

# Model-driven engineering of an openCypher engine: using graph queries to compile graph queries

József Marton<sup>1</sup>, Gábor Szárnyas<sup>2,3</sup>, and Márton Búr<sup>2,3</sup>

<sup>1</sup> Budapest University of Technology and Economics, Database Laboratory  
marton@db.bme.hu

<sup>2</sup> Budapest University of Technology and Economics  
Fault Tolerant Systems Research Group

<sup>3</sup> MTA-BME Lendület Research Group on Cyber-Physical Systems  
{szarnyas, bur}@mit.bme.hu

**Abstract.** Graph database systems are increasingly adapted for storing and processing heterogeneous network-like datasets. Many challenging applications with near real-time requirements—such as financial fraud detection, on-the-fly model validation and root cause analysis—can be formalised as graph problems and tackled with graph databases efficiently. However, as no standard graph query language has yet emerged, users are subjected to the possibility of vendor lock-in. openCypher is an initiative to define an open specification for a declarative graph query language. Still, due to its sophisticated nature, creating an openCypher-compatible query engine requires significant research and engineering efforts. Meanwhile, model-driven language workbenches support the creation of domain-specific languages by providing high-level tools to create parsers, editors and compilers. In this paper, we present a compiler for openCypher, built on model-driven technologies and released under the commercially-friendly Eclipse Public License. We use queries over a cyber-physical system to demonstrate our approach.

## 1 Introduction

**Context.** Graphs are a well-known formalism, widely used for describing and analysing systems. Graphs provide an intuitive formalism for modelling 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*) [31].

The *property graph* data model [34] 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 by matching complex graph patterns, which would otherwise be cumbersome to define and/or inefficient to evaluate on traditional relational databases [39].

Neo4j, a popular NoSQL property graph database, offers the Cypher query language to specify graph patterns. Cypher is a high-level declarative query language, divorced from the query execution plan, which allows the query engine

to use sophisticated optimisation techniques. Neo Technology, the company behind Neo4j initiated the openCypher project [26], which aims to deliver an open specification of Cypher.

**Problem and objectives.** Even though the openCypher specification was released more than 1.5 years ago, there are very few open implementations available and even those offer limited support for the more advanced language constructs. Besides the novelty of the openCypher specification, the primary reason for the lack of open implementations is the complexity of the language. Even with the artifacts provided by the openCypher project—including the specification, the language grammar and a set of test cases—implementing a compiler is a non-trivial task and requires significant research and engineering efforts. Our goal is to deliver a reusable compiler that can be extended with transformation rules for query optimisation.

**Contributions.** In this paper, we use graph queries defined on a cyber-physical system to demonstrate the key challenges in compiling openCypher queries. We present an approach for implementing an openCypher query compiler, built on a modern language workbench. The workbench is based on Eclipse technologies, including a model-based parser generator and a set of model transformation rules. We released the compiler as part of the open-source ingraph project, where it is used as part of an incremental graph query engine.

**Structure of the paper.** We first introduce the concepts of graph queries and model transformations in Sec. 3 and the running example in Sec. 2. We give an overview of the compiler in Sec. 4 and use example queries to elaborate the details of query compilation in Sec. 5. We discuss related research in Sec. 6 and conclude the paper in Sec. 7.

## 2 Running Example

To demonstrate our approach, we use a cyber-physical system demonstrator, MoDeS3 [8,6], which stands for Model-Based Demonstrator for Smart and Safe Systems. It is an educational platform of a model railway system that prevents trains from collision and derailment using runtime verification techniques based on safety monitors. The railway track is instrumented with several sensors, such as *cameras*, *shunt detectors* capable of sensing trains on a particular segment of a track connected to several computing units. In addition to collecting data, these computing units also control the trains to guarantee safe operation. In this paper, we will only introduce a small self-contained fragment of the demonstrator in order to keep the example compact.

Fig. 1(a) depicts a snapshot of the simplified system in operation, where trains are located at different parts of the railway. The railway network itself consists of two types of railway elements: *segments* and *turnouts*. Segments are

selected tracks of the railway network with one entry and exit points individually, they are approximately of same lengths, and they have no intermediate branches between the entry and exit points. As opposed to segments, turnouts allow trains to change tracks. A turnout can either be in *divergent* or *straight* state. A *station* can represent a railway station with an arbitrary purpose, and they can include any number of railway elements.

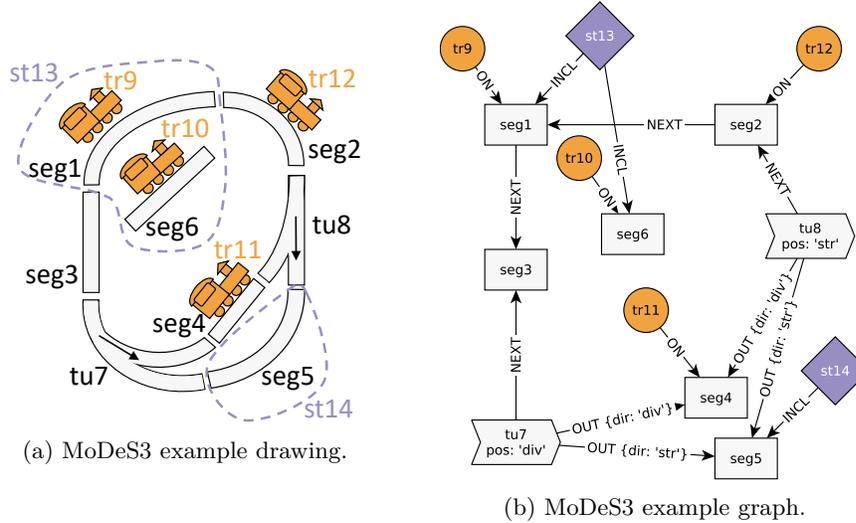


Fig. 1: MoDeS3 example.

