

Logica-TGD: Transforming Graph Databases Logically

Evgeny Skvortsov¹, Yilin Xia², Bertram Ludäscher² and Shawn Bowers³

¹Google LLC, WA, USA

²University of Illinois Urbana-Champaign, School of Information Sciences, IL, USA

³Gonzaga University, Department of Computer Science, Spokane, WA, USA

Abstract

Graph transformations are a powerful computational model for manipulating complex networks, but handling temporal aspects and scalability remain significant challenges. We present a novel approach to implementing these transformations using Logica, an open-source logic programming language and system that operates on parallel databases like DuckDB and BigQuery. Leveraging the parallelism of these engines, our method enhances performance and accessibility, while also offering a practical way to handle time-varying graphs. We illustrate Logica’s graph querying and transformation capabilities with several examples, including the computation of the well-founded solution to the classic “Win-Move” game, a declarative program for pathfinding in a dynamic graph, and the application of Logica to the collection of all current facts of Wikidata for taxonomic relations analysis. We argue that clear declarative syntax, built-in visualization and powerful supported engines make Logica a convenient tool for graph transformations.

Keywords

Logic rules, graph queries, graph transformations

1. Introduction

Graph transformations are a powerful and versatile method for modeling and manipulating complex systems across diverse fields, ranging from software engineering [1, 2] and social network analysis [3] to biology and chemistry [4, 5, 6]. These transformations typically operate by applying rewrite rules to a graph, altering its structure and properties. While these transformations are known for their expressiveness and flexibility, their implementation can often be complex, especially when dealing with time-varying graphs or requiring scalable solutions. Existing graph database systems often provide limited support for such evolution mechanisms, creating a gap between theoretical models and practical implementations.

Logic programming, on the other hand, provides a declarative approach to problem solving by expressing rules and relationships rather than explicitly stating control-flow. Logic-based systems (e.g., Prolog and Answer Set Programming [7]) are well-established in various areas, including (symbolic) AI and knowledge representation and reasoning. More recently, Logica (*Logic + Aggregation*) [8, 9], an open-source logic programming language and system, has emerged, which employs parallel data processing environments such as DuckDB and BigQuery. Logica combines the declarative power of logic programming with the scalability and efficiency of modern databases, offering a promising new path for tackling graph transformation challenges.

This paper explores the application of Logica to the domain of graph transformations, bridging the gap between these two paradigms. Our approach leverages Logica’s ability to process large datasets in a parallel and efficient manner, providing a novel means of implementing and executing graph transformations at scale using logical rules. We demonstrate that with a graph transformation mindset, logic programming can provide a natural, powerful, and intuitive approach for defining transformations, opening new opportunities within both communities. Crucially, Logica also enables a practical approach to addressing issues that are

difficult in classical graph transformation approaches, such as temporal transformations and seamless scalability.

To illustrate the practical nature of our approach, we introduce several example Logica programs implementing different types of graph transformations. Specifically, we show how Logica can be used to define and perform basic graph transformations including message passing, removing transitively implied graph edges, and solving (i.e., labeling or coloring) win-move graphs. We also introduce a novel approach to time-aware pathfinding (e.g., [10, 11, 12]), directly addressing the need for principled and tractable mechanisms to handle evolving graph data, a notoriously difficult problem in classical graph transformation approaches. These examples serve both to show the expressiveness of our approach, and to illustrate its potential benefits.

The rest of this paper is structured as follows: Section 2 provides an overview of Logica. Section 3 demonstrates how to express graph transformations in Logica. We conclude in Section 4 with a discussion of the limitations and potential avenues for future research.

2. Logica Overview

Logica is a freely available, open-source variant of Datalog combining the declarative features of expressive rule-based languages with aggregation. Logica is a descendant of Yedalog [13] and inherits several of its features including support for aggregators, functional predicates, user-defined functions, and complex data types. Logica also supports rules involving recursion and negation. A key feature of the implementation is that it converts programs into SQL. This conversion can be configured by the user and rules can be compiled into: (a) self-contained SQL scripts with fixed recursion depth; or (b) Python-driven pipelines that chain together SQL queries when deep recursion is needed.

