

MillenniumDB: An Open-Source Graph Database System

Domagoj Vrgoč^{1,2†}, Carlos Rojas¹, Renzo Angles^{1,3}, Marcelo Arenas^{1,2}, Diego Arroyuelo^{1,4}, Carlos Buil-Aranda^{1,4}, Aidan Hogan^{1,5}, Gonzalo Navarro^{1,5}, Cristian Riveros^{1,2}, Juan Romero^{1,2}

¹Instituto Milenio Fundamentos de los Datos (IMFD)

²Pontificia Universidad Católica de Chile

³Universidad de Talca

⁴Universidad Técnica Federico Santa María

⁵DCC, Universidad de Chile

Keywords: Graph databases; Knowledge graphs; Worst-case optimal joins; Domain graphs

Citation: Vrgoč, D., Rojas, C., Angles, R., et al.: MillenniumDB: An open-source graph database system. *Data Intelligence* 5, (2023). doi: 10.1162/dint_a_XXXXX

Submitted: XXXX; Revised: XXXX; Accepted: XXXX

ABSTRACT

In this systems paper, we present MillenniumDB: a novel graph database engine that is modular, persistent, and open source. MillenniumDB is based on a graph data model, which we call domain graphs, that provides a simple abstraction upon which a variety of popular graph models can be supported, thus providing a flexible data management engine for diverse types of knowledge graph. The engine itself is founded on a combination of tried and tested techniques from relational data management, state-of-the-art algorithms for worst-case-optimal joins, as well as graph-specific algorithms for evaluating path queries. In this paper, we present the main design principles underlying MillenniumDB, describing the abstract graph model and query semantics supported, the concrete data model and query syntax implemented, as well as the storage, indexing, query planning and query evaluation techniques used. We evaluate MillenniumDB over real-world data and queries from the Wikidata knowledge graph, where we find that it outperforms other popular persistent graph database engines (including both enterprise and open source alternatives) that support similar query features.

[†] Corresponding author: Domagoj Vrgoč (E-mail: dvrhoc@ing.puc.cl; ORCID: 0000-0001-5854-2652).

1. INTRODUCTION

Recent years have seen growing interest in graph databases [1], wherein nodes represent entities of interest, and edges represent relations between those entities. In comparison with alternative data models, graphs offer a flexible and often more intuitive representation of particular domains [2]. Graphs forgo the need to define a fixed (e.g., relational) schema for the domain upfront, and allow for modeling and querying cyclical relations between entities that are not well-supported in other data models (e.g., tree-based models, such as XML and JSON). Graphs have long been used as an intuitive way to model data in domains such as social networks, transport networks, genealogy, biological networks, etc. Graph databases further enable specific forms of querying, such as path queries that find entities related by arbitrary-length paths in the graph. Graph databases have become popular in the context of NoSQL [3], where alternatives to relational databases are sought for specialized scenarios; Linked Data [4], where graph-structured data are published and interlinked on the Web; and recently Knowledge Graphs [5], where diverse data are integrated at large scale into a graph.

Recent years have also seen a growing number of models, languages, techniques and systems emerge for managing and querying graph databases [1, 2]. In the context of NoSQL systems, Neo4j [6], which uses the query language Cypher [7] and the property graph data model [2], is a leading graph database system in practice.[Ⓞ] Other popular graph database systems include ArangoDB [8], JanusGraph [9], OrientDB [10], TigerGraph [11], etc., which support Gremlin [12] and other custom graph query languages. We also find graph database systems supporting the RDF data model and SPARQL query language [13], including Allegrograph [14], Amazon Neptune [15], Blazegraph [16], GraphDB [17], Jena TDB [18], Stardog [19], Virtuoso [20], and (many) more besides [13]. There are now many graph database systems to choose from.

Many of these graph databases implement their own graph model, query language, etc. An open challenge is to design a graph database engine that is both *interoperable*, i.e., able to seamlessly support the diverse graph data models now popular in practice; and *efficient*, i.e., achieving query performance comparable (or ideally better than) systems built with a specific graph data model in mind. These goals are key to use-cases involving diverse, large-scale knowledge graphs. One concrete example is that of the Wikidata knowledge graph [21], which is composed of billions of statements, tens of thousands of node types, thousands of edge types, complex meta-data on edges, etc. Wikidata's query service—currently powered by Blazegraph [16]—receives in the order of millions of queries per day [22]. Interoperability in this setting would allow clients to seamlessly query Wikidata using their preferred syntax. However, as we will argue in Section 2.2, none of the graph models implemented by the aforementioned engines is well-suited for knowledge graphs like Wikidata; for example, for representing meta-data on edges (per Wikidata's *qualifiers* [21]), RDF requires reification [23], which adds bloat and indirection to the graph, while property graphs cannot link edges to nodes [24], requiring syntactic workarounds. Our abstract graph model—which we

[Ⓞ] See, e.g., <https://db-engines.com/en/ranking/graph+dbms>

call *domain graphs*—is designed with such knowledge graphs in mind [24].[Ⓢ] Efficiency then involves supporting more clients posing increasingly complex queries in less time.

Inspired by use-cases such as Wikidata, we propose MillenniumDB: an open-source graph database engine designed from the ground-up in order to achieve both interoperability and efficiency. To achieve *interoperability*, rather than implement the different data models from scratch, we rather adopt a common graph data model, called domain graphs, which can represent popular graph models in practice, and upon which MillenniumDB is founded. We then abstract the key query features common to different graph database engines, capturing them first as a formal query language over which concrete query syntax can be layered. Within this combination, we cover novel features, such as returning shortest paths (like Cypher) that match regular expressions (like SPARQL). In terms of *efficiency*, it is not trivial to know, *a priori*, which techniques are well-suited for evaluating real-world workloads over our model. To achieve efficient query processing, we thus incorporate a mix of both traditional and state-of-the-art techniques for evaluating graph patterns and path queries, adapting them to the domain graph model, and evaluating them empirically in a real-world setting to ascertain which offer the best performance in practice.

Contributions. The contributions of this paper are as follows:

- the domain graph and property domain graph data models, which allow for succinctly representing graph data models popular in practice, including RDF graphs, RDF-star graphs, property graphs, and the Wikidata knowledge graph [21];
- a formal query language based on domain graphs that captures key features of popular query languages for graph databases, along with a concrete query syntax;
- an indexing scheme and query engine designed for domain graphs that incorporates both traditional and state-of-the-art techniques, with optimizations dedicated to the evaluation of graph patterns and path queries;
- experiments over the Wikidata knowledge graph [21], involving real-world graph data and queries, comparing algorithms internal to MillenniumDB as well as other graph database engines.

Our experimental results highlight the benefits, for example, of incorporating worst-case-optimal join algorithms when evaluating complex graph patterns (with many joins) versus a more traditional approach based on applying binary joins with a Selinger-based query engine. We further compare the performance associated with different graph search algorithms in the context of path queries. On a more practical note, we show that MillenniumDB, under optimal configurations, clearly outperforms prominent graph database systems—namely Blazegraph, Neo4j, Jena and Virtuoso—and discuss why. We further publish a first release of MillenniumDB as an open source graph database engine [25], which we plan to extend in future in order to support more query syntax, query features, transactional updates, index structures, and more.

[Ⓢ] Domain graphs provide a concrete data model to use for representing knowledge graphs, and are thus an alternative to RDF graphs, property graphs, etc., but one that—as we will argue in Section 3—encapsulates such models while addressing some of their key limitations in this setting.

Paper structure. The rest of this paper is structured as follows: In Section 2 we describe existing graph data models and their limitations. In Section 3, we propose *domain graphs* as an abstraction of these models used in MillenniumDB. In Section 4, we describe the query language of MillenniumDB, and how it takes advantage of domain graphs. In Section 5, we explain how MillenniumDB stores data and evaluates queries. In Section 6, we provide an experimental evaluation of the proposed methods on a large body of queries over the Wikidata knowledge graph. In Section 7, we provide some concluding remarks and ideas for future research.

Data Availability & Supplementary Material Statement. The source code of MillenniumDB is provided in full at [25]. Experimental data is given at [26].

2. EXISTING GRAPH DATA MODELS AND THEIR LIMITATIONS

In this section we briefly recap the popular graph data models in use today, and discuss their limitations when modeling real-world datasets.

2.1 Graph Data Models

RDF and RDF..* One of the simplest models used for representing knowledge graphs is based on *directed labeled graphs*, composed of a set of edges of the form $\textcircled{a}-b\rightarrow\textcircled{c}$, where a is called the source node, b the edge label, and c the target node. In the context of knowledge graphs, nodes are used to represent entities and edges represent binary relations between pairs of entities. For example, with the edge in Figure 1 we can state that Michelle Bachelet was (or is) the president of Chile.




Figure 1. Information on presidency of Chile.

Such graphs are the basis of RDF [27], where the source node, edge label and target node are called *subject*, *predicate* and *object*, respectively. Given a universe Obj of objects (ids, strings, numbers, IRIs, etc.)[Ⓢ] the RDF data model is defined as follows:

Definition 1. An RDF triple is an element $(s, p, o) \in \text{Obj} \times \text{Obj} \times \text{Obj}$. An RDF graph is a finite set of RDF triples.

Upon analyzing this definition, we can immediately notice that the RDF data model lacks the ability to directly refer to the edge (s, p, o) itself. For instance, if we wanted to add the information about when the presidency represented by the above edge starts and when it ends, we would have to resort to some sort of *reification*, which would introduce an artificial object representing the edge that can then be linked to

[Ⓢ] For simplicity, we will often not distinguish restrictions on different types of objects when not pertinent.

the start and end date information. For example, the (reified) triples representing the duration of this presidency could be represented as shown in Figure 2. The reification is given by the use of the edges labeled as source, label and target.




Figure 2. Reified triples representing the duration of a presidency.

In order to avoid the need for reification, an extension of the RDF data model called *RDF** (or *RDF-star*) was proposed [28]. Intuitively, in *RDF** an entire triple can appear as a subject or an object in another triple. For example, in Figure 3, we are modeling the fact that Michelle Bachelet was the president of Chile from 2014-03-11 to 2018-03-11.




Figure 3. *RDF** triples with another triple as the subject.

