

Formalising openCypher Graph Queries in Relational Algebra

József Marton¹, Gábor Szárnyas^{2,3}, and Dániel Varró^{2,3}

¹ Budapest University of Technology and Economics, Database Laboratory
`marton@db.bme.hu`

² Budapest University of Technology and Economics
Fault Tolerant Systems Research Group
MTA-BME Lendület Research Group on Cyber-Physical Systems
`{szarnyas, varro}@mit.bme.hu`

³ McGill University
Department of Electrical and Computer Engineering

The final publication is available at Springer via
https://doi.org/10.1007/978-3-319-66917-5_13.

Abstract. Graph database systems are increasingly adapted for storing and processing heterogeneous network-like datasets. However, due to the novelty of such systems, no standard data model or query language has yet emerged. Consequently, migrating datasets or applications even between related technologies often requires a large amount of manual work or ad-hoc solutions, thus subjecting the users to the possibility of vendor lock-in. To avoid this threat, vendors are working on supporting existing standard languages (e.g. SQL) or standardising languages.

In this paper, we present a formal specification for openCypher, a high-level declarative graph query language with an ongoing standardisation effort. We introduce relational graph algebra, which extends relational operators by adapting graph-specific operators and define a mapping from core openCypher constructs to this algebra. We propose an algorithm that allows systematic compilation of openCypher queries.

1 Introduction

Context. Graphs are a well-known formalism, widely used for describing and analysing systems. Graphs provide an intuitive formalism for modelling many real-world scenarios, as the human mind tends to interpret the world in terms of objects (*vertices*) and their respective relationships to one another (*edges*) [15].

The *property graph* data model [17] extends graphs by adding labels/types and properties for vertices and edges. This gives a rich set of features for users to model their specific domain in a natural way. Graph databases are able to store property graphs and query their contents with complex graph patterns, which otherwise would be cumbersome to define and/or inefficient to evaluate on traditional relational databases [21].

Neo4j⁴, a popular NoSQL property graph database, offers the Cypher query language to specify graph patterns. Cypher is a high-level declarative query language which allows the query engine to use sophisticated optimisation techniques. Neo Technology, the company behind Neo4j initiated the openCypher project [13], which aims to deliver an open specification of Cypher.

Problem and objectives. The openCypher project provides a formal specification of the *grammar* of the query language and a set of acceptance tests that define the semantics of various language constructs. This allows other parties to develop their own openCypher-compatible query engine. However, there is no mathematical formalisation for the language. In ambiguous cases, developers are advised to consult Neo4j’s Cypher documentation or to experiment with Neo4j’s Cypher query engine and follow its behaviour. Our goal is to provide a formal specification for the core features of openCypher.

Contributions. In this paper, we use a formal definition of the property graph data model [7] and relational graph algebra, which operates on multisets (bags) [6] and is extended with additional graph-specific operators. Using these foundations, we construct a concise formal specification for the core features in the openCypher grammar. This specification is detailed enough to serve as a basis for an openCypher compiler [23].

2 Data Model and Running Example

Data model. A *property graph* is defined as $G = (V, E, st, L, T, \mathcal{L}, \mathcal{T}, P_v, P_e)$, where V is a set of vertices, E is a set of edges and $st : E \rightarrow V \times V$ assigns the source and target vertices to edges. Vertices are labelled and edges are typed:

- L is a set of vertex labels, $\mathcal{L} : V \rightarrow 2^L$ assigns a *set of labels* to each vertex.
- T is a set of edge types, $\mathcal{T} : E \rightarrow T$ assigns a *single type* to each edge.

To define properties, let $D = \cup_i D_i$ be the union of atomic domains D_i and let ε represent the NULL value.

- P_v is a set of vertex properties. A vertex property $p_i \in P_v$ is a partial function $p_i : V \rightarrow D_i \cup \{\varepsilon\}$, which assigns a property value from a domain $D_i \in D$ to a vertex $v \in V$, if v has property p_i , otherwise $p_i(v)$ returns ε .
- P_e is a set of edge properties. An edge property $p_j \in P_e$ is a partial function $p_j : E \rightarrow D_j \cup \{\varepsilon\}$, which assigns a property value from a domain $D_j \in D$ to an edge $e \in E$, if e has property p_j , otherwise $p_j(e)$ returns ε .

