{:\text{HasAddress}} (z : \text{Person}), \langle x \rangle \xrightarrow{:\text{Knows}}^{1..∞} (y) \rangle$$

Such graph patterns generalize conjunctive two-way regular path queries [8] to property graphs.

In GPC, each path pattern occurring in a graph pattern must be qualified with a *restrictor* among the set of simple, trail (used by default if none is given) and shortest. The restrictor's purpose is to ensure a finite result set: simple prevents repetition of nodes along a path; trail prevents repetition of edges; and shortest selects only the paths of minimal length among all the paths between two nodes.

For the ease of exposition, we simplify the semantics of GPC. We assume that a pattern only returns a set of bindings (in [21], a tuple of witnessing paths is also returned with each binding). In GPC, variables used in the scope of a repetition operator, such as $1..∞$, are called *group variables* and are bound to lists of nodes or edges. The remaining variables are called *singleton variables* and are bound to single nodes or edges. For the purpose of our transformation formalism we restrict the output of queries to singleton variables.

For a GPC pattern P , a tuple \bar{x} of singleton variables in P , and a property graph G , we write $\llbracket P \rrbracket_G^{\bar{x}}$ for the set of bindings of \bar{x} returned by P on G . For instance, if P is the first query above and G is the property graph depicted in Figure 1 (iii), we have $\llbracket P \rrbracket_G^y = \{(y \mapsto p2)\}$ when $\$name$ is “Jean”, $\llbracket P \rrbracket_G^y = \{(y \mapsto p1)\}$ when $\$name$ is “Robert”, and $\llbracket P \rrbracket_G^y = \emptyset$ for any other name. (Note that the trail restrictor has been used by default.)

3 PROPERTY GRAPH TRANSFORMATIONS

In this section, we present our declarative formalism for specifying property graph transformations. An example is given in Figure 2. The specification consists of two rules. Each rule collects data from the input graph with a GPC pattern on the left of \implies , and specifies elements of the input graph using the expression on the right. This expression resembles a GPC pattern, but it has specifications of the element’s property values instead of filters and specifications of element identifiers instead of variables to be matched (new variables will reappear on the right-hand side, in a slightly different role). In what follows we discuss how new identifiers are generated using *Skolem functions* (Section 3.1) and how identifiers, labels, and properties of output elements are specified using *content constructors* (Section 3.2). Then, we describe the general form of rules (Section 3.3) and explain their semantics in terms of a procedure that generates an output property graph given an input property graph (Section 3.4). We shall also see if the transformation in Figure 2 fixes the issues discussed in Example 1.1.

3.1 Generating output identifiers

Throughout the paper we assume that all identifiers in input property graphs come from a countable set $\mathcal{S} \subset \mathcal{O}$ of *input identifiers*, and ensure that all identifiers in output property graphs come from a countable set $\mathcal{T} \subset \mathcal{O} \setminus \mathcal{S}$ of *output identifiers*. Following [11], to generate identifiers in the output graph, we use Skolem functions. Specifically, we use a fixed injective Skolem function

$$f : \bigcup_{k \in \mathbb{N}} (\mathcal{O} \cup \text{Const} \cup \mathcal{L})^k \rightarrow \mathcal{T}.$$

In the context of relational schema mappings and data exchange, Skolem functions are used for *value invention* [5], e.g., to generate artificial primary keys of new tuples in a way that makes it possible to refer to them in foreign keys. The way we use Skolem functions is similar, but not the same, because element identifiers are not data values. Rather, they are the property-graph analogue of object identifiers from the object-oriented data model [1, 27]. Most of the time they are invisible to the user, and are not expected to carry any information beyond the identity of the element. Thus, the specific choice of function f is truly irrelevant, as long as f is injective.

Example 3.1. In the rules in Figure 2 the Skolem function is kept implicit, but its arguments are explicitly listed. For example, in the subexpression $((u) : \text{Person})$, on the right-hand side of both rules, (u) indicates that the identifier of the output node is $f(u)$ where u is (the identifier of) a node selected from the input property graph by the left-hand side GPC pattern, such as $u1$. Because the same nodes u are selected in both rules, the subexpressions $((u) : \text{Person})$ in both rules will be referring to the same output nodes. Further, $(\ell.\text{countryName})$ specifies the node identifier as $f(\ell.\text{countryName})$, where $\ell.\text{countryName}$ name refers to the value of the property *countryName* in a node ℓ selected from the input graph, such as “United States”, and similarly for $(a.\text{cityName})$. If $\ell.\text{countryName} = a.\text{cityName}$ for some ℓ and a , which can happen in our example, $(\ell.\text{countryName})$ and $(a.\text{cityName})$ will indicate the same output node. ◀

3.2 Content constructors

A property graph transformation must be able to specify not only the identifiers of output elements, but also their labels and properties. For this purpose, we use content constructors. A *content constructor* is an expression of the form:

$$\begin{aligned} C(\bar{x}) := \{ \\ \text{Id: } (a_1, \dots, a_k) \\ \text{Labels: } L \\ \text{Properties: } \langle k_1 = v_1, \dots, k_n = v_n \rangle \} \end{aligned}$$

where \bar{x} is a tuple of variables, L is a finite set of labels; each k_i is a property name from \mathcal{K} ; each v_i is either a data value $c \in \text{Const}$, or an expression of the form $x.a$ for $x \in \bar{x}$ and $a \in \mathcal{K}$; and each a_i is either a constant $c \in \text{Const}$, or a label $\ell \in \mathcal{L}$, or an expression of the form $x.a$ or x for $x \in \bar{x}$ and $a \in \mathcal{K}$. The field *Id* specifies the identity of the node by listing the arguments to be fed to the Skolem function. The fields *Labels* and *Properties* specify labels and properties present in an element. Importantly, they do not forbid additional labels and properties, which will allow the user to split the description of an element across multiple rules, if the user so desires. We write $C.\text{Id}$ for the content of the *Id* field of C , and similarly for other fields. When \bar{x} is clear from the context, we simply write C instead of $C(\bar{x})$.

Example 3.2. In the first rule in Figure 2, new *Country* nodes are described using the following content constructor:

$$\begin{aligned} C_t(a, u, \ell) := \{ \\ \text{Id: } (\ell.\text{countryName}) \\ \text{Labels: } \{\text{Country}\} \\ \text{Properties: } \langle \text{name} = \ell.\text{countryName}, \text{code} = \ell.\text{countryCode} \rangle \}. \end{aligned}$$

It specifies the identities and the values of properties *name* and *code* of new *Country* nodes in terms of the values of properties *countryName* and *countryCode* retrieved from elements to which variable ℓ is bound in the input graph. Rather than using the abstract syntax introduced above, the rule in Figure 2 presents C_t in GPC-like syntax [21] as

$$\begin{aligned} & ((\ell.\text{countryName}) : \text{Country}) \\ & \langle \text{name} = \ell.\text{countryName}, \text{code} = \ell.\text{countryCode} \rangle \end{aligned}.$$

3.3 Transformations

We describe transformations in terms of property graph transformation rules. Each rule brings together the data retrieved from the input property graph by a GPC pattern and a description of output elements expressed with content constructors.

We recall that the semantics of GPC is defined such that a query returns *tuples*. Each tuple represents a *binding* of singleton variables in that query to elements of the property graph.

We have two kinds of *property graph transformation rules*: node rules and edge rules. A *node rule* is an expression of the form:

$$P(\bar{x}) \implies (C(\bar{x}))$$

where $P(\bar{x})$ is a GPC query with singleton variables \bar{x} and $C(\bar{x})$ is a content constructor. An *edge rule* is an expression of the form:

$$P(\bar{x}) \implies (C_s(\bar{x})) \xrightarrow{C(\bar{x})} (C_t(\bar{x}))$$

where $P(\bar{x})$ is a GPC query with singleton variables \bar{x} and $C_s(\bar{x}), C(\bar{x})$ and $C_t(\bar{x})$ are content constructors. Finally, a *property graph transformation* is a finite set of property graph transformation rules.

Example 3.3. The first edge rule in Figure 2 is built from the content constructor C_t as defined in Example 3.2, and of the following two content constructors C_s and C :

$$\begin{array}{ll} C_s(a, u, \ell) := \{ & C(a, u, \ell) := \{ \\ \text{Id: } (u) & \text{Id: } () \\ \text{Labels: } \{\text{Person}\} & \text{Labels: } \{\text{HasLocation}\} \\ \text{Properties: } \langle \text{name} = u.\text{name} \rangle, & \text{Properties: } \langle \rangle. \quad \blacktriangleleft \end{array}$$

The above definition allows specifying multiple labels with a single constructor as well as specifying the labels of a single element using multiple rules. This feature, illustrated in the following example, is crucial. Without it, in the presence of type hierarchies, one would need negation in the query language to avoid duplicating output elements. In our setting, GPC does not permit negating patterns and it is unlikely for the complexity upper bounds in Section 4.1 to hold when this form of negation is added.

Example 3.4. As discussed in Example 3.1, if for some nodes ℓ and a selected by the GPC patterns in the rules of Figure 2, $\ell.\text{countryName}$ and $a.\text{cityName}$ are equal, then the C_t constructors in both rules refer to the same output node. For instance for, $\ell = l1$ and $a = a1$ in the input graph in Figure 1 (i), both rules refer to the node $f(\text{“Luxemburg”})$ in the output graph in Figure 3. In consequence, this node has two labels, City and Country. This, quite likely, is not what the user actually wants. We will later see how to fix it by adjusting the rules. \blacktriangleleft

Property graphs are *multigraphs* and our rules allow specifying multiple edges with the same endpoints by using different arguments for the Skolem function. We will see an example in Section 5.

We refer to the right-hand side expressions in node (resp. edge) rules as node (resp. edge) constructors. We also allow rules of a more general form, illustrated in Figure 4, where a comma-separated list of node and edge constructors can be used on the right-hand side. We also support aliasing, with scope limited to a single rule. For instance, in the rule in Figure 4, we introduce alias $x = (u)$ in the first edge constructor, and use it in the second edge constructor. Both these extensions are syntactic sugar. To eliminate aliases, we simply substitute them with their definitions: in the example, we replace x in the second edge constructor with (u) . Then, we split the rules: for each node or edge constructor on the right-hand side, we create a separate rule with the same GPC pattern on the left.

3.4 Semantics

In this section, we describe operationally in Algorithm 1 how a transformation given as a set of node and edge rules turns an input property graph into an output property graph. In Section 5, we will see how to implement this efficiently in an existing graph database.

Given a GPC query $P(\bar{x})$, a content constructor $C(\bar{x})$ and a binding \bar{o} for $P(\bar{x})$ over an input property graph G , we define $C.\text{Id}(\bar{o})$ by replacing in $C.\text{Id}$ each x_j with o_j and each $x_j.a$ with $\delta_G(o_j, a)$. Similarly, we define $C.\text{Properties}(\bar{o})$ by replacing in $C.\text{Properties}$ each $x_j.a$ with $\delta_G(o_j, a)$.

Algorithm 1 Semantics of a set of transformation rules.

Input: A property graph G and a set of transformation rules T .
Output: An output of the transformation T over G , a property graph $T(G) = \langle N, E, \lambda, \text{src}, \text{tgt}, \delta \rangle$.

- 1: initialize $T(G)$ to the empty property graph
- 2: **for** each edge rule $P(\bar{x}) \implies (C_s(\bar{x})) \xrightarrow{C(\bar{x})} (C_t(\bar{x})) \in T$ **do**
- 3: add rules $P(\bar{x}) \implies (C_s(\bar{x}))$ and $P(\bar{x}) \implies (C_t(\bar{x}))$ to T
- 4: **for** each node rule $P(\bar{x}) \implies (C(\bar{x})) \in T$ **do**
- 5: **for** each binding $\bar{o} \in \llbracket P \rrbracket_G^{\bar{x}}$ **do**
- 6: $N \leftarrow N \cup \{o := f(C.\text{Id}(\bar{o}))\}$
- 7: $\lambda(o) \leftarrow \lambda(o) \cup C.\text{Labels}$
- 8: set $\delta(o, k)$ to c if $C.\text{Properties}(\bar{o})$ sets property k to c
- 9: **for** each edge rule $P(\bar{x}) \implies (C_s(\bar{x})) \xrightarrow{C(\bar{x})} (C_t(\bar{x})) \in T$ **do**
- 10: **for** each binding $\bar{o} \in \llbracket P \rrbracket_G^{\bar{x}}$ **do**
- 11: $o_s \leftarrow f(C_s.\text{Id}(\bar{o})); o_t \leftarrow f(C_t.\text{Id}(\bar{o}))$
- 12: $E \leftarrow E \cup \{o := f(o_s, C.\text{Id}(\bar{o}), o_t)\}$
- 13: $\text{src}(o) \leftarrow o_s; \text{tgt}(o) \leftarrow o_t$
- 14: $\lambda(o) \leftarrow \lambda(o) \cup C.\text{Labels}$
- 15: set $\delta(o, k)$ to c if $C.\text{Properties}(\bar{o})$ sets property k to c

Example 3.5. We describe, step by step, the operations carried out by Algorithm 1 on the input consisting of the property graph G from Figure 1 (i) and the transformation T_1 which contains only the first of the two rules in Figure 2.