The node representing the edge is called a *quoted triple* [29]. To distinguish edges that originate in a quoted triple, in Figure 3 we denote them with a dotted line. Formally, the *RDF** data model can be defined as follows:

Definition 2 ([28]). An *RDF** triple is defined recursively as follows:

- An *RDF* triple (s, p, o) is an *RDF** triple; and
- If s, o are *RDF** triples or elements of Obj , and $p \in Obj$, then (s, p, o) is an *RDF** triple.

An *RDF** graph is a finite set of *RDF** triples.

Another model extending *RDF* is that of *RDF datasets* [27], which are typically used to represent and manage multiple named *RDF* graphs. This model can be defined in two manners. The first, most general, definition permits empty graphs.

Definition 3. An *RDF dataset* is defined as a pair $D = (G, \{(n_1, G_1), \dots, (n_k, G_k)\})$ where:

- G, G_1, \dots, G_k are *RDF* graphs; and
- n_1, \dots, n_k are objects such that $n_i \neq n_j$ for $1 \leq i < j \leq k$.

The graph G is called the default graph, while each pair (n_i, G_i) is called a named graph, composed of the name n_i and its corresponding RDF graph G_i .

For example, letting G denote the RDF graph of Figure 1, and letting G_1 denote the RDF graph of Figure 2, then we can capture both graphs separately in an RDF dataset of the form $D = (G, \{(n_1, G_1)\})$, where n_1 is a name (e.g., reified) used to reference the graph G_1 . It is common to simply represent RDF datasets as a set of quads of the form $(s, p, o, g) \in \text{Obj} \times \text{Obj} \times \text{Obj} \times \text{Obj}$ [30], which indicates that the RDF triple (s, p, o) is in the RDF graph with name g . In this quad-based view, for example, D would then contain the quad $(e_1, \text{source}, \text{Michelle Bachelet}, \text{reified})$. A special name can be reserved in order to denote the default graph; for example $(\text{Michelle Bachelet}, \text{position held}, \text{President of Chile}, \text{default})$. This quad-based definition cannot directly support naming empty RDF graphs (though it could be extended to incorporate a set of names for empty graphs).

Property graphs. Finally, one of the more popular graph data models is that of *property graphs* [7]. Property graphs extend the simple edge labeled directed graph with two additional features: (i) they assign explicit identifiers to nodes and edges, so that one can refer to them; and (ii) they allow for annotating both nodes and edges with a set of property–value pairs. For example, the information from Figure 8 can be equivalently represented by the property graph in Figure 4.




Figure 4. Property graph representing the information about the presidency of Chile.

Here the nodes have identifiers (n_1, n_2) as well as labels (human, public office). Similarly, edges have both identifiers (e_1) and labels (position held). A node can have multiple labels, while an edge always has a single label (often referred to as its *type*). The edge e_1 has two *properties*, namely start date and end date, each with an associated value. Formally, if Obj is a set of objects, \mathcal{L} is a set of labels, \mathcal{P} a set of properties, and \mathcal{V} a set of values, we define the property graph data model as follows:

Definition 4. A property graph is a tuple $G = (V, E, \text{src}, \text{tgt}, \text{lab}, \text{prop})$, where:

- $V \subset \text{Obj}$ is finite set of node identifiers;
- $E \subset \text{Obj}$ is finite set of edge identifiers disjoint from V ;
- $\text{src} : E \rightarrow V$ assigns a source node to each edge;

- $tgt : E \rightarrow V$ assigns a target node to each edge;
- $lab : (V \cup E) \rightarrow 2^L$ is a function assigning a finite set of labels to nodes and edges, with $|lab(e)| = 1$ for all $e \in E$; and
- $prop : (V \cup E) \times \mathcal{P} \rightarrow \mathcal{V}$ is a partial function assigning a value to a certain property of a node or an edge.

Moreover, we assume that for each object $o \in V \cup E$, there exists a finite number of properties $p \in \mathcal{P}$ such that $prop(o, p)$ is defined.

2.2 Limitations of Existing Models

While all of the described data models have great expressive power, they are sometimes cumbersome to use when representing real-world datasets that contain higher-arity relations. To illustrate this, we will use the Wikidata [21, 31, 5] knowledge graph. Consider the two Wikidata statements shown in Figure 5. Both statements claim that Michelle Bachelet was a president of Chile, and both are associated with nested *qualifiers* that provide additional information: in this case a start date, an end date, who replaced her, and whom she was replaced by. There are two statements for two distinct presidencies. Also the ids for objects (for example, Q320 and P39) are shown; any positional element can have an id and be viewed as a node in the knowledge graph.

Michelle Bachelet [Q320]	
position held [P39]	President of Chile [Q466956]
start date [P580]	2014-03-11
end date [P582]	2018-03-11
replaces [P155]	Sebastián Piñera [Q306]
replaced by [P156]	Sebastián Piñera [Q306]
position held [P39]	President of Chile [Q466956]
start date [P580]	2006-03-11
end date [P582]	2010-03-11
replaces [P155]	Ricardo Lagos [Q331]
replaced by [P156]	Sebastián Piñera [Q306]

Figure 5. Wikidata statement group for Michelle Bachelet.

As aforementioned, representing statements like this in RDF graphs requires reification to decompose n -ary relations into binary relations [23]. Figure 6 shows a graph where e_1 and e_2 are nodes representing two distinct n -ary relationships (an extended version of Figure 2). For greater readability, we use human-readable nodes and labels, where in practice, the node Sebastián Piñera will rather be given as the identifier Q306, and the edge type “replaces” will rather be given as “P155”.

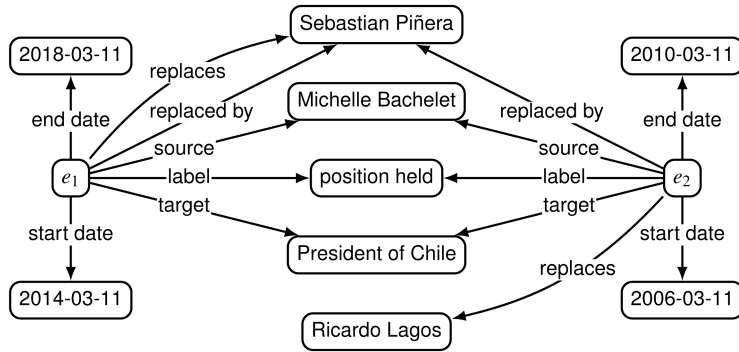


Figure 6. Directed labeled graph reifying the statements of Figure 5.

Since property graphs allow labels and property–value pairs to be associated with both nodes and edges, reification can be avoided in our example. For instance, the statements of Figure 5 can be represented as the property graph in Figure 7. Though more concise than reification, labels, properties and values are considered to be simple strings, which are disjoint with nodes; for example, Ricardo Lagos is neither a node nor a pointer to a node, but a string, which would complicate, for example, querying for the parties of presidents that Michelle Bachelet replaced.

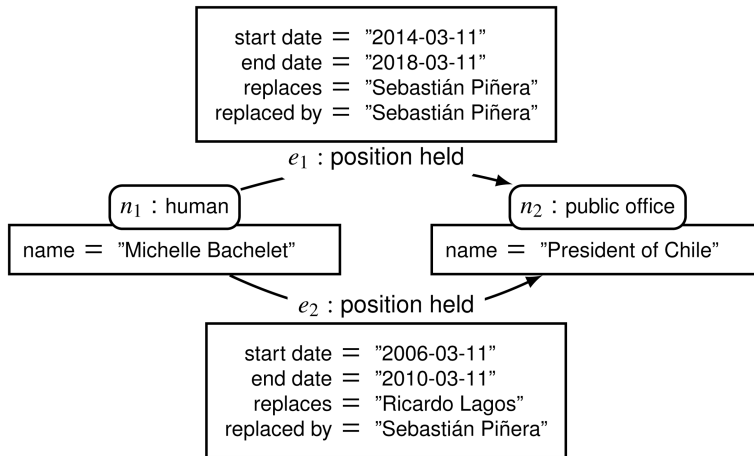


Figure 7. Property graph representing statements of Figure 5.

On the other hand, RDF* allows an edge to be a node. For example, the first statement of Figure 5 can be represented in RDF* as shown in Figure 8. However, we can only represent one of the statements (without reification), as we can only have one distinct node per edge; if we add the qualifiers for both statements, then we would not know which start date pairs with which end date, for example.[®]

[®] A proposed workaround involves adding intermediate nodes to denote different occurrences of quoted triples, but this requires a reserved term [29].




Figure 8. RDF* for one of the statements of Figure 5.

Regarding RDF datasets, we could model both statements by creating two named graphs, each with a copy of the statement that Michele Bachelet has been President of Chile, thereafter defining the start date, end date, replaced by and replaces annotations in another graph using the graph name. The resulting quads could thus be as follows if we define the latter information in the default graph (for example):

(Michelle Bachelet, position held, President of Chile, e_1)
 (e_1 , start date, 2014-03-11, default)
 ...
 (Michelle Bachelet, position held, President of Chile, e_2)
 (e_2 , start date, 2006-03-11, default)
 ...

This is quite a concise way to model the aforementioned Wikidata statements, wherein we effectively use graph names to assign each edge a unique id that serves as a graph node elsewhere. Indeed, the data model we propose follows a similar idea. However, RDF datasets were defined in the context of managing several (named) graphs, where using them to define edge ids gives rise to several complications; for example, SPARQL does not support evaluating path queries that span different named graphs.

3. DATA MODEL UNDERLYING MILLENNIUMDB

In this section, we present the graph data model upon which MillenniumDB is based, called domain graphs, and discuss how it generalizes existing graph data models such as RDF and property graphs. We also show its utility in concisely modeling real-world knowledge graphs that contain higher-arity relations, such as Wikidata [21].

3.1 Domain Graphs

The structure of knowledge graphs is captured in MillenniumDB via *domain graphs*, which follow the natural idea of assigning ids to edges in order to capture higher-arity relations within graphs [23, 32, 33, 24]. Formally, assume a universe Obj of objects (ids, strings, numbers, IRIs, etc.). We define domain graphs as follows:

Definition 5. A domain graph $G = (O, \gamma)$ consists of a finite set of objects $O \subseteq Obj$ and a partial mapping $\gamma : O \rightarrow O \times O \times O$.