In the context of this paper, we define a *relation* as a *bag (multiset)* of tuples: a tuple can occur more than once in a relation [6]. Given a property graph G , relation r is a *graph relation* if the following holds:

$$\forall A \in \text{sch}(r) : \text{dom}(A) \subseteq V \cup E \cup D,$$

⁴ <https://neo4j.com/>

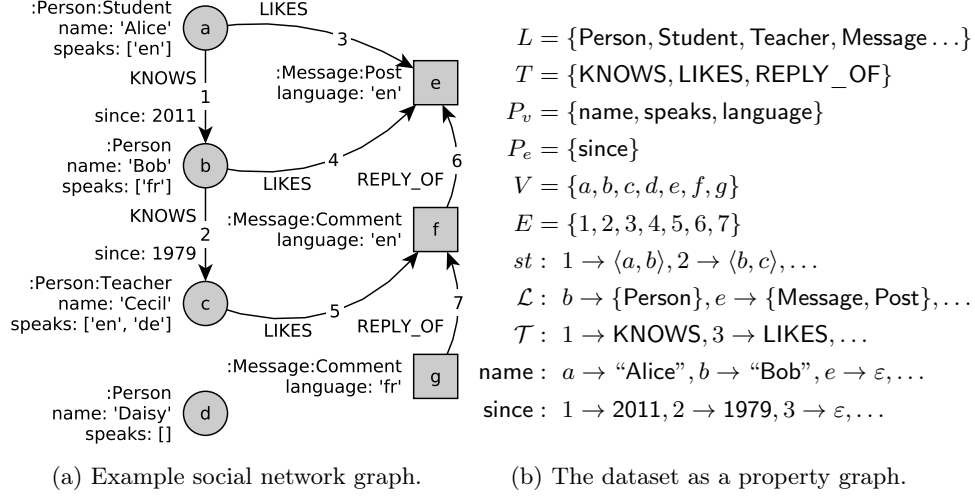


Fig. 1: Social network example represented graphically and formally. To improve readability, we use letters for vertex identifiers and numbers for edge identifiers.

where the schema of r , $\text{sch}(r)$, is a list containing the attribute names, $\text{dom}(A)$ is the domain of attribute A , V is the vertices of G , and E is the edges of G .

Property access. When defining relational algebra expression on graph relations, it is often required (e.g. in projection and selection operators) to access a certain property of a vertex/edge. Following the notation of [7], if x is an attribute of a graph relation, we use $x.p$ to access the corresponding value of property p . Also, $\mathcal{L}(v)$ returns the labels of vertex v and $\mathcal{T}(e)$ returns the type of edge e .

Running example. An example graph inspired by the LDBC Social Network Benchmark [5] is shown on Fig. 1(a), while Fig. 1(b) presents the formalised graph. The graph vertices model four **Persons** and three **Messages**, with edges representing **LIKES**, **REPLY_OF** and **KNOWS** relations. In social networks, the **KNOWS** relation is symmetric, however, the property graph data model does not allow undirected edges. Hence, we use directed edges with an arbitrary direction and model the symmetric semantics of the relation in the queries.

3 The openCypher Query Language

Cypher is the a high-level declarative graph query language of the Neo4j graph database. It allows users to specify graph patterns with a syntax resembling an actual graph, which makes the queries easy to comprehend. The goal of the openCypher project [13] is to provide a standardised specification of the Cypher language. In the following, we introduce features of the language using examples.

3.1 Language Constructs

Inputs and outputs. openCypher queries take a *property graph* as their input, however the result of a query is not a graph, but a *graph relation*.

Vertex and path patterns. The basic building blocks of queries are patterns of vertices and edges. List. 3.1 shows a query that returns all vertices that model a `Person`. The query in List. 3.2 matches `Person` and `Message` pairs connected by a `LIKES` edge and returns the person's name and the message language. List. 3.3 describes person pairs that know each other directly or have a friend in common, i.e. from person `p1`, the other person `p2` can be reached using one or two hops.

```
MATCH (p:Person)
RETURN p
```

List. 3.1: Getting vertices.

```
MATCH (p:Person)-[:LIKES]->(m:Message)
RETURN p.name, m.language
```

List. 3.2: Pattern matching.

```
MATCH
  (p1:Person)-[ks:KNOWS*1..2]-
  (p2:Person)
RETURN p1, p2
```

List. 3.3: Variable length path.

```
MATCH (p:Person) WITH p
UNWIND p.speaks AS language
RETURN language,
  count(DISTINCT p.name) as cnt
```

List. 3.4: Grouping.

Filtering. Pattern matches can be filtered in two ways as illustrated in List. 3.5 and List. 3.6. (1) Vertex and edge patterns in the `MATCH` clause might have vertex label/edge type constraints written in the pattern after a colon, and (2) the optional `WHERE` subclause of `MATCH` might hold predicates.

```
MATCH (p1:Person)-[k:KNOWS]-(p2:Person)
WHERE k.since < 2000
RETURN p1.name, p2.name
```

List. 3.5: Filtering for edge property.

```
MATCH (p:Person)
WHERE p.name = 'Bob'
RETURN p.speaks
```

List. 3.6: Filtering.

Unique and non-unique edges. A `MATCH` clause defines a graph pattern. A query can be composed of multiple patterns spanning multiple `MATCH` clauses. For matches of a pattern within a single `MATCH` clause, edges are required to be unique. However, matches for multiple `MATCH` clauses can share edges. This means that in matches returned by List. 3.7, `k1` and `k2` are required to be different, while in matches returned by List. 3.8, `k1` and `k2` are allowed to be equal. For vertices, this restriction does not apply.⁵ This is illustrated in List. 3.9, which returns adjacent persons who like the same message.

