

EXPRESSIVE PATH QUERIES ON GRAPHS WITH DATA *

PABLO BARCELÓ^a, GAELLE FONTAINE^b, AND ANTHONY WIDJAJA LIN^c

^{a,b} Center for Semantic Web Research & Department of Computer Science, University of Chile
e-mail address: {pbarcelo, gaelle}@dcc.uchile.cl

^c Yale-NUS College, Singapore
e-mail address: anthony.w.to@gmail.com

ABSTRACT. Graph data models have recently become popular owing to their applications, e.g., in social networks and the semantic web. Typical navigational query languages over graph databases — such as Conjunctive Regular Path Queries (CRPQs) — cannot express relevant properties of the interaction between the underlying data and the topology. Two languages have been recently proposed to overcome this problem: *walk logic* (WL) and *regular expressions with memory* (REM). In this paper, we begin by investigating fundamental properties of WL and REM, i.e., complexity of evaluation problems and expressive power. We first show that the data complexity of WL is nonelementary, which rules out its practicality. On the other hand, while REM has low data complexity, we point out that many natural data/topology properties of graphs expressible in WL cannot be expressed in REM. To this end, we propose *register logic*, an extension of REM, which we show to be able to express many natural graph properties expressible in WL, while at the same time preserving the elementariness of data complexity of REMs. It is also incomparable to WL in terms of expressive power.

1. INTRODUCTION

Graph databases have gained renewed interest due to applications, such as the semantic web, social network analysis, crime detection networks, software bug detection, biological networks, and others (e.g., see [1] for a survey). Despite the importance of querying graph

2012 ACM CCS: [Information systems]: Data management systems—Database design and models—Graph-based database models; Data management systems—Query languages—Query languages for non-relational engines; [Theory of computation]: Theory and algorithms for application domains—Database theory—Database query languages (principles); Theory and algorithms for application domains—Database theory—Logic and databases; Formal languages and automata theory—Regular languages.

Key words and phrases: graph databases; graph logics; RPQs; non elementary; register automata.

* This is the full version of the conference paper [3].

^{a,c} Barceló is funded by the Millennium Nucleus Center for Semantic Web Research under Grant NC120004 and Fontaine by Fondecyt postdoctoral grant 3130491. Part of this work done when Lin visited Barceló funded by Fondecyt grant 1130104. This work was also partially done when Lin was at Oxford University supported by EPSRC (H026878).

databases, no general agreement has been reached to date about the kind of features a practical query language for graph databases should support and about what can be considered a reasonable computational cost of query evaluation for the aforementioned applications.

Typical navigational query languages for graph databases — including the conjunctive regular path queries [7] and its many extensions [4] — suffer from a common drawback: they are well-suited for expressing relevant properties about the underlying topology of a graph database, i.e., about the way in which (labeled) nodes are connected via (labeled) edges, but not about how such topology interacts with the node ids or the data. This drawback is shared by common specification languages for verification [6] (e.g. CTL*), which are evaluated over a similar graph data model (a.k.a. transition systems). Examples of important queries that combine graph data and topology, but cannot be expressed in usual navigational languages for graph databases, include the following [8, 13]: (Q1) *Find pairs of people in a social network connected by professional links restricted to people of the same age.* (Q2) *Find pairs of cities x and y in a transportation system, such that y can be reached from x using only services operated by the same company.* In each one of these queries, the connectivity between two nodes (i.e., the topology) is constrained by the data (from an infinite domain, e.g., \mathbb{N}), in the sense that we only consider paths in which all intermediate nodes satisfy a certain condition (e.g. they are people of the same age).

Two languages, *walk logic* and *regular expressions with memory*, have recently been proposed to overcome this problem. These languages have different goals:

(a) Walk logic (WL) was proposed by Hellings et al. [8] as a unifying framework for understanding the expressive power of path queries over graph databases. Its strength is on the expressiveness side. The underlying data model of WL is that of (node or edge)-labeled directed graphs. In this context, WL can be seen as a natural extension of FO with path quantification, plus the ability to check whether positions p and p' in paths π and π' , respectively, have the same data values. In their paper, Hellings et al. assume the restriction that each node carries a distinct data value (and, therefore, that this data value serves as an identifier for the node). However, as we shall see, this makes no difference in terms of the results that we can obtain.

(b) Regular expressions with memory (REMs) were proposed by Libkin and Vrgoč [10] as a formalism for comparing data values along a single path, while retaining a reasonable complexity for query evaluation. The strength of this language is on the side of efficiency. The data model of the class of REMs is that of edge-labeled directed graphs, in which each node is assigned a data value from an infinite domain. REMs define pairs of nodes in the graph database that are linked by a path satisfying a given condition c . Each such condition c is defined in a formalism inspired by the class of *register automata* [9], allowing some data values to be stored in the registers and then compared against other data values. The evaluation problem for REMs is PSPACE-complete (same as for FO over relational databases), and can be solved in polynomial time in *data complexity* [10], i.e., assuming queries to be fixed.¹ This shows that the language is, in fact, well-behaved in terms of the complexity of query evaluation.

The aim of this paper is to investigate the expressiveness and complexity of query evaluation for WL and the class of REMs with the hope of finding a navigational query

¹Recall that data complexity is a reasonable measure of complexity in the database scenario [17], since queries are often much smaller than the underlying data.

language for data graphs that strikes a good balance between these two important aspects of query languages.

Contributions. We start by considering WL, which is known to be a powerful formalism in terms of expressiveness. Little is known about the cost of query evaluation for this language, save for the decidability of the evaluation problem and NP-hardness of its data complexity. Our first main contribution is to pinpoint the exact complexity of the evaluation problem for WL (and thus answering an open problem from [8]): we prove that it is non-elementary, and that this holds even in data complexity, which rules out the practicality of the language.

We thus move to the class of REMs, which suffers from the opposite drawback: Although the complexity of evaluation for queries in this class is reasonable, the expressiveness of the language is too rudimentary for expressing some important path properties due to its inability to (i) compare data values in *different* paths and (ii) express branching properties of the graph database. An example of an interesting query that is not expressible as an REM is the following: (Q) *Find pairs of nodes x and y , such that there is a node z and a path π from x to y in which each node is connected to z .* Notice that this is the query that lies at the basis of the queries (Q1) and (Q2) we presented before.

Our second contribution then is to identify a natural extension of this language, called *register logic* (RL), that closes REMs under Boolean combinations and existential quantification over nodes, paths and register assignments. The latter allows the logic to express comparisons of data values appearing in different paths, as well as branching properties of the data. This logic is incomparable in expressive power to WL. Besides, many natural queries relating data and topology in data graphs can be expressed in RL including: the query (Q), hamiltonicity, the existence of an Eulerian trail, bipartiteness, and connected graphs with an even number of nodes. We then study the complexity of the problem of query evaluation for RL, and show that it can be solved in elementary time (in particular, that it is EXPSpace-complete). This is in contrast to WL, for which even the data complexity is non-elementary. With respect to data complexity, we prove that RL is PSPACE-complete. We then identify a slight extension of its existential-positive fragment, which is tractable (NLOGSPACE) in data complexity and can express many queries of interest (including the query (Q)). The idea behind this extension is that atomic REMs can be enriched with an existential branching operator – in the style of the class of *nested regular expressions* [5] – that increases expressiveness without affecting the cost of evaluation.

Organization of the paper. Section 2 defines our data model. In Section 3, we briefly recall the definition of walk logic and some basic results from [8]. In Section 4, we prove that the data complexity of WL is nonelementary. Section 5 contains our results concerning register logic. We conclude in Section 6 with future work.

2. THE DATA MODEL

We start with a definition of our data model: data graphs.

Definition 2.1 (Data graph). Let Σ be a finite alphabet. A data graph G over Σ is a tuple (V, E, κ) , where V is the finite set of nodes, $E \subseteq V \times \Sigma \times V$ is the set of directed edges labeled in Σ (that is, each triple $(v, a, v') \in E$ is to be seen as an edge from v to v' in G labeled a), and $\kappa : V \rightarrow \mathcal{D}$ is a function that assigns a data value in \mathcal{D} to each node in V .

This is the data model adopted by Libkin and Vrgoč [10] in their definition of REMs. In the case of WL [8], the authors adopted *graph databases* as their data model, i.e., data

graphs $G = (V, E, \kappa)$ such that κ is injective (i.e. each node carries a different data value). In such a case we can think of $\kappa(v)$ as the identifier (id) of v , for each $v \in V$. We shall adopt the general model of [10] since none of our complexity results are affected by the data model: upper bounds hold for data graphs, while all lower bounds are proved in the more restrictive setting of graph databases. However, for the sake of the comparison with the expressiveness of WL, many of our examples are constructed in the scenario of graph databases, that is, when $\kappa(v)$ serves as an id for node v .

There is also the issue of edge-labeled vs node-labeled data graphs. Our data model is edge-labeled, but the original one for WL is node-labeled [8]. We have chosen to use the former because it is the standard in the literature [2]. Again, this choice is inessential, since all the complexity results we present in the paper remains true if the logics are interpreted over node-labeled graph databases or data graphs (applying the expected modifications to the syntax).

Finally, in several of our examples we use logical formulas to express properties of undirected graphs. In each such case we assume that an undirected graph H is represented as a graph database $G = (V, E, \kappa)$ over unary alphabet $\Sigma = \{a\}$, where V is the set of nodes of H and E is a symmetric relation (i.e. $(v, a, v') \in E$ iff $(v', a, v) \in E$). In particular, since $G = (V, E, \kappa)$ is a graph database we have that κ is injective, i.e., each node is uniquely determined by its data value.

3. WALK LOGIC

WL is an elegant and powerful formalism for defining properties of paths in graph databases, which was originally proposed in [8] as a yardstick for measuring the expressiveness of different path logics.

The syntax of WL is defined with respect to countably infinite sets Π of *path variables* (that we denote as π, π_1, π_2, \dots) and $\mathcal{T}(\pi)$, for each $\pi \in \Pi$, of *position variables* of sort π . We assume that different sorts are associated with distinct position variables. We denote position variables by t, t_1, t_2, \dots , and write t^π when we need to emphasize that position variable t is of sort π .

Definition 3.1 (Walk logic (WL)). The set of formulas of WL over finite alphabet Σ is defined by the following grammar, where (i) $a \in \Sigma$, (ii) t, t_1, t_2 are position variables of any sort, (iii) π is a path variable, and (iv) t_1^π, t_2^π are position variables of the same sort π :

$$\phi, \phi' := E_a(t_1^\pi, t_2^\pi) \mid t_1^\pi < t_2^\pi \mid t_1 \sim t_2 \mid \neg\phi \mid \phi \vee \phi' \mid \exists t\phi \mid \exists\pi\phi$$

As usual, WL formulas without free variables are called *Boolean*. □

To define the semantics of WL we need to introduce some terminology. A *path* (a.k.a. *walk* in [8]) in the data graph $G = (V, E, \kappa)$ is a finite, nonempty sequence

$$\rho = v_1 a_1 v_2 \cdots v_{n-1} a_{n-1} v_n,$$

such that $(v_i, a_i, v_{i+1}) \in E$ for each $1 \leq i < n$. The set of *positions* of ρ is $\{1, \dots, n\}$, and v_i is the node in position i of ρ , for $1 \leq i \leq n$. The intuition behind the semantics of WL formulas is as follows. Each path variable π is interpreted as a path $\rho = v_1 a_1 v_2 \cdots v_{n-1} a_{n-1} v_n$ in the data graph G , while each position variable t of sort π is interpreted as a position $1 \leq i \leq n$ in ρ (that is, position variables of sort π are interpreted as positions in the path that interprets π). The atomic formula $E_a(t_1^\pi, t_2^\pi)$ is true iff π is interpreted as path $\rho = v_1 a_1 v_2 \cdots v_{n-1} a_{n-1} v_n$, the position p_2 that interprets t_2 in ρ is the successor of the

position p_1 that interprets t_1 (i.e. $p_2 = p_1 + 1$), and node in position p_1 is linked in ρ by an a -labeled edge to node in position p_2 (that is, $a_{p_1} = a$). In the same way, $t_1^\pi < t_2^\pi$ holds iff in the path ρ that interprets π the position that interprets t_1 is smaller than the one that interprets t_2 . Furthermore, $t_1 \sim t_2$ is the case iff the data value carried by the node in the position assigned to t_1 is the same than the data value carried by the node in the position assigned to t_2 (possibly in different paths). We formalize the semantics of WL below.

Let $G = (V, E, \kappa)$ be a data graph and ϕ a WL formula. Assume that \mathcal{S}_ϕ is the set that consists of (i) all position variables t^π and path variables π such that t^π is a free variable of ϕ , and (ii) all path variables π such that π is a free variable of ϕ . Intuitively, \mathcal{S}_ϕ defines the set of (both path and position) variables that are relevant to define the semantics of ϕ over G . An *assignment* α for ϕ over G is a mapping that associates a path $\rho = v_1 a_1 v_2 \cdots v_{n-1} a_{n-1} v_n$ in G with each path variable $\pi \in \mathcal{S}_\phi$, and a position $1 \leq i \leq n$ with each position variable of the form t^π in \mathcal{S}_ϕ (notice that this is well-defined since $\pi \in \mathcal{S}_\phi$ every time a position variable of the form t^π is in \mathcal{S}_ϕ). As usual, we denote by $\alpha[t \rightarrow i]$ and $\alpha[\pi \rightarrow \rho]$ the assignments that are equal to α except that t is now assigned position i and π the path ρ , respectively.

We say that G *satisfies* ϕ *under* α , denoted $(G, \alpha) \models \phi$, if one of the following holds (we omit Boolean combinations which are standard):

- $\phi = E_a(t_1^\pi, t_2^\pi)$, the path $\alpha(\pi)$ is $v_1 a_1 v_2 \cdots v_{n-1} a_{n-1} v_n$, and it is the case that $\alpha(t_2^\pi) = \alpha(t_1^\pi) + 1$ and $a = a_{\alpha(t_1^\pi)}$.
- $\phi = t_1^\pi < t_2^\pi$ and $\alpha(t_1^\pi) < \alpha(t_2^\pi)$.
- $\phi = (t_1 \sim t_2)$, t_1 is of sort π_1 , t_2 is of sort π_2 , and $\kappa(v_1) = \kappa(v_2)$, where v_i is the node in position $\alpha(t_i)$ of $\alpha(\pi_i)$, for $i = 1, 2$.
- $\phi = \exists t^\pi \psi$ and one of the following holds:
 - (1) t^π does not appear free in ψ , or
 - (2) both t^π and π appear free in ψ , and there is a position i in $\alpha(\pi)$ such that $(G, \alpha[t^\pi \rightarrow i]) \models \psi$, or
 - (3) t_π appears free in ψ , π does not appear free in ψ , and there is a path ρ in G and a position i in ρ such that $(G, \alpha[\pi \rightarrow \rho, t^\pi \rightarrow i]) \models \psi$.
- $\phi = \exists \pi \psi$ and the following holds:
 - (1) π does not appear free in ψ , or
 - (2) there is a path ρ in G such that $(G, \alpha[\pi \rightarrow \rho]) \models \psi$.

Example 3.1. A simple example from [8] that shows that WL expresses NP-complete properties is the following query that checks if a graph G has a Hamiltonian path:

$$\exists \pi \left(\forall t_1^\pi \forall t_2^\pi (t_1^\pi \neq t_2^\pi \rightarrow t_1^\pi \not\sim t_2^\pi) \wedge \forall \pi' \forall t_1^{\pi'} \exists t_2^{\pi'} (t_1^{\pi'} \sim t_2^{\pi'}) \right).$$

In fact, this query expresses that there is a path π in G that does not repeat nodes (because π satisfies $\forall t_1^\pi \forall t_2^\pi (t_1^\pi \neq t_2^\pi \rightarrow t_1^\pi \not\sim t_2^\pi)$), and every node belongs to such path (because π satisfies $\forall \pi' \forall t_1^{\pi'} \exists t_2^{\pi'} (t_1^{\pi'} \sim t_2^{\pi'})$), and, therefore, every node that occurs in some path π' in the graph database also occurs in π). Note that this formula uses in an essential way the fact that G is a graph database, i.e., that each node is uniquely identified by its data value. \square

4. WL EVALUATION IS NON-ELEMENTARY IN DATA COMPLEXITY

In this section we pinpoint the precise complexity of query evaluation for WL. It was proven in [8] that this problem is decidable. Although the precise complexity of this problem was left open in [8], one can prove that this is, in fact, a non-elementary problem by an easy

translation from the satisfiability problem for FO formulas – which is known to be non-elementary [15, 16]. In databases, however, one is often interested in a different measure of complexity – called *data complexity* [17] – that assumes the formula ϕ to be fixed. This is a reasonable assumption since databases are usually much bigger than formulas. Often in the setting of data complexity the cost of evaluating queries is much smaller than in the general setting in which formulas are part of the input. The main result of this section is that the data complexity of evaluating WL formulas is nonelementary even *over graph databases*, which rules out its practicality.