We introduce the following example *monitoring objectives* that are evaluated continuously:

- *Close proximity* identifies trains on consecutive segments with only a limited distance between each other (train `tr9` on `seg9` and `tr12` on `seg2` in the example).
- *Station with free track* monitoring objective finds stations that have at least one free track available (station `st14` in the example).
- *Busy station* identifies stations with at least two trains residing on its corresponding tracks (station `st13` in the example).

### 3 Preliminaries

In this section, we present the theoretical and practical foundations for compiling openCypher queries. This includes the notion of property graphs, a brief description of the openCypher language and the relational algebraic foundations for formalising graph queries. We also discuss model-driven engineering (MDE) along with the MDE tools used in our work.

### 3.1 Property Graphs

The *property graph* data model [33] extends typed graphs with properties on the vertices and edges. This data model is used in NoSQL graph database systems such as Neo4j [25], OrientDB [28] and Titan [40]. Graph databases provide no or weak metamodeling capabilities. Hence, models can either be stored in a weakly typed manner or the metamodel must be included in the graph (on the same metalevel as the instance model).

The property graph of the running example is shown in Fig. 1(b).

### 3.2 The openCypher Query Language

Cypher is a high-level declarative graph query language used in the Neo4j graph database [30]. 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 [26] is to provide a standardised specification of the Cypher language.

List. 3.1 shows a query that returns all `tr`, `seg` pairs, where a particular train `tr` is ON a particular segment `seg`.

```
1 MATCH (tr:Train)-[:ON]->(seg:Segment)
2 RETURN tr, seg
```

List. 3.1: Example openCypher query.

### 3.3 Relational Graph Algebra

We gave a formal specification for the core subset of the openCypher language in [24] using relational graph algebra, which extends relational algebra with graph-specific operators. Here, we give a brief summary of the operators in relational graph algebra, which operates on multisets (bags) [15] of tuples, that form *graph relations*. We refer to named elements of a tuple as *attributes*.

**Nullary operators.** The *get-vertices* [18] 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  $l_1, \dots, l_n$ .

Additionally to our previous works, we introduce `Dual`, which is a relation with no columns and a single (empty) tuple  $\{\langle \rangle\}$ .<sup>4</sup> The `Dual` relation is the identity element of the Cartesian product and the natural join operator. We also introduce `Singular`, which denotes a the empty relation  $\{\}$  and is the zero element of the Cartesian product and the natural join operator.

<sup>4</sup> The `Dual` relation is inspired by the `DUAL` table in the Oracle database [5].

**Unary operators.** The projection operator  $\pi$  keeps the specified set of attributes in the relation:  $\pi_{A_1, \dots, A_n}(r)$ . The projection operator can also rename attributes, e.g.  $\pi_{x_1 \rightarrow x_2}(r)$  renames  $x_1$  to  $x_2$ . Note that tuples are not deduplicated, i.e. the result has the same number of tuples as the input relation  $r$ .

The duplicate-elimination operator  $\delta$  deduplicates tuples from its input.

The selection operator  $\sigma$  filters the incoming relation according to some criteria:  $\sigma_\theta(r)$ , where predicate  $\theta$  is a propositional formula. The operator selects all tuples in  $r$  for which  $\theta$  holds.

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 outgoing edge  $e$  from  $v$  to  $w$ , where  $e$  has *any* of types  $t_1, \dots, t_k$ , while  $w$  has *all* labels  $l_1, \dots, l_n$ . Similarly to the expand-out operator, the *expand-in* operator  $\downarrow$  uses incoming edges, while the *expand-both* operator  $\updownarrow$  uses both incoming and outgoing edges. An extended version of this operator,  $\uparrow_{(v)}^{(w)} [e *_{\min}^{\max}]$  may use between min and max hops.

**Binary operators.** The result of the *natural join* operator  $\bowtie$  is determined by creating the Cartesian product of the relations, then filtering for 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  $s$ :

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

where  $\{A_1, \dots, A_n\} = R \cap S$  is the set of attributes that occur both in  $R$  and  $S$ .

The antijoin operator  $\triangleright$  (also known as *left anti semijoin*) collects the tuples from the left relation  $r$  which have no matching pair in the right relation  $s$ :

$$r \triangleright s = r \setminus \pi_R(r \bowtie s),$$

where  $\pi_R$  denotes a projection operation, which only keeps the attributes of the schema over relation  $r$ . The antijoin operator is not commutative and not associative. The left outer join  $\triangleright\bowtie$  pads tuples from the left relation that did not match any from the right relation with  $\varepsilon$  values and adds them to the result of the natural join [36].

Tab. 1 shows a concise set of rules for mapping openCypher expressions to relational graph algebra [24].