⁵ Requiring uniqueness of edges is called *edge isomorphic matching*. Other query languages and execution engines might use *vertex isomorphic matching* (requiring uniqueness of vertices), *isomorphic matching* (requiring uniqueness of both vertices and edges) or *homomorphic matching* (not requiring uniqueness of either) [8].

```
MATCH (p1)-[k1:KNOWS]-(p2),
      (p2)-[k2:KNOWS]-(p3)
RETURN p1, k1, p2, k2, p3
```

List. 3.7: Different edges.

```
MATCH (p1)-[k1:KNOWS]-(p2)
MATCH (p2)-[k2:KNOWS]-(p3)
RETURN p1, k1, p2, k2, p3
```

List. 3.8: Non-unique edges.

```
MATCH (m:Message)<-[:LIKES]-(p1:Person)--(p2:Person)-[:LIKES]->(m)
RETURN p1, p2, m
```

List. 3.9: Triangle.

Creating the result set. The result set⁶ of a query is basically given in the `RETURN` clause, which can be de-duplicated using the `DISTINCT` modifier, sorted using the `ORDER BY` subclause. Skipping rows after sorting and limiting the result set to a certain number of records can be achieved using `SKIP` and `LIMIT` modifiers.

List. 3.10 illustrates these concepts by returning the name of the persons. The result set is restricted to the second and third names in alphabetical order.

```
MATCH (p:Person)
RETURN DISTINCT p.name
ORDER BY p.name
SKIP 1 LIMIT 2
```

List. 3.10: Deduplicate and sort.

```
1 MATCH
2   ()-[:LIKES]->(m:Message)<-[:LIKES]-( ),
3   (m)<-[:REPLY_OF]-(r)
4 RETURN r
```

List. 3.11: Multiple patterns.

Combining patterns. Multiple patterns (in the same or in different) `MATCH` clauses are combined together based on their common variables. List. 3.11 illustrates this by showing two patterns on lines 2 and 3. The first pattern describes a message `m` that has at least two likes. The second pattern finds replies to `m`.

Aggregation. openCypher specifies aggregation operators for performing calculations on multiple tuples.⁷ Unlike in SQL queries, the *grouping criteria* is determined implicitly in the `RETURN` as well as in `AND` and `WITH` clauses. Each expression of the expression list in `WITH` and `RETURN` are forced to contain either (1) no aggregate functions or (2) a single aggregate function at the outermost level. The grouping key is the tuple built from expressions of type (1).⁸ The query of List. 3.4 counts the number of persons commanding each language.

```
MATCH (p:Person) WITH p
UNWIND p.speaks AS lang
RETURN p.name, lang
```

List. 3.12: Unwind.

p.name	lang
Alice	en
Bob	fr
Cecil	en
Cecil	de

Fig. 2: Output of the unwind query.

⁶ The term *result set* refers to the *result collection*, which can be a set, a bag or a list.

⁷ The `avg`, `count`, `max`, `min`, `percentileCont`, `percentileDisc`, `stdDev`, `stdDevP`, `sum` functions return a single scalar value, while `collect` returns a list.

⁸ Decision on grouping semantics is due after the camera ready submission deadline. The semantics presented in this paper is one of the possible approaches.

Unwinding a list. The `UNWIND` construct takes an attribute and multiplies each tuple by appending the list elements one by one to the tuple, thus modifying the schema of the query part. By applying `UNWIND` to the `speaks` attribute List. 3.12 lists persons along the languages they speak. Fig. 2 shows the output of this query. As Cecil speaks two languages, he appears twice in the output. Note that “Daisy” speaks no languages, thus no tuples belong to her in the output.

3.2 Query Structure

In openCypher a query is composed as the `UNION` of one or more single queries. Each single query must have the same resulting schema, i.e. the resulting tuples must have the same arity and the same name at each position.

Single queries. A single query is composed of one or more query parts written subsequently. Query parts that form a prefix of a single query have one result set with the schema of the last query part’s schema in that prefix.

Query parts. Clause sequence of a query part matches the regular expression as follows: `MATCH*((WITH UNWIND?)|UNWIND|RETURN)`. They begin with an arbitrary number of `MATCH` clauses, followed by either (1) `WITH` and an optional `UNWIND`, (2) a single `UNWIND`, or (3) a `RETURN` in case of the last query part.⁹

The `RETURN` and `WITH` clauses use similar syntax and have the same semantics, the only difference being that `RETURN` should be used in the last query part while `WITH` should only appear in the preceding ones. These clauses list expressions whose value form the tuples, thus they determine the schema of the query parts.

```

1 MATCH (m1:Message)
2 WITH m1.language AS singleLang, count(*) AS cnt
3 WHERE cnt = 1
4 MATCH (m2:Message) WHERE m2.language = singleLang
5 OPTIONAL MATCH (m2)-[:REPLY_OF]->(m3:Message)
6 RETURN m2.language as reply, m3.language as orig

```

reply	orig
fr	en

Fig. 3: Result.

List. 3.13: Single query with multiple query parts.