Let ϕ be a WL formula without free variables. The evaluation problem for ϕ , denoted $\text{EVAL}(\text{WL}, \phi)$, is defined as follows: Given a data graph G , is it the case that $G \models \phi$? We prove the following:

Theorem 4.1. The evaluation problem for WL is non-elementary in data complexity. In particular, for each $k \in \mathbb{Z}_{>0}$, there is a finite alphabet Σ and a Boolean formula ϕ over Σ , such that the problem $\text{EVAL}(\text{WL}, \phi)$ of evaluating the WL formula ϕ is k -EXPSpace-hard. In addition, the latter holds even if the input is restricted to the class of graph databases.

We prove the above result by showing that for all natural numbers k , the data complexity of the model checking problem for WL is k -EXPSpace-hard. For all natural numbers k and f_0 , we provide a reduction to the class of problems solvable by a Turing machine using a tape of size $\text{tower}(k, f_0 n)$ given an input word of size n , where $\text{tower}(1, n) := 2^n$ and $\text{tower}(k + 1, n) = 2^{\text{tower}(k, n)}$.

More precisely, for all natural numbers $k > 0$, there is a Turing machine M and a constant f_0 such that the following problem is k -EXPSpace-hard: given a word w of size n , is there an accepting run of M over w using at most $\text{tower}(k, f_0 n)$ cells? We prove that there is a formula $\phi \in \text{WL}$ such that for all words w of size n , there is a graph G_w such that

$$G_w \models \phi \quad \text{iff} \quad \text{there is an accepting run of } M \text{ over } w \text{ using at most } \text{tower}(k, f_0 n) \text{ cells.} \quad (4.1)$$

Before giving a proof, we sketch the case $k = 1$ here, which illustrates the proof idea. Let M be a Turing machine M such that the following problem is EXPSpace-hard: given a word w of size n , is there an accepting run of M over w using at most $2^{f_0 n}$ cells? The formula ϕ that we will define and satisfying equivalence (4.1) is of the form

$$\exists \pi \psi(\pi),$$


where ψ is a formula that does not contain any quantification over path variables. Given a word w of size n , the label of the path π in the graph G_w will encode an accepting run of M over the word w in the following way.

Given a word w of size n , consider a configuration C of the run of M over w where the head is scanning the cell number i_0 , the machine is in state q and the content of the tape is the word $w' = w'_0 \dots w'_j$ ($j = 2^{f_0 n} - 1$). We may encode the configuration C by the word $e_C = d_0^C \dots d_j^C$ where each d_i^C encodes the information in cell number i and $j = 2^{f_0 n} - 1$. More precisely, we define d_i^C as a word of the form

$$c(i) (q'_i, w'_i), \quad (4.2)$$

where $c(i)$ and q'_i are defined as follows. The word $c(i)$ is the binary encoding of the number i . The letter w'_i is the content of the cell i . The letter q'_i is equal to the dummy symbol $\$$ if the head is not scanning the cell number i ; otherwise, q'_i is equal to the state q . That is, $q'_{i_0} = q$ and for all $i \neq i_0$, $q'_i = \$$. We encode a run $C_0 C_1 \dots$ as the sequence $e_{C_0} e_{C_1} \dots$.

We think of a path π encoding a run as consisting of two parts: the first part contains the encoding e_{C_0} of the initial configuration and is a path through a subgraph I_w of G_w , while the second part contains the encoding $e_{C_1}e_{C_2}\dots$ and is a path through the subgraph H of G_w . If Q is the set of states of M and Σ is the alphabet, we define H as the following graph



where l is equal to $|(Q \cup \{\$\}) \times \Sigma|$, $\{d_i : 1 \leq i \leq l\} = (Q \cup \{\$\}) \times \Sigma$ and the number of nodes with outgoing edges with labels 0 and 1 is equal to f_0n . The label of a path π' from the “left-most” node x to the “right-most” node z with only once occurrence of x is exactly the description of a cell in a configuration: it is the binary encoding of a natural number $< 2^{f_0n}$ followed by a pair of the form (q', a) . We can define a formula $\phi_C \in \text{WL}$ such that for all paths π starting in x and ending in z ,

$$H \models \phi_C(\pi) \quad \text{iff} \quad \text{the label of } \pi \text{ is the encoding of a configuration.}$$

We do not give details; ϕ_C has to express that the encoding of a configuration only has one tape head, that the first number encoded in binary is 0, that the last number is $2^{f_0n} - 1$ and that the encoding of the description of cell number j is followed by the description of cell number $j + 1$. Using the formula ϕ_C , we can define a formula ϕ_1 such that for all paths π ,

$$H \models \phi_1(\pi) \quad \text{iff} \quad \text{the label of } \pi \text{ is the encoding of an accepting run.}$$

The formula ϕ_1 has to ensure that if $e_C e_{C'}$ occurs in the label of π , then C and C' are consecutive configurations according to M . Moreover, ϕ_1 has to express that eventually we reach the final state. In order to express ϕ_C and ϕ_1 , we use the ability of WL to check whether two positions correspond to the same node. For example, in order to define ϕ_1 , since we need to compare consecutive configurations e_C and $e_{C'}$, we need to be able to compare the content of a cell in configuration C and the content of that same cell in C' . In particular, we want to be able to express whether two subpaths π'_0 and π'_1 of π starting in x and ending in y correspond to the binary encoding of the same number. Since the length of such subpaths depends on n , we cannot check node by node whether the two subpaths are equal. However, it is sufficient to check that if $t_0^{\pi'_0}$ and $t_1^{\pi'_1}$ correspond to the same node ($t_0^{\pi'_0} \sim t_1^{\pi'_1}$), then their successors also correspond to the same node ($t_0^{\pi'_0} + 1 \sim t_1^{\pi'_1} + 1$). Note that using the facts that π'_0 and π'_1 are subpaths of π , we will be able to define ϕ_1 such that it only contains quantifications over node variables (and no quantifications over path variables). Similarly, in the formula ϕ_C , we use the operator \sim in order to express that two subpaths correspond to the binary encodings of numbers that are successors of each other.

Similarly to the way we define the graph H , we can introduce a graph I_w and a formula $\phi_0(\pi)$ such that

$$I_w \models \phi_0(\pi) \quad \text{iff} \quad \text{the label of } \pi \text{ is the encoding } e_{C_0},$$

where C_0 is the initial configuration of the run of M over w . By adding an edge from I_w to H , we construct a graph G_w such that for all paths π , $G_w \models \phi_0(\pi) \wedge \phi_1(\pi)$ iff the label of

π is the encoding of an accepting run over w . Hence, the formula $\phi := \exists \pi(\phi_0(\pi) \wedge \phi_1(\pi))$ satisfies (4.1).

For the case where $k > 1$, the problem to adapt the above proof is that we have to consider Turing machine configurations whose size is bounded by a tower of exponentials of height k . If $k > 1$, the binary representation of such a bound is not polynomial. The trick is to represent such exponential towers by k -counters. A 1-counter is the binary representation of a number. If $k > 1$, a k -counter is a word $\sigma_0 l_0 \dots \sigma_j l_j$, where l_j is a $(k-1)$ -counter and $\sigma_j \in \{0, 1\}$.

Definition. For all natural numbers k , we consider the alphabet $\Sigma_k = \{a_k, b_k\}$, where a_k and b_k represent 0 and 1 respectively. We define Γ_k as the alphabet $\Sigma_1 \cup \dots \cup \Sigma_k$.

A 1-counter of length n is a sequence of the form

$$l_0 \dots l_{f_0 n - 1},$$

where for all $0 \leq i < f_0 n$, $l_i \in \Sigma_1$. This 1-counter represents the number $\sum_{i=0}^{f_0 n - 1} l_i 2^i$. Recall that if l_i is equal to a_1 (resp. b_1), then l_i represents 0 (resp. 1).

If $k \geq 2$, a k -counter of length n is a sequence of the form

$$\sigma_0 l_0 \dots \sigma_j l_j,$$

where for all $0 \leq i \leq j$, $l_i \in \Sigma_k$, σ_i is a $(k-1)$ -counter representing the number i and $j = \text{tower}(k-1, f_0 n) - 1$. This k -counter represents the number $\sum_{i=0}^j l_i 2^i$. Again recall that if l_i is equal to a_1 (resp. b_1), then l_i represents 0 (resp. 1).

A $(k, f_0 n, p)$ -description (over an alphabet Δ) is a sequence

$$\sigma_p d_p \dots \sigma_j d_j,$$

where for all $p \leq i \leq j$, $d_i \in \Delta$, σ_i is a $(k-1)$ -counter representing the number i and $j = \text{tower}(k, f_0(n-1)) - 1$. A $(f_0 k, n)$ -description (over an alphabet Δ) is a $(k, f_0 n, 0)$ -description. \square

Note that a $(k, f_0 n)$ -description over the alphabet Σ_k is a k -counter of length n . If Δ is the alphabet $(Q \cup \{\$\}) \times \Sigma$ (where Q is the set of states and Σ is the alphabet of the machine), a $(k, f_0 n)$ -description over Δ is of the form

$$l_0(x_0, y_0) \dots l_j(x_j, y_j)$$

where $j = \text{tower}(k, f_0 n) - 1$. Hence, if we define $c(i)$ in (4.2) as the k -counter encoding the number i , the encoding of a configuration (as defined above) is nothing but a $(k, f_0 n)$ -description.

In particular, if we want to encode a run as the label of a path satisfying some well-chosen formula in a well-chosen graph, we should also be able to encode $(k, f_0 n, p)$ -descriptions as labels of paths. We show how to do so in the following lemma.

Notation. Given a path π in a graph over an alphabet Δ , we denote by $l(\pi)$ the label of π . Given an alphabet $\Delta' \subseteq \Delta$, we denote by $l_{\Delta'}(\pi)$ the trace of $l(\pi)$ over the alphabet Δ' , that is, the subsequence of $l(\pi)$ obtained by deleting the letters that do not belong to Δ' .

Let $G' = (V', E', \kappa')$ be a subgraph of $G = (V, E, \kappa)$ and let π be a path in G and of the form

$$v_1 a_1 v_2 \dots v_{n-1} a_n v_n,$$

where $(v_i, a_i, v_{i+1}) \in E$ for all $1 \leq i < n$. Assume that there are i_0 and i_1 such that $i_0 \leq i_1$ and

$$\{v_i : v_i \in V', 1 \leq i \leq n\} = \{v_{i_0}, \dots, v_{i_1}\},$$

that is, once the path leaves G' , it never goes back to G' . Then we define the *trace* π' of π on G' as the subpath

$$v_{i_0} a_{i_0} v_{i_0+1} \dots v_{i_1-1} a_{i_1} v_{i_1},$$

that is, π' is the longest subpath of ϕ with nodes in G' .

In order to make notation easier, we also abbreviate the formula

$$\exists s E_a(s, t)$$

by $a(t)$.

Given a formula $\phi(\pi, s, t)$ with path variable π and node variables s and t , we denote by $\phi(\pi_{s,t})$ the formula obtained by replacing in $\phi(\pi, s, t)$ each quantification of the form

$$\exists r^\pi$$

by

$$\exists r^\pi \text{ s.t. } (s < r < t).$$

Intuitively, we “restrict” the path π to the nodes occurring between s and t .

Lemma 4.2. For all n and k and for all alphabets Δ , there are formulas $\phi_{k,n,p}^\Delta(\pi)$ ($0 \leq p < n$) and a graph $G_{k,n}^\Delta$ satisfying the following. There is a unique node with an outgoing (resp. incoming) edge with label $i_{k,n}^\Delta$ (resp. $f_{k,n}^\Delta$); moreover, that node has no incoming (resp. outgoing) edge. That node is called the *initial* (resp. *final*) node. Finally, $G_{k,n}^\Delta \models \phi_{k,n,p}^\Delta(\pi)$ iff the label $l(\pi)$ of π satisfies the following conditions:


- only the first edge of π is labeled $i_{k,n}^\Delta$,
- only the last edge of π is labeled $f_{k,n}^\Delta$,
- if $k \geq 2$ and $\Delta' = \Delta \cup \Gamma_{k-1}$, then $l_{\Delta'}(\pi)$ is a $(k, f_0 n, p)$ -description over Δ ;
- if $k = 1$, $l_{\Sigma_1}(\pi)$ is a 1-counter of length n .

We let $\phi_{k,n}^\Delta(\pi)$ be an abbreviation for $\phi_{k,n,0}^\Delta(\pi)$.

Moreover, if $\Delta = \Sigma_k$, then there are formulas $succ_{k,n}(\pi, \pi')$, $number_{k,n}^i$ ($1 \leq i \leq n$), $last_{k,n}$ and $eq_{k,n}(\pi, \pi')$ such that for all paths π and π' satisfying $G_{k,n}^\Delta \models \phi_{k,n}^\Delta(\pi) \wedge \phi_{k,n}^\Delta(\pi')$, we have

- $G_{k,n}^\Delta \models succ_{k,n}(\pi, \pi')$ iff the number encoded by $l_{\Gamma_k}(\pi')$ is the successor of the number encoded by $l_{\Gamma_k}(\pi)$.
- $G_{k,n}^\Delta \models number_{k,n}^i(\pi)$ iff $l_{\Gamma_k}(\pi)$ is the encoding of the number i .
- $G_{k,n}^\Delta \models last_{k,n}(\pi)$ iff $l_{\Gamma_k}(\pi)$ is the encoding of the number $tower(k, f_0 n)$.
- $G_{k,n}^\Delta \models eq_{k,n}(\pi, \pi')$ iff the number encoded by $l_{\Gamma_k}(\pi')$ is equal to the number encoded by $l_{\Gamma_k}(\pi)$.

Proof. The formulas and the graph are defined by induction on k . Suppose first that $k = 1$ and $\Delta = \Sigma_1$. We define G_0 as the following graph



where the number of nodes with outgoing edges with labels a_1 and b_1 , is equal to $f_0 n$. The label N is an additional label that we introduce in order to simplify the notation in the formulas.

We let $G_{1,n}^\Delta$ be the graph G_0 . We define now the formula $\phi_{1,n}^\Delta$. In fact, any path π over G_0 starting with the node with no incoming edge and ending with the node with no outgoing edge, will be such that $l_{\Sigma_1}(\pi)$ is the encoding of a 1-counter. Hence, we can define $\phi_{1,n}^\Delta$ as the conjunction of the formula

$$\exists s^\pi [\neg \exists t^\pi, t < s]$$

and the formula

$$\exists s^\pi [\neg \exists t^\pi, s < t].$$

We show now how to define the formulas $num_{1,n}^i(\pi)$ (by induction on i), $eq_{k,n}(\pi, \pi')$ and $last_{1,n}(\pi)$. For the formula $last_{1,n}(\pi)$, a path π corresponds to the encoding of the number $2^{f_0 n} - 1$ iff we always choose the node with label b_1 . Or equivalently, if we never choose the node with label a_1 . Hence, we may define $last_{1,n}(\pi)$ as the formula

$$\neg \exists s^\pi, a_1(s).$$

For the formula $eq_{1,n}(\pi, \pi')$, two paths π and π' correspond to the same number iff π and π' are equal. Since π and π' are simple paths with the same starting node, this is equivalent over graph databases (where each node carries a different data value) to the fact the the following formula holds

$$\forall t^\pi, (t')^{\pi'} [t \sim t' \rightarrow (t + 1) \sim (t' + 1)].$$

The formulas $num_{1,n}^i(\pi)$ is defined by induction on i . If $i = 0$, the path π encodes the number 0 iff we always choose the node with label a_1 . Or equivalently, if we never choose the node with label b_1 , which is expressed by

$$\neg \exists s^\pi, b_1(s).$$

For the induction case, the path π encodes the number $i + 1$ iff there is a path π'' encoding the number i and the number encoded by π is the successor of the number encoded by π'' . Hence, we can define $num_{1,n}^{i+1}(\pi)$ as the formula

$$\exists \pi'' (num_{1,n}^i(\pi'') \wedge succ_{1,n}(\pi'', \pi)).$$

In order to finish the base case, it remains to define the formula $succ_{1,n}(\pi, \pi')$. Basically, we have to simulate addition in binary. If $x_1 \dots x_{f_0 n}$ is the binary encoding of a number $i < 2^{f_0 n} - 1$, then the binary encoding of the number $i + 1$ is the sequence $x'_1 \dots x'_{f_0 n}$ such that x_m is equal to


- (a) 1 if $x_m = 0$ and all the elements $x_{m+1}, \dots, x_{f_0 n}$ are equal to 1,
- (b) 0 if $x_m = 1$ and all the elements $x_{m+1}, \dots, x_{f_0 n}$ are equal to 1,
- (c) 0 if $x_m = 0$ and there is an element in the sequence $x_{m+1} \dots x_{f_0 n}$ that is equal to 0,
- (d) 1 if $x_m = 1$ and there is an element in the sequence $x_{m+1} \dots x_{f_0 n}$ that is equal to 0.

Case (a) can be expressed by the following formula

$$\forall t^\pi, (t')^{\pi'} [t \sim t' \wedge a_1(t) \wedge \forall s \in \pi [(t < s) \wedge N(s - 1) \rightarrow b_1(s)] \rightarrow b_1(t').$$

The other cases can be treated similarly. This finishes the base case.