An overview of the Logica system is shown in Figure 1. Developers can work with Logica from the command line or via a Jupyter notebook. Programs can import functions and other rules via a module system. Logica parses and analyzes program files and produces a collection of SQL queries in the dialect of the target database engine (currently SQLite, DuckDB, PostgreSQL, or BigQuery). To help manage differences among platforms, Logica employs a type inference engine to create correct SQL for each underlying

Published in the Proceedings of the Workshops of the EDBT/ICDT 2025 Joint Conference (March 25-28, 2025), Barcelona, Spain

✉ evgenys@google.com (E. Skvortsov); yilinx2@illinois.edu (Y. Xia); ludasch@illinois.edu (B. Ludäscher); bowers@gonzaga.edu (S. Bowers)



Copyright © 2025 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

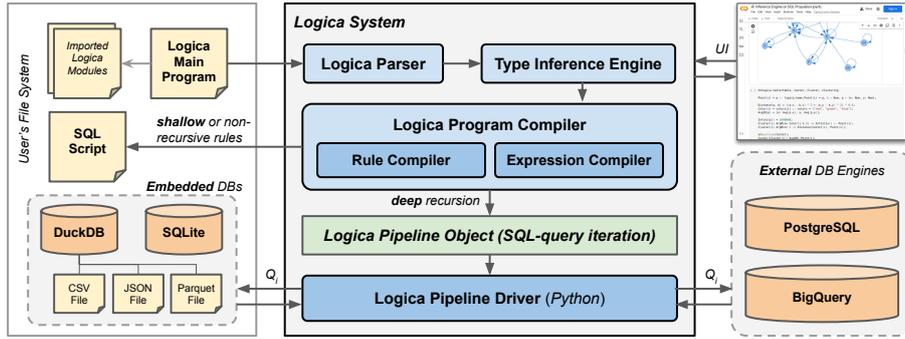


Figure 1: Logica System Architecture (from [9]): Logica supports a module system for importing libraries (top left); users write programs locally (top left) or from within Jupyter notebooks (top right); programs are parsed, type checked (top middle), and compiled to a self-contained SQL script (no/limited recursion) or to an iterative plan of pipelined SQL-queries passed to a pipeline driver (bottom middle). Logica executes and monitors queries as they run on underlying SQL engines, either locally (bottom left) or externally (bottom right).

system. If the set of rules to be evaluated is non-recursive (or if the user has indicated via a directive that a fixed-depth, non-iterative recursion is sufficient), a self-contained SQL script is generated that can be directly executed. For programs requiring deep recursion, Logica generates a pipeline script that iteratively executes the generated SQL queries stage-by-stage until a fixpoint or a user-defined termination condition is reached.

For long-running queries, users can monitor rule execution in the Logica UI: predicate results are rendered as they are being evaluated, so the user knows which (iterated) relations are still running. This information can also be saved and used for logging and profiling program execution.

3. Transforming Graphs with Logica

In Logica, as in Prolog and Answer-Set Programming, graphs can be naturally represented: Nodes are denoted using unique identifiers and edges are represented as facts about the nodes. Graphs can be represented as predicates with two or more *positional* arguments. For example, a simple directed graph can be represented as a binary relation $E(\text{source}, \text{target})$, where *source* and *target* are nodes connected by a directed edge. Additional graph data, e.g., node or edge attributes, can be expressed using separate predicates or extra positional (in Logica also *named*) arguments. For instance, a graph whose edges are assigned colors can be represented using a ternary predicate and facts of the form $E(\text{source}, \text{target}, \text{color})$.

Graph transformations can be implemented by defining new predicates that depend on the original graph: e.g., given a graph with edges E , a new graph $E2$ extending E by adding new edges between nodes that are separated by two “hops” can be derived as follows.¹

```
1 E2(x, z) :- E(x, y), E(y, z);
2 E2(x, y) :- E(x, y);
```

This simple example demonstrates an important aspect of defining graph transformations with logic rules. In native graph transformation languages users define how the graph changes, and edges not involved in the change remain. However, when defining graph transformations logically, we

¹In Logica, variables are lowercase, predicates are uppercase, “~” is used for negation, and semicolons denote the end of rules.

must include rules to *preserve* edges not involved in the transformation (line 2).