Example. An openCypher single query composed of two query parts is shown on List. 3.13 along with its result on Fig. 3. It retrieves the language of messages that were written in a language no other message uses. If that message was a reply, the language of the original message is also retrieved.

The first query part spans lines 1–3 and the second spans lines 4–6. The result of the first query part is a single tuple $\langle \text{“fr”}, 1 \rangle$ with the schema $\langle \text{singleLang}, \text{cnt} \rangle$.

⁹ In openCypher, the filtering `WHERE` operation is a subclause of `MATCH` and `WITH`. When used in `WITH` as illustrated on line 3 of List. 3.13, `WHERE` is similar to the `HAVING` construct of SQL with the major difference that, in openCypher it is also allowed when no aggregation was specified in the query.

The second query part takes this result as an input to retrieve messages of the given languages and in case of a reply the original message in `m3`. The result of these two query parts together produces the final result whose schema is determined by the `RETURN` of the last query part (line 6).

4 Mapping openCypher to Relational Graph Algebra

In this section, we present relational graph algebra using the examples of Sec. 3.1 and provide a mapping that allows compilation from openCypher to this algebra.

#ops.	notation	name	props.	schema
0	$\bigcirc_{(v:L)}$	get-vertices	—	$\langle v \rangle$
1	$\updownarrow_{(v)}^{(w:L)} [e : T] (r)$	expand-both	—	$\text{sch}(r) \parallel \langle e, w \rangle$
	$\neq_{\text{variables}} (r)$	all-different	i	$\text{sch}(r)$
	$\omega_{xs \rightarrow x} (r)$	unwind	—	$\text{sch}(r) \setminus \langle xs \rangle \parallel \langle x \rangle$
	$\sigma_{\text{condition}} (r)$	selection	i	$\text{sch}(r)$
	$\pi_{x_1, x_2, \dots} (r)$	projection	i	$\langle x_1, x_2, \dots \rangle$
	$\gamma_{x_1, x_2, \dots}^{c_1, c_2, \dots} (r)$	grouping	i	$\langle x_1, x_2, \dots \rangle$
	$\delta(r)$	duplicate-elimination	i	$\text{sch}(r)$
	$\tau_{\downarrow x_1, \uparrow x_2, \dots}^{\text{limit}} (r)$	sorting	i	$\text{sch}(r)$
2	$\lambda_{\text{skip}} (r)$	top	—	$\text{sch}(r)$
	$r \cup s$	union	c, a	$\text{sch}(r)$
	$r \uplus s$	bag union	c, a	$\text{sch}(r)$
	$r \bowtie s$	natural join	c, a	$\text{sch}(r) \parallel (\text{sch}(s) \setminus \text{sch}(r))$
	$r \Join s$	left outer join	—	$\text{sch}(r) \parallel (\text{sch}(s) \setminus \text{sch}(r))$

Table 1: Number of operands, properties and result schemas of relational graph algebra operators. A unary operator α is idempotent (i), iff $\alpha(x) = \alpha(\alpha(x))$ for all inputs. A binary operator β is commutative (c), iff $x \beta y = y \beta x$ and associative (a), iff $(x \beta y) \beta z = x \beta (y \beta z)$. For schema transformations, append is denoted by \parallel , while removal is marked by \setminus . L represents a (possibly empty) set of vertex labels and T represents a (possibly empty) set of edge types.

4.1 An Algebra for Formalising Graph Queries

We present both standard operators of relational algebra [4] and operators for graph relations. Tab. 1 provides an overview of the operators of relational graph algebra. We follow the openCypher query language and present a mapping from the language constructs to their algebraic equivalents¹⁰, summarized in Tab. 2. The corresponding rows of the table (e.g. ①) are referred to in the text.

¹⁰ Patterns in the openCypher query might contain anonymous vertices and edges. In the algebraic form, we denote this with names starting with an underscore, such as `_v1` and `_e2`.

Basic operators. The *projection* operator π keeps a specific set of attributes in the relation: $t = \pi_{x_1, \dots, x_n}(r)$. Note that the tuples are not deduplicated (i.e. filtered to sets), thus t will have the same number of tuples as r . The projection operator can also rename the attributes, e.g. $\pi_{x_1 \rightarrow y_1}(r)$ renames x_1 to y_1 .

The *selection* operator σ filters the incoming relation according to some criteria. Formally, $t = \sigma_\theta(r)$, where predicate θ is a propositional formula. Relation t contains all tuples from r for which θ holds.