We turn now to the induction step. If $\Delta = \{d_1, \dots, d_l\}$, we define $G_{k+1,n}^\Delta$ as the following graph



The edge with label $i_{k+1,n}^\Delta$ and the edge with label Δ_f^{k+1} are pointing to the initial node in $G_{k,n}^{\Sigma_k}$. The edge with label Δ^{k+1} is an edge starting from the final node in $G_{k,n}^{\Sigma_k}$.

We define now the formula $\phi_{k+1,n,p}^\Delta(\pi)$. The intuition is as follows. We encode a $(k+1, f_0n, p)$ -description

$$\sigma_p d_p \dots \sigma_j d_j,$$

as a path π starting with the edge with label $i_{k+1,n}^\Delta$ and ending with the edge with label $f_{k+1,n}^\Delta$. Each k -counter σ_i will correspond to a path through the subgraph $G_{k,n}^{\Sigma_k}$, while d_i will correspond to the label of an edge occurring after the edge with label Δ^{k+1} . The formula $\phi_{k+1,n,p}^\Delta(\pi)$ needs to ensure that the following hold:

- (a) The first edge of π is the edge with label $i_{k+1,n}^\Delta$.
- (b) Each “passage” of the path π through the graph $G_{k,n}^{\Sigma_k}$ corresponds to the encoding of a k -counter. To express this, we will use the formula $\phi_{k,n}^{\Sigma_k}(\pi)$ given by the induction hypothesis.
- (c) The first time the path π “goes through” the graph $G_{k,n}^{\Sigma_k}$ corresponds to the encoding of the number p .
- (d) Two successive “passages” of π through the graph $G_{k,n}^{\Sigma_k}$ correspond to two successive k -counters.
- (e) The edge with label $f_{k+1,n}^\Delta$ occurs after the edge with label Δ_f^{k+1} iff the last passage of the path π through the graph $G_{k,n}^{\Sigma_k}$, corresponded to the encoding of the number $tower(k, f_0n)$. This ensures that we fully encode a $(k+1, f_0n)$ -description, and not a subsequence of it.

We only show how to express (b) as this is one of the most difficult cases and the other ones can be treated similarly.

For (b) we have to express that each passage of π through the graph $G_{k,n}^{\Sigma_k}$ corresponds to the encoding of a k -counter. Recall that by the induction hypothesis, since a (k, f_0n) -description over Σ_k is a k -counter of length k , the formula $\phi_{k,n}^{\Sigma_k}(\pi')$ is true in the graph $G_{k,n}^{\Sigma_k}$ iff $l_{\Gamma_k}(\pi')$ is the encoding of a k -counter of length n .

Hence, in order to express (b), it is enough to ensure that if s is the first node of a passage of π through $G_{k,n}^{\Sigma_k}$ and if t is the last node of that same passage, then the formula

$\phi_{k,n}^{\Sigma_k}(\pi_{s,t})$ holds. We introduce a formula $IF_{k,n}(s, t, \pi)$ such that

$$IF_{k,n}(s, t, \pi) \text{ holds} \quad \text{iff} \quad \begin{array}{l} s \text{ is the first node of a passage of } \pi \text{ through } G_{k,n}^{\Sigma_k} \\ \text{and } t \text{ is the last node of that same passage.} \end{array} \quad (4.3)$$

We define $IF_{k,n}(s, t, \pi)$ as the formula

$$i_{k,n}^{\Sigma_k}(s) \wedge E_{f_{k,n}^{\Sigma_k}}(t-1, t) \wedge (s < t) \wedge \neg \exists u^\pi [(s < u < t) \wedge \bigvee_i d_i(u)].$$

This equivalent to saying that s is the initial node of $G_{k,n}^{\Sigma_k}$, that t is the final node of $G_{k,n}^{\Sigma_k}$ and the path “never goes out” of the graph $G_{k,n}^{\Sigma_k}$ (this can be enforced by imposing that we do not go through the edge with label d_i for some i).

We define now the formula $\chi_1(s, t, \pi)$ expressing condition (b), that is, if s is the first node of a passage of π through $G_{k,n}^{\Sigma_k}$ and if t is the last node of that same passage, then the formula $\phi_{k,n}^{\Sigma_k}(\pi_{s,t})$ holds. By (4.3), we may define $\chi_1(s, t, \pi)$ as the formula

$$IF_{k,n}(s, t, \pi) \rightarrow \phi_{k,n}^{\Sigma_k}(\pi_{s,t}).$$

We turn now to the definitions of the formulas $succ_{k+1,n}(\pi)$, $num_{k+1,n}^i(\pi)$, $eq_{k+1,n}(\pi, \pi')$ and $last_{k+1,n}(\pi)$. The formulas $succ_{k+1,n}(\pi)$, $num_{k+1,n}^i(\pi)$ and $last_{k+1,n}(\pi)$ are defined in a similar fashion as the basis case ($k = 1$).

In order to define the formula $eq_{k+1,n}(\pi, \pi')$, let π and π' be two paths satisfying the formula $\phi_{k+1,n}^{\Sigma_{k+1}}$. Recall that π corresponds to the encoding of a $(k+1)$ -counter

$$\sigma_1 d_1 \dots \sigma_j d_j,$$

where each σ_i corresponds to a passage $\pi_{s,t}$ of π through $G_{k,n}^{\Sigma_k}$ and d_i corresponds to the label of an edge occurring right after that passage. Given the structure of the graph $G_{k+1,n}^\Delta$, that edge is the incoming edge of the node $t+2$.

The paths π and π' correspond to the encoding of the same $(k+1)$ -counter if for all passages $\pi_{s,t}$ of π through $G_{k,n}^{\Sigma_k}$ and for all passages $\pi_{s',t'}$ of π' through $G_{k,n}^{\Sigma_k}$ such that $\pi_{s,t}$ and $\pi_{s',t'}$ encode the same k -counter, we have that $t+2$ and $t'+2$ are the same nodes. By (4.3) and by the induction hypothesis, this can be expressed by the following formula $eq_{k+1,n}(\pi, \pi')$ given by

$$\forall s^\pi, t^\pi \forall (s')^{\pi'}, (t')^{\pi'} [IF(s, t, \pi) \wedge IF(s', t', \pi') \wedge eq_{k,n}(\pi_{s,t}, \pi_{s',t'}) \rightarrow (t+2 \sim t'+2)]$$

This finishes the proof of Lemma 4.2. \square

We are now ready to prove Theorem 4.1.

Proof of Theorem 4.1. As explained earlier, we prove that for all Turing machines M and for all k , there is a formula $\phi \in WL$ such that for all words w of size n , there is a graph G_w such that

$$G_w \models \phi \quad \text{iff} \quad \text{there is an accepting run of } M \text{ over } w \text{ using at most } tower(k, n) \text{ cells.}$$

Let $(\Sigma, Q, \delta, q_i, q_f)$ be the Turing machine M , where Σ is the input alphabet together with a blank symbol B , q_0 is the initial state, q_f is the final state and $\delta : Q \times \Sigma \rightarrow (Q \times \Sigma \times \{L, R\})$ is the transition map, where L stands for “left” and R stands for “right”.

The formula ϕ is a formula of the form

$$\exists \pi \psi(\pi),$$

where ψ is a formula that does not contain any quantification over path variables. Given a word w , the label of π in the graph G_w is the encoding of an accepting run of M over the word w . Recall that we encode a configuration of the machine in the following way. Suppose that C is a configuration where the content of the tape is the word $w' = w'_0 \dots w'_j$ ($j = \text{tower}(k, f_0 n) - 1$), the head is scanning the cell number i_0 and the machine is in state q . We may encode C by the word $e_C = d_0^C \dots d_j^C$ where each d_i is a sequence


$$c(i) (q'_i, w'_i),$$

and $c(i), w'_i$ and q'_i are defined as follows. The word $c(i)$ is the k -counter encoding the number i . The letter w'_i is the content of the cell i . The letter q'_i is equal to $\$$ if the head is not scanning the cell number i ; otherwise, q'_i is equal to the state q . This implies that given a configuration C , the word ec_C is a $(k+1, n)$ -counter over the alphabet $\Delta := (Q \cup \{\$\}) \times \Sigma$.

The run of M over the word w is a sequence of configurations of the form $C_0 C_1 \dots$. We encode the run as the word $e_{C_0} e_{C_1} \dots$ (which is a sequence of $(k+1, n)$ -counters). We will define the formula $\psi(\pi)$ and the graph G_w in such a way that a path π satisfies ψ iff the projection of the label of π on the alphabet $\Gamma_k \cup \Delta$ is the encoding of an accepting run of M over w .

We think of a path π encoding a run of M over w as consisting of two parts. The label of the first part contains the encoding e_{C_0} of the initial configuration C_0 . The label of the second part contains the encoding $e_{C_1} e_{C_2} \dots$ of the remaining part of the run. The first part of the path π is a path in a subgraph I_w of G_w , while the second part is a path in the subgraph H (independent of w) of G_w . The graph G_w will be obtained by adding an edge from a node of I_w to a node of H .

We start by defining the graph H . Recall that Δ is the alphabet $(Q \cup \{\$\}) \times \Sigma$. The graph H is defined as the graph $G_{k,n}^\Delta$ with an additional edge from the final node to the initial node. Hence, it follows from the proof of Lemma 4.2 that H is the following graph,




where $\Delta = \{d_1, \dots, d_l\}$ and where the edges with label $i_{k+1,n}^\Delta$ and Δ_f^{k+1} are edges pointing to the initial node of $G_{k,n}^{\Sigma_k}$ and the edge with label Δ^{k+1} is an edge starting from the final node of $G_{k,n}^{\Sigma_k}$. In the above paragraphs, any edge pointing to the graph $G_{k-1,n}^{\Sigma_k}$ is an edge pointing

to the initial node of that graph. Similarly, any edge starting from the graph $G_{k-1,n}^{\Sigma_k}$ will always be referring to an edge starting in the final node of the graph.

Recall that if a path π encodes a run $C_0C_1C_2\dots$, the trace of π on H will encode the part $C_1C_2\dots$ of the run. Each configuration C_i is encoded as a $(k+1, f_0n)$ -description over Δ , which will correspond, as in Lemma 4.2, to a passage of the path π from the initial node to the final node of $G_{k-1,n}^{\Sigma_k}$.


We define now the graph I_w encoding the initial configuration of the tape. Recall that in the initial configuration, the tape contains the word $w = w_0\dots w_{n-1}$, all the cells with number $\geq n$ contain the blank symbol B , the head is scanning the first cell and the state is q_0 . The graph I_w is obtained by “assembling” the subgraphs K_0, \dots, K_{n-1} and K , which we will define next. For each $i \leq n$, the graph K_i is such that the label of its unique maximal path is the encoding of the cell number i in the initial configuration. The trace of the path π on the graph K is the encoding of the contents of the cells with number $\geq n$ in the initial configuration.

More precisely, we define the graphs K_0, \dots, K_{n-1} and K in the following way. The following graph is the graph K_0



The node with label in will be the starting node of the path π . Since the trace of π in K_1 encodes the content of the cell with number 0 in the initial configuration, its label must contain the k -counter encoding the number 0 followed by the letter (q_i, w_0) of the alphabet Δ (indicating that the first cell contains the letter w_0 , the head is scanning the first cell and the current state is q_0). Using Lemma 4.2 and the formula $num_{k,n}^0$, we will impose that the passage of π through the subgraph $G_{k,n}^{\Sigma_k}$ of K_0 corresponds to the encoding of the number 0.

Next, for all $1 \leq i \leq n$, we define K_i as the following graph



Recall that we want to define K_i in such a way that the trace of π on K_i is the encoding of the contents of the cell with number in in the initial configuration (that is, it contains the letter w_i and the head is not scanning the cell since $i \neq 0$). Recall that the encoding of such a cell (and its content) is given by

$$c(i)(\$, w_i),$$

where $c(i)$ is the k -counter encoding the number i . We will use the formula $num_{k,n}^i$ given by Lemma 4.2 to express that the passage of π through the subgraph $G_{k,n}^{\Sigma_k}$ of K_i corresponds to the k -counter encoding i .

Finally we define the graph K as the graph



where Δ_B is the one-letter alphabet containing the blank symbol B . Recall that the trace of π on the graph K will encode the contents of the cells with number $\geq n$ in the initial configuration (that is, the fact that those cells contain the blank symbol and are not scanned by the head). Since the encoding of such a cell with number i is given by


$$c(i)(\$, B)$$

(where $c(i)$ is the k -counter encoding i), the label of the trace of π on K must contain the word

$$c(n+1)(B, \$) \dots c(j)(B, \$),$$

where $j = \text{tower}(k, n) - 1$. That is, the label of the trace of π on K is the unique $(k, f_0n, n+1)$ -description over the alphabet Δ_B . We will express that the passage of π through the graph $G_{k,n}^{\Delta_B}$ corresponds to the $(k, f_0n, n+1)$ -description over the alphabet Δ_B using the formula $\phi_{k,n,n+1}^{\Delta_B}$ provided by Lemma 4.2.

We are now ready to define the graph I_w which is obtained by assembling the graphs previously introduced in the following way.



Each edge between two graphs in the picture above is an edge from the “left-most” node of the first graph to the “right-most” node of the second graph. Finally the graph G_w is the graph obtained by considering the union of the graph I_w and H and adding an edge from the final node of K to the initial node of H .

Now that we have defined the graph G_w , we are ready to define the formula ψ . The formula $\psi(\pi)$ is obtained as the conjunction of the following formulas.

- (A) First we need to express that the path π starts with the edge with label in .
- (B) We need to express that eventually in a configuration, the machine reaches the final state q_f .
- (C) We also have to express that each passage of the path π from the initial node of $G_{k,n}^{\Sigma_k}$ to the final node of $G_{k,n}^{\Sigma_k}$ in the graph H corresponds to the encoding of a $(k+1, f_0n)$ -description.
- (D) We have to express that for all $i < n$, the trace of π on the subgraph $G_{k,n}^{\Sigma_k}$ of the graph K_i corresponds to the k -counter encoding i .
- (E) We need to express that the trace of the path π on the subgraph $G_{k+1,n}^{\Delta_B}$ of K is the unique $(k+1, f_0n, n+1)$ -description over the alphabet Δ_B .
- (F) Finally we need to express how we move from one configuration of the tape to the next one.

Cases (A) and (B) are straightforward. Cases (C), (D) and (E) are similar and we only give details for case (C) and case (F). By Lemma 4.2, case (C) means that

$$\text{if } \pi_{s,t} \text{ is the subpath of } \pi \text{ corresponding to such a passage, then } \phi_{k+1,n,0}^{\Delta}(\pi_{s,t}) \text{ holds.} \quad (4.4)$$

The node s is a node satisfying $i_{k,n}^{\Sigma_k}$, while t is the “closest” node to s with an incoming edge with label $f_{k+1,n}^{\Delta}$. This is expressed by the following formula $IF_{k+1,n}^{\Delta}(s, t, \pi)$ defined by

$$i_{k+1,n}^{\Sigma_k}(s) \wedge E_{f_{k+1,n}^{\Sigma_{k+1}}}(t-1, t) \wedge \neg \exists r^\pi [s < r < t \wedge E_{f_{k+1,n}^{\Sigma_{k+1}}}(r-1, r)]. \quad (4.5)$$

It follows from the definitions of $IF^\Delta(s, t, \pi)$ and the graph $G_{k+1, n}^\Delta$ that

$$G_{k+1, n}^\Delta \models IF_{k+1, n}^\Delta(s, t, \pi) \text{ iff } s \text{ is the first node of a passage of } \pi \text{ through } G_{k+1, n}^\Delta \text{ and } t \text{ is the last node of that same passage.} \quad (4.6)$$

Hence, (4.4) is equivalent to the fact that if $IF_{k+1, n}^\Delta(s, t, \pi)$ holds, so does the formula $\phi_{k+1, n, 0}^\Delta(\pi_{s, t})$. Therefore the formula

$$\forall s^\pi, t^\pi [IF^\Delta(s, t, \pi) \rightarrow \phi_{k+1, n, 0}^\Delta(\pi_{s, t})]$$

expresses case (C).

Finally we treat the most difficult case which is case (F). We need to express how we move from one configuration of the tape to the next one. Recall that the trace of π on the graph H will contain the encoding of the sequence $C_1 C_2 \dots$ of the run, where $C_0 C_1 \dots$ is the full run of the machine on the input w .

Let $\pi_{s, t}$ be the subpath of π corresponding to the encoding of a configuration C_i and let $\pi_{s', t'}$ be the subpath of π corresponding to the configuration C_{i+1} . We need to express how to move from the configuration C_i to the configuration C_{i+1} . Suppose that in the configuration C_i , the current state is q , and the head is scanning the cell c containing the letter u . Suppose also that $\delta(q, u) = (q', v, R)$ (we can treat similarly the case where the head moves to the left). In order to keep our formulas simpler, we use a slightly different definition of a run of a Turing machine, but it would be clear that the notion of run that we use here, can be simulated by a usual Turing machine. Here, we assume that if the machine scans a cell c with content u and $\delta(q, u) = (q', v, R)$, then in the next state, the machine scans the successor c' of c , the content of c' is v , while the content of c is u (in the usual definition, the content of c' is unchanged, while the content of c is v).