### 3.4 Model-Driven Engineering

Model-driven engineering (MDE) is a development paradigm, used in many areas of software and system engineering, such as designing safety-critical systems. MDE focuses on creating and analyzing models at different levels of abstraction during the engineering process. *Model transformations* are used to process models, e.g. to convert models between different modelling languages and to generate code.

Language construct	Relational algebra expression
Vertices and patterns. $\langle p \rangle$ denotes a pattern that contains a vertex $\langle v \rangle$ .	
$\langle v \rangle : \langle l1 \rangle : \dots : \langle ln \rangle$	$\bigcirc_{(v:l1 \wedge \dots \wedge ln)}$
$\langle p \rangle \leftarrow [ \langle e \rangle : \langle t1 \rangle   \dots   \langle tk \rangle ] \rightarrow \langle w \rangle$	$\uparrow \downarrow \binom{w}{v} [ e : t1 \vee \dots \vee tk ] (p)$ , where $e$ is an edge
$\langle p \rangle \leftarrow [ \langle e \rangle * \langle min \rangle .. \langle max \rangle ] \rightarrow \langle w \rangle$	$\uparrow \binom{w}{v} [ e *_{min}^{max} ] (p)$ , where $e$ is a list of edges
Combining and filtering pattern matches	
<code>MATCH</code> $\langle p1 \rangle, \langle p2 \rangle, \dots$	$\neq_{\text{edges of } p1, p2, \dots} (p1 \bowtie p2 \bowtie \dots)$
<code>MATCH</code> $\langle p1 \rangle$ <code>MATCH</code> $\langle p2 \rangle$	$\neq_{\text{edges of } p1} (p1) \bowtie \neq_{\text{edges of } p2} (p2)$
<code>OPTIONAL MATCH</code> $\langle p \rangle$ <code>WHERE</code> $\langle \text{condition} \rangle$	$\{ \langle \rangle \} \bowtie_{\text{condition}} \neq_{\text{edges of } p} (p)$
$\llbracket r \rrbracket$ <code>OPTIONAL MATCH</code> $\langle p \rangle$	$\neq_{\text{edges of } r} (r) \bowtie \neq_{\text{edges of } p} (p)$
$\llbracket r \rrbracket$ <code>WHERE</code> $\langle \text{condition} \rangle$	$\sigma_{\text{condition}}(r)$
$\llbracket r \rrbracket$ <code>WHERE</code> $\langle p \rangle$	$r \bowtie p$
Result and subresult operations. Rules for <code>RETURN</code> also apply to <code>WITH</code> .	
$\llbracket r \rrbracket$ <code>RETURN</code> $\langle x1 \rangle$ <code>AS</code> $\langle y1 \rangle, \dots$	$\pi_{x1 \rightarrow y1, \dots} (r)$
$\llbracket r \rrbracket$ <code>RETURN</code> $\langle x1 \rangle, \langle \text{aggr} \rangle (\langle x2 \rangle)$	$\gamma_{x1, \text{aggr}(x2)}^{x1} (r)$
$\llbracket r \rrbracket$ <code>WITH</code> $\langle x1 \rangle$ $\llbracket s \rrbracket$ <code>RETURN</code> $\langle x2 \rangle$	$\pi_{x2} \left( \left( \pi_{x1} (r) \right) \bowtie s \right)$

Table 1: Mapping from openCypher constructs to relational algebra [24]. 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.

**Domain-specific languages.** While there are some extensible formalisms intended as a general-purpose way of representing models (such as UML), the industrial practice often prefers domain-specific languages (DSLs) for describing models instead. These can be designed and modified to the needs of application domains and actual design processes. On the other hand, developing such a DSL (and providing tool support) is an expensive task.

The Eclipse Modeling Framework (EMF) is a domain-specific modeling technology, built on the Eclipse platform. A DSL development process with EMF starts with the definition of a metamodel, from which several components of the modeling tool can be automatically derived. It is defined using the Ecore formalism, which is defined as part of the EMF. Ecore enables to define metamodels and automates the generation of a wide range of tools. Ecore is discussed in detail in [37].

**Language workbenches.** Model-driven language workbenches [13] support the creation of domain-specific languages by providing high-level tools to cre-

ate parsers, editors and compilers. Xtext [14] is an EMF-based framework for development of programming languages and DSLs. Xtend is a general-purpose programming language (implemented with an Xtext-based parser), which is transpiled to Java source code. Xcore [12] is a textual modeling language for EMF, which provides an Xtext-based language for defining metamodels.

**VIATRA** VIATRA [43] is an open-source Eclipse project written in Java and Xtend [11]. VIATRA builds on the Eclipse Modeling Framework and provides the following main features:

- The VIATRA Query Language, a declarative language for writing queries over models, which are evaluated once or incrementally upon each model change.
- An internal domain-specific language over the Xtend language to specify both batch and event-driven, reactive transformations.
- A rule-based design space exploration framework [17] to explore design candidates with transformation rules where the design candidates must satisfy multiple criteria.