Intuitively, O is the set of database objects and γ models edges between objects. If $\gamma(e) = (n_1, t, n_2)$, this states that the edge (n_1, t, n_2) has id e , type t , and links the source node n_1 to the target node n_2 .⁹ We can analogously define our model as a relation:

$$\text{DOMAINGRAPH}(\text{source}, \text{type}, \text{target}, \underline{\text{eid}})$$

where $\underline{\text{eid}}$ (edge id) is a primary key of the relation.

The domain graph model of MillenniumDB already subsumes the RDF graph model [27]. Recall that an RDF graph is a set of triples of the form (a, b, c) . To show how RDF is modeled in domain graphs, consider again the RDF triple from Figure 1. We can encode this triple in a domain graph by storing the tuple (Michelle Bachelet, position held, President of Chile, e) in the DomainGraph relation, where e denotes a unique (potentially auto-generated) edge id, or equivalently stating that:

$$\gamma(e) = (\text{Michelle Bachelet}, \text{position held}, \text{President of Chile}).$$

The id of the edge itself is not needed in the RDF data model, but it can be used for modeling RDF-star (RDF*) graphs [28, 29]. For example, to represent the RDF* graph from Figure 8, we can extend the function γ with two additional statements:

$$\gamma(e1) = (e, \text{start date}, 2014-03-11)$$
$$\gamma(e2) = (e, \text{end date}, 2018-03-11).$$

Here we use two new edges, $e1$ and $e2$, which have the edge e as their starting node.

For stricter backwards compatibility with legacy property graphs (where desired), MillenniumDB implements a simple extension of the domain graph model, called *property domain graphs*, which allows for *external annotation*, i.e., adding labels and property–value pairs to nodes and edges without creating new nodes and edges. Formally, if \mathcal{L} is a set of labels, \mathcal{P} a set of properties, and \mathcal{V} a set of values, we define a property domain graph as follows:

Definition 6. A property domain graph is defined as a tuple $G = (O, \gamma, \text{lab}, \text{prop})$, where:

- (O, γ) is a domain graph;
- $\text{lab} : O \rightarrow 2^{\mathcal{L}}$ is a function assigning a finite set of labels to an object; and
- $\text{prop} : O \times \mathcal{P} \rightarrow \mathcal{V}$ is a partial function assigning a value to a certain property of an object.

Moreover, we assume that for each object $o \in O$, there exists a finite number of properties $p \in \mathcal{P}$ such that $\text{prop}(o, p)$ is defined.

⁹ Herein, we say “edge type” rather than “edge label” to highlight that the type forms part of the edge, rather than being an annotation on the edge, as in property graphs.

While domain graphs (without properties) can directly capture property graphs—where, for example, the property–value pair (gender, “female”) on node n_1 can be represented by an edge $\gamma(e_3) = (n_1, \text{gender}, \text{female})$, the property–value pair (order, “2”) on an edge e_2 becomes $\gamma(e_4) = (e_2, \text{order}, 2)$, the label on n_1 becomes $\gamma(e_5) = (n_1, \text{label}, \text{human})$, the label on e_1 becomes the type of the edge $\gamma(e_1) = (n_1, \text{father}, n_2)$, etc.[®]—this can generate “incompatibilities” between the legacy property graph and the resulting domain graph; for example, strings like “male”, labels like human, etc., now become nodes in the graph, generating new paths through them that may affect query results. Property domain graphs thus offer an extra layer of flexibility, and interoperability with legacy property graphs, where needed for a given use-case.

To illustrate how property domain graphs work, consider the property graph (as introduced in Definition 4) from Figure 9. To model this information via property domain graphs, we use the domain graph part to capture the graph structure of our model, while property domain graphs also permit annotating that graph structure with labels and property–value pairs. The property graph in Figure 9 can be represented with the following property domain graph $G = (O, \gamma, \text{lab}, \text{prop})$, where the graph structure is as follows:




Figure 9. A property graph with two nodes and two edges. We use the order property on edge e_2 to indicate that Michelle Bachelet is the second child of Alberto Bachelet.

$$O = \{n_1, n_2, e_1, e_2\}, \quad \gamma(e_1) = (n_1, \text{father}, n_2), \\ \gamma(e_2) = (n_2, \text{child}, n_1),$$

and the annotations of the graph structure are as follows:

$$\begin{aligned} \text{lab}(n_1) &= \text{human} & \text{prop}(n_1, \text{last name}) &= \text{"Bachelet"} \\ \text{lab}(n_2) &= \text{human} & \text{prop}(n_2, \text{gender}) &= \text{"male"} \\ \text{prop}(e_2, \text{order}) &= \text{"2"} & \text{prop}(n_2, \text{children}) &= \text{"2"} \\ \text{prop}(n_1, \text{gender}) &= \text{"female"} & \text{prop}(n_2, \text{first name}) &= \text{"Alberto"} \\ \text{prop}(n_1, \text{children}) &= \text{"3"} & \text{prop}(n_2, \text{last name}) &= \text{"Bachelet"} \\ \text{prop}(n_1, \text{first name}) &= \text{"Michelle"} & \text{prop}(n_2, \text{death}) &= \text{"12 March 1974"} \end{aligned}$$

The relational representation of property domain graph then adds two new relations alongside DOMAINGRAPH:

[®] To represent edges in property graphs that permit multiple labels, multiple edges with different types can be added (or the labels can be added on the edge ids).

LABELS(object, label),
 PROPERTIES(object, property, value),

where object, property is a primary key of the second relation, with the first relation allowing multiple labels per object.

3.2 Domain Graphs Compared with Other Graph Data Models

Why did we choose (property) domain graphs as the model of MillenniumDB? As discussed in the previous section, it can be used to model both directed labeled graphs (like RDF) as well as property graphs. It also has a natural relational expression, which facilitates its implementation in a query engine. But it is also heavily inspired by the needs of real-world knowledge graphs like Wikidata [21, 31, 5]. To illustrate its versatility, consider again the Wikidata statements shown in Figure 5. Note that edges that originate in another edge are drawn with a dotted line.

As discussed in Section 2, neither RDF nor RDF* can represent these statements without resorting to reification, while property graphs cannot take nodes as values for properties. The domain graph model allows us to capture higher-arity relations more directly. In Figure 10 we present one possible representation of the statements from Figure 5. We only show edge ids as needed (all edges have ids). We do not use the “property part” of our data model for external annotation, considering that the elements of Wikidata statements shown can form nodes in the graph itself.

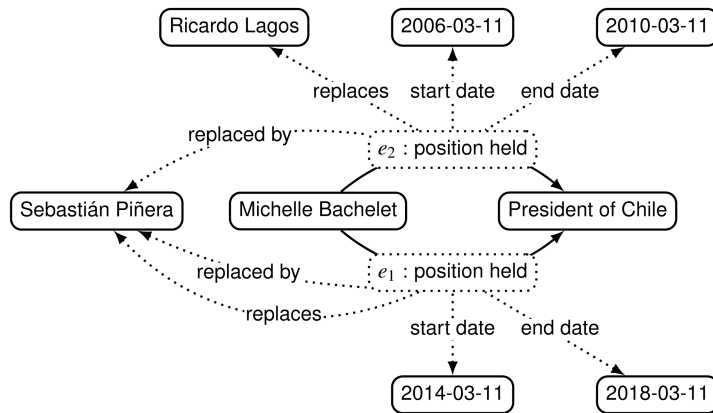


Figure 10. Domain graph for Figure 5.

Domain graphs are similar to named graphs in RDF datasets. Both domain graphs and RDF datasets can be represented as quads. However, the edge ids of domain graphs identify each quad, which, as we will discuss in Section 5, necessitates fewer index permutations. RDF datasets were proposed to represent multiple RDF graphs for publishing and querying. SPARQL thus does not support querying paths that span different named graphs; to support path queries over singleton named graphs, all edges would need to be

duplicated (virtually or physically) into a single graph [23]. Named graphs could be supported in domain graphs using a reserved term *graph*, and edges of the form $\gamma(e_3) = (e_1, \text{graph}, g_1)$, $\gamma(e_4) = (e_2, \text{graph}, g_1)$; optionally, *named domain graphs* could be considered in the future to support multiple domain graphs.

The idea of assigning ids to edges/triples for similar purposes as described here is a natural one, and not new to this work. Hernandez et al. [23] explored using singleton named graphs in order to represent Wikidata qualifiers, placing one triple in each named graph, such that the name acts as an id for the triple. In parallel with our work, recently a data model analogous to domain graphs has been independently proposed for use in Amazon Neptune, which the authors call 1G [33]. Their proposal does not discuss a formal definition for the model, nor a query language, storage and indexing, implementation, etc., but the reasoning and justification that they put forward for the model is similar to ours. Similar such models have been generalized as *multilayer graphs* [24], where the appearance of edge ids within the graph induces different layers of reference. Our work proposes a novel query language, storage and indexing schemes, query planner—and ultimately a fully-fledged graph database engine—built specifically for this model. Furthermore, with property domain graphs, we support annotation external to the graph, which we believe to be a useful extension that enables better compatibility with property graphs.

Table 1 summarizes the features that are directly supported by the respective graph models themselves without requiring *reserved terms*, which would include, for example, source, label and target in Figure 6 (all features except *External annotation* can be supported in all models with reserved vocabulary). Reserved terms can add indirection to modeling (e.g., reification [23]), and can clutter the data, necessitating more tuples or higher-arity tuples to store, leading to more joins and/or index permutations. The features are then defined as follows, considering directed (labeled) edges:

Table 1. Features supported by graph models without reserved terms (RD = RDF Datasets, PG = Property Graphs, DG = Domain Graphs, PDG = Property Domain Graphs).

	RDF	RDF*	RD	PG	DG	PDG
<i>Edge type/label</i>	✓	✓	✓	✓	✓	✓
<i>Node label</i>	✗	✗	✗	✓	✗	✓
<i>Edge annotation</i>	✗	✓	✓	✓	✓	✓
<i>Node annotation</i>	✓	✓	✓	✓	✓	✓
<i>External annotation</i>	✗	✗	✗	✓	✗	✓
<i>Edge as node</i>	✗	✓	✓	✗	✓	✓
<i>Edge as nodes</i>	✗	✗	✓	✗	✓	✓
<i>Nested edge nodes</i>	✗	✓	✓	✗	✓	✓
<i>Graph as node</i>	✗	✗	✓	✗	✗	✗

- *Edge type/label*: assign a type or label to an edge.
- *Node label*: assign labels to nodes.
- *Edge annotation*: assign property–value pairs to an edge.
- *Node annotation*: assign property–value pairs to a node.