Let $\pi_{r, s}$ be the subpath of π corresponding to the encoding of the cell c in the configuration C_i . Let $\pi_{r', s'}$ be the encoding of an arbitrary cell c' in the configuration C_{i+1} . If c' is the successor of the cell c , then the head should scan the cell c' and the content of c' should be the letter v . We express this by the formula $change_{(q, a, q', b)}^R(r, s, r', s', \pi)$ defined by

$$succ_{k, n}(\pi_{r, s}, \pi_{r', s'}) \wedge (q, u)(s + 2) \rightarrow (q', v)(s' + 2).$$

Recall that by Lemma 4.2, $succ_{k, n}(\pi_{r, s}, \pi_{r', s'})$ is the formula expressing that the k -counter associated with $\pi_{r', s'}$ is the successor of the k -counter associated with $\pi_{r, s}$.

If c' is not the successor of the cell c , then the head is not scanning the cell c' and its content remains unchanged. If $\pi_{x, y}$ is the subpath of π corresponding to the content of the cell c' in the configuration C_i , this is expressed by the following formula $stay_{(q, u, q', v)}^R(r, s, x, y, r', s', \pi)$ defined by

$$\neg succ_{k, n}(\pi_{r, s}, \pi_{r', s'}) \wedge (q, u)(s + 2) \wedge eq_{k, n}(\pi_{x, y}, \pi_{r', s'}) \wedge (q'', u_0)(x, y) \rightarrow (\$, u_0)(s + 2)$$

where $q'' \in Q \cup \{\$\}$. Recall that by Lemma 4.2, $eq_{k, n}(\pi_{x, y}, \pi_{r', s'})$ expresses that the k -counters associated with $\pi_{x, y}$ and $\pi_{r', s'}$ are the same.

Now we need to express that the paths $\pi_{r, s}$, $\pi_{s', s'}$ and $\pi_{x, y}$ correspond to the encodings of k -counters. By Lemma 4.2, this means that those paths correspond to passages of π through the graph $G_{k, n}^{\Sigma_k}$. Similarly to (4.5), we introduce a formula $IF_{k, n}^{\Sigma_k}(r, s, \pi)$ defined by

$$i_{k, n}^{\Sigma_k}(r) \wedge E_{f_{k, n}^{\Sigma_k}}(s - 1, s) \wedge \neg \exists t^\pi [r < t < s \wedge E_{f_{k, n}^{\Sigma_k}}(t - 1, t)].$$

This formula expresses that the path $\pi_{r, s}$ corresponds to the encoding of a k -counter.

Next we also need a formula to assert that the paths $\pi_{r,s}$ and $\pi_{r',s'}$ appear in the encodings of successive configurations (and similarly, that the paths $\pi_{x,y}$ and $\pi_{r',s'}$ appear in the encodings of successive configurations). Since the encoding of a configuration starts with the unique edge with label $i_{k+1,n}^\Delta$ (and that edge only occurs at the beginning of the encoding of a configuration), this is equivalent to say that there is a unique edge between s and r' with label $i_{k+1,n}^\Delta$. This is expressed by the formula $config(s, r', \pi)$ defined by

$$\exists t^\pi [(s < t < r') \wedge i_{k+1,n}^\Delta(t) \wedge \neg \exists (t')^\pi [(s < t' < r') \wedge i_{k+1,n}^\Delta(t') \wedge (t' \neq t)]].$$

We are now ready to define $\theta_{q,u,q',v}^R(\pi)$ as the following formula

$$\forall r^\pi, s^\pi, (r')^\pi, (s')^\pi, x^\pi, y^\pi [IF_{k,n}^{\Sigma k}(r, s, \pi) \wedge IF_{k,n}^{\Sigma k}(r', s', \pi) \wedge IF_{k,n}^{\Sigma k}(x, y, \pi) \quad (4.7)$$

$$\wedge config(s, r', \pi) \wedge config(y, r', \pi) \quad (4.8)$$

$$\rightarrow change_{(q,u,q',v)}^R(s, r, s', r', \pi) \wedge stay_{(q,u,q',v)}^R(r, s, x, y, r', s', \pi)]. \quad (4.9)$$

It expresses the following. Suppose that $\pi_{r,s}$, $\pi_{r',s'}$ and $\pi_{x,y}$ are k -counters encoding the numbers of three cells (this corresponds to (4.7)). Suppose that $\pi_{r,s}$ and $\pi_{x,y}$ correspond to cells occurring in the same configuration C and that the cell corresponding to $\pi_{r',s'}$ occurs in the next configuration C' (this is expressed by (4.8)). Then, if we “apply” the transition $\delta(q, u) = (q', v, R)$ to move from C to C' , we move the head to the right and update the content of the cell being scanned (as expressed by the formula $change_{(q,u,q',v)}^R(r, s, r', s', \pi)$) and we leave the other cells unchanged (as expressed by the formula $stay_{(q,u,q',v)}^R(r, s, x, y, r', s', \pi)$).

We define now the formula $\theta_R(\pi)$ as the formula

$$\bigwedge \{\theta_{q,u,q',v}^R(\pi) : \delta(q, u) = (q', v, R)\}.$$

This formula expresses how we move from one configuration to another, when the head moves to the right. Similarly, we can define a formula $\theta_L(\pi)$ expressing how we move from one configuration to another, when the head moves to the left.

This finishes the proof of Theorem 4.1. \square

As a corollary to the proof of Theorem 4.1, we obtain that data complexity is non-elementary even for simple WL formulas that talk about a single path in a graph database.

Corollary 4.3. The evaluation problem for WL over graph databases is non-elementary in data complexity, even if restricted to Boolean WL formulas of the form $\exists \pi \psi$, where ψ uses no path quantification and contains no position variable of sort different than π .

5. REGISTER LOGIC

We saw in the previous section that WL is impractical due to its very high data complexity. In this section, we start by recalling the notion of regular expressions with memory (REM) and their basic results from [10]. In our view, this logic is rather limited in terms of expressive power. For instance, the query (Q) from the introduction cannot be expressed in REM. We then introduce an extension of REM, called regular logic (RL), that remedies this limitation in expressive power (in fact, it can express many natural examples of queries expressible in WL, e.g., those given in [8]) while retaining elementary complexity of query evaluation. Finally, we study which fragments of RL are well-behaved for database applications.

5.1. Regular expressions with memory. REMs define pairs of nodes in data graphs that are linked by a path that satisfies a constraint in the way in which the topology interacts with the underlying data. REMs allow us to remember data values and use them later. Data values are stored in k registers r_1, \dots, r_k . At any point we can compare a data value with one previously stored in the registers. As an example, consider the REM $\downarrow r.a^+[r^-]$. This can be read as follows: Store the current data value in register r (represented by the expression $\downarrow r$), and then check that after reading a word in a^+ we see the same data value again (condition $[r^-]$). We formally define REM next.

Let r_1, \dots, r_k be registers. The set of *conditions* c over $\{r_1, \dots, r_k\}$ is recursively defined as: $c := r_i^- \mid c \wedge c \mid \neg c$, for $1 \leq i \leq k$. Assume that \mathcal{D}_\perp is the extension of the set \mathcal{D} of data values with a new symbol \perp . Satisfaction of conditions is defined with respect to a value $d \in \mathcal{D}$ (the data value that is currently being scanned) and a tuple $\tau = (d_1, \dots, d_k) \in \mathcal{D}_\perp^k$ (the data values stored in the registers, assuming that $d_i = \perp$ represents the fact that register r_i has no value assigned) as follows (Boolean combinations omitted): $(d, \tau) \models r_i^-$ iff $d = d_i$.

Definition 5.1 (REMs). The class of REMs over Σ and $\{r_1, \dots, r_k\}$ is defined by the grammar:

$$e := \varepsilon \mid a \mid e \cup e \mid e \cdot e \mid e^+ \mid e[c] \mid \downarrow \bar{r}.e$$

where a ranges over symbols in Σ , c over conditions over $\{r_1, \dots, r_k\}$, and \bar{r} over tuples of elements in $\{r_1, \dots, r_k\}$. \square

That is, REM extends the class of regular expressions e – which is a popular mechanism for specifying topological properties of paths in graph databases (see, e.g., [18, 2]) – with expressions of the form $e[c]$, for c a condition, and $\downarrow \bar{r}.e$, for \bar{r} a tuple of registers – that define how such topology interacts with the data.

Semantics: To define the evaluation $e(G)$ of an REM e over a data graph $G = (V, E, \kappa)$, we use a relation $\llbracket e \rrbracket_G$ that consists of tuples of the form $(u, \lambda, \rho, v, \lambda')$, for u, v nodes in V , ρ a path in G from u to v , and λ, λ' two k -tuples over \mathcal{D}_\perp . The intuition is the following: the tuple $(u, \lambda, \rho, v, \lambda')$ belongs to $\llbracket e \rrbracket_G$ if and only if the data and topology of ρ can be parsed according to e , with λ being the initial assignment of the registers, in such a way that the final assignment is λ' . We then define $e(G)$ as the pairs (u, v) of nodes in G such that $(u, \perp^k, \rho, v, \lambda) \in \llbracket e \rrbracket_G$, for some path ρ in G from u to v and k -tuple λ over \mathcal{D}_\perp .

We inductively define relation $\llbracket e \rrbracket_G$ below. We assume that $\lambda_{\bar{r}=d}$, for $d \in \mathcal{D}$, is the tuple obtained from λ by setting all registers in \bar{r} to be d . Also, if $\rho_1 = v_1 a_1 v_2 \dots v_{k-1} a_{k-1} v_k$ and $\rho_2 = v_k a_k v_{k+1} \dots v_{n-1} a_{n-1} v_n$ are paths, then:

$$\rho_1 \rho_2 := v_1 a_1 v_2 \dots v_{k-1} a_{k-1} v_k a_k v_{k+1} \dots v_{n-1} a_{n-1} v_n.$$

Then we define:

- $\llbracket \varepsilon \rrbracket_G = \{(u, \lambda, \rho, u, \lambda) : u \in V, \rho = u, \lambda \in \mathcal{D}_\perp^k\}$.
- $\llbracket a \rrbracket_G = \{(u, \lambda, \rho, v, \lambda) : \rho = uav, \lambda \in \mathcal{D}_\perp^k\}$.
- $\llbracket e_1 \cup e_2 \rrbracket_G = \llbracket e_1 \rrbracket_G \cup \llbracket e_2 \rrbracket_G$.
- $\llbracket e_1 \cdot e_2 \rrbracket_G = \llbracket e_1 \rrbracket_G \circ \llbracket e_2 \rrbracket_G$, where $\llbracket e_1 \rrbracket_G \circ \llbracket e_2 \rrbracket_G$ is the set of tuples $(u, \lambda, \rho, v, \lambda')$ such that $(u, \lambda, \rho_1, w, \lambda'') \in \llbracket e_1 \rrbracket_G$ and $(w, \lambda'', \rho_2, v, \lambda') \in \llbracket e_2 \rrbracket_G$, for some $w \in V$, k -tuple λ'' over \mathcal{D}_\perp , and paths ρ_1, ρ_2 such that $\rho = \rho_1 \rho_2$.
- $\llbracket e^+ \rrbracket_G = \llbracket e \rrbracket_G \cup (\llbracket e \rrbracket_G \circ \llbracket e \rrbracket_G) \cup (\llbracket e \rrbracket_G \circ \llbracket e \rrbracket_G \circ \llbracket e \rrbracket_G) \dots$
- $\llbracket e[c] \rrbracket_G = \{(u, \lambda, \rho, v, \lambda') \in \llbracket e \rrbracket_G : (\kappa(v), \lambda') \models c\}$.

- $\llbracket \downarrow \bar{r}.e \rrbracket_G = \{(u, \lambda, \rho, v, \lambda') : (u, \lambda_{\bar{r}=\kappa(u)}, \rho, v, \lambda') \in \llbracket e \rrbracket_G\}$.

For each REM e , we will use the shorthand notation e^* to denote $\varepsilon \cup e^+$.

Example 5.1. The REM $\Sigma^* \cdot (\downarrow r.\Sigma^+[r=]) \cdot \Sigma^*$ defines the pairs of nodes in data graphs that are linked by a path in which two nodes have the same data value. The REM $\downarrow r.(a[\neg r=])^+$ defines the pairs of nodes that are linked by a path ρ with label in a^+ , such that the data value of the first node in the path is different from the data value of all other nodes in ρ . \square

The problem $\text{EVAL}(\text{REM})$ is, given a data graph $G = (V, E, \kappa)$, a pair (v_1, v_2) of nodes in V , and an REM e , is $(v_1, v_2) \in e(G)$? The data complexity of the problem refers again to the case when e is considered to be fixed. REMs are tractable in data complexity and have no worse combined complexity than FO over relational databases:

Proposition 5.1 ([10]). $\text{EVAL}(\text{REM})$ is PSPACE-complete, and NLOGSPACE-complete in data complexity.

5.2. Register logic. REM is well-behaved in terms of the complexity of evaluation, but its expressive power is rather rudimentary for expressing several data/topology properties of interest in data graphs. As an example, the query (Q) from the introduction – which can be easily expressed in WL – cannot be expressed as an REM (we actually prove a stronger result later). The main shortcomings of REM in terms of its expressive power are its inability to (i) compare data values in different paths and (ii) express branching properties of the data.

In this section, we propose register logic (RL) as a natural extension of REM that makes up for this lack of expressiveness. We borrow ideas from the logic CRPQ^\neg , presented in [4], that closes the class of *regular path queries* [7] under Boolean combinations and existential *node* and *path* quantification. In the case of RL we start with REMs and close them not only under Boolean combinations and node and path quantification – which allow to express arbitrary patterns over the data – but also under *register assignment* quantification – which permits comparing data values in different paths. We also prove that the combined complexity of the evaluation problem for RL is elementary (EXPSpace), and, thus, that in this regard RL is in stark contrast to WL.

To define RL we assume the existence of countably infinite sets of *node*, *path* and *register assignment variables*. Node variables are denoted x, y, z, \dots , path variables are denoted $\pi, \pi', \pi_1, \pi_2, \dots$, and register assignment variables are denoted ν, ν_1, ν_2, \dots .

Definition 5.2 (Register logic (RL)). We define the class of RL formulas ϕ over alphabet Σ and $\{r_1, \dots, r_k\}$ using the following grammar:

$$\begin{aligned} \text{atom} &:= x = y \mid \pi = \pi' \mid \nu = \nu' \mid \nu = \bar{\perp} \mid (x, \pi, y) \mid e(\pi, \nu_1, \nu_2) \\ \phi &:= \text{atom} \mid \neg\phi \mid \phi \vee \phi \mid \exists x\phi \mid \exists \pi\phi \mid \exists \nu\phi \end{aligned}$$

Here x, y are node variables, π, π' are path variables, ν, ν' are register assignment variables, and e is an REM over Σ and $\{r_1, \dots, r_k\}$. \square

Intuitively, $\nu = \bar{\perp}$ holds iff ν is the empty register assignment, (x, π, y) checks that π is a path from x to y , and $e(\pi, \nu, \nu')$ checks that π can be parsed according to e starting from register assignment ν and finishing in register assignment ν' . The quantifier $\exists\nu$ is to be read “there exists an assignment of data values in the data graph to the registers”.

Let $G = (V, E, \kappa)$ be a data graph over Σ and ϕ a RL formula over Σ and $\{r_1, \dots, r_k\}$. Assume that D is the set of data values that are mentioned in G , i.e., $D = \{\kappa(v) : v \in V\}$. An *assignment* α for ϕ over G is a mapping that assigns (i) a node in V to each free node variable x in ϕ , (ii) a path ρ in G to each free path variable π in ϕ , and (iii) a tuple λ in \mathcal{D}_\perp^k to each register variable ν that appears free in ϕ . That is, for safety reasons we assume that $\alpha(\nu)$ can only contain data values that appear in the underlying data graph. This represents no restriction for the expressiveness of the logic.

We inductively define $(G, \alpha) \models \phi$, for G a data graph, ϕ an RL formula, and α an assignment for ϕ over G , as follows (we omit equality atoms and Boolean combinations since they are standard):

- $(G, \alpha) \models \nu = \bar{\perp}$ iff $\alpha(\nu) = \perp^k$.
- $(G, \alpha) \models (x, \pi, y)$ iff $\alpha(\pi)$ is a path from $\alpha(x)$ to $\alpha(y)$ in G .
- $(G, \alpha) \models e(\pi, \nu, \nu')$ iff $(u, \alpha(\nu), \alpha(\pi), v, \alpha(\nu')) \in \llbracket e \rrbracket_G$, assuming $\alpha(\pi)$ goes from node u to v .
- $(G, \alpha) \models \exists x \phi$ iff there is node $v \in V$ such that $(G, \alpha[x \rightarrow v]) \models \phi$.
- $(G, \alpha) \models \exists \pi \phi$ iff there is path ρ in G such that $(G, \alpha[\pi \rightarrow \rho]) \models \phi$.
- $(G, \alpha) \models \exists \nu \phi$ iff there is tuple λ in \mathcal{D}_\perp^k such that $(G, \alpha[\nu \rightarrow \lambda]) \models \phi$.