## 4 Overview of the Approach

The high-level workflow of our openCypher query engine is shown in Fig. 2. A domain expert first formulates the *query* using the openCypher language, which serves as the input for our engine. The query is then parsed and transformed into the *query syntax graph* using the openCypher grammar created by the Slizaa project<sup>5</sup>. It is then compiled to our relational graph algebra model. This produces a canonical relational graph algebra representation to keep compiler code simple. The relational graph algebra representation is modified by the **relational algebra optimizer**. The resulting relational algebra model is then passed on to the query execution engine.

**Relational algebra optimizer.** The **relational algebra optimizer** has two main tasks. It removes idempotent operations from the relational graph algebra model and identifies combinations of operations that could be expressed using advanced operations. The relational graph algebra model is also a graph, so both of these tasks are graph manipulation tasks which we have implemented using graph pattern matches using the VIATRA model transformation framework (Sec. 3.4).

**Relational graph algebra metamodel.** The metamodel of the relational graph algebra operators introduced in Sec. 3.3 is shown in Fig. 3. An openCypher query is represented by a *rooted tree* having nullary operators as its leaves and unary or binary operators as its non-leaf vertices.

---

<sup>5</sup> <https://github.com/slizaa/slizaa-opencypher-xtext>, released under EPL v1.0.

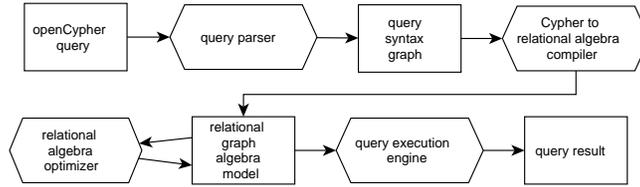


Fig. 2: Workflow of the query engine: compiler and execution engine.

*Nullary operators.* The `GetVertices` and `GetEdges` operators retrieve vertices and edges of the graph, respectively. `SingularObjectSource` and `DualObjectSource` emit the `Singular` and the `Dual` relation, respectively.

*Unary operators.* `Projection` and `Selection` work as given in Sec. 3.3. Exact semantics of the other unary operators are given in [24]. `DuplicateElimination`, `Grouping`, `Sort` and `Top` operators work like their corresponding SQL clauses.<sup>6</sup> `Expand` is a graph-specific operation to traverse one or a sequence of edges from a source to a given target vertex, while `AllDifferent` is specific to openCypher’s edge uniqueness semantics. `Unwind` is the inverse of the list-constructing `collect()` aggregation function.

*Binary operators.* The `Union` operator creates the set or multiset union of its inputs. `Join`, `LeftOuterJoin` and `AntiJoin` operators, based on the `joinVariable` list declared in `AbstractJoin` creates the natural join, antijoin and left outer join operations on their inputs, respectively as given in Sec. 3.3.

## 5 Elaboration

We have shown the overview of our approach in Sec. 4. In this section we present our approach in detail, driven by examples of the MoDeS3 system (Sec. 2).

### 5.1 Compilation of a Multipart Query

In openCypher, queries are composed as a sequence of query parts. Details are given in [24], but essentially a query part contains clauses up to the next `WITH` or `RETURN` clause and defines a result set of the attributes listed, which is then fed into the next query part as its input. For example, the query in List. 5.1 is composed of two query parts: first query part spans lines 1–3 and feeds its result set of the schema  $\langle s, countTrains \rangle$  into the second query part listed in line 4.

*Variable chaining* refers to the fact that attributes of the resulting schema are available in the subsequent query part, i.e. `s` and `countTrains` are available.

<sup>6</sup> In the order of appearance: `DISTINCT`, `GROUP BY`, `ORDER BY` and `SKIP ... LIMIT ...`

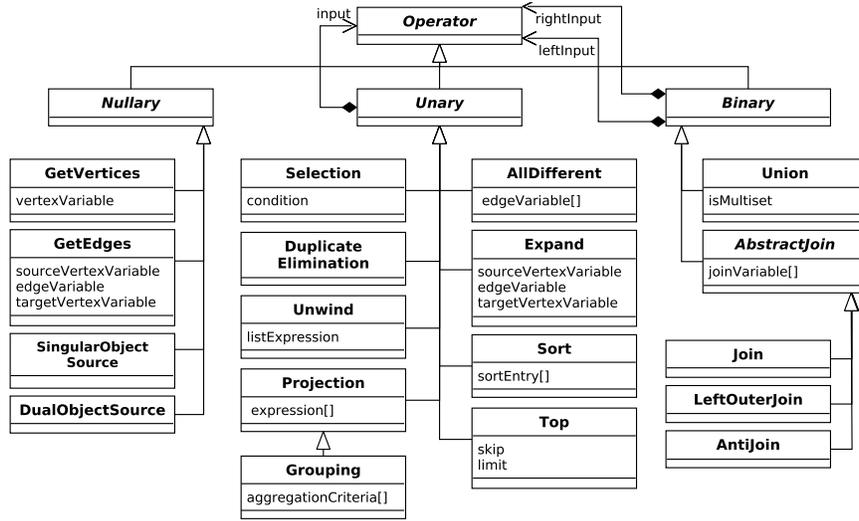


Fig. 3: Operator metamodel of the relational graph algebra.

```

1 MATCH (s:Station)-[:INCL]->(Element)<-[:ON]-(tr:Train)
2 WITH s, count(tr) AS countTrains
3 WHERE countTrains >= 2
4 RETURN s

```

List. 5.1: Busy station.