Vertices and patterns. ①–② The *get-vertices* [7] nullary operator $\bigcirc_{(v:l_1 \wedge \dots \wedge l_n)}$ returns a graph relation of a single attribute v that contains vertices that have all of labels l_1, \dots, l_n . Using this operator, the query in List. 3.1 is compiled to

$$\bigcirc_{(p:\text{Person})}$$

③–⑥ The *expand-out* unary operator $\uparrow_{(v)}^{(w:l_1 \wedge \dots \wedge l_n)} [e : t_1 \vee \dots \vee t_k](r)$ adds new attributes e and w to each tuple iff there is an edge e from v to w , where e has *any* of types t_1, \dots, t_k , while w has *all* labels l_1, \dots, l_n .¹¹ More formally, this operator appends the $\langle e, w \rangle$ to a tuple iff $st(e) = \langle v, w \rangle$, $l_1, \dots, l_n \in \mathcal{L}(w)$ and $\mathcal{T}(e) \in \{t_1, \dots, t_k\}$. Using this operator, the query in List. 3.2 can be formalised as

$$\pi_{p.\text{name}, m.\text{language}} \uparrow_{(p)}^{(m:\text{Message})} [_ e1 : \text{LIKES}] \bigcirc_{(p:\text{Person})}$$

Similarly to the expand-out operator, the *expand-in* operator \downarrow appends $\langle e, w \rangle$ iff $st(e) = \langle w, v \rangle$, while the *expand-both* operator \updownarrow uses edge e iff either $st(e) = \langle v, w \rangle$ or $st(e) = \langle w, v \rangle$. We also propose an extended version of this operator, $\uparrow_{(v)}^{(w)} [e *_{\min}^{\max}]$, which may use between min and max hops. Using this extension, List. 3.3 is compiled to

$$\pi_{p1, p2} \not\equiv_{ks} \updownarrow_{(p1)}^{(p2:\text{Person})} [ks : \text{KNOWS} *_1^2] \bigcirc_{(p1:\text{Person})}$$

Combining and filtering pattern matches. ⑦–⑪ In order to express the uniqueness criterion for edges (illustrated in Sec. 3.1) in a compact way, we propose the *all-different* operator. The all-different operator $\not\equiv_{E_1, \dots, E_n}(r)$ filters r to keep tuples where variables in $\cup_i E_i$ are pairwise different. Note that the operator is actually a shorthand for the following selection:

$$\not\equiv_{E_1, \dots, E_n}(r) = \sigma_{\substack{e_1, e_2 \in \cup_i E_i \\ e_1 \neq e_2}} \bigwedge_{r.e_1 \neq r.e_2} (r)$$

¹¹ Label and type constraints can be omitted for the get-vertices operator and the expand operators. For example, $\bigcirc_{(v)}$ returns all vertices, while $\uparrow_{(v)}^{(w)} [e](r)$ traverses all outgoing edges e from vertices v to w , regardless of their labels/types.

Using the all-different operator, query in List. 3.7 is compiled to

$$\pi_{p1,k1,p2,k2,p3} \not\equiv_{k1,k2} \downarrow_{(p2)}^{(p3)} [k2 : \text{KNOWS}] \downarrow_{(p1)}^{(p2)} [k1 : \text{KNOWS}] \circ_{(p1)}$$

⑦–⑧ The result of the *natural join* operator \bowtie is determined by creating the Cartesian product of the relations, then filtering those tuples which are equal on the attributes that share a common name. The combined tuples are projected: from the attributes present in both of the two input relations, we only keep the ones in r and drop the ones in s . Thus, the join operator is defined as

$$r \bowtie s = \pi_{R \cup S} \left(\sigma_{r.A_1=s.A_1 \wedge \dots \wedge r.A_n=s.A_n} (r \times s) \right),$$

where $\{A_1, \dots, A_n\} = R \cap S$ is the set of attributes that occur both in R and S . In order to allow pattern matches to share the same edge, they must be included in different **MATCH** clauses as shown on List. 3.8 which is compiled to

$$\pi_{p1,p2,p3} \left(\left(\downarrow_{(p1)}^{(p2)} [k1 : \text{KNOWS}] \circ_{(p1)} \right) \bowtie \left(\downarrow_{(p2)}^{(p3)} [k2 : \text{KNOWS}] \circ_{(p2)} \right) \right)$$

The query in List. 3.11 with two patterns in one **MATCH** clause is compiled to:

$$\pi_r \not\equiv_{e1, e2, e3} \left(\left(\downarrow_{(m)}^{(v2)} [_e2 : \text{LIKES}] \uparrow_{(v1)}^{(m:\text{Message})} [_e1 : \text{LIKES}] \circ_{(v1)} \right) \bowtie \left(\downarrow_{(m)}^{(r)} [_e3 : \text{REPLY_OF}] \circ_{(m:\text{Message})} \right) \right)$$

⑨–⑪ The *left outer join* operator produces $t = r \bowtie s$ combining matching tuples of r and s according to a given matching semantics.¹² In case there is no matching tuple in s for a particular tuple $e \in r$, e is still included in the result, with tuple attributes that exclusively belong to relation s having a value of ε .

Result and subresult operations. ⑩ The *duplicate-elimination* operator δ eliminates duplicate tuples in a bag.

⑪ The *grouping* operator γ groups tuples according to their value in one or more attributes and aggregates the remaining attributes.

We generalize the grouping operator to explicitly state the *grouping criteria* and allow for complex aggregate expressions. This is similar to the SQL query language where the grouping criteria is explicitly given in **GROUP BY**.