First, the GPC query

$$P(u, a, \ell) := (u : \text{User}), (a : \text{Address}), (\ell : \text{Location}) \\ \langle u.\text{address}=a.\text{aid}, u.\text{address}=\ell.\text{aid} \rangle$$

is executed on G (only once in the entire process) and outputs the set of bindings $\llbracket P \rrbracket_G^{u,a,\ell} = \{(u \mapsto u1, a \mapsto a1, \ell \mapsto l1), (u \mapsto u2, a \mapsto a2, \ell \mapsto l2)\}$.

In Line 3, the single edge rule of T_1 is split into node rules $P(\bar{x}) \implies (C_s(\bar{x}))$ and $P(\bar{x}) \implies (C_t(\bar{x}))$, where C_s and C_t have been defined in Example 3.3 and 3.2, respectively. These two node rules are added to T_1 , which initially contains no node rules.

Suppose that the node rule $P(\bar{x}) \implies (C_s(\bar{x}))$ is considered first in the loop in Line 4. Two output nodes are created with respective identifiers $f(u1)$ and $f(u2)$ (Line 6), one for each binding. Initially, they have no labels, $\lambda(f(u1)) = \lambda(f(u2)) = \emptyset$, and no properties. Then both get label Person (Line 7) and their property *name* is set to “Jean” and “Robert”, respectively (Line 8).

Next, the algorithm moves to the node rule $P(\bar{x}) \implies (C_t(\bar{x}))$. Two nodes are created in the output with respective identifiers $f(\text{“Luxemburg”})$ and $f(\text{“United States”})$ (Line 6), one for each binding; they both get label Country (Line 7); and their properties *name* and *code* are filled in (Line 8).

Finally, the algorithm steps through the only edge rule in T_1 . For the first binding, the nodes corresponding to the endpoints of the edge that has to be created, namely $o_s := f(u1)$

$$(u : \text{User}), (a : \text{Address}), (\ell : \text{Location}) \Rightarrow ((u) : \text{Person}) \xrightarrow{\text{HasLocation}} ((\ell.\text{countryName}) : \text{Country})$$

$\langle u.\text{address}=a.\text{aid}, u.\text{address}=\ell.\text{aid} \rangle$ $\langle \text{name}=u.\text{name} \rangle$ $\langle \text{name}=\ell.\text{countryName}, \text{code}=\ell.\text{countryCode} \rangle$

$$(u : \text{User}), (a : \text{Address}), (\ell : \text{Location}) \Rightarrow ((u) : \text{Person}) \xrightarrow{\text{HasAddress}} ((a.\text{cityName}) : \text{City})$$

$\langle u.\text{address}=a.\text{aid}, u.\text{address}=\ell.\text{aid} \rangle$ $\langle \text{name}=u.\text{name} \rangle$ $\langle \text{name}=a.\text{cityName}, \text{code}=a.\text{cityCode} \rangle$

Figure 2: Transformation T given as a set of rules.




Figure 3: Output property graph $T(G)$.

and $o_t := f(\text{"Luxemburg"})$ are retrieved (Line 11). They correspond to the nodes that were created, from this binding, by the node rules that were added to T_1 in Line 3. An edge with id $f(f(u1), f(\text{"Luxemburg"}))$ is created (Line 12); its source and target are set to $f(u1)$ and $f(\text{"Luxemburg"})$, respectively (Line 13); it gets label HasLocation (Line 14); and no property is filled (Line 15). For the second binding, an edge is created by the same process between the nodes $f(u2)$ and $f(\text{"United States"})$. ◀

The role of Algorithm 1 is to give semantics to a set of transformation rules: it explains how the outputs of the multiple rules are consolidated into a single output property graph. The following result shows that our transformations are indeed graph-to-graph transformations, offering a way to meet the expected requirements of future versions of standard graph query languages [22].

PROPOSITION 3.6. *Given an input property graph G and a property graph transformation T , Algorithm 1 always returns a valid instance of the property graph data model.*

Although Algorithm 1 always returns a valid property graph (Proposition 3.6), property values may depend on the order in which the rules and bindings are considered in Lines 4–5 and 9–10. Hence, the result of the transformation may be ill-defined on some inputs. We investigate this further in the next section.

4 DETECTING CONFLICTS

As one would expect from any expressive property graph transformation language, our formalism supports manipulating properties of output nodes and edges. Compared to purely structural mechanisms, such as [11], this poses additional challenges.

Example 4.1. Let us continue Example 3.5 by now considering the two-rule transformation T presented in Figure 2. The second rule gets split into two nodes rules, one of which is

$$P(\bar{x}) \Rightarrow ((a.\text{cityName}) : \text{City})$$

$\langle \text{name}=a.\text{cityName}, \text{code}=a.\text{cityCode} \rangle$

Suppose that this node rule is processed in Line 4 after the two node rules discussed in Example 3.5. The algorithm attempts twice

to create a node with identifier $f(\text{"Luxemburg"})$ (Line 6), once for each binding. However, a node with identifier $f(\text{"Luxemburg"})$ has been already created by the second node rule in Example 3.5. In consequence, the label City is added to this node (Line 7) and its properties name and code are set to Luxemburg and 1457, respectively (Line 8), overriding previous values Luxemburg and LUX . This means that one of the two values of property code is lost and it depends on the processing order of rules which one it is. Indeed, the mapping now conflates not only two cities called Luxemburg , as in Example 1.1, but also the country Luxemburg . This time, however, the error is easy to spot: looking at the rules we see immediately that the identity of the output nodes depends exclusively on the name of the city/country, which means that all cities and countries with the same name are conflated. We can fix the transformation easily by including information about the corresponding country in the identity of each City node, for instance by replacing $(a.\text{cityName})$ with $(a.\text{cityName}, \ell.\text{countryName})$ in rule (2) in Figure 2. ◀

Detecting the modelling error in the rules in Figure 2 requires human insight (basic understanding of geography) but we hope to make it easier by insisting on explicit identity specification in transformations. On the other hand, setting an output property to conflicting values is something one can try to capture abstractly and detect automatically. This is what we do next. In the reminder of this section we focus on detecting conflicts statically, by analysing a set of transformation rules to check if it can exhibit this pathological behavior on some input. We come back to handling conflicts dynamically in Section 5 and Section 6. Due to the limited space, most proofs are moved to the appendix [14].

4.1 Consistency

By a *conflict* we mean a situation when Algorithm 1 resets a previously set property to a different value, as illustrated in Example 4.1. A transformation T is *consistent* if for every input property graph G , no execution of Algorithm 1 results in a conflict. Note that even a transformation consisting of a single rule can be inconsistent, because different bindings for the same rule can cause a conflict.

We study the following fundamental static analysis problem, in the setting where there is no source schema constraining the set of possible input property graphs.

Consistency. Given a transformation T , check if T is consistent.

As we show next, consistency of transformations is deeply related to satisfiability of GPC patterns. A GPC pattern is *satisfiable* if it returns a non-empty set of answers on some property graph. Towards the goal of establishing complexity lower bounds for the consistency problem, we provide a polynomial-time reduction from the satisfiability problem for GPC.

Satisfiability. Given a GPC pattern P , check if P is satisfiable.

LEMMA 4.2. *The satisfiability problem for GPC is PTIME-reducible to the transformation consistency problem.*

PROOF. For a GPC pattern P , let T_P be the transformation consisting of the following two rules

$$P() \implies \underbrace{((c) : \ell)}_{\langle k=5 \rangle} \quad \text{and} \quad P() \implies \underbrace{((c) : \ell)}_{\langle k=7 \rangle}$$

for some fixed label ℓ , constant c , and property name k . These rules are not conflicting with themselves, because their node constructors do not depend on the binding. However, they are conflicting with each other on a graph G if P returns at least one answer on G . Hence, T_P is consistent iff P is not satisfiable. \square

For the converse of Lemma 4.2 to hold we need to move to GPC+, a simple extension of GPC with projection and union [21].

LEMMA 4.3. *The transformation consistency problem is PTIME-reducible to the satisfiability problem for GPC+.*

PROOF SKETCH. For a pair of rules R and S , and an attribute a , we can write a Boolean GPC+ query $Q_{R,S,a}()$ that detects if some matches of R and S lead to different values for attribute a in the same element of the output graph. Because there are polynomially many such triples, we can take the union of all such queries to obtain the final GPC+ query to be checked for satisfiability. \square

We now turn to study the complexity of the satisfiability problem for GPC and GPC+. The two lemmas above will allow us to draw conclusions for the consistency problem in Section 4.3.

4.2 The complexity of satisfiability

In Theorem 4.5, we establish that checking if a GPC+ query is satisfiable is a PSPACE-complete problem (modulo certain assumptions on the use of restrictors). We believe that this result is interesting in its own right, beyond the application to transformation consistency we consider this paper. Indeed, deciding whether a query expressed in a given query language is satisfiable is a fundamental problem in database theory. Very little is known to this date about GPC from a theoretical viewpoint, and our work is one of the first to tackle a key static analysis task related to this query language.

LEMMA 4.4. *The satisfiability problem for GPC is PSPACE-hard.*

PROOF SKETCH. We show how to reduce the membership problem for an arbitrary PSPACE language to the satisfiability of a GPC query. Let L be a language in PSPACE and M a deterministic polynomial-space Turing machine that recognizes L in space $c \cdot p(n)$ for a fixed constant c and polynomial p . In the following, n denotes the length of the word w which is an input to M .

We construct the following GPC pattern P :

$$P() := \rho(x)_{\langle \theta_1 \rangle} \left(\left[(u) \rightarrow (v) \right]_{\langle \theta_2 \rangle} \right)^{1.. \infty} (y)_{\langle \theta_3 \rangle}$$

The intuition is the following. We can represent a configuration of M in a single node, using a polynomial number of properties. The pattern $(x)_{\langle \theta_1 \rangle}$ is responsible for encoding the initial configuration of M over the input word w . The pattern $\left[(u) \rightarrow (v) \right]_{\langle \theta_2 \rangle}$ ensures that there exists a valid transition of M between the configurations

represented by nodes u and v . Finally, $(y)_{\langle \theta_3 \rangle}$ specifies that node y represents an accepting configuration.

We can use techniques similar to the proof of the Cook-Levin Theorem [25] to construct in time polynomial in n the formulæ θ_1 , θ_2 , and θ_3 . The size of P is then clearly polynomial in n . This reduction works with any $\rho \in \{\text{shortest, simple, trail}\}$. \square

For completeness, we also provide the matching upper bound (under some assumptions) and obtain Theorem 4.5 as a result. The details of the upper-bound proof are highly technical and the claim depends heavily on the design choices made for GPC.

THEOREM 4.5. *The satisfiability problem for GPC+ queries using only the simple and trail restrictors is PSPACE-complete.*

We prove the upper-bound of Theorem 4.5 by inductively constructing an equality type over all variables in the query. This non-deterministic procedure uses only a polynomial amount of space by avoiding storing the full match of the pattern. Unfortunately, this does not extend to queries using the shortest restrictor: they seem to require storing the full match. We leave open the question of pinpointing the exact complexity of satisfiability for such queries.

Given the high complexity lower bounds, one might wonder whether there are useful subclasses of GPC with tractable satisfiability. In Lemma 4.6 below, we show that even under strong limitations, satisfiability is still intractable.

LEMMA 4.6. *The satisfiability problem is NP-hard even for single-node GPC patterns.*

4.3 Back to consistency

From Theorem 4.5, Lemma 4.2, and Lemma 4.3, we obtain the following fundamental result.

COROLLARY 4.7. *The consistency problem is PSPACE-complete for transformations using only simple and trail restrictors.*

In fact, the PSPACE lower bound holds already for transformations using only two rules and any single restrictor. From Lemma 4.6 and Lemma 4.2 it follows that the problem remains intractable even for transformations using very restricted GPC queries.

In the light of these high complexity lower bounds, it is unlikely that conflict detection can be handled statically in practice. This means that conflicts have to be handled dynamically, when the transformation is executed. In Section 5 we discuss how this can be implemented in practice and in Section 6 we show experimentally that the incurred overhead is affordable.

5 TRANSLATION TO CYPHER

Algorithm 1 can be seen as an abstraction of a transformation engine: it takes a transformation and an input property graph, and produces an output property graph. In this section we show how to compile a transformation to an openCypher script that can be directly executed in any openCypher engine. This is similar in spirit to executable SQL scripts for relational schema mappings, scalable and efficient in producing target solutions [10].

We first discuss the overall complexity of Algorithm 1. Lines 6 and 12 involve a set-theoretic union and, without appropriate optimization, their cost is proportional to the current number of elements in $T(G)$ in each iteration of the loop. Lines 7–8 and 13–15

can be implemented in $O(1)$ provided that Lines 6 and 12 return a pointer to the element $o := f(C.Id(\bar{o})) \in T(G)$. Thus the overall complexity of Algorithm 1 on input G is:

$$O(t_{int} + n_c \cdot Int(G, T) \cdot |T(G)|) \quad (3)$$

where n_c is the total number of content constructors in T , $Int(G, T)$ and t_{int} are respectively the total size of all intermediate results $\llbracket P \rrbracket_G^{\bar{x}}$ and the overall running time for computing $\llbracket P \rrbracket_G^{\bar{x}}$, with $P(\bar{x})$ ranging over all left-hand sides of rules in T .

Thus, the total time taken by Algorithm 1 implemented naively is quadratic in the size of the property graphs, which makes it practically unusable for large input instances. However, the complexity heavily depends on the implementation of the set-theoretic unions.