Compilation of each query part starts from the `Dual` relation. Each pattern given in `MATCH` clause is then compiled and joined to the previous patterns: for `MATCH` clauses we use the natural join operator and for `OPTIONAL MATCH`, we use left outer join. Possible projection, grouping and duplicate-elimination operators are appended above as required by the `WITH` or `RETURN` clauses.

Query parts are compiled one by one and combined together using the natural join operator as follows. The natural join is injected into the compiled form of the current query part just below the possible projection, grouping and duplicate-elimination operators populating its right input with the descendants. Its left input is the compiled form of the query parts processed so far.

Each query part that begins with a non-optional `MATCH` clause, like the first query part in List. 5.1 is joined with `Dual`. As the second query part has no patterns, its inputs are the first query part's result set and the `Dual`. Raw compiled form of this query is shown in Fig. 4(a), which contains two joins having `Dual`, its identity operand as one of its operands. Thus these natural join operations along with `Dual` should be removed which we implemented using a VIATRA graph transformation rule. Applying this transformation, we get the simplified form shown in Fig. 4(b).

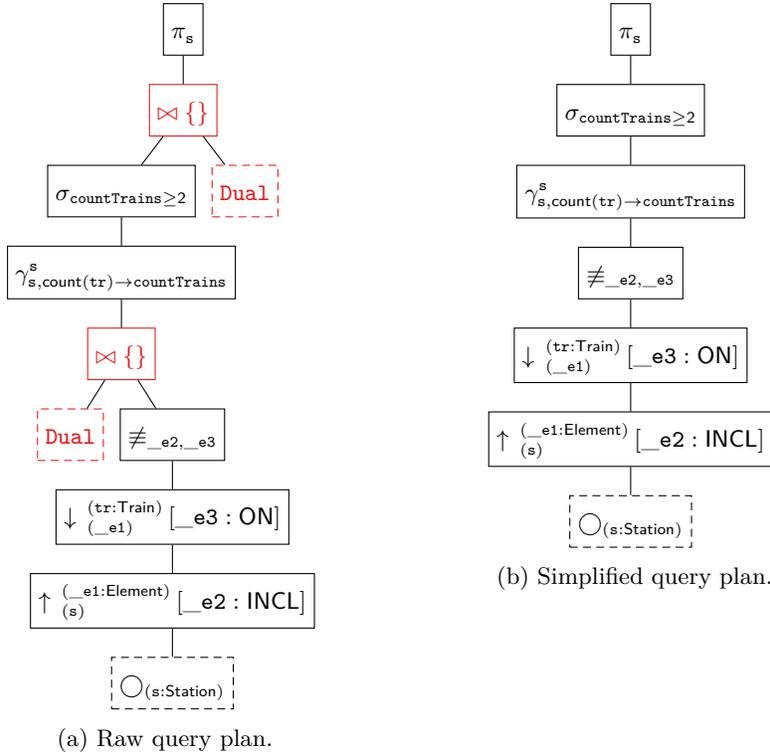


Fig. 4: Query plans for *Busy station*.

## 5.2 Compilation of Variable Length Path Patterns

Query in List. 5.2 features a variable length path pattern stating that two segments, `seg1` and `seg2` are connected through one to two edges of type `NEXT`. A variable length path pattern is compiled to the expand-both operator given in Sec. 3.3. Raw compiled form of this query is shown in Fig. 5(a), which is simplified to Fig. 5(b) using the transformation rule described in Sec. 5.1.

```

1 MATCH
2   (t1:Train)-[:ON]->(seg1:Element)-[:NEXT*1..2]-
3   (seg2:Element)<-[:ON]-(t2:Train)
4 RETURN t1, t2, seg1, seg2

```

List. 5.2: Close proximity.

## 5.3 Identifying Antijoin Operators

Query in List. 5.3 uses negative pattern match on line 2 to express that track element `re` does not have a train on it. This is essentially an antijoin operation.

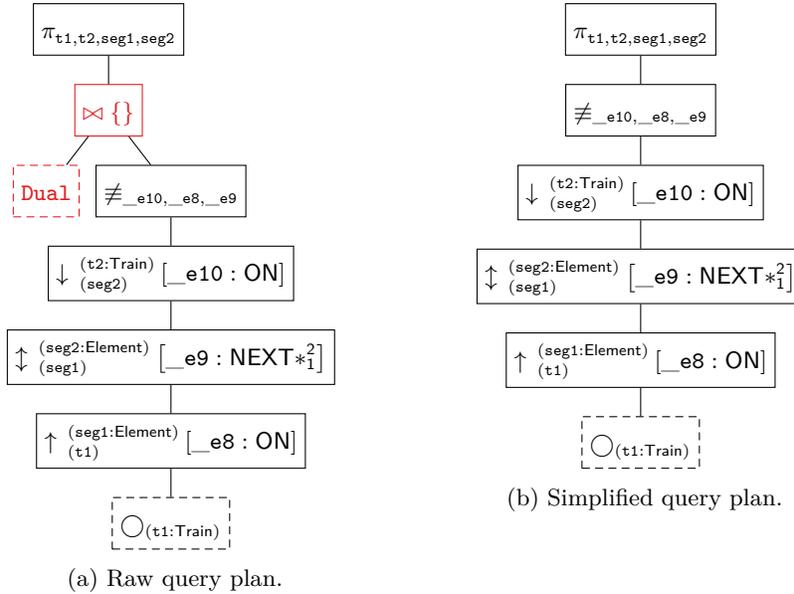


Fig. 5: Query plans for *Close proximity*.