3.1. Message Passing

Our first example involves the simple passing of a message along the directed edges of a graph, as described in [14].

```
1 M0(0); # start node
2 # Rule 1: Message initialization.
3 M(x) :- M=nil, M0(x);
4 # Rule 2: Message passing.
5 M(y) :- M(x), E(x, y);
6 # Rule 3: Message retention.
7 M(x) :- M(x), ~E(x, y);
```

Fact $M0(0)$ states that initially the message is at node 0. Rule 1 initializes the message relation M . It states that node x has a message M if it was specified as the initial node in $M0(x)$. Condition $M=nil$ makes the rule fire (trigger) only at the start of the recursive iteration. The message propagates from node x to its neighbors y through the Rule 2. Finally, Rule 3 requires the message of a node x to be retained if x has no outgoing edges.

3.2. Computing Distances in a Graph

Logica has native support for aggregation. Predicate level aggregation is computed over the specified values of the fields of the defined predicate. The following Logica program computes the minimum distance $D(x)$ of a node x from a given $Start$ node.

```
1 # Rule 1: Distance from the Start node is 0.
2 D(Start()) Min= 0;
3 # Rule 2: Triangle inequality.
4 D(y) Min= D(x) + 1 :- E(x,y);
```

$Start()$ is used as a *functional predicate*: All Logica relations have an additional “special attribute” (named `Logica_value`) to store and access a relation’s functional value, i.e., its value when used syntactically as a function. Here, $Start()$ simply returns (a predefined) starting value. Rule 1 assigns the distance 0 to the starting node, where $D(x)$ is also a functional predicate that returns the (minimum) distance of a node x . Rule 2 computes the distance to a node y as *at most* one more than the distance to its prede-

cessor x . The Min= aggregation operator specifies that the minimum distance will be taken among all possible paths.

3.3. Solving Win-Move Games

Win-Move is a two-player combinatorial game played on a finite directed graph (representing a board) where nodes represent positions and edges represent moves [15, 16]. After choosing a starting position (node), the two players take turns moving a game pebble along the edges of the graph. A player who cannot move loses the game. We propose to compute the solution to Win-Move games in a new way, inspired by graph transformations. A concise formulation of the game is given by the logic rule:

$$\text{Win}(x) \text{ :- Move}(x, y), \sim\text{Win}(y).$$

A solution to the game can be computed using the 3-valued well-founded semantics [17]. The moves of the game are given by the predicate `Move`. A game starting at node x is objectively won if there exists a (*winning*) move to some y that is lost. Then, no matter how an opponent (Player II) moves from y (assuming there is an outgoing move to begin with), Player I can always force a win after the opponent has moved. The well-founded model assigns $\text{win}(x)$ the value *true*, *false*, and *undefined* if and only if x is objectively *won*, *lost*, and *drawn*, respectively. In contrast, the 2-valued *stable model* semantics does not compute the desired game solution: There can be 0, 1, or more stable models, none of which may agree with the correct, well-founded solution. Similarly, in Logica, the above rule cannot be used as-is.

Looking at the problem through a graph-transformation lens, however, we can define a new predicate $w(x, y)$, indicating that the move from x to y is a winning move. Computation of w , from a graph transformation perspective, amounts to selecting a collection of winning moves. In Logica this can be specified with a single rule:

```
1 W(x,y) :- Move(x,y), (Move(y,z1) => W(z1,z2));
```

The rule says: A move from x to y is *winning* iff for *every* opponent move from y to some $z1$ there *exists* a move from $z1$ to $z2$ for Player I that is again winning! The logical implication $\text{Move}(y, z1) \Rightarrow W(z1, z2)$ is a shorthand for the nested negation $\sim(\text{Move}(y, z1), \sim W(z1, z2))$.

The position values can now be derived from the winning moves w . The source and target nodes x and y of a winning move are won and lost positions, respectively. Positions that are neither won nor lost are drawn:

```
1 Won(x), Lost(y) :- W(x,y);
2 Drawn(x) :- Position(x), ~Won(x), ~Lost(x);
```

Here, the unary predicate `position` is defined as the union of the domain and range of the `Move` relation:

```
1 Position(x) :- x in [a,b], Move(a,b);
```

3.4. Finding Paths in Evolving Graphs

Graphs that change over time [10, 11, 12] can be naturally modeled and computed over in Logica, where time can be represented using extra positional arguments, named arguments, or as functional values. As a simple example, assume a graph that changes over time is represented by the relation