Plain implementation. In Figure 5 we showcase the result of our translation strategy for the variant T_r of T , presented in Figure 4. This transformation has only one rule and is translated into a single executable script. For transformations with several rules, each rule of the transformation is independently translated into a script.

Cypher’s built-in `elementId` primitive provides access to the identifier of an element, which is unique among all elements in the database. It plays a crucial role in our implementation as we actively use these identifiers as arguments to the Skolem function generating output identifiers. To the best of our knowledge, there is no explicit control of the creation of new identifiers in Neo4j, so we equip nodes and edges in the output graph with a special property `_id` that plays the role of controllable element identifier.

Lines 1–3 correspond to the left part of the rule and are responsible for retrieving the necessary information from the input property graph. Recall that, in Line 3 of Algorithm 1, a node rule is added for each endpoint of every edge constructor in the transformation. Accordingly, in the openCypher script, each node constructor used on the right-hand side of the rule is considered separately (Lines 4–12). Similarly to how Skolem functions are usually implemented in relational data exchange for schema mapping tasks [10], we implement them with string operations, e.g., `_id: "(" + elementId(u) + ")"`. We rely on the semantics of Cypher’s `MERGE` clause, described in [26], to implement the set-theoretic union: in Lines 4, 7, and 10, `MERGE` checks whether an element with this identifier already exists in the graph; either one exists and is retrieved, or a new element is created. Adding the corresponding label(s) to the retrieved node (Line 7 of Algorithm 1) is implemented with the native Cypher’s `SET` clause in Lines 5, 8, and 11. Similarly, the properties of the nodes (Line 8 of Algorithm 1) are set in Lines 6, 9, and 12.

Finally, the relationships are created (Lines 13–18). To keep the value of `_id` unique among all elements in the output, and given the restriction that relationships hold a single label in Neo4j, the edge labels have been provided as arguments to the Skolem functions in Figure 4. Note that, when we merge an edge pattern, we are sure that the endpoints already exist in the database.

We point out that the `_id` property and the `_dummy` label are internal data; they are of no interest to the end user and can be dropped after the transformation with Cypher’s `REMOVE` command.

Optimizations. Optimizing the `MERGE` clauses in Lines 4, 7, 10, 13, and 16 which implement the set-theoretic unions is crucial in reducing the overall execution time of the transformation.

As is the case in most database management systems, Neo4j provides facilities for query optimization. The two that are relevant in this context are indexes and uniqueness constraints. An *index* permits to retrieve efficiently nodes with a given label that have a specific value at a given property. When we know in advance that all these values are unique, we can make further use of *uniqueness constraints* (UCs). Note that in our implementation, we maintain the invariant that each `_id` is unique across all elements in the output.

In the version of Neo4j Community Edition that we use for running the experiments, indexes are implemented using b-trees, which means that the cost of testing if an index with n elements contains a given key is $O(\log n)$. That is, by using indexes we can improve the worst-case complexity of Algorithm 1 to:

$$O(t_{int} + n_c \cdot Int(G, T) \cdot \log |T(G)|) \quad (4)$$

In the next section we comprehensively evaluate the advantages and disadvantages of using indexes and uniqueness constraints on nodes and relationships, defined on the label/property pair `_dummy/_id`.

Conflict detection. The consistency problem is unfortunately PSPACE-complete by Corollary 4.7, so we cannot efficiently check the declarative specification at compile time. Instead, we need to be ready for potential inconsistencies at run time.

Figure 6 illustrates how one can detect conflicts on the property code when creating a new `City` node. We use the `ON MATCH` sub-clause of the `MERGE` clause to perform a comparison when we set a property for an existing node. Notice that a different rule could have led to the creation of this node and, consequently, `z`. code may be empty; in this case the operator `<>` returns `false` and the correct specification is reached.

6 EXPERIMENTS

Our experimental study has three main objectives: (i) evaluate the benefits of using this formalism for transforming property graphs in practical use-cases over a large amount of data, (ii) evaluate the involved overhead of detecting potential inconsistencies at run-time, and (iii) compare with the native openCypher approach such as the one presented in Figure 1 (ii).

Experimental setting. We have implemented our property graph transformations in openCypher 9 using a local Neo4j Community Edition instance in version 5.9.0. For monitoring the results and performing the database management tasks required in our methodology, we have used Python 3.11 and the official Neo4j Python Driver 5.9.0. The source code, datasets, and configuration files are available on the public `GITHUB` repository of the project. We performed the experiments on an HP EliteBook 840 G3 with an Intel Core i7-6600U CPU and 32GiB of system memory (2133 MHz).

Datasets. Due to the lack of benchmarks for property graph transformations, in order to build realistic scenarios we have adapted the mappings from several relational data integration scenarios from the iBench suite [4]. In particular, we encode relational input instances as property graphs by creating a node for each tuple (no edges), and we let the target instances be property graphs as well, thus simulating graph-to-graph transformations. Each mapping in a scenario corresponds to a rule of our formalism. Following the method described in Section 5, we compute an openCypher script implementing each rule.




Figure 4: Refined property graph transformation T_r .

```

1 MATCH (u:User)
2 MATCH (a:Address) WHERE a.aid = u.address
3 MATCH (l:Location) WHERE l.aid = u.address
4 MERGE (x:_dummy { _id: "(" + elementId(u) + ")" })
5 SET x:Person,
6   x.name = u.name
7 MERGE (y:_dummy { _id: "(" + l.countryName + ")" })
8 SET y:Country,
9   y.name = l.countryName, y.code = l.countryCode
10 MERGE (z:_dummy { _id: "(" + a.cityName + ")" })
11 SET z:City,
12   z.name = a.cityName, z.code = a.cityCode
13 MERGE (x)-[hl:HasLocation] {
14   _id: "(" + elementId(x) + "," + "HasLocation" + "," +
15   + elementId(y) + ")" }->(y)
16 MERGE (x)-[ha:HasAddress] {
17   _id: "(" + elementId(x) + "," + "HasAddress" + "," +
18   + elementId(z) + ")" }->(z)

```

Figure 5: openCypher script corresponding to T_r (Figure 4).

```

1 MERGE (z:_dummy { _id: "(" + a.cityName + ")" })
2 ON CREATE SET z:City, z.code = a.cityCode
3 ON MATCH SET z:City, z.code = CASE WHEN z.code <> a.cityCode
4   THEN "Conflict detected!" ELSE a.cityCode END

```

Figure 6: Detecting conflicts on the property code.

Table 1: Scenarios characteristics.

Scenario	Labels / Properties			Rules	
	$ \mathcal{L}_{in} $	$ \mathcal{L}_{out}^{node} $	$ \mathcal{L}_{out}^{edge} $	$ \mathcal{K} $	$ T $ n_c
PersonAddress	2	2	1	7	2 6
FlightHotel	2	3	2	5	1 7
PersonData	3	3	2	3	1 5
GUSToBIOSQL	7	5	4	80	8 18
DBLPToAmalgam1	7	5	4	140	10 22
Amalgam1ToAmalgam3	15	2	1	128	8 22

The middle part of Table 1 reports the number $|\mathcal{L}_{in}|$ of input labels in each scenario (corresponding to the number of different relations in the original iBench scenario), the number $|\mathcal{L}_{out}^{node}|$ of output node labels, the number $|\mathcal{L}_{out}^{edge}|$ of output edge labels, and the number $|\mathcal{K}|$ of properties. The right part provides information about the number of rules in the scenario $|T|$ and the total number n_c of content constructors. In each scenario, for each of the $|\mathcal{L}_{in}|$ input node labels, we generated up to 10^5 nodes.

Methodology. The main abstraction in our implementation is a *Scenario* which describes an input property graph database that contains some data of a given size stored in specific node and relationship properties. As previously shown in Figure 1 (i), given the iBench output, we create a node for each tuple, having as properties (key/value pairs) the columns names and column values. We also add the Cypher specification of a set of indexes and constraints on the output side, that are created before executing the transformation




Figure 7: Comparison between uniqueness constraints and indexes for computing $T(G)$.




Figure 8: Impact of indexing strategies and implementation variants on the computation of $T(G)$.




Figure 9: Average computation time for different orders of execution of the rules.

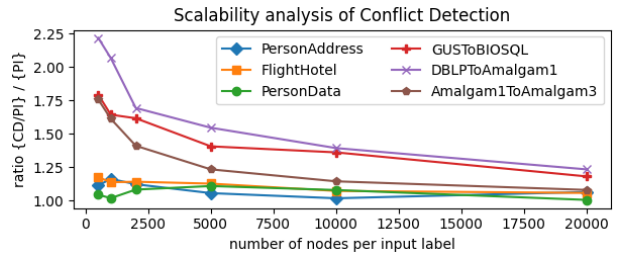


Figure 10: Ratio between the time for computing $T(G)$ with and without conflict detection (using PI_NI).

when the output data is still empty. This step is not time consuming and takes on average less than one millisecond per index.

A scenario includes several Cypher queries—one for each transformation rule—that are successively applied. To simulate the process of transforming one graph into another, and to distinguish between input and output data, we have used disjoint sets of labels in the input and output instance. Thus, a single database instance holds both input and output data at a time, but contains initially no output data. As a final step, a scenario is responsible for flushing the database and removing the indexes and constraints in order to have a fresh database instance before executing the next scenario. Note that the query cache (execution plans) is cleared when one of them is dropped. We monitored the total amount of time spend by Neo4j in applying the transformation rules. Each experiment generally represents the average taken over 5 runs of a scenario.

Alternative implementation using separate indexes. In Section 5, we discussed an implementation of the framework, the *Plain implementation* (PI), which uses a single index on the output side to speed up the retrieval of already existing nodes by Cypher’s **MERGE** clause. Using a single index for all nodes in the output may severely impact the performance of the implementation as the cost of index maintenance may become prohibitive. To quantify this, we compare with an alternative implementation, the *Separate indexes implementation* (SI), where the label is part of the argument list, similar to the case of relationships. The goal here is to mitigate the cost of maintaining a very large index by splitting the data into many smaller ones. Note that it is still possible to detect conflicts in this variant with a slight modification of the code from Figure 6.

Impact of indexes and uniqueness constraints. We start by comparing the advantages of using uniqueness constraints on nodes (NUC) and indexes (NI) on the two alternative implementations, NI and SI. Figure 7 reports the results for our FlightHotel scenario, showing that for large input data, indexes tend to outperform UCs.

We next investigate the impact of using combinations of indexes on nodes and relationships. We compared variants with *indexes on nodes and relationships* (NI_RI), *indexes on nodes only* (NI), *indexes on relationships only* (RI), and *without indexes* (WI) for the previous PI and SI implementations and their respective variants with conflict detection enabled: *Conflict Detection over Plain implementation* (CD/PI), *Conflict Detection over Separate indexes* (CD/SI). We showcase in Figure 8, on a logarithmic scale, the results that were obtained for the DBLPToAmalgam1 scenario. Other scenarios show similar trends and they are reported in the appendix [14]. It is clear from the figure that the choice of indexes to use is crucial. Using indexes only on nodes is more efficient than using a combination of indexes on nodes and relationships, which is in turn more efficient than using indexes only on relationships or using no index at all. The key reason of this behavior is that indexes on nodes already allow accessing the endpoints of edges, along with the edges themselves, efficiently. Additional indexes on edges do not help, but do incur additional overhead.

The positive point that emerges from this study is that the implementation does not require fine tuning to be efficient in a specific scenario; using indexes only on nodes is consistently the best approach to use. Additionally, when using indexes only on nodes (NI), the Plain implementation (PI) is negligibly slower than Separate indexes implementation (SI), whereas for other combinations of indexes it is noticeably slower. We discussed in Example 3.4 that PI allows for more flexible use of labels compared to the SI (which

corresponds to having a dedicated Skolem function for each set of labels). In view of the above results, in the remaining experiments we focus on the Plain implementation with node indexes (PI_NI).

Impact of rule order. Our formalism is declarative and does not specify the order for the execution of the rules. Hence, we have investigated the impact of different orders on the computation time of the transformation. We compare the minimum, average and maximum running times using random orders with the (fixed) order provided in iBench as baseline. Figure 9 reports the results for the DBLPToAmalgam1 scenario; error bars indicate minimum and maximum computation times observed over 20 independent runs. For space reasons, GUSToBIOSQL and Amalgam1ToAmalgam3, exhibiting similar results, are deferred to the appendix [14].

We can observe that the impact of the order in which the rules are applied on the execution time of the transformation is not substantial; randomized orders have a variance similar to that of a fixed order. It is fair to say that the performance of our implementation does not rely on any specific execution order.

Overhead of detecting potential inconsistencies. We evaluated the impact of turning on conflict detection (over PI_NI) by investigating the ratio between computation time with and without conflict detection. The theoretical complexity of our implementation of Algorithm 1 with conflict detection is:

$$O(t_{int} + n_c \cdot \text{Int}(G, T) \cdot (\log |T(G)| + c)) \quad (5)$$

for c a constant modeling the cost of the conditional statement. Thus the overhead incurred by detecting conflicts is $1 + \frac{c}{\log |T(G)|}$, which tends to 1 in larger scenarios.

The results presented in Figure 10 experimentally validate that the incurred overhead of conflict detection is reasonably low for large input instances, and stays within a constant factor, roughly between 1 and 1.3, depending on the scenario.