In order to keep compiler simple, it is compiled in the raw form to left outer join of the two pattern matches and a negated selection stating that edge and vertex variables of the pattern condition are all not null. We highlighted using blue boxes the corresponding operator nodes in the raw compiled form of this query shown in Fig. 6(a). It is transformed by an other VIATRA rule to the antijoin operator also highlighted using blue in Fig. 6(b).

Simplification of this query again shows the removal of unused join (highlighted with red). Green box in Fig. 6(a) shows the all-different operator which states that the listed edge variables match unique edges. This is specified by openCypher's edge uniqueness semantics. As one edge is always unique, we added another transformation rule to remove this operator from the operation tree.

```

1 MATCH (s:Station)-[:INCL]->(re:Element)
2 WHERE NOT (re)<-[:ON]-(:Train)
3 RETURN DISTINCT s

```

List. 5.3: Station with free track.

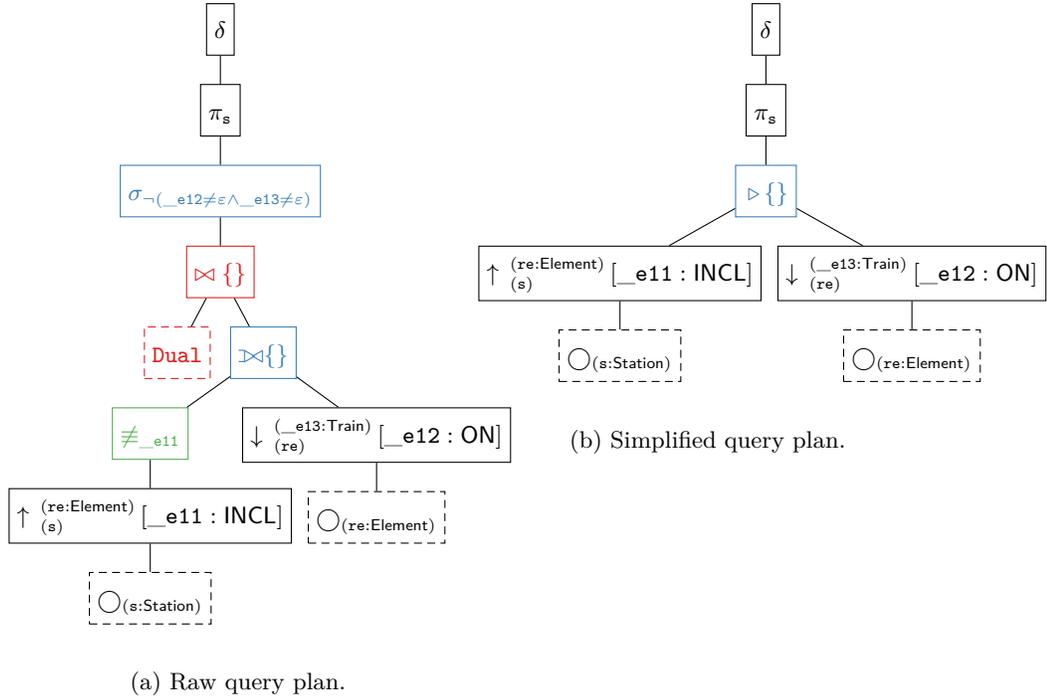


Fig. 6: Query plans for *Station with free track*.

## 6 Related Work

### 6.1 Industrial Graph Query Languages

As graph queries are increasingly used in industry, graph query languages are available across different technological spaces. Here, we discuss related query languages and compilers.

**Property graphs.** The Cypher language was originally designed as the primary query language of the Neo4j graph database system [25,30]. The grammar specification and the language behaviour of openCypher was defined to match those of Neo4j. Consequently, the compiler and query engine of Neo4j form the most complete openCypher implementation available, and is dual licensed (GPLv3/AGPLv3 for compatible projects and custom licensing for commercial applications).

The authors of [18] studied the Cypher query language and defined graph-specific relational algebra operators, such as get-vertices and expand-out (Sec. 3.3). While their work focused on optimisation, our work aims to provide a mapping and compilation steps for transforming openCypher to relational graph algebra.

In [19], graph queries were defined in a Cypher-like language and evaluated on Apache Flink and evaluated in the Apache Flink-based GRADOOP framework. However, formalisation and compilation of the queries was not discussed in detail.

**TinkerPop.** The TinkerPop framework aims to define a standard data model for property graphs. For graph queries, it provides the Gremlin Structure API, a low-level programming interface and Gremlin language, a high-level imperative graph traversal language [32]. The latter is implemented as a Groovy DSL [20].

**EMF.** The Eclipse Modeling Framework (Sec. 3.4) is an object-oriented modelling framework widely used in model-driven engineering. Henshin [3] provides a visual language for defining patterns, while Epsilon [21] and VIATRA Query [4] provide high-level declarative (textual) query languages, the Epsilon Pattern Language and the VIATRA Query Language (Sec. 3.4), respectively. VIATRA Query supports both incremental and search-based queries [9].