We use the notation $\gamma_{e_1, e_2, \dots}^{c_1, c_2, \dots}$, where c_1, c_2, \dots in the superscript form the *grouping criteria*, i.e. the list of expressions whose values partition the incoming tuples into groups. For each and every group this aggregation operator emits a single tuple of expressions listed in the subscript, i.e. $\langle e_1, e_2, \dots \rangle$. Given attributes

¹² Matching semantics might use value equality of attributes that share a common name (similarly to natural join) or use an arbitrary condition (similarly to θ -join).

Combining results. The \cup operator produces the *set union* of two relations, while the \uplus operator produces the *bag union* of two operators, e.g. $\{\langle 1, 2 \rangle, \langle 3, 4 \rangle\} \uplus \{\langle 1, 2 \rangle\} = \{\langle 1, 2 \rangle, \langle 1, 2 \rangle, \langle 3, 4 \rangle\}$. For both the union and bag union operators, the schema of the operands must have the same attributes.

4.2 Mapping openCypher Queries to Relational Graph Algebra

In this section, we give the mapping algorithm of openCypher queries to relational graph algebra and also give a more detailed listing of the compilation rules for the query language constructs in Tab. 2. We follow a bottom-up approach to build the relational graph algebra expression.

1. Process each single query as follows and combine their result using the union operation. As the union operator is technically a binary operator, the union of more than two single queries are represented as a left-deep tree of **UNION** operators.
2. For each query part of a single query, denoted by t , the relational graph algebra tree built from the prefix of query parts up to—but not including—the current query part, process the current query part as follows.
 1. A single pattern is turned left-to-right to a get-vertices for the first vertex and a chain of expand-in, expand-out or expand-both operators for inbound, outbound or undirected relationships, respectively.
 2. Comma-separated patterns in a single **MATCH** are connected by natural join.
 3. Append an all-different operator for all edge variables that appear in the **MATCH** clause because of the non-repeating edges language rule.
 4. Process the **WHERE** subclause of a single **MATCH** clause.
 5. Several **MATCH** clauses are connected to a left-deep tree of natural join. For **OPTIONAL MATCH**, left outer join is used instead of natural join. In case there is a **WHERE** subclause, its condition becomes part of the join condition, i.e. it will never filter on the input from previous **MATCH** clauses.
 6. If there is a positive or negative pattern deferred from **WHERE** processing, append it as a natural join or a combination of left outer join and selection operator filtering on no matches were found, respectively.
 7. If this is not the first query part, combine the current query part with the relational graph algebra tree of the preceding query parts by appending a natural join here. Its left operand will be t and its right operand will be the relational graph algebra tree built so far from the current subquery.
 8. Append grouping, if **RETURN** or **WITH** clause has grouping functions inside.
 9. Append a projection operator based on the **RETURN** or **WITH** clause. This operator will also handle the renaming (i.e. **AS**).
 10. Append a duplicate-elimination operator, if the **RETURN** or **WITH** clause has the **DISTINCT** modifier.
 11. Append a selection operator if **WITH** had the optional **WHERE** subclause.