- *External annotation*: nodes/edges can be annotated without adding new nodes or edges.
- *Edge as node*: an edge can be referenced as a node (this allows edges to be connected to nodes of the graph).
- *Edge as nodes*: a single unique edge can be referenced as multiple nodes.
- *Nested edge nodes*: an edge involving an edge node can itself be referenced as a node, and so on, recursively.
- *Graph as node*: a graph can be referenced as a node.

Some unsupported features in Table 1 are more benign than others; for example, *Node label* requires a reserved term (e.g., `rdf:type`), but no extra tuples; on the other hand, *Edge as node* requires reification, using at least one extra tuple, and also a reserved term.

Wikidata requires *Edge as nodes* as per Figure 5, where values like Ricardo Lagos are themselves nodes. Only RDF datasets, domain graphs and property domain graphs can model such examples without reserved terms; however, the use of RDF datasets requires co-opting graph names, which are typically used to manage multiple graphs, to rather serve as edge ids. Comparing RDF datasets and domain graphs, the latter sacrifices the “*Graph as node*” feature without reserved vocabulary to reduce indexing permutations (discussed in Section 5). Property domain graphs further support external annotation, and better compatibility with legacy property graphs.

4. QUERY LANGUAGE

Per our goal of supporting multiple graph models, MillenniumDB aims to support a number of graph query languages. However, no existing query language would take full advantage of the property domain graph model defined in the previous section. We have thus implemented a base query language, called DGQL, which closely resembles Cypher [7], but is designed for the property domain graph model, and adds features of other query languages, such as SPARQL, that are commonly used for querying knowledge graphs [34, 5]. Herein we provide a guided tour of the syntax of DGQL. A full formal specification of the language can be found in the appendix of this paper.

To introduce the features of the query language, in Figure 11 we present (a snippet of) a bibliographical knowledge graph representing data about publications, authors, institutions, etc. The knowledge graph is represented as a property domain graph, where, for authorship relations, we use properties on the edge to indicate the author order, but directly link the edges (via their ids) to the organization node with which the author was affiliated for that particular paper (something not directly possible in property graphs). We further use abstract node and edge ids ($n_1, \dots, n_{15}, e_1, \dots, e_{21}$) for brevity, though these may be instantiated with application ids; for example, in Wikidata, the node n_{15} denoting the U.S. might rather have the id Q30.

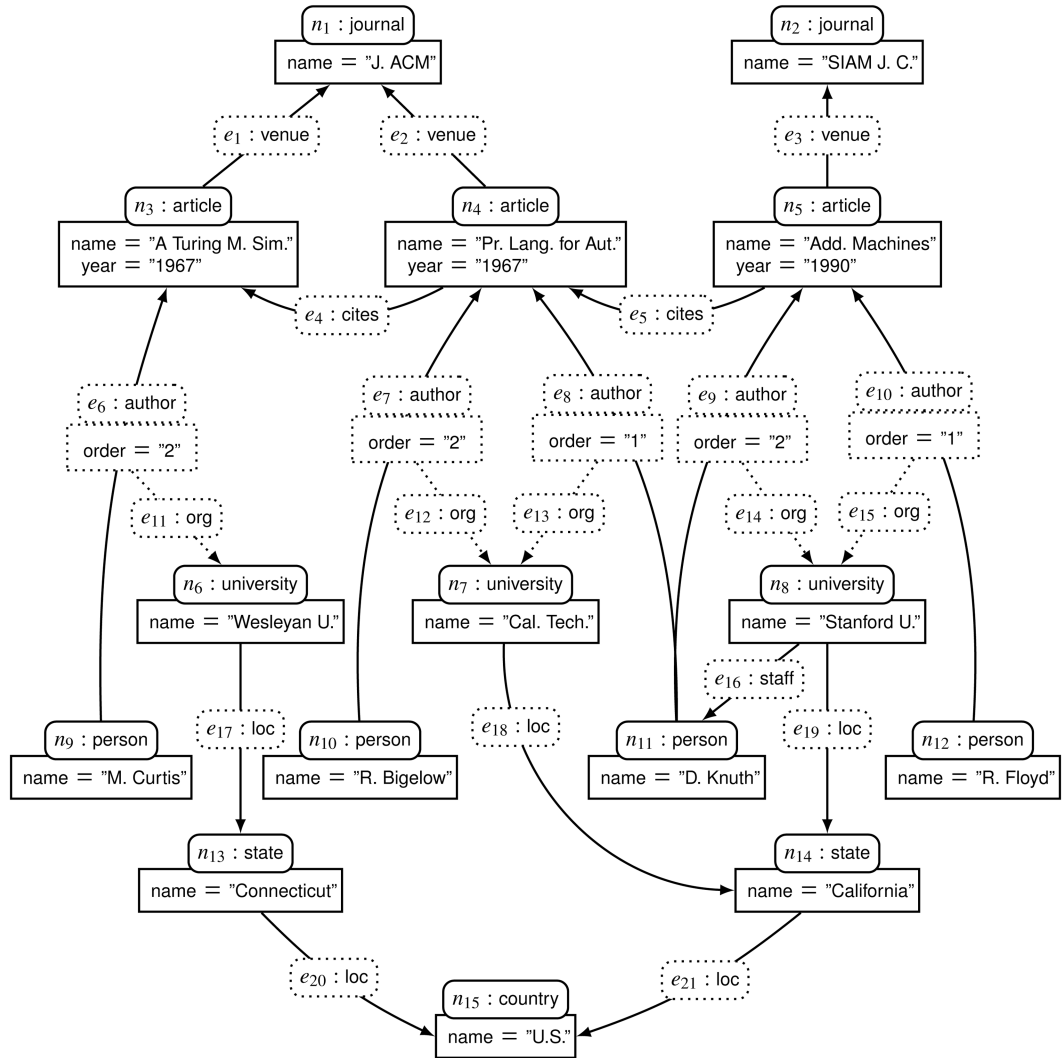


Figure 11. A property domain graph describing venues, papers, universities, authors, locations, and their relations.

We will use this knowledge graph as a running example in order to illustrate the MillenniumDB query language in the context of a bibliographical use-case, where we wish to analyze citations, find possible collaborators, etc.

4.1 Domain Graph Queries

A DGQL query takes the following high-level form:

MATCH Pattern
WHERE Filters
RETURN Variables

When evaluated over a property domain graph, such a query will return a multiset of mappings binding Variables to database objects (or values) that satisfy the Pattern specified in the MATCH clause and the Filters specified in the WHERE clause.

Querying objects. The most basic query will return all the objects (or more precisely, their ids) in our property domain graph. In MillenniumDB we can achieve this via the following query:

```
MATCH (?x)
RETURN ?x
```

Over the knowledge graph of Figure 11, this would return a table of all node and edge ids: $n_1, \dots, n_{15}, e_1, \dots, e_{21}$. Of course, one usually wants to select objects with a certain label, or a certain value in a specific property, as illustrated in the following example.

EXAMPLE 4.1. *The following DGQL query returns articles published in 1967 from Figure 11:*

```
MATCH (?x :article { year : 1967 })
RETURN ?x, ?x.name
```

This returns the ids of nodes with label article and value 1967 for the property year, along with their value for the property name, i.e., we return two results as follows:

?x	?x.name
n_3	"A Turing M. Sim."
n_4	"Pr. Lang. for Aut."

If, for example, n_3 did not have a name, we would still return n_3 as a result, leaving the corresponding value for ?x.name blank.

If we wish to specify a range, we can rather use the WHERE clause, which allows us to specify conditions on the results returned.

EXAMPLE 4.2. *If we want to find articles published before 1990, we can use the following query:*

```
MATCH (?x :article)
WHERE ?x.year < 1990
RETURN ?x, ?x.name
```

This returns the same solutions over Figure 11 as in Example 4.1. If we were to replace "<" with "<=", we would receive a third result for n_5 and "Add. Machines".

Querying edges.. In order to query over edges, we can write the following query, which returns γ , i.e., the relation `DOMAINGRAPH`:

```
MATCH (?x)-[?e :?t]->(?y)
RETURN *
```

The `RETURN *` operation projects all variables specified in the `MATCH` pattern, while the construct `(?x)-[?e :?t]->(?y)` specifies that we want to connect the object in `?x` with an object in `?y`, via an edge with type `?t` and id `?e`. This is akin to a query `DOMAINGRAPH(?x,?t,?y,?e)` over the domain graph relation. Variable or constant edge types (e.g. `?t` above) are prefixed by a colon.

EXAMPLE 4.3. Over the graph of Figure 11, the aforementioned query would return results of the following form, with 21 results in total (one for each edge):

<code>?x</code>	<code>?e</code>	<code>?t</code>	<code>?y</code>
n_3	e_1	<i>venue</i>	n_1
n_4	e_2	<i>venue</i>	n_1
n_5	e_3	<i>venue</i>	n_2
...
e_6	e_{11}	<i>org</i>	n_6
...
n_{14}	e_{21}	<i>loc</i>	n_{15}

This is akin to returning the `DOMAINGRAPH` relation.

We can also restrict which edges are matched, as shown in the following example.

EXAMPLE 4.4. The following query in `DGQL` will return the ids and names of articles that cite an article of the same year:

```
MATCH (?x)-[:cite]->(?y)
WHERE ?x.year == ?y.year
RETURN ?x, ?x.name
```

Here we choose to omit the edge id variable as we do not need it (e.g., in the `WHERE` or `RETURN` clause). Over the knowledge graph of the running example, this returns a single result:

<code>?x</code>	<code>?x.name</code>
n_4	"Pr. Lang. for Aut."

In the next example, we illustrate two features together: the ability to return and specify conditions on edge properties, and the ability to query known objects.

EXAMPLE 4.5. We look for the names of papers where Donald Knuth (n_{11}) is second or third author, returning his position in the author list:

```
MATCH (n11)-[?e :author]->(?y)
WHERE ?e.order > 1 AND ?e.order < 4
RETURN ?y.name ?e.order
```

This query returns a single result:

?y.name	?e.order
"Add. Machines"	2

As per the previous example, the WHERE clause may use Boolean combinations. We recall that the running example uses abstract node ids for brevity. In practice, the node id n_{11} could rather be an id such as Q17457, which identifies Donald Knuth on Wikidata.