$E(x, y, t_0, t_1)$ stating that there exists an edge from node x to y added to the graph between time moments t_0 and t_1 . For simplicity we assume that time starts at moment 0, edges are transitioned instantly, and the start node is given as `Start()`. The following Logica program computes the earliest possible arrival time for each node.

```
1 # Rule 1: Starting condition.
2 Arrival(Start()) Min= 0;
3 # Rule 2: Traversal of an edge when edge exists.
4 Arrival(y) Min= Greatest(Arrival(x), t0) :-
5   E(x,y,t0,t1), Arrival(x) <= t1;
```

Rule 1 sets the arrival time of the start node to 0 (the assumed initial time). Rule 2 considers edge transitions: if we arrive at the source x of an edge before the edge expires (given by t_1) then the arrival time of the edge's target y is set to the larger (latest) of x 's arrival time and the time t_0 that the edge was added to the graph. Once the arrival times are computed, additional rules can be easily added to select specific (time aware) paths of the graph. Figure 2 gives an example of path finding over an evolving graph. An initial graph (with nodes A, B, etc., and colored blue) is given with edges labeled by their start and end times. The start node A is shown in green. The computed arrival times are displayed as additional nodes (in yellow). The graph in Figure 2 was created within Logica as further described below in Section 3.6.

3.5. Transitive Reduction of DAGs

The *transitive reduction* [18] of a directed graph is a graph with the fewest possible edges that has the same reachability as the original graph. For directed acyclic graphs (DAGs), the transitive reduction coincides with finding the (unique) *minimum equivalent subgraph*, i.e., a proper subgraph with the fewest possible edges needed to maintain the same reachability relation as the original graph. While finding a minimum equivalent subgraph for (arbitrary) graphs with cycles is NP-complete, for DAGs, the problem can be solved in PTIME. The following Logica program computes the transitive reduction $\text{TR}(x, y)$ of a DAG given by the edge relation $E(x, y)$. The program first computes the transitive closure and then uses the transitive closure to remove redundant edges.

```
1 # Rule 1: Transitive closure base case.
2 TC(x,y) distinct:- E(x,y);
3 # Rule 2: Transitive closure inductive step.
4 TC(x,y) distinct:- TC(x,z), TC(z,y);
5 # Rule 3: Transitive reduction.
6 TR(x,y) :- E(x,y), ~(E(x,z), TC(z,y));
```

Rules 1 and 2 compute the transitive closure. Rule 3 identifies edges $E(x, y)$ in the original graph that cannot be bypassed by going to some other node and taking a transitive path from there. These edges are the essential edges needed to maintain reachability and make up the edges of the resulting transitive reduction graph TR .

3.6. Rendering Graphs with Logica

One of the benefits of Logica is the convenience of rendering graphs directly from predicate definitions. With just a few lines of Logica code and a minimal Python wrapper, it is possible to generate visually appealing and informative