Figure 11: Run-time comparison for different likelihood of conflicts (using CD/PI_NI with 10^5 nodes for each input type).

Robustness against incidence of conflicts. iBench’s scenarios have very few or no conflicts. To investigate the generalizability of these results to more conflict-prone scenarios, we designed an experiment using an additional randomization step: when a rule attempts to set a value for an attribute, the value is changed randomly. This allows us to control the average number of conflicts in the output. Figure 11 reports, on a logarithmic scale, the results for all our scenarios, with varying likelihood of conflicts, ranging from 0% to 100%. Note that, the size of the output is preserved because only the attributes are affected, not the topology of the graph.

Table 2: Running times and size of intermediate data (ICIJ).

Rules	t_{int}	t	$Int(G, T)$	$ T(G) $	O/I	t_{int}^e	t^e
R1 – R4	2,757	11,192	374,955	748,524	1.996	0.007	0.015
R5 – R9	3,553	5,946	62,242	82,616	1.327	0.057	0.072
R10 – R13	15,509	36,775	1,906,686	1,905,547	0.999	0.008	0.019
R14 – R17	9,667	21,006	493,556	1,173,720	2.378	0.020	0.018
R18 only	8,407	25,640	785,124	1,570,470	2.000	0.011	0.016

We observe that the prevalence of conflicts has no impact on the execution time, suggesting that our framework’s stability is preserved, even with a large proportion of conflicts in the output.




Figure 12: Horizontal scaling, with varying number of independent copies of the scenario.

Horizontal scalability. We have investigated how well our framework scales with the number of rules and input labels. We built larger scenarios by taking an increasing number of independent copies of the scenarios from Table 1. The resulting transformations reach over one hundred rules and input labels, and over 1.5 million input nodes (in total). Figure 12 reports the results for the PersonData and Amalgam1ToAmalgam3 scenarios.

We observe the running time scales smoothly (almost linearly) as the number of copies increases. Results on other scenarios follow similar trends and are deferred to the appendix [14].

Improvement over handcrafted scripts: a user study. To compare empirically the readability and usability of the script-based approach and our framework, we ran an ad-hoc user study involving 12 participants that were all already familiar with openCypher.

We compared the ability of the participants to understand the behavior of some provided openCypher scripts and transformations in clearly defined scenarios. Only 25% of the participants have been able to fully understand the behaviour of the openCypher scripts, whereas 67% of them succeeded with transformations. In average, participants have scored 50% on openCypher scripts and 90% on our framework. Participants were also asked to compare openCypher scripts and our framework in terms of understandability, intuitiveness, and flexibility; they all have favored our framework by a great margin. For space reasons, the questionnaire, the participant’s answers, and the full discussion of the obtained results are deferred to the appendix [14].

6.1 Use-case Study: Improving Data Integration

In this section, we want to compare the cost of running the whole transformation compared to the cost of querying the source property graph to extract the bindings (intermediate data). To this end, we use a real-world dataset, the *Offshore Leaks Database and guide from the International Consortium of Investigative Journalists (ICIJ)* [20], a property graph with 1,908,466 nodes and 3,193,390 edges taken from [31]. This dataset consolidates data from several leaks (*Panama Papers*, *Bahamas Leaks*, etc.) collected by ICIJ over a

period of ten years, but still presents the consolidated data in a heterogeneous manner. The dataset contains information about *entities* (off-shore companies), *officers* of those, *intermediaries* (middlemen who help set up off-shore companies), and *jurisdictions* (countries or territories where off-shore companies are registered). We have designed a modular 18-rule transformation aiming to uniformize the presentation of the information contained in the graph. The rules are grouped into 5 subsets, each addressing a specific refactoring goal motivated below. For space reasons, we have deferred the rules themselves to the appendix [14].

Refactoring registered addresses (R1 – R4). The ICIJ database contains the registered *addresses* of the *officers*, and *entities*. These rules are responsible for creating nodes representing *countries* and linking to the *addresses*. Because the data is semi-structured and collected from multiple sources, information may be stored in attributes that have different names, or may even not be available at all. All these cases are covered by these four rules.

Uniformizing address information for intermediaries (R5 – R9). After careful investigation, we found that the registered *address* of *intermediaries* can be stored in three different ways in the database: (i) an *intermediary* can have a direct relationship with an *address*, (ii) the *address* can be stored in the properties of the node itself, and (iii) when neither of the two previous cases applies, it is necessary to retrieve the *address* of an *entity* linked to this *intermediary*. These rules permit to consistently store address information.

Exporting the nodes (R10 – R13). These rules copy the node information from the source to the target; they are necessary to preserve all the information from the original graph.

Improving similarity detection (R14 – R17). Because the dataset consolidates multiple leaks, certain specific relationships, such as *similar* and *same_as*, are used to indicate that some *officers* (resp. *addresses*) are likely to represent the same real life entity. These rules focus on exporting this data and improving the similarity detection. This is illustrated by Rule 15 shown in Figure 13 which composes the relationships *similar* and *same_as* to ensure that both its endpoints correspond to *officers* having the same address. (This is because *similar* encompasses address similarity.) Then, it checks whether their names are also similar. If both conditions hold, it safely adds a similarity edge between the endpoints in the output.

Refactoring jurisdictions (R18). The last rule is responsible for connecting the *jurisdictions* with their associated *countries*; this information is not explicitly stored in the initial database.

Results. Our experimental results are reported in Table 2. We report the time t_{int} (in *ms*) the database takes to retrieve the intermediate data; the total time t of running the transformation (extracting the bindings and constructing the output); the size $Int(G, T)$ of intermediate data; the size of the output $T(G)$; the ratio O/I of the size of the output to the size of intermediate data. To account for the differences in the sizes of the outputs of the respective tasks, we also report the average time t_{int}^e taken to produce each binding of the intermediate result, and the average time t^e taken to construct each element of the output. We break down the reported values into groups of rules corresponding to the aforementioned integration tasks.

There are several things that we can learn from Table 2. First, the overhead t^e/t_{int}^e of turning the intermediate results into a proper property graph is reasonable. For Rules R14 – R17, it is even comparatively more efficient to compute the output property graph

$$\begin{array}{c}
(o : \text{Officer}) \xrightarrow{\text{:similar}^{0..∞}} (: \text{Officer}) \xrightarrow{\text{:registered_address}} (: \text{Address}) \xrightarrow{\text{:similar}} (: \text{Address}) \xleftarrow{\text{:registered_address}} (: \text{Officer}) \xleftarrow{\text{:similar}^{0..∞}} (p : \text{Officer}) \\
\text{(toLower(o.name)=toLower(p.name))} \\
\implies ((o) : \text{T_Officer}) \xrightarrow{\text{():T_Similar} / \langle \text{link}=\text{"similar name and address as"} \rangle} ((p) : \text{T_Officer})
\end{array}$$

Figure 13: Improved similarity detection (R15).

(overhead is 0.9). The worst case is for rules *R10* – *R13* exhibiting an overhead of 2.4. Second, the ratio O/I is also reasonable, ranging from 1 to 2.4. This shows that, in practical contexts, $Int(G, T)$ can be assumed to have a size comparable to $|T(G)|$.

Thus, we have demonstrated that the overhead incurred by producing a property graph rather than a set of bindings is acceptable for a realistic transformation in a real-life integration scenario.

7 RELATED WORK

Schema mapping and data exchange. Specifying the relationship between two relational (or XML) schemas using a set of declarative assertions is a task known as *schema mapping* [9, 17, 28]. This relation, is usually *non functional*, i.e. given an input instance I , several target instances satisfying the mapping constraints exist.

Schema mappings and data exchange have been studied in [6, 24] for graph databases. The mapping languages considered are based on classical graph database queries such as regular path queries [7], limited in their expressivity by not supporting data values. Moreover, answering queries on the target is already intractable in data complexity for RPQs [6] and undecidable for data RPQs [24]. In comparison, our transformation framework provides more flexibility by including the support for data values, and any target query can be answered by simple execution on the produced property graph.

Graph transformations. Graph database transformations defined using Datalog-like rules based on acyclic conjunctive two-way regular path queries have been investigated in [11]. They study three fundamental static analysis problems: *type checking*, *equivalence of transformations under graph schemas*, and *schema elicitation*. They show all these problems to be in EXPTIME.

A key difference with our work lies in the graph database model they consider, which does not have data values. We have seen that dealing with data values gives rise to the consistency checking problem, which is key to understanding if a property graph transformation is well-defined. Moreover, their query language – a fragment of Datalog, is not practical for querying property graphs [13]. Another difference is that they are using a single dedicated node constructor for each label. In Section 3 and 5, we have seen that this approach is too rigid for dealing with multiple labels.

Object-creating functions. The Skolem functions we use in our constructors resemble to the object creating functions that are used in the object-oriented database model [1, 27]. Among transformation languages based on oid generation, StruQL [19] specifically operates on object-oriented semi-structured instances where nodes can either be data values or contain an oid and labeled edges can connect oid nodes to oid or value nodes. The major difference with our work is that they have *multi-valued* attributes: i.e., an oid node may be connected via a -edges to several value nodes. Hence, additional integrity constraints are necessary to ensure a correct

modeling of property graphs in their model. Therefore, they did not take into account the problem of consistency.

Interoperability of graph data. Although RDF, RDF-star and the property graph data model share striking similarities, both being based on elementary graph concepts, like nodes and edges, intricate interoperability issues arise when attempting to exchange data between them. RDF-star notably allows for annotating RDF triples with metadata annotations, which are notoriously difficult to capture within the property graph data model as witnessed in [2].

The main concern of transformation languages between graph data models is thus primarily focused on solving the well-known impedance mismatch problem [10], which does not arise in our setting because we have property graphs for both input and output. Our transformation language can be thus more expressive, and can be executed by the graph database management system itself.

Mining the identities of nodes across networks. Network alignment is a technique for finding node correspondences between two or more networks. It can be used, for example, to associate nodes from different social networks with the same user [34]. Nodes are identified based on their similarities with respect to both their features (i.e., their properties) and their neighborhood.

While these methods are not part of graph transformation formalisms, they can be used to guide the construction of graph transformations. For instance, in Section 6.1, the results of network alignment (the similarity edges in the Offshore Leaks Database) were leveraged to better integrate data coming from multiple leaks.

8 CONCLUSION

Our research is the first to lay the theoretical foundations for declarative property graph transformations, and facilitate practical solutions for turning such specifications into executable scripts in modern property graph query languages. New challenges arise from the specification of property-aware transformations, notably the task of checking if a transformation is *consistent*. Using a proof-of-concept implementation of our formalism in openCypher, we showcase the efficiency of our approach for transforming property graphs for both real-world and synthetic datasets.

This work paves the way for obtaining compositional semantics for graph query languages. As a future direction, we will investigate the model extensions needed for the above semantics, by addressing label and path variables, and aggregates. Meanwhile, our framework can already seamlessly support the group variables of GPC because those are list of identifiers that can be flattened into the identifier lists of the constructors. Finally, we will investigate how to assist users in the design process of their transformation rules; for instance by lifting *schema matching* techniques [9, 10] from relational to property graph schemas.