Path queries. A key feature of graph databases is their ability to explore paths of arbitrary length. DGQL supports two-way regular path queries (2RPQs), which specify regular expressions over edge types, including concatenation (/), disjunction (|), inverses (^), optional (?), Kleene star (*) and Kleene plus (+). We use [= [] => (rather than - [] ->) to signal a path query in DGQL.

EXAMPLE 4.6. If we wish to find all of the citations of the article named "Add. Machines", and their respective citations, and so on transitively, we can use the regular expression `:cites+` in the following way, further returning the name and year of the articles where available:

```
MATCH (?x :article { name : "Add. Machines" })=[:cites+]=>(?y)
RETURN ?y, ?y.name, ?y.year
```

This returns:

?y	?y.name	?y.year
n_3	"A Turing M. Sim."	1967
n_4	"Pr. Lang. for Aut."	1967

DGQL can also return a single shortest path witnessing the query result by binding such a path to a variable.

EXAMPLE 4.7. The following query extends that of Example 4.6 by binding paths witnessing each result to a variable `?p`:

```
MATCH (?x :article { name : "Add. Machines" })=[?p :cites+]=>(?y)
RETURN ?y, ?y.name, ?y.year, ?p
```

This returns a string representation of a shortest path for each result, as follows:

?y	?y.name	?y.year	?p
n ₃	"A Turing M. Sim."	1967	" n ₅ $\xrightarrow{\text{cites}}$ n ₄ $\xrightarrow{\text{cites}}$ n ₃ "
n ₄	"Pr. Lang. for Aut."	1967	" n ₅ $\xrightarrow{\text{cites}}$ n ₄ "

We can also combine different features to capture more complex paths.

EXAMPLE 4.8. The following query looks for publications of staff and students of U.S. institutions, further including the direct citations of these publications:

```
MATCH (?x { name : "U.S." })=[?p (^:loc)*/(:staff|:student)/:author/:cites?]=>(?y)
RETURN ?y, ?y.name, ?p
```

This query returns:

?y	?y.name	?p
n ₃	"A Turing M. Sim."	" n ₁₅ $\xleftarrow{\text{loc}}$ n ₁₄ $\xleftarrow{\text{loc}}$ n ₈ $\xrightarrow{\text{staff}}$ n ₁₁ $\xrightarrow{\text{author}}$ n ₄ $\xrightarrow{\text{cites}}$ n ₃ "
n ₄	"Pr. Lang. for Aut."	" n ₁₅ $\xleftarrow{\text{loc}}$ n ₁₄ $\xleftarrow{\text{loc}}$ n ₈ $\xrightarrow{\text{staff}}$ n ₁₁ $\xrightarrow{\text{author}}$ n ₄ "
n ₅	"Add. Machines"	" n ₁₅ $\xleftarrow{\text{loc}}$ n ₁₄ $\xleftarrow{\text{loc}}$ n ₈ $\xrightarrow{\text{staff}}$ n ₁₁ $\xrightarrow{\text{author}}$ n ₅ "

Notice that a shortest path to each node is returned, so an additional path to e.g. n₄ using cites is not returned.

The final example for paths illustrates operators nested inside a Kleene star.

EXAMPLE 4.9. In order to find potential collaborators for a researcher, the following query finds shortest paths from that researcher (in this case Donald Knuth: n₁₁) to other authors via (transitive) citation or coauthorship relations:

```
MATCH (n11)=[?p :author/(:cites|(^:author/:author))*/^:author]=>(?y)
RETURN ?y, ?y.name, ?p
```

This query returns:

?y	?y.name	?p
n ₉	"M. Curtis"	" n ₁₁ $\xrightarrow{\text{author}}$ n ₄ $\xrightarrow{\text{cites}}$ n ₃ $\xleftarrow{\text{author}}$ n ₉ "
n ₁₀	"R. Bigelow"	" n ₁₁ $\xrightarrow{\text{author}}$ n ₄ $\xleftarrow{\text{author}}$ n ₁₀ "
n ₁₁	"D. Knuth"	" n ₁₁ $\xrightarrow{\text{author}}$ n ₄ $\xleftarrow{\text{author}}$ n ₁₁ "
n ₁₂	"R. Floyd"	" n ₁₁ $\xrightarrow{\text{author}}$ n ₅ $\xleftarrow{\text{author}}$ n ₁₂ "

A path that cycles back to Donald Knuth (n_{11}) is included. If we wished to filter such results, we could use *WHERE* to require the inequality $?y \neq n_{11}$. Also there are two possible shortest paths for the n_{11} result (via n_4 or via n_5), where the first such path to be found is returned.

If we wished to return *all shortest paths*, we could use the DGQL keyword *ALL* before the path variable. For instance, in the previous example, we can write `[ALL ?p ...]`, which returns a second result for n_{11} indicating the other shortest path.

Unlike Cypher, we can return paths matching 2RPQs, not just Kleene star. Unlike SPARQL, we can return paths, not just pairs of nodes. No manipulation of path variables, apart from outputting the result, is currently supported in MillenniumDB, but a full path algebra will be supported in future versions.

Basic graph patterns. Basic graph patterns [2] lie at the core of many graph query languages, including DGQL. Such graphs following the same structure as the data model, but allowing variables in any position. They can be seen as expressing natural (multi)joins over sets of atomic edge patterns. In DGQL, they are given in the *MATCH* clause. Basic graph patterns are evaluated under *homomorphism-based semantics* [2], which allows multiple variables in a result to map to the same element of the data.

EXAMPLE 4.10. *Illustrating basic graph patterns, the following query finds pairs of co-authors:*

```
MATCH (?x)-[:author]->(?z),
      (?y)-[:author]->(?z)
RETURN ?x, ?y
```

If we evaluate this query over the running example, we get the following results:

?x	?y
n_9	n_9
n_{10}	n_{10}
n_{10}	n_{11}
n_{11}	n_{10}
n_{11}	n_{11}
n_{11}	n_{12}
n_{12}	n_{11}
n_{12}	n_{12}

*Given the homomorphism-based semantics, results are returned that map both variables to the same author. If we wished to filter such results, we could stipulate the desired inequalities with *WHERE* $?x \neq ?y$, which would filter the first, second, fifth and eighth result.*

The previous example could equivalently be expressed as a path of the form `?x=[:author/^ :author]=>?y`. However, with basic graph patterns, we can also capture branches and cycles, as illustrated in the following example.

EXAMPLE 4.11. If we wanted to detect self-citations among journals in the same year, we could use the following DGQL query:

```
MATCH (?x)-[:cites]->(?y),
      (?x)-[:venue]->(?z),
      (?y)-[:venue]->(?z)
WHERE ?x.year == ?y.year
RETURN ?x, ?x.name, ?y, ?y.name, ?z, ?z.name
```

This would return:

?x	?x.name	?y	?y.name	?z	?z.name
n ₄	"Pr. Lang. for Aut."	n ₃	"A Turning M. Sim."	n ₁	"J. ACM"

The DGQL query language allows us to take full advantage of domain graphs by allowing joins between edges, types, etc., as illustrated by the following example.

EXAMPLE 4.12. The following query looks for articles with an affiliation that is current, i.e., where an author is still staff at the indicated organization:

```
MATCH (?x)-[?e :author]->(?y :article),
      (?e)-[:org]->(?z),
      (?z)-[:staff]->(?x)
RETURN ?y, ?y.name
```

The variable *?e* invokes a join between an edge and a node, returning:

?y	?y.name
n ₅	"Add. Machines"

Navigational graph patterns. If we further allow path queries within basic graph patterns, we arrive at *navigational graph patterns* [2].

EXAMPLE 4.13. Suppose we wish to find, for example, instances of self-citation of staff at U.S. institutions; we could write this as follows:

```
MATCH (?w)=[^:staff/loc+]=>(?x :country { name : "U.S." }),
      (?w)-[:author]->(?y),
      (?w)-[:author]->(?z),
      (?y)-[:cites]->(?z)
RETURN ?w, ?w.name, ?y, ?y.name, ?z, ?z.name
```

This query would return:

?w	?w.name	?y	?y.name	?z	?z.name
n ₁₁	"D. Knuth"	n ₅	"Add. Machines"	n ₄	"Pr. Lang. for Aut."

Optional graph patterns. Considering that the graph patterns considered previously allow for extracting tables from graphs, a way to enrich a graph query language is to support relational operators over these tables [2]; this gives rise to the notion of *relational graph patterns*. Thus far we have seen the ability to *project* results with SELECT, and apply *selections* over results with WHERE. We have also spoken about how basic and navigational graph patterns can be interpreted as natural joins in the relational algebra. DGQL also supports optional graph patterns, which behave akin to left outer joins in the relational algebra, i.e., they allow for extending solutions with data that may or may not be available; in case the data of the optional pattern are not available, the solution is still returned and the optional data are left blank.

EXAMPLE 4.14. Assume we want to find the authors who have published articles in the Journal of the ACM, their affiliation in those articles, and, if available, the organization at which they are currently staff. The following DGQL query achieves this:

```
MATCH (?v)-[?e :author]->(?w),
      (?w)-[:venue]->(?x { name = "J. ACM" }),
      (?e)-[:org]->(?y)
OPTIONAL {
  (?v)-[:staff]->(?z)
}
RETURN ?v, ?v.name, ?y, ?y.name, ?z, ?z.name
```

This query would return:

?v	?v.name	?y	?y.name	?z	?z.name
n ₉	"M. Curtis"	n ₆	"Wesleyan U."		
n ₁₀	"R. Bigelow"	n ₇	"Cal. Tech."		
n ₁₁	"D. Knuth"	n ₇	"Cal. Tech."	n ₈	"Stanford U."

Results for n₉ and n₁₀ are still returned though they are not currently staff at any organization; the corresponding variables are left blank.

Nested optional patterns are also supported. However, optional patterns must form *well-designed patterns* [35].

Limits and ordering. Some additional operators that MillenniumDB supports are LIMIT and ORDER BY. These allow us to limit the number of output mappings, and sort the obtained results, as illustrated by the following example.

EXAMPLE 4.15. *In order to find the most recent paper by Donald Knuth, we can use the following DGQL query:*

```
MATCH (?x { name : "D. Knuth" })-[:author]->(?y)
ORDER BY DESC ?y.year
RETURN ?x, ?x.name
LIMIT 1
```

The result returned is as follows:

?x	?x.name
n_5	"Add. Machines"

Ordering is always applied before limiting results.