Thus, each REM e is expressible in RL using the formula:

$$\exists \pi \exists \nu \exists \nu' (\nu = \bar{\perp} \wedge e(\pi, \nu, \nu')).$$

Example 5.2. Recall query (Q) from the introduction: *Find pairs of nodes x and y in a graph database, such that there is a node z and a path π from x to y in which each node is connected to z .* This query can be expressed in RL over $\Sigma = \{a\}$ and a single register r as follows:

$$\exists \pi \left((x, \pi, y) \wedge \exists z \forall \nu (e_1(\pi, \nu, \nu) \rightarrow \exists z' \exists \pi' ((z', \pi', z) \wedge e_2(\pi', \nu, \nu))) \right),$$

where $e_1 := a^*[r^=] \cdot a^*$ is the REM that checks whether the node (i.e. data) stored in register r appears in a path, and $e_2 := \varepsilon[r^=] \cdot a^*$ is the REM that checks if the first node of a path is the one that is stored in register r .

In fact, this formula defines the pairs of nodes x and y such that there exists a path π that goes from x to y and a node z for which the following holds: for every register value ν (i.e., for every node ν) such that $e_1(\pi, \nu, \nu)$ (i.e. node ν is in π), it is the case that there is a path π' from some node z' to z such that $e_2(\pi', \nu, \nu)$ (i.e., $z' = \nu$ and π' connects ν to z). Notice that this uses the fact that the underlying data model is that of graph databases, in which each node is uniquely identified by its data value. \square

The limitations in expressive power of RL have also been independently recognized by Libkin, Martens and Vrgoc [12]. In order to allow for interesting data value comparisons while retaining reasonable complexity of evaluation, they propose to use query languages based on the XML language XPath. These languages are not comparable in terms of expressive power to the ones we study here.

Complexity of evaluation for RL: The evaluation problem for RL, denoted EVAL(RL), is as follows: Given a data graph G , an RL formula ϕ , and an assignment α for ϕ over G , is

it the case that $(G, \alpha) \models \phi$? As before, we denote by $\text{EVAL}(\text{RL}, \phi)$ the evaluation problem for the fixed RL formula ϕ .

We show next that, unlike WL, register logic RL can be evaluated in elementary time, and, actually, with only one exponential jump over the complexity of evaluation of REMs:

Theorem 5.2. $\text{EVAL}(\text{RL})$ is EXPSPACE-complete. The lower bound holds even if the input is restricted to graph databases.

Proof. We start by proving the upper bound, that is, $\text{EVAL}(\text{RL})$ is in EXPSPACE. The structure of the proof is quite similar to the one that proves that CRPQ^\top queries can be evaluated in PSPACE in combined complexity [4]. The difference is that now we have to accommodate the extra expressive power of RL, that allows to express properties of register values and check acceptance of data walks by REMs.

Let $G = (V, E, \lambda)$ be a Σ -labeled data graph and ϕ a RL formula over Σ and $\{x_1, \dots, x_k\}$. Let us denote by $D = \{\lambda(v) : v \in V\}$ and define D_\perp to be $D \cup \{\perp\}$. Further, let α be an assignment for ϕ over G . We define

- $\bar{v} = (v_1, \dots, v_{k_1})$ as a tuple of nodes in G such that

$$\{v_1, \dots, v_{k_1}\} = \{\alpha(x) : x \text{ is a free node variable}\},$$

- $\bar{\rho} = (\rho_1, \dots, \rho_{k_2})$ as a tuple of paths in G such that

$$\{\rho_1, \dots, \rho_{k_2}\} = \{\alpha(\rho) : \rho \text{ is a free path variable}\},$$

- $\bar{\lambda} = (\lambda_1, \dots, \lambda_{k_3})$ as a tuple of register values for $\{x_1, \dots, x_k\}$ over G (i.e., tuples in $(D \cup \{\perp\})^k$) such that $\{\lambda_1, \dots, \lambda_{k_3}\} = \{\alpha(\lambda) : \lambda \text{ is a free register variable}\}$.

Further, assume that e_1, \dots, e_m are all the REMs mentioned in ϕ . Our goal is to define an EXPSPACE procedure that checks whether $(G, \alpha) \models \phi$. In order to do that, we first have to introduce some new terminology.

Let τ be a first-order (FO) vocabulary

$$\langle \text{Nodes}, \text{Paths}, \text{Registers}, \text{Endpoints}, e_1, \dots, e_m, \bar{\perp} \rangle,$$

where (a) **Nodes**, **Paths** and **Registers** are unary relation symbols, (b) **Endpoints** and e_i ($1 \leq i \leq m$) are ternary relation symbols, and (c) $\bar{\perp}$ is a constant. We define, from G , an FO structure \mathcal{M}_G over τ as follows: The domain of \mathcal{M}_G is the disjoint union of V , all the paths that belong to G , and all k -tuples over D_\perp . (Notice that each node in V is also a path in G , but here we consider them to be different objects. That is, each $v \in V$ appears separately as a node and as a path in the domain of \mathcal{M}_G). The constant $\bar{\perp}$ is interpreted in \mathcal{M}_G as the tuple \perp^k . The interpretation of **Nodes** in \mathcal{M}_G contains all those elements of the domain that are nodes. The interpretation of **Paths** in \mathcal{M}_G contains all those elements of the domain that are paths. The interpretation of **Registers** in \mathcal{M}_G contains all those elements of the domain that are k -tuples over D_\perp . The interpretation of the ternary relation **Endpoints** contains all tuples (v, ρ, v') such that ρ is a path in G from node v to node v' . Finally, the interpretation of the symbol e_i ($1 \leq i \leq m$) contains all tuples $(\lambda, \rho, \lambda')$ such that ρ is a path in G , λ, λ' are k -tuples over D_\perp , and $e_i(\rho, \lambda, \lambda')$.

Let ϕ_τ be the FO formula over vocabulary τ obtained from ϕ by simultaneously replacing (1) each subformula of the form $\exists x\theta$ (for x a node variable) with $\exists x(\text{Nodes}(x) \wedge \theta)$, (2) each subformula of the form $\exists \pi\theta$ (for π a path variable) with $\exists \pi(\text{Paths}(\pi) \wedge \theta)$, (3) each subformula of the form $\exists \nu\theta$ (for ν a register variable) with $\exists \nu(\text{Registers}(\nu) \wedge \theta)$, and (4) each atomic formula of the form (x, π, y) with $\text{Endpoints}(x, \pi, y)$.

Clearly, $G, \alpha \models \phi$ iff $\mathcal{M}_G, \alpha \models \phi_\tau$.

Of course, \mathcal{M}_G cannot be effectively constructed from G since the set of paths in G is potentially infinite, and, thus, \mathcal{M}_G is also potentially infinite. However, it is possible to prove that there exists a finite structure $\mathcal{M}'_{G, \bar{\rho}}$ such that $G, \alpha \models \phi$ iff $\mathcal{M}'_{G, \bar{\rho}}, \alpha \models \phi_\tau$. We show how to define $\mathcal{M}'_{G, \bar{\rho}}$ next.

Assume that the *quantifier rank* of ϕ_τ is $k \geq 0$, where the quantifier rank of an FO formula θ is the depth of nested quantification in θ . Let $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$. A path ρ in G *satisfies* \mathcal{E} if the following holds: For each triple $(e_i, \lambda, \lambda') \in \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$, it is the case that $G, \alpha \models e_i(\rho, \lambda, \lambda')$ iff $(e_i, \lambda, \lambda') \in \mathcal{E}$. (Notice that for each path in G there is one, and only one, subset \mathcal{E} of $\{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$ that it satisfies.) For each pair (v, v') of nodes in V , and for every $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$, let $c_{\mathcal{E}, v, v'} \geq 0$ be the minimum between $k + |\bar{\rho}|$ and the number of paths in G that go from v to v' and satisfy \mathcal{E} . We arbitrarily pick, for each pair (v, v') of nodes in V and for each $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$, $c_{\mathcal{E}, v, v'}$ distinct paths $\rho_{\mathcal{E}, v, v'}^1, \dots, \rho_{\mathcal{E}, v, v'}^{c_{\mathcal{E}, v, v'}}$ from v to v' that satisfy \mathcal{E} .

We define the structure $\mathcal{M}'_{G, \bar{\rho}}$ as follows: Its domain contains all the nodes of V , each path ρ that belongs to the tuple $\bar{\rho}$, every path of the form $\rho_{\mathcal{E}, v, v'}^i$, where $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$ ($v, v' \in V$ and $1 \leq i \leq c_{\mathcal{E}, v, v'}$), and every tuple in D_\perp^k . The constant \perp is interpreted in $\mathcal{M}'_{G, \bar{\rho}}$ as the tuple \perp^k . The interpretation of **Nodes** in $\mathcal{M}'_{G, \bar{\rho}}$ contains all nodes in the domain. The interpretation of **Paths** in $\mathcal{M}'_{G, \bar{\rho}}$ contains all those elements of the domain that are paths. The interpretation of **Registers** in $\mathcal{M}'_{G, \bar{\rho}}$ contains all those elements of the domain that are k -tuples over D_\perp . The interpretation of the ternary relation **Endpoints** contain all tuples of the form (v, ρ, v') , where $v, v' \in V$ and ρ is a path in the domain that goes from v to v' in G . Finally, the interpretation of e_i ($1 \leq i \leq m$) in $\mathcal{M}'_{G, \bar{\rho}}$ contains all tuples $(\lambda, \rho, \lambda')$ such that ρ is a path in the domain, λ, λ' are k -tuples over D_\perp , and $e_i(\rho, \lambda, \lambda')$.

By using a standard Ehrenfeucht-Fraïssé argument it is possible to prove the following:

Claim 5.2.1. The structures $(\mathcal{M}_G, \bar{v}, \bar{\rho}, \lambda)$ and $(\mathcal{M}'_{G, \bar{\rho}}, \bar{v}, \bar{\rho}, \lambda)$ are indistinguishable by FO sentences of quantifier rank $\leq k$.

Proof. We show that the duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game played on $(\mathcal{M}_G, \bar{v}, \bar{\rho}, \lambda)$ and $(\mathcal{M}'_{G, \bar{\rho}}, \bar{v}, \bar{\rho}, \lambda)$. The duplicator's response to a spoiler move in round $i \leq k$ is (inductively) defined as follows (we assume without loss of generality that the spoiler never repeats moves, i.e. in no round does the spoiler choose an element that has already been chosen by either player in previous rounds):

- If the spoiler's move in round i is a node in either of the two structures, then the duplicator responds by mimicking the spoiler's move on the other structure;
- if the spoiler's move in round i is a k -tuple over $D \cup \{\perp\}$ in either of the two structures, then the duplicator responds by mimicking the spoiler's move on the other structure;
- if the spoiler's move in round i is a path ρ in $\bar{\rho}$ in either of the two structures, then again the duplicator responds by mimicking the spoiler's move on the other structure;
- if the spoiler plays a path ρ from node v to v' , in either of the two structures, such that ρ satisfies $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$ and ρ is not a path in $\bar{\rho}$, then the duplicator responds with any path from v to v' in the other structure that (1) satisfies \mathcal{E} , (2) does not belong to $\bar{\rho}$, and (3) has not been previously chosen in the game. Notice that it is always possible for the duplicator to choose such a path, since for each pair of nodes

$v, v' \in V$ and for each $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_{\perp}^k \times D_{\perp}^k$, the number of paths from v to v' that satisfy \mathcal{E} and do not belong to $\bar{\rho}$ is the same up to k .

It is easy to see that duplicator's response defined in this way always preserves a partial isomorphism between the two structures. This implies that the duplicator has a winning strategy in the k -round Ehrenfeucht-Fraïssé game played on $(\mathcal{M}_G, \bar{v}, \bar{\rho}, \bar{\lambda})$ and $(\mathcal{M}'_{G, \bar{\rho}}, \bar{v}, \bar{\rho}, \bar{\lambda})$, and, thus, by well-known results, that the structures are indistinguishable by FO sentences of quantifier rank $\leq k$. \square

The previous claim shows that $(G, \alpha) \models \phi$ iff $(\mathcal{M}'_{G, \bar{\rho}}, \alpha) \models \phi_{\tau}$. Thus, a straightforward approach to check whether $(G, \alpha) \models \phi$ would be to construct $\mathcal{M}'_{G, \bar{\rho}}$ and then evaluate ϕ_{τ} over it. The problem with this approach is that $\mathcal{M}'_{G, \bar{\rho}}$ could be of double exponential size (because there is a double exponential number of different subsets \mathcal{E} of $\{e_1, \dots, e_m\} \times D_{\perp}^k \times D_{\perp}^k$), and, thus, impossible to construct in exponential space. It will be necessary to follow a different approach.

Assume that ϕ_{τ} is given in prenex normal form, i.e. ϕ_{τ} is of the form

$$Q_1 y_1 \cdots Q_m y_m \psi(\bar{x}, \bar{\pi}, \bar{\nu}, y_1, \dots, y_m),$$

where each Q_i is either \exists or \forall , each y_i is a node, path or register variable, and ψ is quantifier-free (if ϕ_{τ} is not in prenex normal form, we can convert it in polynomial time into an equivalent formula in prenex normal form). We follow a usual argument to evaluate FO formulas on structures. The main problem with this is that some of the elements in $\mathcal{M}'_{G, \bar{\rho}}$ are paths and register values, and have to be treated as such. Therefore, we define a way of encoding paths (in exponential space) and register values (in polynomial space).

Clearly, each register value can be codified with a tuple of length $k \cdot \log_2(|V|)$. In order to denote that this tuple is the address of a register value (and not, say, of a path), we add an extra bit at the beginning of the tuple which is labeled with a new symbol r . Codification of paths requires a bit of extra work. Each path ρ is encoded with an *address*, that is, a string that satisfies the following:

- It starts with a new symbol p , that states that this is the address of a path;
- the address continues with the encodings of the two endpoints v and v' of the path (separated with some delimiter); this part of the address uses $O(\log_2 |V|)$ space;
- then the address encodes the subset \mathcal{E} of $\{e_1, \dots, e_m\} \times D_{\perp}^k \times D_{\perp}^k$ that ρ satisfies; this encoding can be easily expressed with a string of length $m \times |V|^k \times |V|^k$ over alphabet $\{0, 1\}$ that flags with a 1 those elements of $\{e_1, \dots, e_m\} \times D_{\perp}^k \times D_{\perp}^k$ that belong to \mathcal{E} ;
- finally, the address contains an encoding of the integer $i \leq \bar{k} + |\bar{\rho}|$ such that $\rho = \rho_{\mathcal{E}, v, v'}^i$; this encoding uses $O(\log_2(|\phi| + |\bar{\rho}|))$ space.

Clearly, the address of a path defined in this way can be specified using at most exponential space.

We show next how the problem of checking whether $(\bar{v}, \bar{\rho}, \bar{\lambda})$ belongs to the evaluation of ϕ_{τ} over $\mathcal{M}'_{G, \bar{\rho}}$ can be solved in exponential time by an alternating Turing machine. This will finish the proof of the theorem, since the class of problems that can be solved in exponential time by alternating Turing machines coincides with the class of problems that can be solved in EXPSpace.

The alternating machine proceeds as follows. It first replaces in ϕ_{τ} each node variable x in \bar{x} with the encoding of the corresponding node v of \bar{v} , each path variable π in $\bar{\pi}$ with the encoding (address) of the corresponding path ρ in $\bar{\rho}$, and each register variable ν in $\bar{\nu}$ with

the encoding of the corresponding tuple λ in $\bar{\lambda}$. Then the machine reads the formula ϕ_τ from left-to-right. Each time it encounters an existential quantifier $\exists y_i$ it enters an existential state, and each time it encounters a universal quantifier $\forall y_i$ it enters a universal state. In each case, the machine “guesses” the interpretation of y_i as the encoding of a node, a path or a register value $c(y_i)$ in the domain. (Since encodings of paths are of exponential size, this alternating machine requires at least exponential time to work). Finally, the machine verifies that $\psi(\bar{v}, \bar{\rho}, \bar{\lambda}, c(y_1), \dots, c(y_m))$ holds, and if that is the case it accepts. We show next that the latter can be done in exponential time. Notice that this implies that the whole process can be performed in exponential time.

We start with the case of the atomic formulas in ψ . In order to check whether the element assigned to a variable belongs to the interpretation of **Nodes** in $\mathcal{M}'_{G, \bar{\rho}}$, we only have to check that the encoding of this element does not start with a p or an r . In order to check whether the element belongs to the interpretation of **Paths** (resp., **Registers**), it is sufficient to check that its encoding starts with a p (resp., r). In order to check whether the elements a, b, c assigned to variables x, π, y , respectively, are such that (a, b, c) belongs to the interpretation of **Endpoints**, we only have to check that b is the encoding of a path, a and c are encodings of nodes, and that b is a path from a to c . Finally, in order to check whether the elements (a, b, c) assigned to a variable belongs to the interpretation of e_i ($1 \leq i \leq m$), we only have to check that a, c are register values (i.e. their encodings start with symbol r), that b encodes a path ρ (i.e. its encoding starts with p), and that the bit that corresponds to tuple (e_i, a, c) in the part of the address b that encodes the set $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$ that ρ satisfies is set to 1.