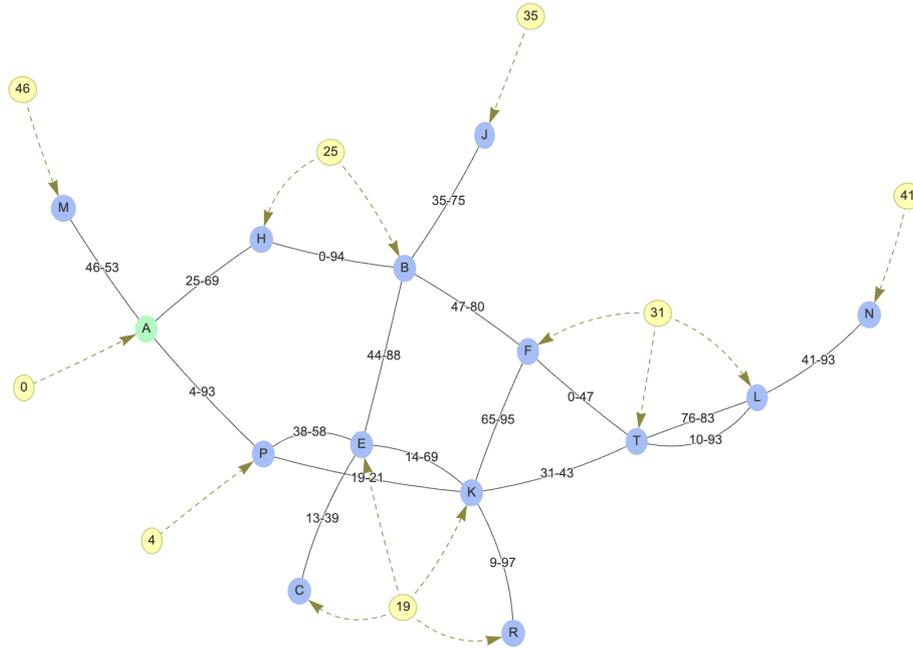


Figure 2: Visualizing the solution of pathfinding in a dynamic graph. The time of existence is shown as labels on edges; the earliest possible time of arrival is shown via yellow nodes. Section 3.6 shows how such visualizations can be created in Logica.

graph representations that highlight case-specific attributes. For example, consider the transitive reduction computed in the previous section. Instead of exporting the TR and E predicates and configuring their rendering in a separate graphing library, Logica makes it possible to define a predicate that *directly* specifies the visual attributes of the graph. The following relations define graph visualization properties in Logica highlighting the edges in both the original graph and its transitive reduction:

```

1 R(x, y,
2   arrows:"to" ,
3   color?Max= "rgba (40, 40, 40, 0.5)",
4   dashes?Min= true,
5   width?Max= 2,
6   physics?Max= false,
7   smooth?Max= false)distinct:- E(x, y);
8 R(x, y,
9   arrows:"to" ,
10  color?Max= "rgba (90, 30, 30, 1.0)",
11  dashes?Min= false,
12  width?Max= 4,
13  physics?Max= true,
14  smooth?Max= true)distinct:- TR(x, y);

```

The $R(x, y, \dots)$ relation defines the edges of the graph and their visual properties. The first rule draws the original edges from $E(x, y)$ in a light gray, dashed style with a thin line. The second rule draws the transitive reduction edges from $TR(x, y)$ in bold red, solid lines. The `distinct` keyword triggers aggregation, so that each edge occurs only once and values of attributes such as `color`, `dashes`, etc., are chosen based on whether this edge belongs to the transitive reduction. The visual attributes, e.g., `arrows`, `color`, `dashes`, `width`, `physics`, and `smooth`, are directly embedded within the Logica rules. The question mark in “`color? Max=`” indicates that the `color` can have

different values from different rules. The `Max=` indicates that if there are multiple rules that apply to the same edge, the maximum value will be used. The actual rendering is then handled with a Python function call:

```

1 # from logica.common import graph
2 graph.SimpleGraph(
3   R, extra_edges_columns=[
4     "arrows" ,"physics" ,"dashes" ,"smooth" ],
5   edge_color_column="color" ,
6   edge_width_column="width" )

```

This code leverages Logica’s graph module. The `SimpleGraph` function takes the R predicate, specifies which columns in R represent edge attributes (using `extra_edges_columns`, `edge_color_column`, and `edge_width_column`), and provides overall graph options. The resulting visualization is shown in Figure 3 (where the transitive reduction graph is overlaid over the original graph). The tight integration between logical definitions and graph rendering makes Logica a powerful tool for analyzing and understanding complex relationships in graph data. The ability to define graph properties directly in Logica, rather than relying on external tools, can significantly streamline the data exploration and visualization process.

3.7. Graph Condensation

Graph condensation [19] is a technique for simplifying a graph by collapsing strongly connected components (SCCs) into single nodes. This can be useful for visualizing the overall structure of a graph and for performing analyses that are less sensitive to the details within SCCs. In the condensed graph, each node represents an SCC, and an edge exists between two nodes if there is an edge in the original graph between nodes in the corresponding SCCs.