Similarly to our approach, the VIATRA compiler is also built using model-driven technologies (EMF, Xtend). The build process of VIATRA uses VIATRA to build itself, thus leading to the bootstrapping problem, well-known in the field of compiler construction [23]. Meanwhile, our openCypher compiler builds on the VIATRA system as a third-party component, so it avoids this problem.

**RDF.** Widely used in semantic technologies, SPARQL is a standardised declarative graph pattern language for querying RDF [47] 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. A formal definition of the language is given in [29]. Apache Jena ARQ [2] and Eclipse RDF4J [10] are open-source compilers and query engines for the SPARQL language.

**Comparing graph query engines.** To evaluate different graph query approaches, G. Szárnyas et al. designed the Train Benchmark, a framework for comparing graph query frameworks across different technological spaces, such as property graphs, EMF, RDF and SQL [39].

## 6.2 Query Compilation in Graph Transformation Systems

The authors of [7] adapted the Rete algorithm originally developed in the domain of production rule systems (see the following section for details) for pattern matching in a GT engine. The presented solution supported a simple core graph pattern language.

The Fujaba [27] graph transformation tool fixes a single, breadth-first traversal strategy at compile-time, using simple heuristics, e.g. that navigation along an edge with an at most one multiplicity constraint precedes navigations along edges with arbitrary multiplicity. PROGRES [35] uses a sophisticated cost model

for basic operations and generates the search plan at compile-time by a greedy algorithm.

An algorithm to produce a high-quality (e.g. compact) Rete network from a pattern specification was proposed in [44]. Paper [45] presented an algorithm to define efficient search plans on EMF models. These approaches are used in the eMoflon system [22].

The approach of [46] uses both metamodel- and instance model-level information to adaptively optimize graph queries based on statistical data collected from the current instance model. GrGen.NET provides a dynamic, runtime optimization engine, which uses a mix of heuristical and cost-based techniques [16].

The first VIATRA prototype, which was capable of generating Prolog code from metamodels and model transformations defined in XMI (XML Metadata Interchange) format, was presented in [42].

G. Szárnyas et al. designed the INCQUERY-D [38] system, an incremental graph query engine, built on top of the components of the VIATRA Query framework [43] (then known as EMF-INCQUERY [41]). INCQUERY-D reused the query parser and compiler of EMF-INCQUERY, but used a different query engine, tailored for scalable query evaluation and operating on RDF data sets.

## 7 Conclusion and Future Work

In this paper, we presented an approach to design and implement a query engine for the openCypher graph query language. We implemented this approach based on a language workbench built on EMF-based technologies, such as Xcore, Xtext, Xtend and VIATRA. The resulting prototype is part of the ingraph project, an openCypher-compatible incremental graph query engine.<sup>7</sup>

In the future, we plan to implement a query optimizer, which relies on *search-based optimization techniques using model transformations*, also known as *planning by rewriting* [1]. As our approach already utilizes the VIATRA query engine, the optimizer can be based on the VIATRA-DSE design-space exploration framework [17] without a significant integration overhead.

## Acknowledgements

This work was partially supported by the MTA-BME Lendület Research Group on Cyber-Physical Systems. We would like to thank János Magincz and Dávid Szakállas for their contributions to the relational graph algebra model. We are also grateful to András Vörös and Gábor Bergmann for their suggestions and comments on the draft of this paper.

## References

1. J. L. Ambite and C. A. Knoblock. Planning by rewriting. *J. Artif. Intell. Res.*, 15:207–261, 2001.

<sup>7</sup> Available at <http://docs.inf.mit.bme.hu/ingraph/>.