REFERENCES

- [1] Serge Abiteboul and Paris C. Kanellakis. 1998. Object Identity as a Query Language Primitive. *J. ACM* 45, 5 (1998), 798–842.
- [2] Ghadeer Abuoda, Daniele Dell’Aglio, Arthur Keen, and Katja Hose. 2022. Transforming RDF-star to Property Graphs: A Preliminary Analysis of Transformation Approaches. In *QuWeDa 2022*. 17–32.
- [3] Marcelo Arenas, Pablo Barcelo, Leonid Libkin, and Filip Murlak. 2010. *Relational and XML Data Exchange* (1st ed.). Morgan and Claypool Publishers.
- [4] Patricia C. Arocena, Boris Glavic, Radu Ciucanu, and Renée J. Miller. 2015. The IBench Integration Metadata Generator. *VLDB* 9, 3 (2015), 108–119.
- [5] Patricia C. Arocena, Boris Glavic, and Renee J. Miller. 2013. Value Invention in Data Exchange. In *SIGMOD*. 157–168.
- [6] Pablo Barceló, Jorge Pérez, and Juan Reutter. 2013. Schema Mappings and Data Exchange for Graph Databases. In *ICDT*. 189–200.
- [7] Pablo Barceló, Jorge Pérez, and Juan L. Reutter. 2012. Relative Expressiveness of Nested Regular Expressions. In *AMW*. 180–195.
- [8] Pablo Barceló Baeza. 2013. Querying Graph Databases. In *PODS*. 175–188.
- [9] Z. Bellahsene, A. Bonifati, and E. Rahm. 2011. *Schema Matching and Mapping*.
- [10] Philip A. Bernstein and Sergey Melnik. 2007. Model Management 2.0: Manipulating Richer Mappings. In *SIGMOD*. 1–12.
- [11] Iovka Boneva, Benoît Groz, Jan Hidders, Filip Murlak, and Slawek Staworko. 2023. Static Analysis of Graph Database Transformations. In *PODS*. 251–261.
- [12] Angela Bonifati, Ugo Comignani, Emmanuel Coquery, and Romuald Thion. 2017. Interactive Mapping Specification with Exemplar Tuples. In *SIGMOD*. 667–682.
- [13] Angela Bonifati, G.H.L. Fletcher, Hannes Voigt, and N. Yakovets. 2018. *Querying graphs*. Morgan & Claypool Publishers.
- [14] Angela Bonifati, Filip Murlak, and Yann Ramusat. 2024. Transforming Property Graphs (Appendix). <https://github.com/yannramusat/TPG/blob/main/Appendix.pdf>
- [15] Laura Chiticariu and Wang-Chiew Tan. 2006. Debugging schema mappings with routes. In *PVLDB*. 79–90.
- [16] Alin Deutsch, Nadime Francis, Alastair Green, Keith Hare, Bei Li, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Wim Martens, Jan Michels, Filip Murlak, Stefan Plantikow, Petra Selmer, Oskar van Rest, Hannes Voigt, Domagoj Vrgoc, Mingxi Wu, and Fred Zemke. 2022. Graph Pattern Matching in GQL and SQL/PGQ. In *SIGMOD*. 2246–2258.
- [17] Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller, and Lucian Popa. 2005. Data Exchange: Semantics and Query Answering. *TCS* 336, 1 (2005), 89–124.
- [18] Ronald Fagin, Phokion G. Kolaitis, and Lucian Popa. 2005. Data Exchange: Getting to the Core. *TODS* 30, 1 (2005), 174–210.
- [19] Mary Fernandez, Daniela Florescu, Alon Levy, and Dan Suciu. 1997. A Query Language for a Web-Site Management System. *SIGMOD* 26, 3 (1997), 4–11.
- [20] Miguel Fiandor and Michael Hunger. [n.d.]. Offshoreleaks Data Packages. Retrieved March 1, 2024 from <https://github.com/ICIJ/offshoreleaks-data-packages>
- [21] Nadime Francis, Amélie Gheerbrant, Paolo Guagliardo, Leonid Libkin, Victor Marsault, Wim Martens, Filip Murlak, Liat Peterfreund, Alexandra Rogova, and Domagoj Vrgoc. 2023. GPC: A Pattern Calculus for Property Graphs. In *PODS*. 241–250.
- [22] Nadime Francis, Amélie Gheerbrant, Paolo Guagliardo, Leonid Libkin, Victor Marsault, Wim Martens, Filip Murlak, Liat Peterfreund, Alexandra Rogova, and Domagoj Vrgoc. 2023. A Researcher’s Digest of GQL. In *ICDT*, Vol. 255. 1–22.
- [23] Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Stefan Plantikow, Mats Rydberg, Petra Selmer, and Andrés Taylor. 2018. Cypher: An Evolving Query Language for Property Graphs. In *SIGMOD*. 1433–1445.
- [24] Nadime Francis and Leonid Libkin. 2017. Schema Mappings for Data Graphs. In *PODS’17*. 389–401.
- [25] Michael R. Garey and David S. Johnson. 1990. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.
- [26] Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Stefan Plantikow, Martin Schuster, Petra Selmer, and Hannes Voigt. 2019. Updating graph databases with Cypher. *VLDB* 12, 12 (2019), 2242–2254.
- [27] Richard Hull and Masatoshi Yoshikawa. 1990. ILOG: Declarative Creation and Manipulation of Object Identifiers. In *VLDB*. 455–468.
- [28] Phokion G. Kolaitis. 2005. Schema Mappings, Data Exchange, and Metadata Management. In *PODS*. 61–75.
- [29] Neo4j. 2023. APOC user guide for Neo4j 5. Retrieved November 9, 2023 from <https://neo4j.com/docs/apoc/current/>
- [30] Neo4j. 2023. Graph Data Modeling Fundamentals. Retrieved November 9, 2023 from <https://graphacademy.neo4j.com/courses/modeling-fundamentals/>
- [31] Philipp Skavantzios and Sebastian Link. 2023. Normalizing Property Graphs. *Proc. VLDB Endow.* 16, 11 (2023), 3031–3043.
- [32] Oskar van Rest, Sungpack Hong, Jinha Kim, Xuming Meng, and Hassan Chafi. 2016. PGQL: A Property Graph Query Language. In *GRADES*. 1–6.
- [33] Moshe Y. Vardi. 2016. A Theory of Regular Queries. In *PODS*. 1–9.
- [34] Si Zhang and Hanghang Tong. 2020. Network Alignment: Recent Advances and Future Directions. In *CIKM*. 3521–3522.

A PRELIMINARIES

A.1 Data model

We introduce thereafter further basic notation around the data model of a property graph.

A *path* $p := (n_0, e_1, n_1, \dots, e_k, n_k)$ is an alternating sequence of node and edges, which starts and ends with nodes. Given a path p we denote by $\text{src}(p)$ and $\text{tgt}(p)$ the first and last node of p ; in this case, $\text{src}(p) = n_0$ and $\text{tgt}(p) = n_k$. $\text{len}(p) \in \mathbb{N}$ is the *length* of p , i.e, the number of edges in the path; and if $\text{len}(p) = 0$ then the path consists of a single node which is both the source and the target. We can define as usual the *concatenation* $p \cdot p'$ of two paths p and p' whenever $\text{tgt}(p) = \text{src}(p')$.

A.2 Graph Pattern Calculus

We make a brief summary on the syntax and semantics of GPC [21], focusing only on the concepts we need to formally define our property graph transformation rules.

The *atomic* GPC patterns are node and edge patterns. A *node pattern* is of the form $(x : \ell)$ and an *edge pattern* is of the form $\xrightarrow{x:\ell}$. In both cases x is an optional variable (picked from a countably infinite set \mathcal{X} of variables) which bounds, if present, to the matched element and ℓ is an optional label indicating that we want to restrict to ℓ -elements. In an edge pattern \rightarrow may indicate one of the two possible directions:

forward \rightarrow and backward \leftarrow . A GPC *pattern* denoted π is inductively constructed on top of the atomic patterns by using arbitrarily many *union* ($\pi + \pi$), *concatenation* ($\pi \cdot \pi$), *conditioning* ($\pi_{(\theta)}$), and *repetition* ($\pi^{n..m}$) constructs.

A GPC *query* is of the form $\rho \pi$ with ρ a *restrictor* among the set of simple, trail, shortest, which purpose is to ensure a finite result set. simple prevents repetition of nodes along a path, trail prevents repetition of edges and shortest selects only the paths of minimal length among all the paths between two nodes.

The structure of a GPC query can be inspected using a *type system*, a set of typing rules [21]. A query is *well-typed* if the rules permit to deduce a unique type to every variable appearing in the query. When an expression Q is well-typed, the schema $\text{sch}(Q)$ of this expression associates a type to each variable.

The *answer* of a GPC query $Q(\bar{x})$ on a property graph G , denoted $\llbracket Q \rrbracket_G^{\bar{x}}$ is a set of assignments. An *assignment* binds the variables \bar{x} , present in the query, to values. *Values* to be associated to variables are dependent upon the deduced type of the variable for that query. Hence, for each type τ , there is a set of values \mathcal{V}_τ . Values may be references to elements in the graph, e.g., for Node and Edge types.

All answers to queries we need to define our property graph transformations will have assignments of the variables among the types Node and Edge. In our framework, we will use GPC queries extended with the capability to use conditioning on top of joins. This is not part of the specification in [21], but this is planned to be in GQL [22].

B PROPERTY GRAPH TRANSFORMATIONS

We provide additional notation and definitions that are used in the main proofs of this paper.

B.1 Generating output identifiers

Given a tuple of m variables $\bar{x} = (x_1, \dots, x_m)$, we define the sets of *value arguments* and *arguments* for \bar{x} as

$$\mathcal{V}_{\bar{x}} ::= c \mid x_i.a$$

$$\mathcal{A}_{\bar{x}} ::= x_i \mid c \mid \ell \mid x_i.a$$

where $x_i \in \bar{x}$, $c \in \text{Const}$, $\ell \in \mathcal{L}$ and $a \in \mathcal{K}$. We denote by $\mathcal{A}_{\bar{x}}^k$ the set of all tuples of arguments for \bar{x} of length k .

For a given property graph G , $A = (a_1, \dots, a_k)$ a tuple of arguments for \bar{x} defines a function $\mathcal{O}^m \rightarrow (\mathcal{O} \cup \text{Const} \cup \mathcal{L})^k$ defined as $(o_1, \dots, o_m) \mapsto (v_1, \dots, v_k)$ where, for all $1 \leq i \leq k$:

- $v_i := o_j$ if $a_i = x_j$;
- $v_i := c$ if $a_i = c$;
- $v_i := \ell$ if $a_i = \ell$;
- $v_i := c$ if $a_i = x_j.a$ and $\delta_G(o_j, a) = c$.

B.2 Semantics

PROPOSITION 3.6. *Given an input property graph G and a property graph transformation T , Algorithm 1 always returns a valid instance of the property graph data model.*

PROOF. Given T and G with identifiers from \mathcal{S} , let $T(G) := \langle N, E, \lambda, \text{src}, \text{tgt}, \delta \rangle$ be a property graph returned by Algorithm 1. We have to check that (i) both N and E are finite and disjoint, (ii) all elements have a finite number of labels, and (iii) every edge has exactly one source and one target.

(i) The set of bindings resulting from querying a property graph G with the query $P(\bar{x})$, $\llbracket P \rrbracket_G^{\bar{x}}$, is assumed to be finite, moreover, we have a finite number of rules in T , hence the finiteness of $N \cup E$. A similar reasoning shows the finiteness of the label set for each element in $T(G)$; this is because each rule can mention at most a finite number of labels.

(ii) We now show that $N \cap E = \emptyset$. Let us assume that $o \in \mathcal{T}$ is both a node and an edge id in $T(G)$ – respectively resulting from a node rule $R := P(\bar{x}) \implies (D)$ for \bar{o} and an edge rule $S := Q(\bar{y}) \implies (C_s) \xrightarrow{C} (C_t)$ for \bar{p} . The *injectivity* of the Skolem function f enforces that o has been generated, in both cases, by using the same arguments. Moreover, by *injectivity* again, we necessarily have $D.Id(\bar{o}) = (o_s, C.Id(\bar{p}), o_t)$ for some $o_s, o_t \in N$, with $C.Id \in \mathcal{A}_{\bar{x}}^{k-2}$. (Note that o_s and o_t have been respectively obtained from the source and target rules of S for \bar{p} .) By definition of the range of the Skolem function f , o_s and o_t belong to \mathcal{T} ; thus, they could not be equal to the first and last values of $D.Id$. We conclude that $N \cap E = \emptyset$.

(iii) Finally, by *injectivity* of f , for an $o \in \mathcal{O}$ which is an edge id in $T(G)$, there are, by definition, exactly one $o_s \in N$ and one $o_t \in N$ which correspond to the source and the target of this edge. \square

C DETECTING CONFLICTS

We formally define the notion of *conflicts* and provide the proofs for the results in Section 4.

C.1 Consistency checking

We now formalize the notion of *node conflict*. For any $k \geq 0$, let

$$R := P(\bar{x}) \implies \begin{matrix} (A : L) \\ \langle \dots, a=v, \dots \rangle \end{matrix} \quad \text{and} \quad S := Q(\bar{y}) \implies \begin{matrix} (B : M) \\ \langle \dots, a=w, \dots \rangle \end{matrix}$$

be two node rules in T with $L, M \in \mathcal{P}_{fin}(\mathcal{L})$. These two *node* rules are *potentially conflicting* on the property a when:

- their respective argument lists have same length, i.e., $A = (a_1, \dots, a_k) \in \mathcal{A}_{\bar{x}}^k$ and $B = (b_1, \dots, b_k) \in \mathcal{A}_{\bar{y}}^k$ for a $k \geq 0$;
- their argument lists are *compatible*, which means that for each $1 \leq i \leq k$, a_i is a value argument if and only if b_i is a value argument;
- if a_i and b_i are respectively x_i and y_j , then $\text{sch}(P)(x_i)$ should be equal to $\text{sch}(P)(y_j)$;
- they have *potentially conflicting properties*, which means that they both define the same property to (possibly) different values, i.e., $a \in \mathcal{K}, v \in \mathcal{V}_{\bar{x}}$ and $w \in \mathcal{V}_{\bar{y}}$.

A *node conflict* for a pair of possibly conflicting rules on the property a occurs in a property graph G whenever it exists $\bar{o} \in \llbracket P \rrbracket_G^{\bar{x}}$ and $\bar{p} \in \llbracket Q \rrbracket_G^{\bar{y}}$ with $A(\bar{o}) = B(\bar{p})$ and $v(\bar{o}) \neq w(\bar{p})$.

We now formalize the notion of *edge conflict*. For any $k \geq 0$, let

$$R := P(\bar{x}) \implies \begin{matrix} (A_s :) \\ \langle \dots, a=v, \dots \rangle \end{matrix} \left[\begin{matrix} A : L \\ \longrightarrow \end{matrix} \right] \begin{matrix} (A_t :) \\ \langle \dots, a=w, \dots \rangle \end{matrix} \quad \text{and} \quad S := Q(\bar{y}) \implies \begin{matrix} (B_s :) \\ \langle \dots, a=w, \dots \rangle \end{matrix} \left[\begin{matrix} B : M \\ \longrightarrow \end{matrix} \right] \begin{matrix} (B_t :) \\ \langle \dots, a=w, \dots \rangle \end{matrix}$$

be two edge rules in T with $L, M \in \mathcal{P}_{fin}(\mathcal{L})$. These two *edge* rules are *potentially conflicting* on the property a when:

- their respective argument lists have same length, i.e., $A = (a_1, \dots, a_k) \in \mathcal{A}_{\bar{x}}^k$ and $B = (b_1, \dots, b_k) \in \mathcal{A}_{\bar{y}}^k$ for a $k \geq 0$;
- their argument lists are *compatible*, which means that for each $1 \leq i \leq k$, a_i is a value argument if and only if b_i is a value argument;
- if a_i and b_i are respectively x_i and y_j , then $\text{sch}(P)(x_i)$ should be equal to $\text{sch}(P)(y_j)$;
- the three previous points also apply to the pairs (A_s, A_t) and (B_s, B_t) ;
- they have *potentially conflicting properties*, which means that they both define the same property to (possibly) different values, i.e., $a \in \mathcal{K}, v \in \mathcal{V}_{\bar{x}}$ and $w \in \mathcal{V}_{\bar{y}}$.