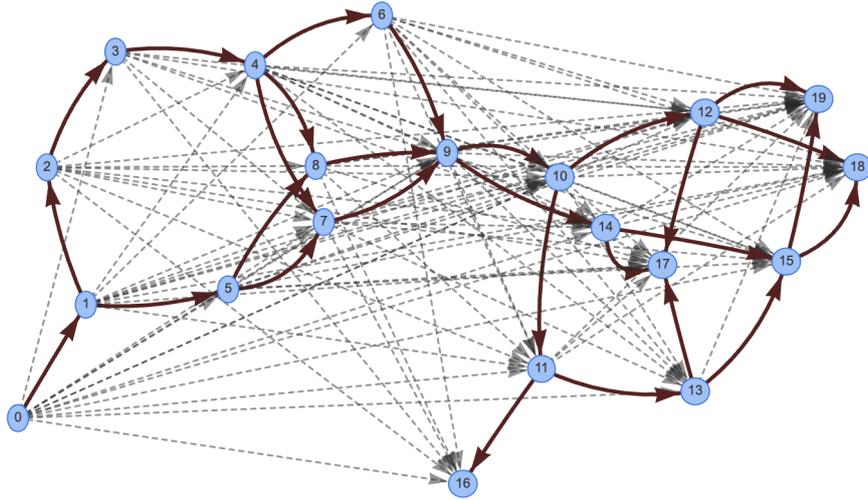


Figure 3: Visualizing a transitive reduction with Logica: Original edges are displayed with dashed lines, while edges present in the transitive reduction are highlighted with solid lines. This visualization, generated directly from Logica predicates, shows how the transitive reduction simplifies the graph while preserving reachability.

The following Logica program computes the condensation of a graph given by edges defined by the $E(x, y)$ relation, and where all nodes are defined by a $\text{Node}(x)$ predicate. We assume that transitive closure is already computed as in Subsection 3.5.

```

1 # Minimal node ID of the component
2 # ..is used as the component ID.
3 CC(x) Min= x :- Node(x);
4 CC(x) Min= y :- TC(x,y), TC(y,x);
5 # Compute condensation graph edges.
6 ECC(CC(x),CC(y)) distinct:-
7 E(x,y), CC(x)!=CC(y);

```

Once the condensation is computed, it can be easily rendered using Logica's graph visualization capabilities. To enhance readability, we use the following naming conventions: nodes in the original graph are named directly with their IDs (e.g., "1", "2", "3"), while the condensed components are prefixed with "c-" (e.g., "c-1", "c-2"). The predicate definitions for rendering both the original and condensed graphs simultaneously are as follows:

```

1 NodeName(x) = ToString(ToInt64(x));
2 CompName(x) = "c-"++ ToString(ToInt64(x));
3
4 # Edges of the original graph.
5 Render(NodeName(a), NodeName(b),
6 physics:true,arrows:"to" ,
7 dashes:false,smooth:true,
8 color:"#33e"):- E(a, b);
9 # Edges of the condensation.
10 Render(CompName(x), CompName(y),
11 physics:true,arrows:"to" ,
12 dashes:false,smooth:true,
13 color:"#33e"):- ECC(x, y);
14 # Mapping from the graph to
15 # the condensation.
16 Render(NodeName(ToInt64(a)), CompName(CC(a)),
17 physics:false,arrows:"to" ,
18 dashes:true,smooth:false,
19 color:"#888");

```

The `NodeName` and `CompName` functions convert node and component IDs to strings with appropriate prefixes. The `Render` predicate then defines the appearance of the graph:

- Edges in the original graph (defined by E) are rendered as solid blue lines.
- Edges between components (defined by ECC) are also rendered as solid blue lines.
- Dashed gray lines connect each original node to its corresponding component. The physics is disabled between nodes and their condensation components to help with readability.

Figure 4 shows the result of using the above Logica code and the appropriate Python wrapper as described in Subsection 3.6 to view the original and corresponding condensed graph. By combining the power of Logica for defining graph structures and relationships with its seamless graph rendering capabilities, it is possible to gain valuable insights into complex systems through clear and concise visualizations.

3.8. Inferring a Taxonomic Tree

Another application of Logica lies in its ability to infer relationships from large datasets and represent them as graphs. Consider the problem of constructing a simplified taxonomic tree for a set of species, showing their evolutionary relationships leading back to common ancestors. Logica can perform this task on massive amounts of data. Here, we aim to build a taxonomic tree for *Homo sapiens* (humans), *Crocodylidae* (crocodiles), *Tyrannosaurus* (T-Rex), and *Columbidae* (pigeons). The input is a knowledge graph of triples $T(a, b, c)$ and a functional predicate $L(a) = \ell$ that maps objects a to human readable labels ℓ . We define $\text{SuperTaxon}(item, parent)$ as the result of the query $T(item, "P171", parent)$, which states that $parent$ is a direct supertaxon of $item$, and $\text{TaxonLabel}(x) = L(x)$ to give labels restricted to taxons. The following Logica program builds the taxonomic tree.

```

1 # Use unbounded depth (-1) to find ancestors.
2 @Recursive(E, -1, stop: FoundCommonAncestor);
3 E(x, item,
4 TaxonLabel(x), TaxonLabel(item)) distinct:-
5 SuperTaxon(item,x),
6 ItemOfInterest(item) | E(item);