Thus, the value of the atomic formulas involved in $\psi(\bar{v}, \bar{\rho}, \bar{\lambda}, c(y_1), \dots, c(y_m))$ can be computed in polynomial time (in the size of $\psi(\bar{v}, \bar{\rho}, \bar{\lambda}, c(y_1), \dots, c(y_m))$). But since ψ is a polysize Boolean combination of atomic formulas, the value of $\alpha(\bar{v}, \bar{\rho}, \bar{\lambda}, c(y_1), \dots, c(y_m))$ can be computed in polynomial time from the values of the atomic formulas. We conclude that computing the value of $\alpha(\bar{v}, \bar{\rho}, \bar{\lambda}, c(y_1), \dots, c(y_m))$ can be done in polynomial time.

There is, however, one small issue that requires explanation in order for the previous procedure to work properly. Assume that the procedure “guesses” the interpretation of a variable y_i in ϕ_τ to be the encoding of a path in G from v to v' that satisfies $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$. Then it is necessary to check that, if the encoding implies that this path is $\rho_{\mathcal{E}, v, v'}^i$, then $i \leq c_{\mathcal{E}, v, v'}$. In order to do so, the procedure needs to check, in a subroutine, whether there exist i different paths from v to v' that satisfy \mathcal{E} . The next claim shows that this can be done in exponential space, which finishes the proof of the theorem.

Claim 5.2.2. For each pair $v, v' \in V$, $\mathcal{E} \subseteq \{e_1, \dots, e_m\} \times D_\perp^k \times D_\perp^k$, and $i \leq k + |\bar{\rho}|$, one can check in EXPSPACE whether there are i distinct paths in G from v to v' that satisfy \mathcal{E} .

Proof. Let $\#$ be a symbol not in Σ and denote by $\Sigma_\#$ the alphabet $\Sigma \cup \{\#\}$. Let $\mathcal{A}_{v, v'}$ be the automaton over alphabet $\{v\} \cup (\Sigma_\# \times V)$ defined as follows. The set of states is the disjoint union of V with a new state s . The initial state of \mathcal{A} is s and the final state is v' . Further, the transition relation of \mathcal{A} is defined as follows: (1) For every edge $(v_1, a, v_2) \in E$ there is a transition in \mathcal{A} from v_1 to v_2 labeled (a, v_2) , (2) for every node $v_1 \in V$ there is a transition from v_1 to v_1 in \mathcal{A} labeled $(\#, v_1)$, and (3) there is a transition in \mathcal{A} from s to v labeled v . Intuitively, $\mathcal{A}_{v, v'}$ accepts exactly those strings of the form $v(a_1, v_1)(a_2, v_2) \cdots (a, v')$ such that $va_1v_1a_2v_2 \cdots av'$ is a path in G from v to v' , when we allow paths to loop arbitrarily many times on $\#$ -labeled nodes.

Let $\mathcal{A}_{v,v'}^i$ be the automaton over alphabet $\{v^i\} \cup (\Sigma_{\#} \times V)^i$ defined as follows: The set of states is $V^i \cup \{s^i\}$, the initial state is s^i and the final state is $(v')^i$. There is a transition in \mathcal{A}^i from $\bar{u} = (u_1, \dots, u_i)$ to $\bar{w} = (w_1, \dots, w_i)$ labeled $\bar{t} = (t_1, \dots, t_i)$ iff there is a transition labeled t_ℓ from u_ℓ to w_ℓ in $\mathcal{A}_{v,v'}$, for each $1 \leq \ell \leq i$. (Notice that $\mathcal{A}_{v,v'}^i$ is not exactly the i -th product of $\mathcal{A}_{v,v'}$ with itself, as $\mathcal{A}_{v,v'}^i$ does not contain all states in such a product). Clearly, $\mathcal{A}_{v,v'}^i$ is of exponential size but the size of each one of its states is polynomial. Furthermore, it is decidable in polynomial time whether there exists a transition labeled \bar{t} from state \bar{u} to \bar{w} in $\mathcal{A}_{v,v'}^i$.

Define now an automaton $\mathcal{A}'_{v,v'}$ that is the restriction of $\mathcal{A}_{v,v'}^i$ to those strings

$$v^i(w_1^1, \dots, w_i^1) \cdots (w_1^p, \dots, w_i^p)$$

over alphabet $\{v^i\} \cup (\Sigma_{\#} \times V)^i$ that satisfy the following:

- For each $1 \leq \ell \leq i$, if for some $1 \leq j < p$ it is the case that $w_\ell^j = (\#, v')$, for some $v' \in V$, then for each $j < k \leq p$ it is the case that $w_\ell^k = (\#, v')$.
- For each $1 \leq \ell, t \leq i$, if $\ell \neq t$ then the strings $vw_\ell^1 \cdots w_\ell^p$ and $vw_t^1 \cdots w_t^p$ over alphabet $\{v\} \cup (\Sigma_{\#} \times V)$ are different.

The first condition says that each projection of a string accepted by $\mathcal{A}'_{v,v'}$ represents a path in G from v to v' that loops only on v' and only at the end of the path. The second condition ensures that any two distinct projections of a path accepted by $\mathcal{A}'_{v,v'}$ represent different strings.

It is not hard to prove that the language accepted by $\mathcal{A}'_{v,v'}$ is nonempty iff there exist i distinct paths in G from v to v' . Further, it is also not hard to see that $\mathcal{A}'_{v,v'}$ is of exponential size but the size of each one of its states is polynomial; and it is decidable in polynomial time whether there exists a transition labeled \bar{t} from state \bar{q} to state \bar{q}' in $\mathcal{A}'_{v,v'}$.

Using techniques in [10], it is also possible to construct in exponential time an NFA $\mathcal{A}_{v,v',e_i,\lambda,\lambda'}$ ($e_i \in \{1, \dots, m\}$, $\lambda, \lambda' \in D_{\perp}^k$) over alphabet $v \cup (\Sigma_{\#} \times V)$, that accepts precisely the strings w accepted by $\mathcal{A}_{v,v'}$ such that the path ρ from v to v' in G that is represented by w satisfies $e_i(\rho, \lambda, \lambda')$. (The main idea is to construct $\mathcal{A}_{v,v',e_i,\lambda,\lambda'}$ in such a way that, at each position while reading w , it keeps in its state the k -tuple of data values that is stored in the registers of e_i). The set of states of $\mathcal{A}_{v,v',e_i,\lambda,\lambda'}$ is of exponential size, but each particular state can be represented using polynomial space. Further, deciding if there is a transition between two states of $\mathcal{A}_{v,v',e_i,\lambda,\lambda'}$ can be done in polynomial time. This means that for each $\mathcal{A}_{v,v',e_i,\lambda,\lambda'}$, its complement can be constructed in double exponential time, with each state using only exponential space.

It is not hard to see, then, that one can construct in double exponential time an automaton $\mathcal{A}_{\mathcal{E},v,v'}$ over alphabet $\{v^i\} \cup (\Sigma_{\#} \times V)^i$ that does the following: It starts from $\mathcal{A}'_{v,v'}$, and restricts acceptance to strings $v^i(w_1^1, \dots, w_i^1) \cdots (w_1^p, \dots, w_i^p)$ over alphabet $\{v^i\} \cup (\Sigma_{\#} \times V)^i$ such that:

- For each $1 \leq j \leq i$ and tuple $(e_\ell, \lambda, \lambda') \in E$, it is the case that $vw_j^1 \cdots w_j^p$ is accepted by $\mathcal{A}_{v,v',e_\ell,\lambda,\lambda'}$.
- For each $1 \leq j \leq i$ and tuple $(e_\ell, \lambda, \lambda') \notin E$, it is the case that $vw_j^1 \cdots w_j^p$ is accepted by the complement $\mathcal{A}_{v,v',e_\ell,\lambda,\lambda'}$.

Further, each state in $\mathcal{A}_{\mathcal{E},v,v'}$ can be represented using exponential space and checking whether there is a transition between two given states of $\mathcal{A}_{\mathcal{E},v,v'}$ can be done in polynomial time.

It is clear that there exist i distinct paths in G from v to v' that satisfy \mathcal{E} if and only if $\mathcal{A}_{\mathcal{E},v,v'}$ accepts at least one string. But we can check $\mathcal{A}_{\mathcal{E},v,v'}$ for nonemptiness in EXPSPACE using a standard “on-the-fly” argument. This finishes the proof of the claim. \square

This also finishes the proof that $\text{EVAL}(RL)$ is in EXPSPACE. Now we show that $\text{EVAL}(RL)$ is EXPSPACE-hard.

For all constants f_0 , we provide a reduction from the class of problems solvable by a Turing machine using a tape of size $2^{f_0 n}$ given an input word of size n . There are a Turing machine M and a constant f_0 such that the following problem is EXPSPACE-hard: given a word w of size n , is there an accepting run of M over w using at most $2^{f_0 n}$ cells? We prove that there is a formula $\phi \in RL$ such that for all words w of size n , there are a formula ϕ_w and a graph G_w such that

$$G_w \models \phi_w \quad \text{iff} \quad \text{there is an accepting run of } M \text{ over } w \text{ using at most } 2^{f_0 n} \text{ cells.}$$

Let $(\Sigma, Q, \delta, q_0, q_f)$ be the Turing machine M , where Σ is the alphabet consisting of the input alphabet and the blank symbol B , q_0 is the initial state, q_f is the final state and $\delta : Q \times \Sigma \rightarrow (Q \times \Sigma \times \{L, R\})$ is the transition map, where L stands for “left” and R stands for “right”.

The formula ϕ_w that we associate with the machine M and a word w is a formula of the form

$$\exists \pi \psi_w(\pi),$$


where ψ_w is a formula that does not contain any quantification over path variables. The formula $\psi_w(\pi)$ expresses that the path π in the graph G_w encodes an accepting run of M over the word w .

As in the proof of Theorem 4.1, we encode a configuration C of a run of M in the following way. Suppose that the content of the tape is the word $w' = w'_1 \dots w'_{2^{f_0 n}}$, the head is scanning the cell number i_0 and the machine is in state q . We encode the configuration C by the word $e_C = d_1^C \& \dots d_{2^{f_0 n}}^C \&$ where $\&$ plays the role of a delimiter and each d_i^C encodes the information in cell number i . More precisely, if Δ is the alphabet $(Q \cup \{\$\}) \times \Sigma$, we define d_i^C as the word

$$c(i) (q'_i, w'_i),$$


where $c(i)$ and q'_i are defined as follows. The word $c(i)$ is the binary representation of the number i . The letter w'_i is the content of the cell i . The letter q'_i is equal to $\$$ if the head is not scanning the cell number i ; otherwise, q'_i is equal to the state q . That is, $q'_{i_0} = q$ and for all $i \neq i_0$, $q'_i = \$$. We encode a run $C_0 C_1 \dots$ as the word $e_{C_0} \# e_{C_1} \# \dots$. We define the formula $\psi_w(\pi)$ and the graph G_w in such a way that a path π satisfies ψ_w iff the label of π is the encoding of an accepting run of M over w .

The graph G_w is (almost) the same graph as the graph in the proof of Theorem 4.1 in the case where $k = 1$. That is, G_w is obtained by “linking” two graphs I_w and H . If $C_0 C_1 \dots$ is an accepting run with associated path π , the trace of π on I_w will correspond to the encoding of the initial configuration C_0 , while the trace of π on H is the encoding of the run $C_1 C_2 \dots$. The graph H is given by




Recall that the set $\{d_i : 1 \leq i \leq l\}$ is defined as $(Q \cup \{\$\}) \times \Sigma$. Consider a simple path π' starting from the node with data value 1 and ending in the node with outgoing edges with labels $\&$ and $\#$. Its label is of the form $c(i)(q', a)$; that is, it is the encoding of the information in a cell with number i . Hence, the label of a path π starting and ending in the node with data value 1 satisfies $c(i_1)d_1(\& \vee \#)c(i_2)d_2(\& \vee \#) \dots c(i_k)d_k(\& \vee \#)$, where each d_j is the encoding of the information of a cell and each $c(i_j)$ is the encoding of the number i_j . We define the formula ψ in such a way that if $\psi(\pi)$ holds, the succession of encodings of cells describe a run of M .


Next we define the graph I_w where we encode the initial configuration. Suppose that $w = w_1 \dots w_n$. For all $1 \leq i \leq n$, we introduce a graph K_i describing the cell number i in the initial configuration. If $b_1 \dots b_n$ is the binary encoding of the number i , the graph K_i is given by




where $q'_1 = q_0$ and $q'_i = \$$ if $i \neq 1$. The label of the longest path of K_i is exactly the encoding of cell number i in the initial configuration. Next we define the graph K which allows us to encode the cells with number $> n$ in the initial configuration. The graph K is given by



The label of a simple path from the node with data value 1 to the node with outgoing edge with label $\&$ is of the form $c(i)(\$, B)$; that is, it is the encoding of an unscanned cell with a blank symbol. In particular, it is the encoding of all the cells with number $> n$ in the initial configuration. The graph I_w is obtained by linking together the graph K_1, \dots, K_n, K in the following way



The arrow with a label s is an arrow pointing to the node of K_1 with data value 1. Each arrow with label $\#$ between two graphs is from the “right-most” node of the first graph to the “left-most” node of the second graph. Finally, the graph G_w is defined by



where the edge with label $\#$ is an edge from the “right-most” node of K to the node with data value 1 in H . We define now a formula ψ_π such that

$$G_w \models \psi(\pi) \quad \text{iff}$$

$l(\pi)$ is the encoding of an accepting run of M over w using at most $2^{f_{0n}}$ cells,

where $l(\pi)$ is the label of π . In fact it is easier to define a formula $\chi(\pi)$ such that

$$G_w \models \chi(\pi) \quad \text{iff}$$

$l(\pi)$ is not the encoding of an accepting run of M over w using at most $2^{f_{0n}}$ cells. (5.1)

Suppose that π is a path through G_w . The label of π is not the encoding of an accepting run of M over w using at most $2^{f_{0n}}$ cells iff at least one of the following conditions holds.

- (i) The first letter of $l(\pi)$ is not the initial letter s or the run never reaches a final state, i.e. there is no pair of the form (q_f, a) for some a , occurring in $l(\pi)$.
- (ii) The symbol $\#$ is not at the “right place”:
 - either after we reach the symbol $\#$ (i.e. we are going to enter the encoding of a new configuration), the label contains the binary encoding of a number $\neq 1$,
 - or after the binary encoding of the number $2^{f_{0n}}$ (that is, after encoding the information of the last cell), the symbol $\#$ does not occur in the label of π . Since $\#$ is used as a delimiter between encodings of configurations, this means that although we finished encoding the last cell of a configuration, we do not move to a new configuration.
- (iii) There is a substring $c(i)d\&c(j)d'$ (where $i < 2^n$) such that j is not the successor of i . That is, after encoding the information of cell number i , we do not encode the information of cell number $i + 1$.
- (iv) Finally, there is a string $e_C\#e_{C'}$ of D_π such that C and C' are not successive configurations.

Expressing cases (i) and (ii) is fairly easy. We concentrate on (iii) and (iv). We start by showing how to express case (iv). Suppose that $c(i)$ and $c(j)$ are two successive binary encodings occurring in the label of π . Suppose $c(i) = b_1 \dots b_{f_{0n}}$ and $c(j) = b'_1 \dots b'_{f_{0n}}$. Then j is not the successor of i iff one of the following holds.

- (a) For some k , $b_k \dots b_{f_{0n}}$ is equal to $1 \dots 1$ and b'_k is not equal to 0.
- (b) For some k , $b_k \dots b_{f_{0n}}$ is equal to $01 \dots 1$ and b'_k is not equal to 1.
- (c) For some k , $b_k = 1$, 0 occurs in $b_{k+1} \dots b_{f_{0n}}$ and b'_k is not equal to 1.
- (d) For some k , $b_k = 0$, 0 occurs in $b_{k+1} \dots b_{f_{0n}}$ and b'_k is not equal to 0.

We show how we can express (a); the other cases are similar. Case (a) is expressed by the following REM

$$(\Lambda \setminus \Delta_f)^* \downarrow r_1.1^* \Delta \& \{0, 1\}^* [r_1^-] 1 \Lambda^*.$$

where Δ_f is the set $\{q_f\} \times \Sigma$ and Λ is the alphabet of the graph G_w . In the register r_1 , we store the number k such that $b_k \dots b_{f_0n} = 1 \dots 1$ (that is, we only go through edges with label 1 until we reach an edge with label in Δ). When we reach again the node with data value k , the label of the outgoing edge is 1, expressing that $b'_k = 1$.