An *edge conflict* for a pair of possibly conflicting rules on the property a occurs in a property graph G whenever it exists $\bar{o} \in \llbracket P \rrbracket_G^{\bar{x}}$ and $\bar{p} \in \llbracket Q \rrbracket_G^{\bar{y}}$ with $A(\bar{o}) = B(\bar{p})$, $A_s(\bar{o}) = B_s(\bar{p})$, $A_t(\bar{o}) = B_t(\bar{p})$ and $v(\bar{o}) \neq w(\bar{p})$.

LEMMA 4.3. *The transformation consistency problem is PTIME-reducible to the satisfiability problem for GPC+.*

PROOF. Recall that a node conflict for a pair of possibly conflicting rules R and S on a property a occurs on a property graph G whenever it exists $\bar{o} \in \llbracket P \rrbracket_G^{\bar{x}}$ and $\bar{p} \in \llbracket Q \rrbracket_G^{\bar{y}}$ with $A(\bar{o}) = B(\bar{p})$ and $v(\bar{o}) \neq w(\bar{p})$. We can rewrite all of those conditions in a single boolean GPC query:

$$Q_{(R,S,a)}() := P(\bar{x}), Q(\bar{y}), A = B, v \neq w$$

which is satisfiable on a property graph H iff this specific conflict occurs on H . Note that we assume w.l.o.g. in the following construction that \bar{x} and \bar{y} are disjoint sets of variables.

Similarly, for an edge conflict, we obtain a single boolean GPC query with the same properties:

$$Q_{(R,S,a)}() := P(\bar{x}), Q(\bar{y}), A = B, A_s = B_s, A_t = B_t, v \neq w$$

We provide an example of the GPC pattern encoding $v \neq w$. Let assume that $v := x_i.b$ and $w := y_j.c$ and $\text{sch}(P)(x_i) = \text{Node}$ and $\text{sch}(Q)(y_j) = \text{Edge}$, the following join query encodes $v \neq w$:

$$\left[(x_i), () \xrightarrow{y_j} () \right].$$

$$\langle \neg(x_i.b=y_j.c) \rangle$$

Similarly, we can apply point-wise this construction and take their join to encode $A = B, A_s = B_s$ and $A_t = B_t$.

Finally, to wrap-up the proof, it is easy to see that given a property graph transformation T , there are at most polynomially many such triplets (R, S, a) satisfying this criteria, so we can take the union of all the $Q_{(R,S,a)}()$ as the final GPC+ query on which to check for satisfiability. \square

C.2 GPC satisfiability

LEMMA 4.4. *The satisfiability problem for GPC is PSPACE-hard.*

PROOF. Let $M = (Q, \Sigma, s, F, \delta)$ be the TM that recognizes L in deterministic polynomial-space. Let w be an input word of length n . Assume that M works over w using at most $c \cdot p(n)$ for a fixed constant c and polynomial p tape cells.

We build a GPC query using the following set of properties $\mathcal{K}_0 \subset \mathcal{K}$ which contains all the following elements:

- $\text{pos}_{(i,\sigma)}$; the tape contains symbol $\sigma \in \Sigma$ at position $1 \leq i \leq c \cdot p(n)$;
- $\text{head}_{(i)}$; the head of the TM is in position $1 \leq i \leq c \cdot p(n)$;
- q ; the TM is in state $q \in Q$.

Notice that we will be using only two constants values: 0 and 1.

To encode the consistency of a state represented by the set of properties of an element x , we will use formula $\theta_c(x)$ defined as the conjunction of the following formulas:

- $\neg(x.\text{pos}_{(i,a)} = 1) \vee \neg(x.\text{pos}_{(i,b)} = 1)$; for $1 \leq i \leq c \cdot p(n)$, $a, b \in \Sigma$, $a \neq b$;
- $\neg(x.\text{head}_{(i)} = 1) \vee \neg(x.\text{head}_{(j)} = 1)$; for $1 \leq i \neq j \leq c \cdot p(n)$;
- $\neg(x.q = 1) \vee \neg(x.q' = 1)$; for $q, q' \in Q$, $q \neq q'$
- $(x.k = 1) \vee (x.k = 0)$; for $k \in \mathcal{K}_0$.

$\theta_1(x)$ is a conjunction of the following formulas; it ensures that the set of properties of the element pointed to by x encodes the initial configuration of the TM M over w :

- $x.\text{pos}_{(i,a)} = 1$; if $w_i = a$, for $1 \leq i \leq n$; w is stored at the beginning of the tape;
- $x.\text{pos}_{(i,\square)} = 1$; for $n < i \leq c \cdot p(n)$, where \square is the blank symbol; the rest of the tape is filled with blank symbols;
- $(x.\text{head}_{(1)} = 1) \wedge (x.s = 1)$; initialisation of both head and state;
- $\theta_c(x)$; consistency check.

$\theta_2(u, v)$ checks that the configuration stored in the record of v can be obtained from u in a single computation step of M ; it consists in the conjunction of the following formulas:

- $(u.\text{head}_{(i)} = 0 \wedge u.\text{pos}_{(i,a)} = 1) \implies v.\text{pos}_{(i,a)} = 1$; for $1 \leq i \leq c \cdot p(n)$, $a \in \Sigma$; the tape remains unchanged unless written by head;
- $(u.\text{head}_{(i)} = 1 \wedge u.\text{pos}_{(i,a)} = 1 \wedge u.q = 1) \implies (v.\text{pos}_{(i,b)} = 1 \wedge v.\text{head}_{(i+d)} = 1 \wedge v.q' = 1)$; for $1 \leq i \leq c \cdot p(n)$, $a, b \in \Sigma$, $q, q' \in Q$, $d \in \{-1, 0, 1\}$, when $(q', b, d) \in \delta(q, a)$; there is a valid transition;
- $\theta_c(u) \wedge \theta_c(v)$; consistency checks.

Finally, θ_3 checks that we have reached an accepting configuration:

$$\theta_3(y) := \bigvee_{q \in F} y.q = 1$$

It is clear that the pattern P – which is constructed in polynomial time given an entry word w – is satisfiable if and only if there exists an accepting run for M over w using at most a polynomial amount of space, i.e., iff $w \in L \in \text{PSPACE}$. \square

THEOREM 4.5. *The satisfiability problem for GPC+ queries using only the simple and trail restrictors is PSPACE-complete.*

PROOF. By Lemma 4.4, it is only left to prove the upper-bound. Let Q be a GPC query. We assume a mild syntactic restriction that, all edge and node patterns must mention a variable in their descriptor. This can be achieved by picking a fresh variable when none is specified in a descriptor.

We note respectively Const_0 , \mathcal{K}_0 and \mathcal{L}_0 , the set of constants, keys and labels mentioned in Q . Additionally, let d be the number of occurrences of conditions of the form $x.a = c$ or $x.a = y.b$ in the formula. We extend the set Const_0 with $d + 1$ fresh distinct constants. Let also \mathcal{X} and \mathcal{Y} be respectively the sets of variables of type Node and Edge in the schema of Q .

Preliminary remarks. We start with some key observations:

- It is clear that we cannot always guess an answer made up of a path and an assignment because some patterns are only satisfiable by paths of exponential length: e.g., we can simulate a counter to count up to 2^n with n properties and a polynomial-sized formula similar to that used in Lemma 4.4;
- The concatenation of patterns implies an implicit equality over the endpoints; thus, we need to book-keep the endpoints of a pattern alongside an assignment for it;
- The semantics and the typing rules of GPC isolate the variables under a repetition pattern; this means that, in a repetition pattern, we cannot refer to externally defined variables (i.e., non-local occurrences are not permitted); moreover, because conditions are only defined over *singleton* variables (i.e., variables of type Node or of type Edge in the schema of Q), we cannot refer in conditions to variables appearing under a repetition sub-pattern. Thus, we can assume that, if a query is satisfiable, an answer path for a pattern inside a repetition pattern can be disjoint (except for its endpoints) from the answer paths for the outside of the repetition pattern.

This shows that we can store in polynomial space an *extended* assignment corresponding to a valid answer to a path pattern π – by extended, we mean that we additionally store the two endpoints of the answer path in variables named src_π and tgt_π (that we assume to belong to \mathcal{X}); additionally, Group and Maybe variables will not be tracked because they can no longer be mentioned in conditions. In this case, an *assignment* $\eta(\cdot)$ stores for each variable the description of an element which is made of its label set (referred to by $\lambda(\eta(\cdot))$) and its record (referred to by $\delta(\eta(\cdot), \cdot)$); it does not contain an identifier for this element.

Saturation procedure and consistency check. In the body of the main algorithm we make use of a procedure called *saturation*. The main idea of this procedure is to propagate equality and inequality constraints between variables whenever new ones are found. We illustrate why this procedure is needed and how it works on the following pattern:

$$(u)_{\langle u.a=1 \rangle} \xrightarrow{z} (v) \xrightarrow{z} (w)_{\langle w.a=2 \rangle}$$

Because both occurrences of z must map to the same edge, the constraints over the endpoints of z enforce equalities between the variables u , v and w . Successively, these equalities imply that $u.a = w.a$, which is in conflict with the requirements of the pattern. Note that if we remove the conditions $u.a = 1$ and $w.a = 2$, this pattern remains unsatisfiable w.r.t. the simple semantics; the forced repetition of nodes in bindings of this pattern was not explicit before applying the saturation procedure because no variable was reused. However, reusing the variable z makes the pattern explicitly unsatisfiable under trail semantics.

To formalize this, we introduce the notion of an *equality graph* G for a query Q (or a pattern π), which is a 2-layer undirected edge-labeled graph with nodes in each layer that respectively belong to the sets \mathcal{Y} and \mathcal{X} . There are two edge labels, $=$ and \neq . Edges can only connect nodes from the same layer. If G is an equality graph for a pattern π , there are two distinguished nodes $\text{src}_\pi, \text{tgt}_\pi \in \mathcal{X}$ which are respectively called the source and the target, and abbreviated src, tgt if clear from the context. This structure supports the following set of operations:

Saturation: The first step of the *saturation procedure over a pattern π* is to equate the endpoints of edge variables. Let e be an edge variable in G , if G contains say x and y , two distinct node variables, which are both at the source or the target of an occurrence of e in π , then add $x = y$ in G .

In a second step, it performs the transitive closure on the $=$ -edges of the graph at layer \mathcal{X} .

After obtaining a $=$ -transitive graph, we pass on the inequalities: if there is an \neq -edge between say x and y , and if both $x' = x$ and $y' = y$ in G , then we add $x' \neq y'$ in G .

Finally, it does a backward step to pull-up the inequalities: if there is an \neq -edge between two elements in \mathcal{X} and if both are either the source or the target of some edges in Q , then it adds an inequality edge between these two edge variables.

Check consistency: Given an assignment for all the variables mentioned in the equality graph, check if all equalities are satisfied in this assignment: if x and y are two variables with $x = y$ in G , check if their assignments are strictly the same. Moreover, check if there is no conflict in the graph: a *conflict* is when there are both an $=$ -edge and an \neq -edge between the same pair of nodes of G ; or if there is a \neq -loop.

Merge: Given two equality graphs G_1 and G_2 , *merging both on nodes x and y for policy ρ* consists in:

- (1) taking the union of their vertices and edges;
- (2) adding an $=$ -edge between x and y ; if x and y are given as parameters;
- (3) if $\rho = \text{simple}$, adding an \neq edge from each node in the layer \mathcal{X} of G_1 not $=$ -connected to tgt , to each node of the layer \mathcal{X} of G_2 not $=$ -connected to src ; (This is consistent with the semantics of the concatenation of $\rho_1 \cdot \rho_2$, where the target of ρ_1 must be equal to the source of ρ_2);
- (4) if $\rho = \text{trail}$, adding an \neq edge from each node of the layer \mathcal{Y} of G_1 , to each node of the layer \mathcal{Y} of G_2 ;
- (5) applying the saturation procedure.

Note that the result of this operation is also a valid equality graph and that x, y and ρ are optional parameters.

The maximum number of nodes and edges in an equality graph for Q is polynomial in Q ; and the three procedures can be implemented in polynomial time.