4.2 Formal Definitions for DGQL

For readers interested in a formal specification, we provide the full definition of DGQL, together with the associated semantics, in the appendix to this paper. Specifically, in Appendix A we provide a grammar for DGQL queries, and in Appendix B we define (an equivalent) abstract syntax of DGQL and formal semantics of the language.

4.3 Comparing Graph Query Languages

A variety of query languages for graphs have been proposed in recent years [1, 2]. This section compares DGQL with six prominent query languages: Cypher [7] (Neo4j), SPARQL [30] (the standard query language for RDF Triple Stores), G-CORE [36] (LDBC), GSQL [11] (TigerGraph), Gremlin [12] (supported by several systems like Amazon Neptune and JanusGraph) and nGQL [37] (NebulaGraph). For each query language, we evaluate its support (total or partial) for six query features, namely: basic graph patterns, relational graph patterns, querying edges, regular path queries, navigational graph patterns, and full path recovery. Our comparison is shown in Table 2.

Table 2. Query features supported by graph query languages (BGP = basic graph patterns, RGP = relational graph patterns, QE = querying edges, RPQ = regular path queries, NGP = navigational graph patterns, FPR = full path recovery). The symbol \sim is used to indicate partial support of a feature.

Query language	<i>BGP</i>	<i>RGP</i>	<i>QE</i>	<i>RPQ</i>	<i>NGP</i>	<i>FPR</i>
DGQL	✓	~	✓	✓	✓	~
Cypher	✓	✓	~	~	~	~
SPARQL	✓	✓	~	✓	✓	✗
G-CORE	✓	✓	~	✓	✓	✓
GSQL	✓	~	~	✓	~	✗
Gremlin	✓	~	~	~	~	~
nGQL	✓	✓	~	~	~	~

Every query language considered in Table 2 supports the notion of a basic graph pattern (*BGP*), which, in its most general form, is a graph pattern structured like the data, but allowing variables to replace constants. In most cases, the result of a basic graph pattern is a relation (or table) consisting of results, and in some cases it is possible to construct/return a graph (like in G-CORE and SPARQL).

Considering that a graph pattern extracts a table from a graph (as seen in the examples of Section 4.1), *relational graph patterns (RGPs)* allow the use of relational-based operators to combine the results of one or more graph patterns into a single relation. Full support of this feature in Table 2 indicates that a language provides join, optional, union and negation of graph patterns. Partial support indicates that a language supports some of these operators, usually join and optional graph patterns, as is the case for DGQL (we plan to extend this in future to support more relational operators).

Querying edges (QE) is a particular feature of DGQL, allowing for querying relationships involving edges. Notably, DGQL allows an id to be extracted as an edge in one part of the query, and then used as a node in another part. Other query languages provide partial support for querying edges, as they are restricted to query the labels and properties of the edges, require reserved vocabulary (reification), or have other restrictions (e.g., using named graphs in SPARQL over which paths cannot be resolved).

The *regular path queries (RPQs)* feature refers to matching paths based on (2-way) regular expressions, with concatenation, disjunction, inverse, optional and Kleene star. Partial support indicates that a language offers a restricted group of such operators, such as in the case of Cypher, which supports only Kleene star on top of a single edge type, and not over a subexpression, thus supporting an expression such as `cites+`, but not a more complex expression, such as `(author/^\^ author)+`.

We use the term *navigational graph patterns (NGPs)* to represent the combination of basic graph patterns and regular path queries. These queries are akin to conjunctive (2-way) regular path queries. NGPs are supported by DGQL, SPARQL and G-Core.

Finally, a query language with *full path recovery* allows not only to search for some paths, but also to return such paths as objects that can be manipulated (with the nodes and edges in a path). This is a particular feature of G-CORE as it supports path construction operations, and the data model permits storing paths. In Cypher, the resulting paths can be assigned to a variable, so the elements of each path can be accessed by using ad-hoc functions, although with reduced facilities. SPARQL does not support this feature as the output of a path expression is only the start and end nodes of each path. In GSQL and Gremlin, the result of a path query is a set of objects, so the resulting paths must be processed by using a programming language. Currently, DGQL partially supports this feature by returning a path as a string; however, MillenniumDB has been designed to support path manipulation in the future.

Table 2 focuses on core features for querying graphs [2], and thus omits features (e.g., borrowed from SQL) that are supported by some of the languages, and that are potentially very useful in practice, such as aggregations, solution modifiers, federation, etc. Such features can be layered atop the features mentioned.

5. SYSTEM ARCHITECTURE

In this section, we describe the internals of the MillenniumDB engine, which have been designed to efficiently support the domain graph model. The overall architecture of the system is presented in Figure 12, and will be explained in the following.

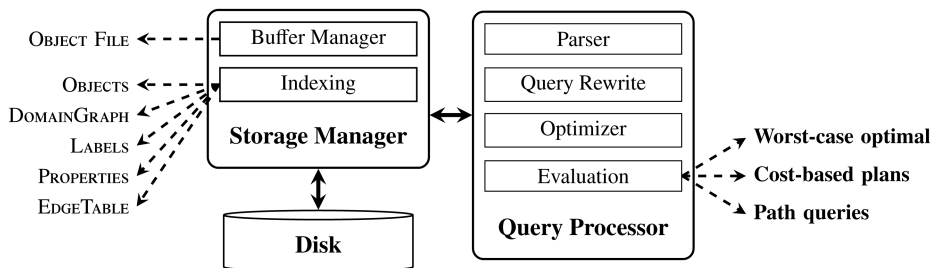


Figure 12. MillenniumDB Architecture.

MillenniumDB is founded on tried and tested relational techniques: it stores the (property) domain graph model as several relations indexed in B+ trees, loading parts into main memory as needed using a fixed-size buffer. It also uses algorithmic techniques recently suggested in the theoretical literature for evaluating queries [38, 39]—techniques not typically implemented in graph database systems—for supporting the domain graphs model in practice. Specifically, we combine three different techniques that are new to the architecture of graph database systems when used in conjunction. First, the data model is encoded as basic relations, indexed following different attributes orders, wherein data objects (e.g., nodes, strings) are represented by ids. Second, we translate the evaluation of any query to several joins between basic relations, which we manage using *worst-case optimal join algorithms* [40]: an evaluation technique recently proposed for relational database systems. Last, we combine join algorithms with the evaluation of path queries by

compiling the path pattern into an automaton and running the query on the fly. These techniques, together, are at the heart of how MillenniumDB optimizes queries over the domain graphs model in practice.

In what follows, we explain how one can store (property) domain graphs and index them. We then outline the query evaluation process and the algorithmic techniques it uses, like the worst-case optimal query plan and the evaluation of path queries.

Storage and indexing. Let us start by explaining the Disk and Storage Manager part of the MillenniumDB architecture from Figure 12. The main component of the domain graph data model are objects. Objects are represented internally as 8-byte identifiers. To optimize query execution, identifiers are divided into classes and the first byte of the identifier specifies a class it belongs to. The main classes in a property domain graph $G = (O, \gamma, \text{lab}, \text{prop})$ are:

- *Nodes*, which are objects in the range of γ . They are divided into two subclasses: *named nodes*, which are objects in the domain graph for which an explicit name is available (e.g. Q320 in Wikidata), and *anonymous nodes*, which are internally generated objects without an explicit name available to the user (similar to blank nodes in RDF [41]).
- *Edges*, which are objects in the domain and range of γ , and are always anonymous, internally generated objects.
- *Values*, which are data objects like strings, integers, etc. These values are classified in two subclasses: *inlined values*, which are values that fit into 7 bytes of the identifier after the mask (e.g. 7 byte strings, integers, etc.), and *external values*, which are values longer than 7 bytes (e.g. long strings).

All records stored in MillenniumDB are composed of these identifiers. We will explain later how long strings for external values are handled.

To store property domain graphs, MillenniumDB deploys B+ trees [42]. For this purpose, we build a B+ tree template for fixed sized records, which store all classes of identifiers. To store a property domain graph $G = (O, \gamma, \text{lab}, \text{prop})$, we simply store and index in B+ trees the four components defining it:

- `OBJECTS(id)` stores the identifiers of all the objects in the database (i.e., O).
- `DOMAINGRAPH(source,type,target,eid)` contains all information on edges in the graph (i.e., γ), where `eid` is an edge identifier, and `source`, `type`, and `target` can be ids of any class (i.e., node, edge, or value). By default, four permutations of the attributes are indexed in order to aid query evaluation. These are: `source-target-type-eid`, `target-type-source-eid`, `type-source-target-eid` and `type-target-source-eid`.
- `LABELS(object,label)` stores object labels (i.e., `lab`). The value of object can be any identifier, and the values of label are stored as ids. Both permutations are indexed.
- `PROPERTIES(object,property,value)` stores the property–value pairs associated with each object (i.e., `prop`). The object column can contain any id, and property and value are value ids. Aside from indexing the primary key, an additional permutation is added to search objects by property–value pairs.

All the B+ trees are created through a bulk-import phase, which loads multiple tuples of sorted data, rather than inserting records one by one. In order to enable fast lookups by edge identifier, we use the fact that this attribute is the key for the relation. Therefore, we also store a table called `EDGETABLE`, which contains triples of the form $(source, type, target)$, such that the position in the table equals to the identifier of the object e such that $\gamma(e) = (source, type, target)$. This implies that edge identifiers must be assigned consecutive ids starting from zero internally by MillenniumDB (they are not specified by the user). In total, we use ten B+ trees for storing the data.

To transform external strings and values (longer than 7 bytes) to database object ids and values, we have a single binary file called `OBJECTFILE`, which contains all such strings concatenated together. The internal id of an external value is then equal to the position where it is written in the `OBJECTFILE`, thus allowing efficient lookups of a value via its id. The identifiers are generated upon loading, and an additional hash table is kept to map a string to its identifier; we use this to ensure that no value is inserted twice, and to transform explicit values given in a query to their internal ids. Only strings are currently supported, but the implementation interface allows for adding support for different value types in a relatively simple manner.

All of the stored relations are accessed through linear iterators which provide access to one tuple at a time. All of the data is stored on pages of fixed (but parametrized) size (currently 4kB). The data from disk is loaded into a shared main memory buffer, whose size can be specified upon initializing the MillenniumDB server. The buffer uses the standard clock page replacement policy [42]. Additionally, for improved performance, upon initializing the server, it can be specified that the `ObjectFile` be loaded into main memory in order to quickly convert internal identifiers to string and integer values that do not fit into 7 bytes.