Finally we look at case (iv), i.e. how to express that there is a string $e_C \# e_{C'}$ of D_π such that C and C' are not successive configurations. This might happen for several reasons: (A) either we did not modify properly the content of a cell or move properly the head or (B) we modified the content that was supposed to remain constant. We only treat case (A), as the other case can be handled in a similar way (note that the proof of Theorem 5.3 is very similar to this proof and there, we will treat case (B)). In case (A), we also only consider the case of a transition moving the head to the right, the other case being symmetric.

So suppose that $\delta(q, a) = (q', b, R)$. As in the proof of Theorem 4.1, we use a slightly different definition of a run of a Turing machine (which is equivalent to the usual definition, but helps us to keep our formulas simpler). We assume that if $\delta(q, a) = (q', b, R)$ and the machine scans a cell c with content a , then in the next state, the machine scans the successor c' of c , the content of c' is b , while the content of c is a .

Let $(\Lambda_{\#})^*$ be the set of words over Λ that contain at most one occurrence of $\#$. We define $e_{(q,a)}^R$ as the following REM

$$\Lambda^* \{0, 1\} \downarrow r_1 \dots \{0, 1\} \downarrow r_n.(q, a)(\Lambda_{\#})^* \{0, 1\} [r_1^-] \dots \{0, 1\} [r_{f_0n}^-] \Delta \& \{0, 1\}^* (\Delta \setminus \{(q', b)\}) \Lambda^*.$$

We store in the registers r_1, \dots, r_{f_0n} the binary encoding of a number i . That number is the number of a scanned cell with content a and the current state is q . After the next occurrence of $\#$ (after reading a word in $(\Lambda_{\#})^*$), we enter a new configuration. In that new configuration, we reach the cell number i when we read a sequence matching the contents of the registers. The encoding of the next cell (after reading a sequence in $\Delta \&$) must consist of the binary encoding of a number followed by a symbol in Δ that is not (q', b) .

If Δ_R is the set $\{(q, a) : \delta(q, a) = (q', b, R) \text{ for some } q', b\}$ and λ is a register of size f_0n , we define the formula

$$\bigvee \{ \exists \nu e_{(q,a)}^R(\pi, \perp, \nu) : (q, a) \in \Delta_R \},$$

taking care of case (iv)(A). □

The increase in expressiveness of RL over REM has an important cost in data complexity, which becomes intractable:

Theorem 5.3. $\text{EVAL}(RL)$ is in PSPACE in data complexity. Furthermore, there is a finite alphabet Σ and a RL formula ϕ over Σ and a single register r , such that $\text{EVAL}(RL, \phi)$ is PSPACE-hard. In addition, the latter holds even if the input is restricted to graph databases.

Proof. The upper bound follows as a corollary to the proof of the upper bound in Theorem 5.2. In fact, it is clear that the whole process can be carried in PSPACE if we assume a fixed RL query (in fact, to obtain a PSPACE upper bound we do not need more than to fix the number of registers used in the query).

For the lower bound, we define a formula ϕ in RL such that for all constants f_0 , there is a reduction from the class of problems solvable by a Turing machine using a tape of size f_0n given an input word of size n , to the evaluation problem of ϕ . More precisely, there are

a Turing machine M and a constant f_0 such that the following problem is PSPACE)-hard: given a word w of size n , is there an accepting run of M over w using at most $f_0 n$ cells?

We prove that the formula ϕ is such that for all words w of size n , there is a graph G_w such that

$$G_w \models \phi \quad \text{iff} \quad \text{there is an accepting run of } M \text{ over } w \text{ using at most } f_0 n \text{ cells.}$$

Let $(\Sigma, Q, \delta, q_i, q_f)$ be the Turing machine M , where Σ is the input alphabet together with a blank symbol B , q_0 is the initial state, q_f is the final state and $\delta : Q \times \Sigma \rightarrow (Q \times \Sigma \times \{L, R\})$ is the transition map, where L stands for “left” and R stands for “right”.

The formula ϕ that we associate with the machine M is a formula of the form

$$\exists \pi \psi(\pi),$$

where ψ is a formula that does not contain any quantification over path variables. Given a word w , the path π in the graph G_w will encode an accepting run of M over the word w .

Given a word w of size n , consider a configuration C of the run of M over w where the contents of the tape is the word $w' = w'_1 \dots w'_{f_0 n}$, the head is scanning the cell number i_0 and the machine is in state q . Similarly to the proof of Theorem 4.1, we encode the configuration C by the word


$$e_C = Nd_1^C Nd_2^C \dots Nd_{f_0 n}^C$$

where each d_i^C encodes the information in cell number i in the configuration C . We define d_i^C as the pair (q'_i, w'_i) , where q'_i is defined as follows. The letter w'_i is the contents of the cell i . The letter q'_i is equal to $\$$ if the head is not scanning the cell number i ; otherwise, q'_i is equal to the state q . That is, $q'_{i_0} = q$ and for all $i \neq i_0$, $q'_i = \$$.

The run of M over the word w is a (possibly infinite) sequence of configurations of the form $C_0 C_1 \dots$. We encode the run as the word $e_{C_0} \# e_{C_1} \# \dots$, where $\#$ plays the role of a delimiter. We will define the formula $\psi(\pi)$ and the graph G_w in such a way that a path π satisfies ψ iff the label of π is the encoding (as defined above) of an accepting run of M over w .

We think of a path π encoding a run of M over w as consisting of two parts. The label of the first part contains the encoding e_{C_0} of the initial configuration C_0 . The label of the second part contains the encoding $e_{C_1} \# e_{C_2} \# \dots$ of the remaining part of the run. The first part of the path π is a path in a subgraph I_w of G_w , while the second part is a path in the subgraph H (independent of w) of G_w . The graph G_w will be obtained by adding an edge from a node in I_w to a node in H .

The graph I_w is given by




In the graph above the data value i carried by a node v indicates that the label of the outgoing edge of v is $d_i^{C_0}$ (where C_0 is the initial configuration and $d_i^{C_0}$ is defined as above). Recall that $d_i^{C_0}$ indicates the contents of the cell number i and whether or not the head is

scanning that cell. There is a unique path π_0 from the node with no incoming edge to the unique node with no outgoing edge. The label of π_0 is

$$sN(q_i, w_0) (\$, w_1)N(\$, w_2) \dots N(\$, w_n)N(\$, B) \dots N(\$, B),$$

that is, the word $s.e_{C_0}$, where e_{C_0} is the encoding of the initial configuration C_0 .

We define now the graph H encoding the remaining part of the run of the machine.



For all $1 \leq i \leq f_0n$, all $q \in Q$ and all $a \in \Sigma$, the node with data value i admits outgoing edges with label (q, a) and $(\$, a)$. A path from the “left-most” node to the “right-most” node that does not go through the edge with label $\#$ has a label of the form

$$Nd_1Nd_2 \dots Nd_{f_0n}N,$$


where each d_i belongs to $(Q \cup \{\$\}) \times \Sigma$, that is, an encoding of a configuration of the machine.

Hence, a path π' from the “left-most” node to the “right-most” node of H (possibly going through the edge with label $\#$) has a label of the form

$$e_{C_1}\#e_{C_2}\# \dots e_{C_l},$$

where each e_{C_i} is the encoding of a configuration of the machine.

We are now ready to define the graph G_w .



The edge with label $\#$ is an edge from the unique node in I_w with no outgoing edge to the “left-most” node in H . We define now the formula ψ . Let Λ be the alphabet $[(Q \cup \{\$\}) \times \Sigma] \cup \{\#, s\}$. The formula ψ must be such that

$$G_w \models \psi(\pi) \quad \text{iff} \quad l(\pi) \text{ is of the form } se_{C_0}\#e_{C_1} \dots e_{C_p}\Lambda^*,$$

where $l(\pi)$ is the label of π and $C_0 \dots C_p$ is an accepting run of the machine over w . In fact, it will be intuitively easier to first define a formula $\chi(\pi)$ such that

$$G_w \models \chi(\pi) \quad \text{iff} \quad l(\pi) \text{ is not of the form } se_{C_0} \# e_{C_1} \dots e_{C_p} \Lambda^*, \quad (5.2)$$

and define ψ as $\neg\chi$. The formula χ is obtained as a disjunction of the following subformulas.

- First $l(\pi)$ might not satisfy $se_{C_0} \# e_{C_1} \dots e_{C_p} \Lambda^*$ because: (i) it does not start with the letter s or (ii) it does not contain the encoding of a final configuration, i.e. it does not contain any occurrence of a pair of the form (q_f, a) for some $a \in \Sigma$. Case (i) is expressed by the formula

$$\bigvee \{e_\gamma(\pi, \perp, \perp) : \gamma \in \Lambda, \gamma \neq s\},$$

where e_γ is the REM $\gamma \Lambda^*$. Case (ii) is expressed by the formula

$$(\Lambda \setminus \Delta_f)^*(\pi, \perp, \perp),$$

where Δ_f is the set $\{q_f\} \times \Sigma$. This formula express that there is no pair of the form (q_f, a) , for some $a \in \Sigma$.

- Next $l(\pi)$ might not be of the “right form” because it contains a substring of the form $e_C \# e_{C'}$ occurring before a pair of the form (q_f, a) and such that C and C' are not consecutive configurations. This might happen for several reasons: (A) either we did not modify properly the contents of a cell or move properly the head or (B) we modified the contents that was supposed to remain constant.

We will only treat one case. As we treated case (A) in the proof of Theorem 5.2, here we treat case (B). As in the proof of Theorem 5.2, we also consider a slightly different definition of a run of a Turing machine (which is equivalent to the usual definition, but helps us to keep our formulas simpler). We assume that if $\delta(q, a) = (q', b, R)$ and the machine scans a cell c with content a , then in the next state, the machine scans the successor c' of c , the content of c' is b , while the content of c is a .

In case (B), we make the following case distinction.

- Suppose first that case (B) happened because we modified the contents of the cell that was scanned by the head in configuration C . Note that where when moving from C to C' , by definition of a run, we cannot modify the contents of the cell scanned in configuration C .

Let (q, a) be a pair in $Q \times \Sigma$ and let Δ be the set $(Q \cup \{\$\}) \times \Sigma$. We define also $\Lambda_{! \#}^*$ as the set of words over Λ that contain at most one occurrence of the symbol $\#$. We let $e_{(q,a)}(\pi)$ be the formula

$$(\Lambda \setminus \Delta_f)^* \downarrow r.(q, a) \Lambda_{! \#}^* [r^=](\Delta \setminus \{(q, a) : q \in Q \cup \{\$\}\}) \Lambda^*.$$

Before we reach a final state, we store in register r the number of a cell that is scanned in the current configuration and with contents a . In the next configuration (after reading a word in $\Lambda_{! \#}^*$), when we read the same cell, it does not contain a . That is, the label of the edge is not a pair of the form (q, a) .

The formula

$$\bigvee \{\exists \nu e_{(q,a)}(\pi, \perp, \nu) : (q, a) \in Q \times \Sigma\},$$

takes care of the cases where from moving from one configuration to the next, we modified the contents of the cell scanned in the first configuration.

- Next suppose that when moving from C to C' , we modified the contents of a cell that was not scanned in C . Suppose also that according to the machine, we were not supposed to modified that contents.

Let a be a letter in Σ . Let Δ_L be the set of pairs (q, a) such that $\delta(q, a) = (q', b, L)$ and let Δ_R be the set of pairs (q, a) such that $\delta(q, a) = (q', b, R)$. We define $e_a(\pi)$ as the REM

$$(\Lambda \setminus \Delta_f)^*(\Delta \setminus \Delta_R) \downarrow r.(\$, a)N(\Delta \setminus \Delta_L)\Lambda_{! \#}^*[r^-](\Delta \setminus (\$, a))\Lambda^*.$$

Before we reach a final state, in a configuration C , we store in register r the number i of an unscanned cell with contents a . We assume that if the cell with number $(i - 1)$ (if it exists) is scanned in C , then the head is not moving to the right (i.e. to cell number i) in the next configuration. That is, the pair describing the cell number $(i - 1)$ in C is not a pair in Δ_R . We also assume that if cell number $(i + 1)$ (if it exists) is scanned, then the head is not moving to cell number i in the next configuration, i.e. cell number $(i + 1)$ is not described by a pair in Δ_L .

Next we express that in the next configuration, it is not the case that the cell with number i is an unscanned cell with contents a . That is, after reading a word in $\Lambda_{! \#}^*$, when we see again the node stored in the register, the label of the edge is not the pair $(\$, a)$.

The formula

$$\bigvee \{ \exists \nu e_a(\pi, \perp, \nu) : a \in \Delta \},$$

takes care of case (B) in the case where the cell modified is a cell unscanned by the head.

This finishes the proof of the theorem. \square

In the next section we introduce an interesting language, based on a restriction of RL, that is tractable in data complexity, and thus better suited for database applications. This language is a proper extension of REM. But before that, we make some important remarks about the expressive power of RL.

Expressive power of RL. We now look at the expressive power of the logic RL. It was proven in [8] that CRPQ is not subsumed by WL. Since RL subsumes CRPQ $^\neg$, it follows that RL is not subsumed by WL. On the other hand, WL is also not subsumed by RL due to Theorem 4.1, Theorem 5.2, and the standard time/space hierarchy theorem from complexity theory. Therefore, we have the following proposition:

Proposition 5.4. The expressive power of WL and RL are incomparable.

On the other hand, we shall argue now that many natural queries about the interaction between data and topology are also expressible in RL. The aforementioned query (Q) is one such example. We shall now mention other examples: hamiltonicity (H), the existence of an eulerian trail (E), bipartiteness (B), and connected graphs with an even number of nodes (C2). The first two are expressible in WL, while (B) and (C2) are not known to be expressible in WL. We conjecture that they are not.

We now show how to express in RL the existence of a hamiltonian path in a graph; the query (E) can be expressed in the same way but with two registers (to remember edges, i.e., consisting of two nodes). This is done with the following formula over $\Sigma = \{a\}$ and a single register r :

$$\exists \pi (\forall \lambda \forall \lambda' \neg e_1(\pi, \lambda, \lambda') \wedge \forall \lambda (\lambda \neq \perp \rightarrow e_2(\pi, \lambda, \lambda))),$$

where $e_1 := a^* \cdot (\downarrow r.a^+[r^-]) \cdot a^*$ is the REM that checks whether in a path some node is repeated (i.e., that it is not a simple path), and $e_2 := a^*[r^-]a^*$ is the REM that checks that

the node stored in register r appears in a path. In fact, this query expresses that there is a path π that it is simple (as expressed by the formula $\forall\lambda\forall\lambda'\neg e_1(\pi, \lambda, \lambda')$), and every node of the graph database is mentioned in π (as expressed by the formula $\forall\lambda(\lambda \neq \perp \rightarrow e_2(\pi, \lambda, \lambda))$).

We now show how to express in RL the bipartiteness property from graph theory. An undirected graph $G = (V, E)$ is *bipartite* if its set of nodes can be partitioned into two sets S_1 and S_2 such that, for each edge $(v, w) \in E$, either (i) $v \in S_1$ and $w \in S_2$, or (ii) $v \in S_2$ and $w \in S_1$. It is well-known that a graph database is bipartite if and only if it does not have cycles of odd length. The latter is expressible in RL since the existence of an odd-length cycle can be expressed as $\exists\pi\exists\lambda\exists\lambda'e(\pi, \lambda, \lambda')$, where $e = \downarrow r.a(aa)^*[r=]$.

We now show how to express in RL that a graph database is a connected graph with an odd number of nodes. To this end, it is sufficient and necessary to express the existence of a hamiltonian path π with an even number of edges in the graph. But this is a simple modification of our formula for expressing hamiltonicity: we add the check that π has an even number of edges by adding the conjunct $e(\pi, \nu, \nu')$, where $e = (aa)^*$, and close the entire formula under existential quantification of ν and ν' .


5.3. Tractability in data complexity. Let RL^+ be the *positive* fragment of RL, i.e., the logic obtained from RL by forbidding negation but adding conjunctions (as they were not explicitly present in RL). It is easy to prove that the data complexity of the evaluation problem for RL^+ is tractable (NLOGSPACE). This fragment contains the class of *conjunctive* REMs, that has been previously identified as tractable in data complexity [10]. However, the expressive power of RL^+ is limited as the following proposition shows.

Proposition 5.5. The query (Q) from the introduction is not expressible in RL^+ .

Proof. Recall that Q is the following query: Find pairs of nodes x and y such that there is a node z and a path π from x to y in which each node is connected to z . Suppose for contradiction that there is a formula ϕ in RL^+ over an alphabet Σ and registers r_1, \dots, r_k , expressing $\exists x\exists y Q$. We may assume that ϕ is of the form

$$\exists x_1 \dots \exists x_{n_1} \exists \pi_1 \dots \exists \pi_{n_2} \exists \nu_1 \dots \exists \nu_{n_3} \psi$$

where ψ is a disjunction of conjunctions of atoms. Let $G = (V, E, \kappa)$ be the following graph



where each edge is labeled with a . The query $\exists x\exists y Q$ is true in G ; hence, the formula ϕ must be true in G . That is, there is an assignment α mapping each variable x_i to a node in G , each path variable π_i to a path ρ_i in G and each variable ν_i to a tuple in $\{\perp, \$, 1, \dots, n_2+1\}^k$ such that

$$(G, \alpha) \models \psi.$$

Let G' be the graph (V, E', κ) where E' is the set


$$\{(i, a, i+1) : 1 \leq i < n_2\} \cup \{(i, a, \$) : \text{for some } 1 \leq j \leq n_2, (i, \$) \text{ occurs in } \rho_j\}$$

That is, we delete the edges $(i, \$)$ that α “does not use”. By definition of E' , the formula ψ remains true in G' under the assignment α . In particular, ϕ is true in G' . This implies that $\exists x, y Q$ holds in G' .