Language construct	Relational algebra expression	
Vertices and patterns. $\langle p \rangle$ denotes a pattern that contains a vertex $\langle v \rangle$.		
$\langle v \rangle$	$\bigcirc_{(v)}$	①
$\langle v \rangle : \langle l1 \rangle : \dots : \langle ln \rangle$	$\bigcirc_{(v:l1 \wedge \dots \wedge ln)}$	②
$\langle p \rangle - [\langle e \rangle : \langle t1 \rangle \dots \langle tk \rangle] \rightarrow \langle w \rangle$	$\uparrow_{(v)}^{(w)} [e : t1 \vee \dots \vee tk] (p)$, where e is an edge	③
$\langle p \rangle \leftarrow [\langle e \rangle : \langle t1 \rangle \dots \langle tk \rangle] - \langle w \rangle$	$\downarrow_{(v)}^{(w)} [e : t1 \vee \dots \vee tk] (p)$, where e is an edge	④
$\langle p \rangle \leftarrow [\langle e \rangle : \langle t1 \rangle \dots \langle tk \rangle] \rightarrow \langle w \rangle$	$\updownarrow_{(v)}^{(w)} [e : t1 \vee \dots \vee tk] (p)$, where e is an edge	⑤
$\langle p \rangle - [\langle e \rangle * \langle min \rangle \dots \langle max \rangle] \rightarrow \langle w \rangle$	$\uparrow_{(v)}^{(w)} [e *_{min}^{max}] (p)$, where e is a list of edges	⑥
Combining and filtering pattern matches		
MATCH $\langle p1 \rangle, \langle p2 \rangle, \dots$	$\neq_{\text{edges of } p1, p2, \dots} (p1 \bowtie p2 \bowtie \dots)$	⑦
MATCH $\langle p1 \rangle$ MATCH $\langle p2 \rangle$	$\neq_{\text{edges of } p1} (p1) \bowtie \neq_{\text{edges of } p2} (p2)$	⑧
OPTIONAL MATCH $\langle p \rangle$	$\{\langle \rangle\} \bowtie \neq_{\text{edges of } p} (p)$	⑨
OPTIONAL MATCH $\langle p \rangle$ WHERE $\langle condition \rangle$	$\{\langle \rangle\} \bowtie_{\text{condition}} \neq_{\text{edges of } p} (p)$	⑩
$\llbracket r \rrbracket$ OPTIONAL MATCH $\langle p \rangle$	$r \bowtie \neq_{\text{edges of } p} (p)$	⑪
$\llbracket r \rrbracket$ WHERE $\langle condition \rangle$	$\sigma_{\text{condition}}(r)$	⑫
$\llbracket r \rrbracket$ WHERE $\langle v \rangle : \langle l1 \rangle : \dots : \langle ln \rangle$	$\sigma_{\mathcal{L}(v)=l1 \wedge \dots \wedge \mathcal{L}(v)=ln}(r)$	⑬
$\llbracket r \rrbracket$ WHERE $\langle p \rangle$	$r \bowtie p$	⑭
Result and subresult operations. Rules for RETURN also apply to WITH .		
$\llbracket r \rrbracket$ RETURN $\langle x1 \rangle$ AS $\langle y1 \rangle, \dots$	$\pi_{x1 \rightarrow y1, \dots}(r)$	⑮
$\llbracket r \rrbracket$ RETURN DISTINCT $\langle x1 \rangle$ AS $\langle y1 \rangle, \dots$	$\delta(\pi_{x1 \rightarrow y1, \dots}(r))$	⑯
$\llbracket r \rrbracket$ RETURN $\langle x1 \rangle, \langle \text{aggr} \rangle(\langle x2 \rangle)$	$\gamma_{x1, \text{aggr}(x2)}^{x1}(r)$ (see Sec. 3.1)	⑰
$\llbracket r \rrbracket$ WITH $\langle x1 \rangle$ $\llbracket s \rrbracket$ RETURN $\langle x2 \rangle$	$\pi_{x2}(\pi_{x1}(r) \bowtie s)$	⑱
Unwinding and list operations		
$\llbracket r \rrbracket$ UNWIND $\langle xs \rangle$ AS $\langle x \rangle$	$\omega_{xs \rightarrow x}(r)$	⑲
$\llbracket r \rrbracket$ ORDER BY $\langle x1 \rangle$ ASC , $\langle x2 \rangle$ DESC , \dots	$\tau_{\uparrow x1, \downarrow x2, \dots}(r)$	⑳
$\llbracket r \rrbracket$ SKIP $\langle s \rangle$ LIMIT $\langle l \rangle$	$\lambda_l^s(r)$	㉑
Combining results		
$\llbracket r \rrbracket$ UNION $\llbracket s \rrbracket$	$r \cup s$	㉒
$\llbracket r \rrbracket$ UNION ALL $\llbracket s \rrbracket$	$r \uplus s$	㉓

Table 2: Mapping from openCypher constructs to relational algebra. Variables, labels, types and literals are typeset as $\langle v \rangle$. The notation $\langle p \rangle$ represents patterns resulting in a relation p , while $\llbracket r \rrbracket$ denotes previous query fragment resulting in a relation r . To avoid confusion with the “.” language construct (used for ranges), we use \dots to denote omitted query fragments.

4.3 Summary and Limitations

In this section, we presented a mapping that allows us to express the example queries of Sec. 3.1 in graph relational algebra. We extended relational algebra by adapting operators (\bigcirc , \uparrow , τ , λ), precisely specifying grouping semantics (γ) and defining the all-different operator (\neq). Finally, we proposed an algorithm for compiling openCypher graph queries to graph relational algebra.

Our mapping does not completely cover the openCypher language. As discussed in Sec. 3, some constructs are defined as legacy and thus were omitted. The current formalisation does not include expressions (e.g. conditions in selections) and maps. Compiling data manipulation operations (such as `CREATE`, `DELETE`, `SET`, and `MERGE`) to relational algebra is also subject of future work.

5 Related Work

Property graph data models. The TinkerPop framework aims to provide a standard data model for property graphs, along with Gremlin, a high-level imperative graph traversal language [16] and the Gremlin Structure API, a low-level programming interface.

EMF. The Eclipse Modeling Framework is an object-oriented modelling framework widely used in model-driven engineering. Henshin [1] provides a visual language for defining patterns, while Epsilon [9] and VIATRA Query [2] provide high-level declarative (textual) query languages, the Epsilon Pattern Language and the VIATRA Query Language.

SPARQL. Widely used in semantic technologies, SPARQL is a standardised declarative graph pattern language for querying RDF [24] graphs. SPARQL bears close similarity to Cypher queries, but targets a different data model and requires users to specify the query as triples instead of graph vertices/edges [14]. G-SPARQL [19] extended the SPARQL language for attributed graphs, resulting in a language with an expressive power similar to openCypher.