Inductive procedure for patterns. In the following, we describe a non-deterministic polynomial space procedure to check if a GPC pattern query $\rho\pi$ is satisfiable. This inductive procedure over the structure of π succeeds if and only if π is satisfiable under policy ρ . It returns a pair consisting of an extended assignment and an equality graph over the variables in the assignment:

- Case $\pi := (x : \ell)$. Guess a node element with a label set $L \subseteq \mathcal{L}_0$ s.t. $\ell \in L$ and a record consisting in a partial assignment from the keys in \mathcal{K}_0 to Const_0 . Return the pair consisting of the extended assignment binding src_π , tgt_π and x to this node; and of the equality graph containing three nodes: src_π , tgt_π and x , and two edges: $\text{src}_\pi = x$ and $\text{tgt}_\pi = x$.
- Case $\pi := \xrightarrow{y:\ell}$. Guess an edge element with a label set $L \subseteq \mathcal{L}_0$ s.t. $\ell \in L$ and a record consisting in a partial assignment from the keys in \mathcal{K}_0 to Const_0 . Return one of the following two possibilities:
 - Return a pair consisting of the extended assignment binding y to this edge, and src_π and tgt_π to two arbitrarily guessed endpoint nodes (which act as the source and target of y); the equality graph contains these three elements and an \neq -edge between src_π and tgt_π .
 - (Loop; if ρ is not simple) Return a pair consisting of the extended assignment binding this edge to y and src_π and tgt_π to the same arbitrarily guessed endpoint node; the equality graph contains these three elements and an $=$ -edge between src_π and tgt_π .
- Case $\pi := \pi_1 + \pi_2$. Guess $i \in \{1, 2\}$ and return the extended assignment and equality graph obtained by a recursive call to this procedure on π_i , after removing the variables that do not appear in π_{3-i} . (This is because those variables are of type $\text{Maybe}(\cdot)$ in π .)
- Case $\pi := \pi_1\pi_2$. Perform a recursive call to this procedure on both π_1 and π_2 to obtain an extended assignment and an equality graph for π_i , $i \in \{1, 2\}$. Check whether they unify (i.e., check whether the assignments of π_1 and π_2 are strictly the same on their common variables). Then, merge the two equality graphs on tgt_{π_1} and src_{π_2} under the policy ρ ; and check consistency w.r.t. the unified extended assignment. Return the pair consisting of the unified extended assignment and the merged equality graph after removing tgt_{π_1} and src_{π_2} and renaming src_{π_1} to src_π and tgt_{π_2} to tgt_π .
- Case $\pi := \pi_1(\theta)$. Perform a recursive call to this procedure on π_1 to obtain an extended assignment and an equality graph for π_1 . Check the validity of θ (as defined in the *Semantics of conditioned patterns* in [21]) over the extended assignment of π_1 and return the same extended assignment and equality graph.
- Case $\pi := \pi_1^{n..m}$. Guess a length k between n and m . (Note that k can be assumed to be at most exponential in the size of the whole query if m is ∞ by Lemma C.1; hence, it can be written in binary using a polynomial amount of space.) Perform k successive recursive calls to this procedure on π_1 . Each time drop from the obtained extended assignment and equality set all but the src and tgt variables and the equality or inequality edge between them. Check if they concatenate with the previous block obtained so far (i.e., perform the case $\pi := \pi_1\pi_2$). Return a pair consisting of an extended assignment binding the src of the very first guess to src_π and the tgt of the very last guess to tgt_π ; and of the equality graph containing only the src_π and tgt_π nodes, possibly with an $=$ or an \neq edge between them if should be.

Example. We illustrate how our procedure works for repetition patterns with the following example:

$$\pi := \text{simple} \left(\left[(u) \rightarrow (v) \right]_{\langle \neg(u.a=v.a) \rangle} \right)^{0..∞}$$

We note $\pi_1 := \left[(u) \rightarrow (v) \right]_{\langle \neg(u.a=v.a) \rangle}$. There are three different types of behavior depending on k :

- If $k = 0$, it returns a pair consisting of an extended assignment binding src_π and tgt_π to an arbitrarily guessed node element; and of an equality graph containing the edge $\text{src}_\pi = \text{tgt}_\pi$.
- If $k = 1$, it returns a pair consisting of an extended assignment binding src_π and tgt_π to two node elements having a different value for a (if both set); and of the quality graph containing $\text{src}_\pi \neq \text{tgt}_\pi$.
- If $k \geq 2$, it returns a pair consisting of an extended assignment binding src_π and tgt_π to two arbitrary node elements; and of the quality graph containing $\text{src}_\pi \neq \text{tgt}_\pi$.

Invariant. Let π be a pattern matched under a restrictor $\rho \in \{\text{simple}, \text{trail}\}$. The pair (η, G_π) is an output of the procedure consisting of an extended assignment and an equality graph for π iff there exists a property graph P such that $(\rho, \mu) \in \llbracket \pi \rrbracket_{P^2}$ and if, for all x s.t. $\text{sch}(\pi)(x) \in \{\text{Node}, \text{Edge}\}$, we have:

- $\lambda(\mu(x)) = \lambda(\eta(x))$;
- for all key (property) a , $\delta(\mu(x), a) = \delta(\eta(x), a)$, or both are not defined;
- if $x = y$ in G then $\mu(x) = \mu(y)$; similarly, if $x \neq y$ in G then $\mu(x) \neq \mu(y)$.

In the following, we provide the key ideas for proving this invariant by induction:

- Case $\pi := (x : \ell)$ and case $\pi := \xrightarrow{y:\ell}$ are trivial as P can be obtained directly from η .
- Case $\pi := \pi_1 + \pi_2$ by applying the hypothesis on either side. Note that π may contain fewer singleton variables than π_i . (This is because the variables of π_i that do not appear in π_{3-i} are of type $\text{Maybe}(\cdot)$ in π .)

²The set of answers to π on P [21].

- Case $\pi := \pi_1\pi_2$ relies on the merge procedure to propagate the *forced* equalities and inequalities.
- Case $\pi := \pi_1\langle\theta\rangle$ by noticing that we only need to check whether the condition holds over the extended assignment because conditions only apply on singleton variables. (This is enforced by the *Typing rules for the GPC type system* presented in Figure 2 of [21].) Notice that $\neg(u.a = v.a)$ does not lead to $u \neq v$ because either $u.a$ or $v.a$ may be undefined; thus, we don't need to update G . (Again, this is because of the *Semantics of conditioned patterns* of [21] where $\mu \models (x.a = x.b)$ iff $\delta(\mu(x), a)$ and $\delta(\mu(y), b)$ are defined and equal.)
- Case $\pi := \pi_1^{n..m}$ because π does not contain any singleton variables. Hence, only a potential equality or inequality between its endpoints is tracked throughout the k iteration steps over π_1 .

Extension to queries (with join and conditioning). Let $Q := Q_1, Q_2$ be a join query with Q_1 and Q_2 matched under restrictors simple or trail; and let the pairs $(\eta_i, G_{Q_i}), i \in \{1, 2\}$ be returned by the previous procedure on each Q_i . The procedure (for Q) checks if the η_i 's unify (i.e., check whether the η_i 's agree on their common variables), and if the result of merging the G_i 's remains consistent.

Note that this procedure supports *conditioning over join queries* (i.e., we add the following query expression $Q := Q\langle\theta\rangle$) by checking if the condition is valid over η , similar to what is done for conditioning in patterns.

Extension to GPC+. Simply guess a GPC query among all the disjuncts, and check its satisfiability using the previous procedure.

Shortest restrictor. Note that the invariant is not valid if the shortest restrictor is used. For instance, consider the following pattern, which is not supported by the above procedure:

$$\begin{aligned} & \text{shortest simple } (x) \left[\begin{array}{c} () \rightarrow (u) \rightarrow () + () \rightarrow () \rightarrow (u) \rightarrow () \rightarrow () \\ \langle a=1 \rangle \end{array} \right] (y), \\ & \text{simple } (x) \rightarrow (w) \rightarrow (y), \\ & \qquad \qquad \qquad \langle a=1 \rangle \\ & \text{simple } (u) \\ & \qquad \qquad \qquad \langle a=2 \rangle \end{aligned}$$

Nevertheless, we can easily prove that the above procedure works when all patterns in a GPC query Q that are matched under the shortest restrictor have only their endpoints for singleton variables. \square

LEMMA C.1. *Let $\pi := \pi_1^\infty$ be a sub-pattern of a GPC query $\rho\pi_0$. If there is an answer path for π , then there is another answer path for π consisting of at most k repetitions of π_1 , with k exponential in the size of π_0 , and with all answer paths for π_1 being inner disjoint.*

PROOF. There are at most exponentially many distinct records for nodes with values in Const_0 , keys in \mathcal{K}_0 and labels in \mathcal{L}_0 . Given the *Semantics of repeated patterns* in [21], only the target node of an iteration has an impact over the next iteration by being its source; only this information is transferred across successive repetitions of π_1 . Thus, we can reduce the number of repetitions of π_1 in the initial answer path because one target node necessarily gets repeated. We can moreover assume w.l.o.g. that all answer paths to π are inner disjoint, by taking disjoint copies of the initial answer paths. \square

LEMMA 4.6. *The satisfiability problem is NP-hard even for single-node GPC patterns.*

PROOF. We reduce 3-SAT to the satisfiability of a GPC query. Let F be the following 3-SAT formula over n variables and m clauses:

$$\bigwedge_{1 \leq i \leq m} (l_{i1} \vee l_{i2} \vee l_{i3})$$

where for all $1 \leq i \leq m$ and $j \in \{1, 2, 3\}$, l_{ij} is a *literal* which is either x_k or \bar{x}_k for a $k \in \{1, \dots, n\}$.

We construct the following GPC query $Q(x) := \rho \pi\langle\theta\rangle$ with $\pi := (x : \ell)$ and

$$\begin{aligned} \theta := & \bigwedge_{1 \leq i \leq n} (x_i = 1 \wedge \bar{x}_i = 0) \vee (x_i = 0 \wedge \bar{x}_i = 1) \\ & \wedge \bigwedge_{1 \leq j \leq m} \bigvee_{(b_1, b_2, b_3) \in C^3} (l_{j1} = b_1 \wedge l_{j2} = b_2 \wedge l_{j3} = b_3) \end{aligned}$$

where $C^3 = \{(0, 0, 1), \dots, (1, 1, 1)\}$; x and ℓ are optional in π ; the literals of F are used as property names in Q ; and size of Q is clearly polynomial in the size of F .

We now show that F is satisfiable if and only if there exists a property graph G on which $Q(x)$ returns at least a node.

(\Rightarrow) Assume that F is satisfied by an assignment ν to \bar{x} . We construct a property graph containing a unique ℓ -labeled node with identifier o , having the following record:

$$\forall i \in \{1, \dots, n\} \left\{ \begin{array}{ll} x_i \mapsto 1, & \text{if } \nu_i = \top \\ \bar{x}_i \mapsto 0, & \\ x_i \mapsto 0, & \text{if } \nu_i = \perp \\ \bar{x}_i \mapsto 1 & \end{array} \right\}.$$

By design, the top-most conjunct of θ is satisfied. Let C_i for any $1 \leq i \leq m$ be a clause in F ; by hypothesis c_i is satisfied by v , so there exists $(b_1, b_2, b_3) \in C^3$ such that $\delta(o, l_{ij}) = b_j$ for all $1 \leq j \leq 3$.

(\Leftarrow) Conversely, if $Q(x)$ is satisfied in a property graph G for an element o ; we have that $o \in N$ and each x_i, \bar{x}_i are defined in the record of o . The restrictions enforced by the top-most conjunct of θ ensure that we can retrieve a well-defined assignment for F ; the last conjunct ensures that this is a valid assignment for F . \square

D EXPERIMENTS

User study. The User study consists of four parts, respectively aiming to:

- Evaluate the Understandability of openCypher scripts; we asked four yes/no questions asking to the participants whether an assertion is true or not w.r.t. the behavior of a provided script in a concrete transformation scenario (e.g. Does this script create as many Director nodes as there are Person nodes that have an outgoing relationship of type DIRECTED to a Movie node?).
- Evaluate the Understandability of Property Graph Transformations; with four similar yes/no questions on a slightly different transformation to avoid biases;
- The third part required participants to modify some provided openCypher scripts and transformations to adapt to a new requirement. We collected the participants' answers and checked them.
- The last part required the participants to give their opinion on the following questions and to indicate, in a range for 1 to 5 (3 is neutral) whether they found openCypher scripts and/or transformation rules better on:
 - Which one of the two methods do you find easier to understand?
 - Which one of the two methods do you find more intuitive? (Better for describing the desired output.)
 - Which one of the two methods do you find more flexible? (Easier to adapt to a new specification.)

We collected the answers of 12 participants, that were asked to self report their level of expertise in a range from (1 - Novice) to (5 - Expert) on the following topics:

- How would you rate your knowledge about databases?
The answers filled in by the participants ranged from 3 to 5, included.
- How would you rate your knowledge about openCypher?
The answers filled in by the participants ranged from 2 to 4, included.
- How would you rate your knowledge about the MERGE clause of openCypher?
The answers filled in by the participants ranged from 1 to 5, included.
- How would you rate your knowledge about property graph transformations?
The answers filled in by the participants ranged from 1 to 3, included.

We have a pool of people that all have prior exposure to openCypher (2-4) but a great diversity w.r.t. the knowledge of the MERGE clause of openCypher (the basic tool for updates in openCypher), i.e. from novice to expert.

The results on the first two parts are as follow:

- The average number of correct answers is 50% (2.0 out of 4) for the understandability of the openCypher scripts, and 90% (3.6 out of 4) for the understandability of the transformation rules.
- 25.0%, resp. 67% of participants checked all the correct answers in the first, resp. second part.
- All participants scored higher on their individual understanding of transformation rules compared to openCypher scripts.