```

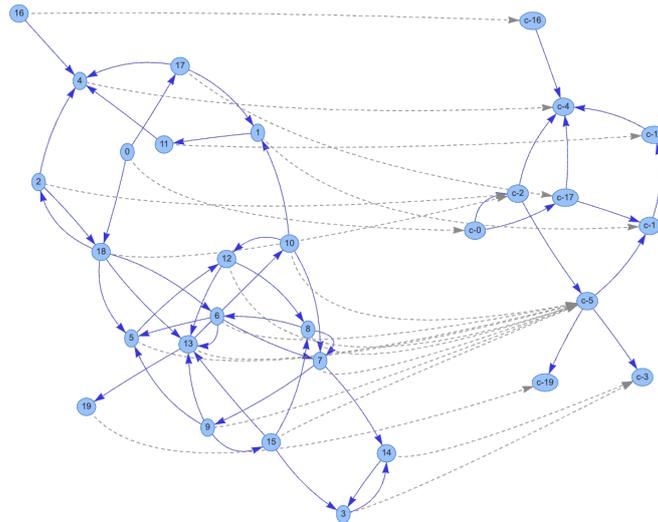


Figure 4: Graph Condensation visualized with Logica. This figure displays both the original graph and its condensed representation. Solid blue lines represent edges within the original graph and between condensed components. Dashed gray lines connect each node to its corresponding condensed component, illustrating how strongly connected components are collapsed into single nodes. This condensation simplifies the graph while preserving high-level connectivity information.

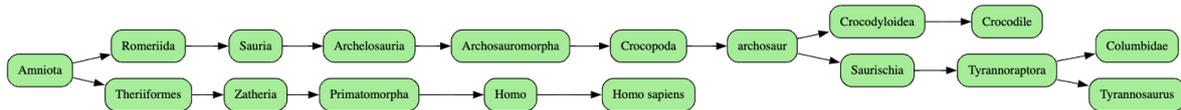


Figure 5: Inferred Taxonomic Tree for humans, crocodiles, T-Rex, and pigeons, generated by a Logica program from a large Wikidata dump. Edges represent evolutionary relationships, leading from an ancestral species to descendant species. This demonstrates that Logica can operate on very large, real-world knowledge graphs.

```

7 NumRoots() += 1 :- E(x,y), ~E(z,x);
8 # Stop when common ancestor is found.
9 FoundCommonAncestor() :- NumRoots() = 1;

```

Figure 5 shows an example output from running the program, where the GraphViz² library is used to render the resulting tree. The input for the example consisted of the set of statements obtained by a dump of Wikidata, which contained 806M facts defined over 89M objects. When stored using DuckDB, the entire input is 13GB in size. The full recursive search was run on a Google Cloud c2d-standard-32 instance in under 7 seconds via DuckDB with no custom index setup. The majority of the execution time was spent selecting the taxonomy edges from all possible relations in Wikidata. Note that the number of taxons returned by the original program is large. The result shown in Figure 5 is only a sample of the obtained taxonomic tree (where the sampling is also performed by Logica; not shown above).

By expressing the problem in Logica, we leverage its ability to efficiently perform (recursive) query processing (e.g., through its compilation to DuckDB). In this case, Logica is used to quickly navigate the complex relationships within a large taxonomic database and generate output that can then be displayed visually through its Python integration (here, using the GraphViz library).

4. Conclusion

We have presented a novel approach to implementing graph transformations using Logica, a free and open-source logic programming language designed for parallel databases. While Logica is typically used for data processing and analysis, we have shown that it can also effectively be used for graph transformations. By leveraging Logica’s ability to process large datasets in a parallel and efficient manner, graph transformations can be performed efficiently. By showcasing several examples, including the Win-Move problem, dynamic path finding, transitive reduction, graph condensation, and graph querying, we demonstrated that graph transformation problems can be expressed in an intuitive and efficient manner. The examples show how describing graph transformations using logic rules can lead to concise declarative specifications that also improve readability. For example, our approach models time-varying graphs in a principled manner, an issue that can be challenging in classical graph transformation languages.