SQL. In general, relational databases offer limited support for graph queries: recursive queries are supported by PostgreSQL using the `WITH RECURSIVE` keyword and by the Oracle Database using the `CONNECT BY` keyword. Graph queries are supported in the SAP HANA prototype [18], through a SQL-based language [10].

Cypher. Due to its novelty, there are only a few research works on the formalisation of (open)Cypher. The authors of [7] defined *graph relations* and introduced the `GETNODES`, `EXPANDIN` and `EXPANDOUT` operators. While their work focused on optimisation transformations, this paper aims to provide a more complete and systematic mapping from openCypher to relational algebra.

In [8], graph queries were defined in a Cypher-like language and evaluated on Apache Flink. However, formalisation of the queries was not discussed in detail.

Comparison of graph query frameworks. Previously, we published the Train Benchmark, a framework for comparing graph query frameworks across different technological spaces, such as property graphs, EMF, RDF and SQL [21].

6 Conclusion and Future Work

In this paper, we presented a formal specification for a subset of the openCypher query language. This provides the theoretical foundations to use openCypher as a language for graph query engines.

As future work, we plan to provide a formalisation based on graph-specific theoretical query frameworks, such as [12]. We will also give the formal specification of the operators for incremental query evaluation, which requires the definition of *maintenance operations* that keep the result in sync with the latest set of changes [22]. Our long-term research objective is to design an openCypher-compatible *distributed, incremental graph query engine* [20].¹⁴

Acknowledgements

Gábor Szárnyas and Dániel Varró were supported by the MTA-BME Lendület Research Group on Cyber-Physical Systems and the NSERC RGPIN-04573-16 project. The authors would like to thank Gábor Bergmann and János Maginecz for their comments on the draft of this paper.

References

1. T. Arendt et al. Henshin: Advanced concepts and tools for in-place EMF model transformations. In *MODELS*, pages 121–135, 2010.
2. G. Bergmann et al. Incremental evaluation of model queries over EMF models. In *MODELS*, pages 76–90, 2010.
3. E. Botoeva et al. OBDA beyond relational DBs: A study for MongoDB. In *Proceedings of the 29th International Workshop on Description Logics*, 2016.
4. R. Elmasri and S. B. Navathe. *Fundamentals of Database Systems*. Addison-Wesley-Longman, 3rd edition, 2000.
5. O. Erling et al. The LDBC Social Network Benchmark: Interactive workload. In *SIGMOD*, pages 619–630, 2015.
6. H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database systems – The complete book*. Pearson Education, 2nd edition, 2009.
7. J. Hölsch and M. Grossniklaus. An algebra and equivalences to transform graph patterns in Neo4j. In *GraphQ at EDBT/ICDT*, 2016.
8. M. Junghanns et al. Cypher-based graph pattern matching in Gradoop. In *GRADES at SIGMOD*, 2017.
9. D. S. Kolovos et al. The Epsilon transformation language. In *ICMT*, 2008.
10. C. Krause et al. An SQL-based query language and engine for graph pattern matching. In *ICGT*, pages 153–169, 2016.
11. C. Li, K. C. Chang, I. F. Ilyas, and S. Song. RankSQL: Query algebra and optimization for relational top-k queries. In *SIGMOD*, pages 131–142, 2005.

¹⁴ Our prototype, *ingraph*, is available at: <http://docs.inf.mit.bme.hu/ingraph/>

12. L. Libkin et al. Querying graphs with data. *J. ACM*, 63(2):14:1–14:53, 2016.
13. Neo Technology. openCypher project. <http://www.opencypher.org/>, 2017.
14. J. Pérez et al. Semantics and complexity of SPARQL. *ACM TODS*, 34(3), 2009.
15. M. A. Rodriguez. A collectively generated model of the world. In *Collective intelligence: creating a prosperous world at peace*, pages 261–264. 2008.
16. M. A. Rodriguez. The Gremlin graph traversal machine and language (invited talk). In *DBPL*, pages 1–10, 2015.
17. M. A. Rodriguez and P. Neubauer. The graph traversal pattern. In *Graph Data Management: Techniques and Applications*, pages 29–46. 2011.
18. M. Rudolf et al. The graph story of the SAP HANA database. In *BTW*, 2013.
19. S. Sakr, S. Elnikety, and Y. He. G-SPARQL: a hybrid engine for querying large attributed graphs. In *CIKM*, pages 335–344, 2012.
20. G. Szárnyas et al. IncQuery-D: A distributed incremental model query framework in the cloud. In *MODELS*, pages 653–669, 2014.
21. G. Szárnyas et al. The Train Benchmark: Cross-technology performance evaluation of continuous model validation. *Softw. Syst. Model.*, 2017.
22. G. Szárnyas, J. Maginecz, and D. Varró. Evaluation of optimization strategies for incremental graph queries. *Periodica Polytechnica, EECS*, 2017.
23. G. Szárnyas and J. Marton. Formalisation of openCypher queries in relational algebra. Technical report, Budapest University of Technology and Economics, 2017. <http://hdl.handle.net/10890/5395>.
24. W3C. Resource Description Framework. <https://www.w3.org/RDF/>, 2014.