Now, for all $1 \leq j \leq n_2$, there is at most one natural number i such that $(i, a, \$)$ occurs in ρ_j . This is simply because there is no path going through edges $(i, a, \$)$ and $(i', a, \$)$ if $i \neq i'$. This implies that the set

$$\{(i, a, \$) : \text{for some } 1 \leq j \leq n_2, (i, \$) \text{ occurs in } \rho_j\}$$

contains at most n_2 edges. Since G admits $n_2 + 1$ edges of the form $(i, a, \$)$, there must be an edge $(i_0, a, \$)$ occurring in G , but not in G' . This means that G' is a graph of the form



In particular, $\exists x, y Q$ is not true in G' , which contradicts the fact that ϕ is true in G' . \square

On the other hand, increasing the expressive power of RL^+ with some simple forms of negation leads to intractability of query evaluation in data complexity:

Proposition 5.6. There is a finite alphabet Σ and REMs e_1, e_2, e_3, e_4 over Σ and a single register r , such that $\text{EVAL}(\text{RL}, \phi)$ is PSPACE-complete, where ϕ is either $\exists \pi \exists \lambda \neg (e_1(\pi, \perp, \lambda) \vee e_2(\pi, \perp, \perp))$ or $\exists \pi \forall \lambda \neg (e_3(\pi, \perp, \lambda) \vee e_4(\pi, \perp, \perp))$.

Proof. By the proof of Theorem 5.3 (and using the same notation), we know that for every Turing machine M , there is a formula $\chi(\pi)$ such that for all words w of size n , there is a graph G_w (of size polynomial in n) such that

$$G_w \models \exists \pi \neg \chi(\pi) \quad \text{iff} \quad \text{there is an accepting run of } M \text{ over } w \text{ using at most } cn \text{ cells.}$$

Moreover, the formula $\chi(\pi)$ is a formula of the form

$$\bigvee \{e_i(\pi, \perp, \perp) : i \in I\} \vee \bigvee \{\exists \lambda f_j(\pi, \perp, \lambda) : j \in J\},$$

where $\{e_i, f_j : i \in I, j \in J\}$ is a set of REMs that do not contain any \cup . Since REMs are closed under union, $\chi(\pi)$ is equivalent to a formula of the form

$$\exists \lambda (e(\pi, \perp, \perp) \vee f(\pi, \perp, \lambda)).$$

Hence, the formula $\exists \pi \neg \chi(\pi)$ is equivalent to

$$\exists \pi \forall \lambda \neg (e(\pi, \perp, \perp) \vee f(\pi, \perp, \lambda)).$$

This proves that $\text{EVAL}(\text{RL}, \phi)$ is PSPACE-complete, where $\phi = \exists \pi \forall \lambda \neg (e(\pi, \perp, \perp) \vee f(\pi, \perp, \lambda))$.

Now we prove that $\text{EVAL}(\text{RL}, \phi')$ is PSPACE-complete, where ϕ' is a formula of the form $\exists \pi \exists \lambda \neg (e'(\pi, \perp, \perp) \vee f'(\pi, \perp, \lambda))$, for some REMs e' and f' . The intuition is as follows. The difference between ϕ and ϕ' is that in ϕ , we may choose the data value that is in the register after checking that f is true. However, in ϕ' , we must be able to store any value in the register after checking that f' is true. We will make two changes to make this possible.

First, we modify the graph G_w in such a way that two arbitrary nodes are always reachable. This can be easily achieved by adding an edge from the “right-most node” of the graph H to the “left-most node” of the graph I_w (allowing to encode the initial configuration of a run). Second, we modify the REMs of ϕ in such a way that the label of a path satisfying

those REMs, encodes an accepting run and after reaching the final state, it goes through all the nodes of G_w . Hence, once we checked that the run reaches the final state, we can simply store any value in the register. We leave out the details, as the intuition is pretty simple and the details a bit tedious. \square

In the case of basic navigational languages for graph databases, it is possible to increase the expressive power – without affecting the cost of evaluation – by extending formulas with a branching operator (in the style of the class of *nested regular expressions* [5]). The same idea can be applied in our scenario, by extending atomic REM formulas in RL^+ with such a branching operator. The resulting language is more expressive than RL^+ (in particular, this extension can express query (Q)), yet remains tractable in data complexity. We formalize this idea below.

The class of *nested* REMs (NREM) extends REM with a nesting operator $\langle \cdot \rangle$ defined as follows: If e is an NREM then $\langle e \rangle$ is also an NREM. Intuitively, the formula $\langle e \rangle$ filters those nodes in a data graph that are the origin of a path that can be parsed according to e . Formally, if e is an NREM over k registers and G is a data graph, then $\llbracket \langle e \rangle \rrbracket_G$ consists of all tuples of the form $(u, \lambda, \rho = u, u, \lambda)$ such that $(u, \lambda, \rho', v, \lambda') \in \llbracket e \rrbracket_G$, for some node v in G , path ρ' in G , and k -tuple λ' over \mathcal{D}_\perp .

Let NRL^+ be the logic that is obtained from RL^+ by allowing atomic formulas of the form $e(\pi, \lambda, \lambda')$, for e an NREM. Given a data graph G and an assignment α for π, λ and λ' over G , we write as before $(G, \alpha) \models e(\pi, \lambda, \lambda')$ if and only if $\alpha(\pi)$ goes from u to v and $(u, \alpha(\lambda), \alpha(\pi), v, \alpha(\lambda')) \in \llbracket e \rrbracket_G$. The semantics of NRL^+ is thus obtained from the semantics of these atomic formulas in the expected way. The following example shows that query (Q) is expressible in NRL^+ , and, therefore, that NRL^+ increases the expressiveness of RL^+ .

Example 5.3. Over graph databases, the query (Q) from the introduction is expressible in NRL^+ using the following formula over $\Sigma = \{a\}$ and register r :

$$\phi = \exists \pi \exists \lambda ((x, \pi, y) \wedge e(\pi, \lambda, \lambda)),$$

where $e := (\langle e_1 \rangle \cdot a)^* \langle e_1 \rangle$, for $e_1 = a^*[r=]$. Intuitively, e_1 checks in a path whether its last node is precisely the node stored in register r , and thus e checks whether every node in a path can reach the node stored in register r . Therefore, the formula ϕ defines the set of pairs (x, y) of nodes, such that there is a path π that goes from x to y and a register value λ (i.e., a node λ) that satisfies that every node in π is connected to λ . \square

The extra expressive power of NRL^+ over RL^+ does not affect the data complexity of query evaluation:

Theorem 5.7. Evaluation of NRL^+ formulas is in NLOGSPACE in terms of data complexity.

Proof. Let $G = (V, E, \kappa)$ be a data graph and ϕ an NRL^+ formula. Also, let $D = \{\kappa(v) \mid v \in V\}$. We assume without loss of generality that ϕ is Boolean, that is, we study the complexity of deciding whether $G \models \phi$. In the case when ϕ is not Boolean, that is, when the input consists of G and an assignment α for ϕ over G , we simply replace each free variable η in ϕ by $\alpha(\eta)$, and then use the evaluation algorithm we describe below for the resulting formula.

Assume without loss of generality that ϕ is of the form $\exists \bar{x} \exists \bar{v} \exists \bar{\pi} \psi$, where \bar{x} is a tuple of node variables, \bar{v} is a tuple of register assignment variables, $\bar{\pi}$ is a tuple of path variables, and ψ is quantifier-free. Assume also that $\{e_1, \dots, e_m\}$ is the set of NREMs mentioned in ϕ , and

that each such NREM is over $\{r_1, \dots, r_k\}$. The evaluation algorithm does the following: It first guesses witnesses for the existentially quantified node and register assignment variables. Assume that the guess for \bar{x} is \bar{v} , where \bar{v} is a tuple of nodes of the same arity than \bar{x} , and that the guess for $\bar{\nu}$ is $\bar{\lambda}$, where $\bar{\lambda}$ is a tuple of elements in $(D \cup \{\perp\})^k$ of the same arity than $\bar{\nu}$. Clearly, both \bar{v} and $\bar{\lambda}$ can be represented using only logarithmic space (since ϕ , and therefore k , is fixed).

The algorithm then guesses a witness ξ for each existentially quantified path variable π in $\bar{\pi}$. This witness codifies all the information that we need to know about the actual path ρ that interprets $\bar{\pi}$. In our case, ξ is a string that lists (in a precise order) the endpoints u and v of ρ and the tuples (e, λ, λ') , for $e \in \{e_1, \dots, e_m\}$ and $\lambda, \lambda' \in \bar{\lambda}$, such that $(u, \lambda, \rho, v, \lambda') \in \llbracket e \rrbracket_G$. Let $\bar{\xi}$ be the witnesses for the tuple $\bar{\pi}$. It is not hard to see that $\bar{\xi}$ can be represented using logarithmic space.

Finally, the algorithm checks that the guess $(\bar{v}, \bar{\lambda}, \bar{\xi})$ satisfies ψ . Since ψ is a Boolean combination of atomic formulas, we only have to explain how to do this in logarithmic space for each atomic formula of the logic. Atomic formulas of the form $x = y$, $\nu = \nu'$ and $\nu = \perp$ are self-evident. Atomic formulas of the form $\pi = \pi'$ only require checking whether the witness of π is equal to the witness of π' . Atomic formulas of the form (x, π, y) require checking in the witness ξ of π whether its endpoints correspond to the witnesses of x and y . Finally, formulas of the form $e(\pi, \nu, \nu')$ require checking in ξ whether $(u, \lambda, \rho, v, \lambda') \in \llbracket e \rrbracket_G$, where λ and λ' are the witnesses of ν and ν' , respectively. Clearly, any of this can be done in logarithmic space.

The only thing that remains to be done is checking that $\bar{\xi}$ is consistent with G . That is, we need to check the following for each path variable π that is mentioned in $\bar{\pi}$ whose witness in $\bar{\xi}$ is ξ : There is a path ρ in G whose codification corresponds to ξ . In other words, if ξ tells us that the endpoints of the path are u and v , we need to check in G if there is a path ρ from u to v such that, for each $e \in \{e_1, \dots, e_m\}$ and $\lambda, \lambda' \in \bar{\lambda}$, it is the case that $(u, \lambda, \rho, v, \lambda') \in \llbracket e \rrbracket_G$ each time that ξ tells us so. We explain next how this can be done in NLOGSPACE combining techniques from REM and nested regular expressions evaluation.

The algorithm starts computing, for each expression of the form $\langle e \rangle$ that appears in any of the e_i 's, the set $U(e)$ of pairs of the form (w, λ) , for w a node in G and λ a k -tuple over $D \cup \{\perp\}$, such that $(w, \lambda, \rho' = w, w, \lambda) \in \llbracket \langle e \rangle \rrbracket_G$. In order to do so it proceeds recursively depending on the *nesting depth* of the expression $\langle e \rangle$, which is 1 if e contains no nested subexpression, and it is 1 more than the maximum nesting depth of any subexpression of e of the form $\langle e' \rangle$ otherwise. The algorithm starts with those expressions $\langle e \rangle$ of nesting depth one, i.e., when e is an REM (no nesting). Using techniques from [10] it is possible to compute in NLOGSPACE the set $U(e)$, for each such expression $\langle e \rangle$. Then it continues with expressions of nesting depth two. In such case, it uses the same aforementioned techniques, but each time the procedure is asked to check whether $(w, \lambda, \rho' = w, w, \lambda) \in \llbracket \langle e' \rangle \rrbracket_G$, for a subexpression $\langle e' \rangle$ of nesting depth one, it simply checks whether $(w, \lambda) \in U(e')$. The process continues iteratively in this way until all sets $U(e)$, for $\langle e \rangle$ a subexpression of any of the e_i 's, are computed. Clearly, this iterative process can be performed in NLOGSPACE.

Once the previous step is finished, the algorithm checks whether there is a path ρ from u to v such that, for each $e \in \{e_1, \dots, e_m\}$ and $\lambda, \lambda' \in \bar{\lambda}$, it is the case that $(u, \lambda, \rho, v, \lambda') \in \llbracket e \rrbracket_G$ each time that ξ tells us so. This can be done in NLOGSPACE applying the same techniques used in the previous paragraph and the knowledge provided by the sets $U(e)$. The result follows from the fact that NLOGSPACE functions are closed under composition. \square

From the proof of Theorem 5.7 it also follows that NRL^+ formulas can be evaluated in PSPACE in combined complexity.

6. CONCLUSIONS AND FUTURE WORK

We have proven that the data complexity of walk logic (WL) is nonelementary, which rules out the practicality of the logic. We have proposed register logic (RL), which is an extension of regular expressions with memory. Our results in this paper suggest that register logic is capable of expressing natural queries about interactions between data and topology in data graphs, while still preserving the elementary data complexity of query evaluation (PSPACE). Finally, we showed how to make register logic more tractable in data complexity (NLOGSPACE) through the logic NRL^+ , while at the same time preserving some level of expressiveness of RL.

We leave open several problems for future work. One interesting question is to study the expressive power of extensions of walk logic, in comparison to RL and ECRPQ^\neg from [4]. For example, we can consider extensions with regularity tests (i.e. an atomic formula testing whether a path belongs to a regular language). Even in this simple case, the expressive power of the resulting logic, compared to RL and ECRPQ^\neg , is already not obvious. Secondly, we do not know whether NRL^+ is strictly more expressive than RL. Finally, we will also mention that expressibility of bipartiteness in WL is still open (an open question from [8]). We also leave open whether the query that a graph database is a connected graph with an even number of nodes is expressible in WL.

REFERENCES

- [1] R. Angles, C. Gutiérrez. Survey of graph database models. *ACM Computing Surveys* 40(1), 2008.
- [2] P. Barceló. Querying graph databases. In *32nd ACM Symposium on Principles of Database Systems (PODS)*, pages 175-188, 2013.
- [3] P. Barceló, G. Fontaine, A. W. Lin. Expressive languages for path queries on graphs with data. In *19th Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, pages 71-85, 2013.
- [4] P. Barceló, L. Libkin, A. W. Lin, P. T. Wood. Expressive languages for path queries over graph-structured data. *ACM Transactions on Database Systems* 37(4): 31, 2012.
- [5] P. Barceló, J. Pérez, J. L. Reutter. Relative expressiveness of nested regular expressions. In *6th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW)*, pages 180-195, 2012.
- [6] E. M. Clarke, O. Grumberg, D.A. Peled. *Model checking*. MIT Press, 2000.
- [7] I. Cruz, A. O. Mendelzon, P. T. Wood. A graphical query language supporting recursion. In *ACM SIGMOD Conference*, pages 323-330, 1987.
- [8] J. Hellings, B. Kuijpers, J. van den Bussche, X. Zhang. Walk logic as a framework for path query languages on graph databases. In *15th International Conference on Database Theory (ICDT)*, pages 117-128, 2013.
- [9] M. Kaminski, N. Francez. Finite memory automata. *Theoretical Computer Science*, 134(2), pages 329-363, 1994.
- [10] L. Libkin, D. Vrgoč. Regular path queries on graphs with data. In *15th International Conference on Database Theory (ICDT)*, pages 74-85, 2012.
- [11] L. Libkin, D. Vrgoč. Regular expressions for data words. In *18th Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR)*, pages 274-288, 2012.
- [12] L. Libkin, W. Martens, D. Vrgoč. Querying graph databases with XPath. In *16th International Conference on Database Theory (ICDT)*, pages 129-140, 2013.
- [13] L. Libkin, J. L. Reutter, D. Vrgoč. Trial for RDF: adapting graph query languages for RDF data. In *32nd ACM Symposium on Principles of Database Systems (PODS)*, pages 201-212, 2013.

- [14] A. O. Mendelzon, P. T. Wood. Finding regular simple paths in graph databases. *SIAM Journal on Computing*, 24(6), pages 1235-1258, 1995.
- [15] L.E. Robertson. Structure of complexity in the weak monadic second-order theories of the natural numbers. In *6th Annual ACM Symposium on Theory of Computing (STOC)*, pages 161-171, 1974.
- [16] L. Stockmeyer. The complexity of decision problems in automata theory and logic. Ph.D. thesis, MIT, 1974.
- [17] M. Y. Vardi. The complexity of relational query languages. In *14th Annual ACM Symposium on Theory of Computing (STOC)*, pages 137-146, 1982.
- [18] P. T. Wood. Query languages for graph databases. *SIGMOD Record* 41(1), pages 50-60, 2012.