2. Apache Software Foundation. Apache Jena. <https://jena.apache.org/>.
3. T. Arendt et al. Henshin: Advanced concepts and tools for in-place EMF model transformations. In *MODELS*, pages 121–135, 2010.
4. G. Bergmann et al. Incremental evaluation of model queries over EMF models. In *MODELS*, pages 76–90, 2010.
5. B. Bryla and K. Loney. *Oracle Database 12C The Complete Reference*. McGraw-Hill Osborne Media, 1st edition, 2013.
6. Budapest University of Technology and Economics, Department of Measurement and Information Systems. Model-based Demonstrator for Smart and Safe Systems. <https://inf.mit.bme.hu/en/research/projects/modes3>, 2015.
7. H. Bunke, T. Glauser, and T.-H. Tran. An efficient implementation of graph grammars based on the RETE matching algorithm. In *Graph-Grammars and Their Application to Computer Science*, pages 174–189, 1990.
8. M. Búr et al. Distributed query-based runtime verification of critical cyber-physical systems. In *Runtime Verification*, 2017. Submitted. Available for reviewers at <http://docs.inf.mit.bme.hu/preprints/rv2017-modes3.pdf>.
9. M. Búr, Z. Ujhelyi, Á. Horváth, and D. Varró. Local search-based pattern matching features in EMF-IncQuery. In *ICGT*, pages 275–282, 2015.
10. Eclipse Foundation. RDF4J. <http://rdf4j.org/>.
11. Eclipse Foundation. Xtend – Modernized Java. <https://www.eclipse.org/xtend/>.
12. Eclipse Foundation. Xcore, 2017. <http://wiki.eclipse.org/Xcore>.
13. S. Erdweg et al. The state of the art in language workbenches - conclusions from the language workbench challenge. In *SLE*, pages 197–217, 2013.
14. M. Eysholdt and H. Behrens. Xtext: Implement your language faster than the quick and dirty way. In *SIGPLAN, SPLASH/OOPSLA*, pages 307–309, 2010.
15. H. Garcia-Molina, J. D. Ullman, and J. Widom. *Database systems – The complete book*. Pearson Education, 2nd edition, 2009.
16. R. Geiß, G. V. Batz, D. Grund, S. Hack, and A. Szalkowski. GrGen: A fast SPO-based graph rewriting tool. In *ICGT*, pages 383–397, 2006.
17. Á. Hegedüs, Á. Horváth, and D. Varró. A model-driven framework for guided design space exploration. *Autom. Softw. Eng.*, 22(3):399–436, 2015.
18. J. Hölsch and M. Grossniklaus. An algebra and equivalences to transform graph patterns in Neo4j. In *GraphQ at EDBT/ICDT*, 2016.
19. M. Junghanns et al. Cypher-based graph pattern matching in Gradoop. In *GRADES at SIGMOD*, 2017.
20. D. Koenig, A. Glover, P. King, G. Laforge, and J. Skeet. *Groovy in Action*. Manning Publications Co., Greenwich, CT, USA, 2007.
21. D. S. Kolovos et al. The Epsilon transformation language. In *ICMT*, pages 46–60, 2008.
22. E. Leblebici, A. Anjorin, and A. Schürr. Developing eMoflon with eMoflon. In *ICMT*, pages 138–145, 2014.
23. K. C. Loudon. *Compiler Construction: Principles and Practice*. PWS Publishing Co., Boston, MA, USA, 1997.
24. J. Marton, G. Szárnyas, and D. Varró. Formalising openCypher graph queries in relational algebra. In *ADBIS*, Lecture Notes in Computer Science. Springer, 2017.
25. Neo Technology. Neo4j. <http://neo4j.org/>.
26. Neo Technology. openCypher project. <http://www.opencypher.org/>, 2017.
27. U. Nickel, J. Niere, and A. Zündorf. The FUJABA environment. In *ICSE*, pages 742–745. ACM, 2000.

28. OrientDB LTD. OrientDB graph-document NoSQL DBMS. <http://www.orientdb.org/>.
29. J. Pérez et al. Semantics and complexity of SPARQL. *ACM TODS*, 34(3), 2009.
30. I. Robinson, J. Webber, and E. Eifré. *Graph Databases*. O'Reilly Media, 2nd edition, 2015.
31. M. A. Rodriguez. A collectively generated model of the world. In *Collective intelligence: creating a prosperous world at peace*, pages 261–264. 2008.
32. M. A. Rodriguez. The Gremlin graph traversal machine and language (invited talk). In *DBPL*, pages 1–10, 2015.
33. M. A. Rodriguez and P. Neubauer. Constructions from dots and lines. *Bulletin of the American Society for Information Science and Technology*, 36(6):35–41, 2010.
34. M. A. Rodriguez and P. Neubauer. The graph traversal pattern. In *Graph Data Management: Techniques and Applications*, pages 29–46. 2011.
35. A. Schürr, A. J. Winter, and A. Zündorf. Handbook of graph grammars and computing by graph transformation. pages 487–550. World Scientific Publishing Co., Inc., 1999.
36. A. Silberschatz, H. F. Korth, and S. Sudarshan. *Database System Concepts, 5th Edition*. McGraw-Hill Book Company, 2005.
37. D. Steinberg, F. Budinsky, M. Paternostro, and E. Merks. *EMF: Eclipse Modeling Framework 2.0*. Addison-Wesley Professional, 2nd edition, 2009.
38. G. Szárnyas et al. IncQuery-D: A distributed incremental model query framework in the cloud. In *MODELS*, pages 653–669, 2014.
39. G. Szárnyas et al. The Train Benchmark: Cross-technology performance evaluation of continuous model validation. *Softw. Syst. Model.*, 2017.
40. ThinkAurelius. Titan. <https://github.com/thinkaurelius/titan>.
41. Z. Ujhelyi et al. EMF-IncQuery: An integrated development environment for live model queries. *Sci. Comput. Program.*, 98:80–99, 2015.
42. D. Varró. Automated program generation for and by model transformation systems. In *AGT*, pages 161–174, 2002.
43. D. Varró et al. Road to a reactive and incremental model transformation platform: three generations of the VIATRA framework. *Softw. Syst. Model.*, 15(3):609–629, 2016.
44. G. Varró and F. Deckwerth. A rete network construction algorithm for incremental pattern matching. In *ICMT*, pages 125–140, 2013.
45. G. Varró et al. An algorithm for generating model-sensitive search plans for pattern matching on EMF models. *Softw. and Syst. Model.*, 14(2):597–621, 2015.
46. G. Varró, K. Friedl, and D. Varró. Adaptive graph pattern matching for model transformations using model-sensitive search plans. *Electronic Notes in Theoretical Computer Science*, 152:191–205, 2006.
47. W3C. Resource Description Framework. <https://www.w3.org/RDF/>, 2014.