Evaluating a query. In MillenniumDB, the execution pipeline follows the standard database template where the string of the query is parsed and translated into a logical plan, which is then analyzed and converted into a physical plan, and finally evaluated, as illustrated by the Query Processor component of Figure 12.

A key part is in how the *patterns* and *filters* of a DGQL query (see Section 4.1) are evaluated. Specifically, patterns and filters are grouped together into a list of relations that can be edges, labels, properties, or path queries, forming a large *multi-way join* query. In essence, evaluating these joins is analogous to selecting an appropriate join plan for the relations representing the different elements. This also goes in hand with selecting the appropriate join algorithm for each of the joins. Given that edges, labels, and properties are all indexed, this will most commonly be index nested-loop join. Paths on the other hand are not directly indexed. For this reason, they are pushed to the end of the join plan and joined via nested-loop with the rest of the multi-way join[Ⓢ].

MillenniumDB supports different mechanisms for evaluating the multi-way join formed by the pattern and filter of DGQL query.

[Ⓢ] While not the best option, based on empirical evidence (see Section 6), this solution seems to be adequate in practice.

- A worst-case optimal query plan as described in [43] is used whenever possible. This approach implements a modified leapfrog algorithm [38] in order to minimize the number of intermediate results that are generated.
- The classical relational optimizer, which is based on cost estimation, and tries to order base relations in such a way as to minimize the amount of (intermediate) results. We currently support two modes of execution here:
 - (i) Selinger-style join plans [44] which use dynamic programming to determine the optimal order of relations.
 - (ii) In the presence of a large number of relations, a greedy planner [45] is used which simply determines the cheapest relation to use in each step.

Two particular points of interest are the worst-case optimal query planner, and the way that paths are evaluated. Both of these deploy state-of-the-art research ideas that are usually not implemented in practical graph database systems (though some prototypes exist [43, 46, 47]). We provide some additional details on these next.

Worst-case optimal query plan.. Evaluating multiple joins in a worst-case optimal way is done using a modified leapfrog algorithm [43]. While a classical join plan does a nested for-loop over relations, leapfrog performs a nested for-loop over variables [38]. Specifically, the algorithm first selects a variable order for the query, say ($?x$, $?y$, $?z$). It then intersects all relations where the first variable $?x$ appears, and over each solution for $?x$ returned, it intersects all relations where $?y$ appears (replacing $?x$ in its current solution), and so on to $?z$, until all variables are processed and the final solutions are generated. We refer the reader to [38] and [43] for a detailed explanation. Two critical aspects for supporting this approach are indexes and variable ordering, explained next.

To support the leapfrog algorithm over traditional relational indexes such as B+ trees, we should index all relations in all possible orders of their attributes in order to ensure efficient intersections, which greatly increases disk storage [43]. In MillenniumDB, we include four orders for `DOMAINGRAPH`, and all orders for `LABELS` and `PROPERTIES`. With these orders we can cover the most common join-types that appear in practice [34] by a worst-case optimal query plan. We use the classical relational optimizer if the plan needs an unsupported order or one of the relations uses a path query.

The leapfrog algorithm further requires choosing a variable ordering, which is crucial for its performance. The heuristic we deploy for selecting the variable ordering mixes a greedy approach, and the ideas of the Graham–Yu–Özsoyoglu (GYO) reduction [48]. More precisely, we first order the variables based on the minimal cost of the relations they appear in and resolve ties by selecting the variable that appears in more distinct relations. The variables “connected” to the first one chosen are then processed in the same manner (where connected means appearing in the same relation) until the process can not continue. The isolated variables are then treated last.

Evaluating path queries. For evaluating a path query (2RPQ), the path pattern is compiled into an automaton. Then a “virtual” cross-product of this automaton and the graph is constructed on-the-fly, and navigated via breadth-first search, as commonly suggested in the theoretical literature [49, 50, 39] (our experiments in Section 6 will also test a depth-first search (DFS) variant). Our assumption is that each path pattern will have at least one of the endpoints assigned before evaluation. This can be done either explicitly in the pattern, or via the remainder of the query. For instance, a path pattern $(Q1)=[P31^*]=>(?x)$ has the starting point of our search assigned to Q1. On the other hand, $(?x)=[P31^*]=>(?y :person)$ does not have any of the endpoints assigned, however, the $(?y :Person)$ allows us to instantiate $?y$ with any node with the label $:person$.

Intuitively, from a starting node (tagged with the initial state of the automaton), all edges with the type specified by the outgoing transitions from this state are followed. The process is repeated until reaching an end state of the automaton, upon which a result can be returned. This allows a fully pipelined evaluation of path queries, while only requiring at most a fixed amount of memory (the neighbors of the node on the top of the BFS queue). Additionally, the BFS algorithm also allows us to return a single shortest path between each pair of endpoints (see Section 4 for an example). Returning a single shortest path comes almost for free, given that it can be reconstructed using the set of visited nodes as used for bookkeeping in the BFS algorithm. The algorithm can also be extended to return all shortest paths (as supported by DGQL) by keeping a list of predecessors that reach the node via a path of shortest length.

The implemented algorithm only requires two permutations of the `DOMAINGRAPH` relation: one for retrieving all of a node’s successors via an edge of a specified type; and another for retrieving all such predecessors of a given node.

6. BENCHMARKING

In this section, we provide an experimental evaluation of the core graph querying features of MillenniumDB addressing two key questions: (Q1) *Which join and path algorithms provide the best performance over domain graphs?* (Q2) *How does MillenniumDB’s performance compare with existing graph database engines?*

We base our experiments on the Wikidata knowledge graph [21], which is one of the largest and most diverse real-world knowledge graphs that is publicly available, and also provides a public log of real-world queries posted by Wikidata users that we can use for experiments [22, 34]. The experiments focus on two fundamental query features: (i) basic graph patterns (BGPs); and (ii) path queries. Regarding (Q1), we compare the performance of different join and path algorithms within MillenniumDB. Regarding (Q2), we also provide a side by side comparison with several popular persistent graph database engines that support BGPs and at least the Kleene star feature for paths. We publish the data, queries, scripts, and configuration files for each engine online, together with the scripts used to load the data and run the experiments [26].

Internal baselines. The base of our comparison is the MillenniumDB implementation available at [25]. For comparing the performance of different join and path algorithms in MillenniumDB (per Q1), we include internal baselines, where we test: (i) MillenniumDB LF, which is the default version implementing the leapfrog triejoin algorithm; (ii) MillenniumDB GR, which implements the greedy algorithm for selecting the join order; and (iii) MillenniumDB SL, implementing the Sellinger join planner. Similarly, for path queries, we test (a) MillenniumDB BFS, the default version of the engine; and (b) MillenniumDB DFS, which evaluates path queries using the depth-first traversal.

Other engines. We also compare the performance of MillenniumDB with five persistent graph query engines (per Q2). First, we include three popular RDF engines: Jena TDB version 4.1.0 [18], Blazegraph (BlazeG for short) version 2.1.6 [16], and Virtuoso version 7.2.6 [20]. We further include a property graph engine: Neo4j community edition 4.3.5 [6].[®] Finally, we also compare with Jena Leapfrog (Jena LF, for short)—a version of Jena TDB implementing a leapfrog-style algorithm [43]—in order to compare with an external graph database using a worst case optimal algorithm.

The machine. All experiments described were run on a single commodity server with an Intel®Xeon®Silver 4110 CPU, and 128GB of DDR4/2666MHz RAM, running the Linux Debian 10 operating system with the kernel version 5.10. The hard disk used to store the data was a SEAGATE ST14000NM001G with 14TB of storage.

The data. The base for our experiments is the Wikidata dataset. In particular, we used the truthy dump version 20210623-truthy-BETA [51], keeping only triples in which (i) the subject position is a Wikidata entity, and (ii) the predicate is a direct property. We call this dataset *Wikidata Truthy*. The size of the dataset after this process was 1,257,169,959 triples. The simplification of the dataset is done to facilitate comparison across multiple engines, specifically to keep data loading times across all engines manageable while keeping the nodes and edges necessary for testing the performance of BGPs and property paths. The size of the *Wikidata Truthy* dataset, when loaded into the respective systems, is summarized in Table 3. Default indices were used on Jena TDB, Blazegraph and Virtuoso. Jena LF stores three additional permutations of the stored triples to efficiently support the leapfrog algorithm for any join query, thus using more space. Neo4j by default creates an index for edge types (as of version 4.3.5). To speed up searches for particular entities and properties, we also created an index linking a Wikidata identifier (such as, e.g., Q510) to its internal id in Neo4j. We also tried to index literal values in Neo4j, but the process failed (the literals are still stored). MillenniumDB uses extra disk space because of the additional indices needed to support worst-case optimal join over domain graphs (similar to the case of Jena LF).

Table 3. Wikidata Truthy sizes when loaded into each engine. The base dataset consists of roughly 1.25 billion triples.

MillenniumDB	BlazeG	Jena	Jena LF	Virtuoso	Neo4j
203GB	70GB	110GB	195GB	70GB	112GB

[®] Though TigerGraph meets the technical requirements, its license currently restricts benchmarking and thus it is excluded.

How we ran the queries. We detail the query sets used for the experiments in their respective subsections. To simulate a realistic database load, we do not split queries into cold/hot run segments. Rather, we run them in succession, one after another, after a cold start of each system (and after cleaning the OS cache). This simulates the fact that query performance can vary significantly based on the state of the system buffer, or even on the state of the hard drive, or the state of OS's virtual memory. For each system, queries were run in the same order. We record the execution time of each individual query, which includes iterating over all results. We set a limit of 100,000 distinct results for each query, again in order to enable comparability as some engines showed instability when returning larger results.

Memory usage. Blazegraph, Jena and Virtuoso were assigned 64GB of RAM, as is recommended. Neo4J was run with default settings[®], while MillenniumDB had access to 32GB for main-memory buffer, and it uses an additional 10GB for in-memory dictionaries. Since the systems tested are buffer-based—i.e., since they reserve a fixed amount of scrap space (the buffer) in main memory for their operation, and do not exceed this memory (except perhaps modulo a small amount used for internal operations)—and since they tend to use the buffer available, their maximum memory usage corresponds to these settings. Thus, in the rest of this section we focus on comparing runtimes.