Given those results, it is extremely clear that it is very difficult for people to understand even the basic openCypher scripts used to transform property graphs, whereas our framework – despite being absolutely new to the participants, has been widely understood.

The results on the last part are as follow (recall that 3 is neutral, 1 is strong preference for openCypher scripts and 5 is strong preference for transformation rules):

- Which one of the two methods do you find easier to understand?
Collected answers range from 1 to 5 with an average of 3.3.
- Which one of the two methods do you find more intuitive? (Better for describing the desired output.)
Collected answers range from 3 to 5 with an average of 3.8.
- Which one of the two methods do you find more flexible? (Easier to adapt to a new specification.)
Collected answers range from 3 to 5 with an average of 4.1.

Moreover, we have noticed that the participants having a low understanding of openCypher scripts (scored 2 or less out of 4 in the first part) have been more inclined to provide less credit to the transformation rules than other people. So we decided to split the participants in two groups to investigate this more.

The results on the last part are as follow only for the 8 people that have score 2 or lower out of 4 in their understanding of openCypher scripts (recall that 3 is neutral, 1 is strong preference for openCypher scripts and 5 is strong preference for transformation rules):

- Which one of the two methods do you find easier to understand?
Collected answers range from 1 to 5 with an average of 2.75.

- Which one of the two methods do you find more intuitive? (Better for describing the desired output.)
Collected answers range from 3 to 5 with an average of 3.9.
- Which one of the two methods do you find more flexible? (Easier to adapt to a new specification.)
Collected answers range from 3 to 5 with an average of 3.9.

The results on the last part are as follow only for the 4 people that have score 3 or higher out of 4 in their understanding of openCypher scripts (recall that 3 is neutral, 1 is strong preference for openCypher scripts and 5 is strong preference for transformation rules):

- Which one of the two methods do you find easier to understand?
Collected answers range from 1 to 5 with an average of 4.7.
- Which one of the two methods do you find more intuitive? (Better for describing the desired output.)
Collected answers range from 3 to 5 with an average of 3.7.
- Which one of the two methods do you find more flexible? (Easier to adapt to a new specification.)
Collected answers range from 3 to 5 with an average of 4.7.

It is therefore clear that those who have been less convinced that openCypher scripts are more error-prone and harder to interpret and analyze have not figured out for themselves that openCypher scripts are difficult to understand and manipulate. Moreover people with a good understanding of openCypher are clearly in favor that our framework is easier to understand.

With this study, we empirically and experimentally demonstrated that the script-based approach can be error prone, hard to interpret and analyze (i.e., less usable) and that the improvement of usability and accuracy over handcrafted, script-based solutions have clearly been attested by a majority of the participants.

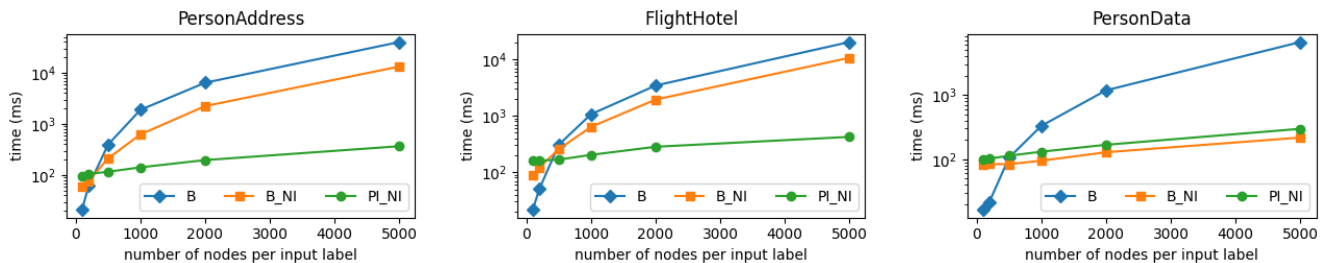


Figure 14: Run-time comparison with the baseline approach, depending on the number of input nodes of each type in G .

Comparison with native Cypher approach. Finally, we compared our framework (using PI_NI) with ad-hoc transformation scripts (B-NI, B; respectively with and without node indexes), such as the one presented in Figure 1 (ii). The result over PersonAddress, FlightHotel and PersonData are presented in Figure 14. For larger scenarios, such as GUSToBIOSQL, DBLPToAmalgam1 and Amalgam1ToAmalgam3, handcrafted transformation scripts are exceedingly large due to the number of rules and properties involved.

We can observe that our solution clearly outperforms the handcrafted solutions in most of the cases. The only exception occurs when using the PersonData scenario, for which the B-NI baseline is slightly better than our solution, while the B baseline is still outperformed. The underlying reason is due to the nature of this scenario, for which the collect clause contains only one element.

D.1 Offshore Leaks Dataset




Figure 15: Horizontal scaling, with varying number of independent copies of the scenario.

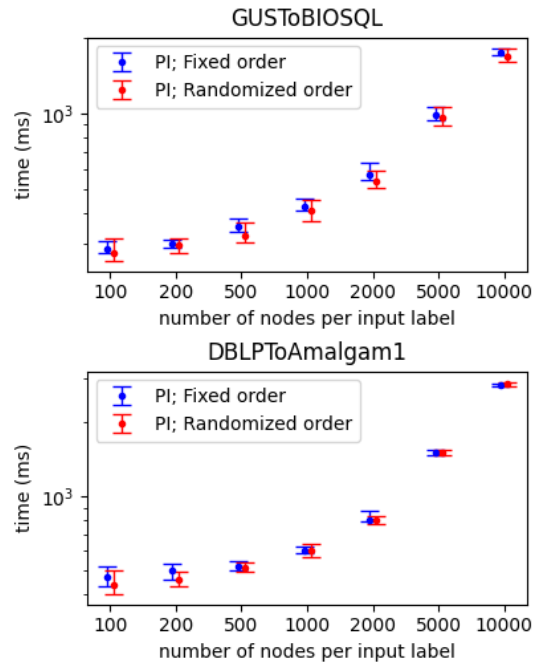


Figure 16: Average computation time for different orders of execution of the rules; error bars indicate minimum and maximum computation times observed over 20 independent runs.




Figure 17: Impact of indexing strategies and implementation variants on the computation of $T(G)$.

$$\begin{aligned}
& \langle (a : \text{Address}) \mid \langle \text{sourceID}=\text{"Malta registry"}, a.\text{address}=a.\text{address} \rangle \rangle \implies (x = (a) : \text{T_Address}) \xrightarrow{\text{:T_LOCATED}} (y = (a.\text{country_code}) : \text{T_Country}) \quad (R_1) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=a.\text{sourceID} \rangle \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{name}=a.\text{country} \rangle \\
& \langle (a : \text{Address}) \mid \langle \text{sourceID}=\text{"Malta registry"}, \neg(a.\text{address}=a.\text{address}) \rangle \rangle \implies (x = (a) : \text{T_Address}) \quad (R_2) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=a.\text{sourceID} \rangle \\
& \langle (a : \text{Address}) \mid \langle \neg(\text{sourceID}=\text{"Malta registry"}), a.\text{address}=a.\text{address} \rangle \rangle \implies (x = (a) : \text{T_Address}) \xrightarrow{\text{:T_LOCATED}} (y = (a.\text{country_codes}) : \text{T_Country}) \quad (R_3) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=a.\text{sourceID} \rangle \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{name}=a.\text{country} \rangle \\
& \langle (a : \text{Address}) \mid \langle \neg(\text{sourceID}=\text{"Malta registry"}), \neg(a.\text{address}=a.\text{address}) \rangle \rangle \implies (x = (a) : \text{T_Address}) \quad (R_4) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=a.\text{sourceID} \rangle
\end{aligned}$$

(i) Refactoring registered addresses (R1 – R4).

$$(i : \text{Intermediary}) \xrightarrow{\text{:registered_address}} (a : \text{Address}) \implies (x = (i) : \text{T_Intermediary}) \xrightarrow{\text{:T_REG_ADDRESS}} (y = (a) : \text{T_Address}) \quad (R_5)$$

$$\begin{aligned}
& \langle (i : \text{Intermediary}) \mid \langle \neg(\text{address}=\text{""}), \text{address}=\text{address}, \text{country_code}=\text{country_code} \rangle \rangle \implies (x = (i) : \text{T_Intermediary}) \xrightarrow{\text{:T_REG_ADDRESS}} (y = (i.\text{address}) : \text{T_Address}), \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=i.\text{sourceID} \rangle \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad (y) \xrightarrow{\text{:T_LOCATED}} (z = (i.\text{country_codes}) : \text{T_Country}) \quad (R_6) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{name}=i.\text{countries} \rangle
\end{aligned}$$

$$\langle (i : \text{Intermediary}) \mid \langle \neg(\text{address}=\text{""}), \text{address}=\text{address}, \neg(\text{country_code}=\text{country_code}) \rangle \rangle \implies (x = (i) : \text{T_Intermediary}) \xrightarrow{\text{:T_REG_ADDRESS}} (y = (i.\text{address}) : \text{T_Address}) \quad (R_7)$$

$\langle \text{source}=i.\text{sourceID} \rangle$

$$\langle (i : \text{Intermediary}) \mid \langle \neg(\text{address}=\text{""}), \text{address}=\text{address}, \neg(\text{country_code}=\text{country_code}) \rangle \rangle \xrightarrow{\text{:intermediary_of}} (e : \text{Entity}) \implies (x = (i) : \text{T_Intermediary}) \xrightarrow{\text{:T_REG_ADDRESS}} (y = (e.\text{address}) : \text{T_Address}) \quad (R_8)$$

$\langle \text{source}=i.\text{sourceID} \rangle$

$$\begin{aligned}
& \langle (i : \text{Intermediary}) \mid \langle \neg(\text{address}=\text{""}), \text{address}=\text{address}, \text{country_code}=\text{country_code} \rangle \rangle \xrightarrow{\text{:intermediary_of}} (e : \text{Entity}) \implies (x = (i) : \text{T_Intermediary}) \xrightarrow{\text{:T_REG_ADDRESS}} (y = (e.\text{address}) : \text{T_Address}), \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{source}=i.\text{sourceID} \rangle \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad (y) \xrightarrow{\text{:T_LOCATED}} (z = (e.\text{country_codes}) : \text{T_Country}) \quad (R_9) \\
& \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \qquad \langle \text{name}=i.\text{countries} \rangle
\end{aligned}$$

(ii) Uniformizing address information for intermediaries (R5 – R9).

$$(i : \text{Intermediary}) \implies (x = (i) : \text{T_Intermediary}) \quad (R_{10})$$

$\langle \text{name}=i.\text{name}, \text{status}=i.\text{status}, \text{valid_until}=i.\text{valid_until}, \text{source}=i.\text{sourceID} \rangle$

$$(a : \text{Address}) \implies (x = (a) : \text{T_Address}) \quad (R_{11})$$

$\langle \text{address}=a.\text{address}, \text{orig_addr}=a.\text{orig_addr}, \text{valid_until}=a.\text{valid_until} \rangle$

$$(e : \text{Entity}) \implies (x = (e) : \text{T_Entity}) \quad (R_{12})$$

$\langle \text{name}=e.\text{name}, \text{orig_name}=e.\text{orig_name}, \text{inact_date}=e.\text{inact_date}, \text{inc_date}=e.\text{inc_date}, \dots, \text{source}=e.\text{sourceID} \rangle$

$$(o : \text{Officer}) \implies (x = (o) : \text{T_Officer}) \quad (R_{13})$$

$\langle \text{name}=o.\text{name}, \text{status}=o.\text{status}, \text{source}=o.\text{sourceID} \rangle$

(iii) Exporting the nodes (R10 – R13).

Figure 18: Rules of the ICIJ database transformation.

$$(o : \text{Officer}) \xrightarrow{\text{:similar}} (p : \text{Officer}) \implies (x = (o) : \text{T_Officer}) \xrightarrow{\langle \text{link}=\text{"similar name and address as"} \rangle} (y = (p) : \text{T_Officer}) \quad (R_{14})$$

$$(a : \text{Address}) \xrightarrow{\text{:same_as}} (b : \text{Address}) \implies (x = (a) : \text{T_Address}) \xrightarrow{\text{:T_SAME_AS}} (y = (b) : \text{T_Address}) \quad (R_{16})$$

$$(o : \text{Officer}) \xrightarrow{\text{:registered_address}} (a : \text{Address}) \implies (x = (o) : \text{T_Officer}) \xrightarrow{\text{:T_REGISTERED_ADDRESS}} (y = (a) : \text{T_Address}) \quad (R_{17})$$

(i) Improving similarity detection (R14 – R17).

$$\begin{aligned} (e : \text{Entity}) \xrightarrow{\langle \text{jurisdiction_desc}=\text{jurisdiction_desc} \rangle} (x = (e) : \text{T_Entity}) &\xrightarrow{\text{:T_IN_JURIS}} (y = (e.\text{jurisdiction_desc}) : \text{T_Jurisdiction}), \\ &\langle \text{juris}=\text{e.jurisdiction_desc} \rangle \\ (y) \xrightarrow{\text{:T_RELATED}} (z = (e.\text{jurisdiction}) : \text{T_Country}) &\quad (R_{18}) \end{aligned}$$

(ii) Refactoring jurisdictions (R18).

Figure 19: Rules of the ICIJ database transformation. (Bis)