In future work, we plan to explore more complex graph transformation patterns, including rewritings that may require solving NP-hard problems. We also plan to benchmark our approach against other graph transformation tools, which would be valuable for demonstrating Logica’s performance advantages on real-world datasets.

We hope that Logica, with its ability to leverage underlying SQL engines, can make graph transformations a more widely accessible and practical approach for managing complex and dynamic data.

²<https://graphviz.org/>

References

- [1] F. de la Parra, T. Dean, Survey of graph rewriting applied to model transformations, in: *International Conference on Model-Driven Engineering and Software Development*, 2014, pp. 431–441.
- [2] T. Kräuter, A. Rutle, H. König, Y. Lamo, Formalization and Analysis of BPMN Using Graph Transformation Systems, in: *International Conference on Graph Transformation*, 2023, p. 204–222.
- [3] M. Fernández, H. Kirchner, B. Pinaud, J. Vallet, Labeled Graph Rewriting Meets Social Networks, in: *International Workshop on Rewriting Logic and Its Applications*, volume 9942 of *LNCS*, Springer, 2016, pp. 1–25.
- [4] F. Rosselló, G. Valiente, *Graph Transformation in Molecular Biology*, Springer, 2005, pp. 116–133.
- [5] M. Nagl, Graph rewriting systems and their application in biology, in: *Mathematical Models in Medicine*, Springer, 1976, pp. 135–156.
- [6] F. Rosselló, G. Valiente, Chemical Graphs, Chemical Reaction Graphs, and Chemical Graph Transformation, *Electronic Notes in Theoretical Computer Science* 127 (2005) 157–166. *Proceedings of the International Workshop on Graph-Based Tools*.
- [7] M. Gelfond, Answer Sets, in: F. van Harmelen, V. Lifschitz, B. W. Porter (Eds.), *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, Elsevier, 2008, pp. 285–316.
- [8] E. S. Skvortsov, Y. Xia, B. Ludäscher, Logica: Declarative Data Science for Mere Mortals, in: *International Conference on Extending Database Technology*, 2024, pp. 842–845.
- [9] E. S. Skvortsov, Y. Xia, S. Bowers, B. Ludäscher, The Logica System: Elevating SQL Databases to Declarative Data Science Engines, in: *International Workshop on the Resurgence of Datalog in Academia and Industry (Datalog-2.0)*, volume 3801 of *CEUR Workshop Proceedings*, 2024, pp. 69–73.
- [10] D. Kempe, J. Kleinberg, A. Kumar, Connectivity and inference problems for temporal networks, *J. Comput. Syst. Sci.* 64 (2002) 820–842.
- [11] F. Rosselló, G. Valiente, *Graph Transformation in Molecular Biology*, Springer, 2005, pp. 116–133. URL: https://doi.org/10.1007/978-3-540-31847-7_7.
- [12] A. Casteigts, A.-S. Himmel, H. Molter, P. Zschoche, Finding Temporal Paths Under Waiting Time Constraints, *Algorithmica* 83 (2021) 2754–2802.
- [13] B. Chin, D. von Dincklage, V. Ercegovic, P. Hawkins, M. Miller, et al., Yedalog: Exploring knowledge at scale, *Summit on Advances in Programming Languages* (2015).
- [14] B. König, *Graph Transformation Meets Logic, GReTA - Graph Transformation Theory and Applications Symposium*, 2020.
- [15] C. A. Smith, Graphs and composite games, *Journal of Combinatorial Theory* 1 (1966) 51–81.
- [16] J. Flum, M. Kubierschky, B. Ludäscher, Total and Partial Well-Founded Datalog Coincide, in: *Intl. Conf. on Database Theory (ICDT)*, LNCS 1186, Springer, 1997, pp. 113–124.
- [17] A. Van Gelder, K. A. Ross, J. S. Schlipf, The Well-founded Semantics for General Logic Programs, *Journal of the ACM* 38 (1991) 619–649.
- [18] A. V. Aho, M. R. Garey, J. D. Ullman, The Transitive Reduction of a Directed Graph, *SIAM Journal on Computing* 1 (1972) 131–137.
- [19] R. Tarjan, Depth-First Search and Linear Graph Algorithms, *SIAM Journal on Computing* 1 (1972) 146–160.