Handling timeouts. We defined a timeout of 10 minutes per query for each system. Apart from that, we note that most systems had to be restarted upon a timeout as they often showed instability, particularly while evaluating path queries. This was done without cleaning the OS cache in order to preserve some of the virtual memory mapping that the OS built up to that point. In comparison, MillenniumDB managed to return a non-trivial amount of query results on each query, and did not need to be restarted, thus handling timeouts gracefully.

6.1 Basic Graph Patterns

We focus first on basic graph pattern queries. To test different query execution strategies of MillenniumDB, we use two benchmarks: *Real-world BGPs* and *Complex BGPs*, which are described next.

Real-world BGPs. The Wikidata SPARQL query log contains millions of queries [22], but many are trivial to evaluate. We thus generate our benchmark from more challenging cases, i.e., a smaller log of queries that timed-out on the Wikidata public endpoint [22]. From these queries we extracted their BGPs, removing duplicates (modulo isomorphism on query variables). We distinguish queries consisting of a single triple pattern (*Single*) from those containing more than one triple pattern (*Multiple*). The former set tests the triple matching capabilities of the systems, whereas the latter set tests join performance. *Single* contains 399 queries, whereas *Multiple* has 436 queries.

[®] We also tried increasing the `dbms.memory.pagecache.size` parameter manually to 64GB, and setting `dbms.memory.heap.initial_size` and `dbms.memory.heap.max_size` to 30GB each, but the variation in the runtimes between the two settings was less than 0.5%. We believe that this is because in both cases, Neo4J manages to run in main memory without swapping, so varying these configurations has little effect.

Real-world Single. Table 4 (top) summarizes the query times on this set, whereas Figure 13 (left) shows boxplots with more detailed statistics on the distributions of runtimes. Since these queries do not require joins, we show one variant for MillenniumDB. MillenniumDB is the fastest overall (median of 0.05 s), followed by Blazegraph (median of 0.09 s). In terms of average times and higher percentiles, MillenniumDB more clearly outperforms other engines, being able to enumerate up to 100,000 results (the limit) more quickly due to decoding internal ids more quickly. In MillenniumDB, values such as P12 or Q10 that fit within 7 bytes are inlined, and do not need to be dictionary decoded. The remaining dictionary fits entirely in available memory (~8 GB of RAM). In the other systems, dictionary decoding generates random accesses to the disk. The four SPARQL engines tested must store IDs as IRIs within the RDF model, which include relatively long prefixes. However, since RDF datasets typically have few prefixes repeated often, we could support full IRIs within MillenniumDB with minimal overhead by encoding a prefix id for the top k prefixes in $\lceil \log_2(k) \rceil$ bits within the object identifier, keeping the small mapping from prefix id to string in memory. The Wikidata query service lists 32 prefixes, which would require 5 bits that would fit “for free” in the class byte (essentially considering each prefix to be a class).

Table 4. Summary of runtimes (in seconds) for BGPs.

Engine	Supported	Error	Timeouts	Average	Median
<i>Real-world Single (399 queries)</i>					
MillenniumDB	399	0	0	0.07	0.05
Blazegraph	399	0	0	2.21	0.09
Jena	399	0	0	14.10	0.34
Jena LF	395	4	0	10.08	0.44
Virtuoso	399	0	0	2.22	0.32
Neo4j	394	5	0	28.00	1.33
<i>Real-world Multiple (436 queries)</i>					
MillenniumDB LF	436	0	0	4.84	0.24
MillenniumDB GR	436	0	1	10.19	0.30
MillenniumDB SL	436	0	1	10.04	0.27
Blazegraph	436	0	3	31.79	2.42
Jena	426	10	0	35.43	4.90
Jena LF	418	18	0	16.78	3.39
Virtuoso	436	0	0	7.87	5.11
Neo4j	405	31	0	75.55	6.84
<i>Complex (850 queries)</i>					
MillenniumDB LF	850	0	0	0.38	0.10
MillenniumDB GR	850	0	1	3.30	0.17
MillenniumDB SL	850	0	1	3.51	0.17
Blazegraph	850	0	2	4.63	0.34
Jena	850	2	0	3.37	0.16
Jena LF	850	0	0	0.88	0.14
Virtuoso	850	0	0	1.00	0.19
Neo4j	850	10	0	17.92	0.66

Real-world Multiple. Table 4 (middle) and Figure 13 (middle) show the results for this set. Comparing different join execution strategies within MillenniumDB, we can see the superiority of the leapfrog triejoin variant (particularly on average, i.e., for more complex queries). The Selinger variant of MillenniumDB outperforms the greedy algorithm for join selection, but only marginally. Compared with existing graph engines, MillenniumDB clearly outperforms other systems on this query set. Its medians are an order of magnitude faster than those of Blazegraph, the next best contender. The difference is less sharp for averages, but MillenniumDB LF still takes 60% of the time of Virtuoso, the next best contender.

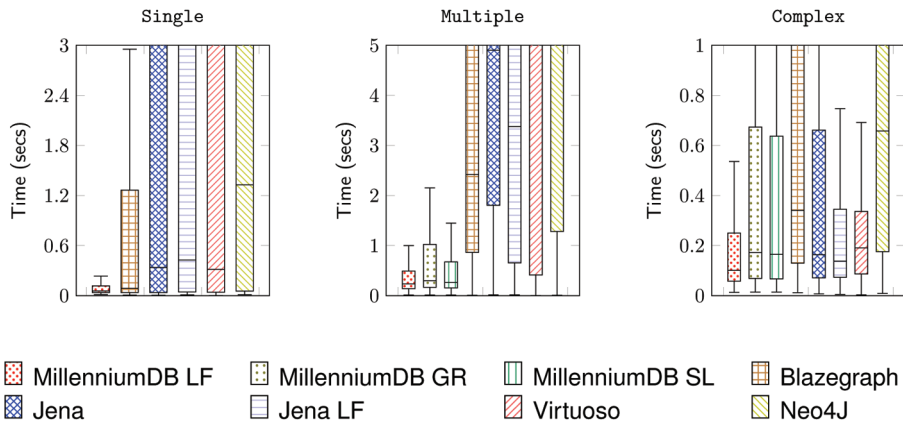


Figure 13. Boxplots showing distributions for BGP runtimes.

Complex BGPs. This is a benchmark used to test the performance of worst-case optimal joins [43]. Here, 17 different complex join patterns were selected, and 50 different queries generated for each pattern, resulting in a total of 850 queries. Figure 13 (right), and Table 4 (bottom), show the resulting query times. In this case, the difference between the join algorithms of MillenniumDB is more clear. The worst-case-optimal version (MillenniumDB LF) is not only considerably more stable than the other two versions, but also twice as fast in the median. We can also observe that MillenniumDB GR wins out over MillenniumDB SL on average (but not the median case). When comparing with other engines, the next-best competitor after MillenniumDB LF is Jena LF, showing the benefits of worst-case optimal joins. Virtuoso follows not far behind, while MillenniumDB GR, Jena, Blazegraph and Neo4j are considerably slower. Overall, MillenniumDB LF offers the best performance for every statistic shown in the plot.

6.2 Path Queries

To test the performance of path queries, we extracted 2RPQ expressions from a log of queries that timed out on the Wikidata endpoint [22]. The original log has 2110 queries. After removing queries that do not use direct properties (which are absent in the *Wikidata Truthy* dataset), we ended up with 1683 queries. These were run in succession, each restricted to return at most 100,000 results. In the case of SPARQL engines, we added the `DISTINCT` keyword to remove duplicates caused by the rewriting of fixed-length path queries to unions of BGPs that are then evaluated under bag semantics. To make the comparison fair,

the DISTINCT keyword was also added in MillenniumDB queries. Each system was started after cleaning the system cache, and with a timeout of 10 minutes. Since these are originally SPARQL queries, not all of them were supported by Neo4J given the restricted regular-expression syntax it supports. MillenniumDB and Neo4J were the only systems able to handle timeouts without being restarted.[®] In this comparison we do not include Jena LF since it uses the same execution strategy as Jena for property paths. Likewise, for MillenniumDB, we introduce two internal baselines for breadth-first search (BFS) and depth-first search (DFS). The experimental results for these path queries are summarized in Table 5 and Figure 14.

Table 5. Summary of runtimes (in seconds) for path queries.

Engine	Supported	Error	Timeouts	Average	Median
MillenniumDB BFS	1683	0	0	1.1	0.095
MillenniumDB DFS	1683	0	0	1.1	0.072
Blazegraph	1683	2	44	27.6	0.396
Jena	1683	14	46	22.8	0.207
Virtuoso	1683	55	4	5.8	0.325
Neo4J	1622	0	42	23.3	0.328

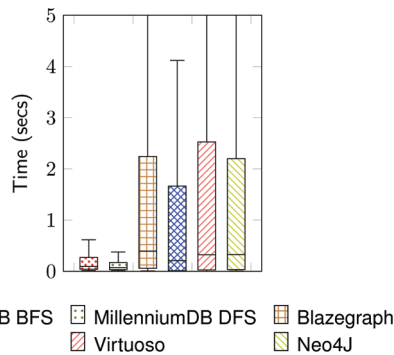


Figure 14. Boxplots of query times on property paths.

In terms of our internal comparison, we can see that the DFS algorithm slightly outperforms BFS. The reason for keeping BFS as the default algorithm is twofold: (i) it significantly outperforms DFS when paths are also returned; and (ii) it supports returning all shortest paths between any pair of nodes. To illustrate point (i), we ran our experiments again, but now also returning a single path witnessing each query answer. In this case, the average for BFS is 5.9 sec, and median is 0.086 sec. On the other hand, when paths are returned, DFS takes 7.9 sec on average, and 0.1 sec median time.

[®] In fact, MillenniumDB did not give any timeouts. However, we re-ran the experiments with a lower timeout, and observed that the system could recover from interrupting the query gracefully and was able to return the results found before being interrupted.

Compared with other engines, MillenniumDB is generally the fastest, and has the most stable performance. Its average is near a second, i.e., five times faster than the next best contender (Virtuoso). Its median, below 0.1 seconds, is half the next one (Jena's). Even after removing the queries that timed-out on the other systems, they are considerably slower than MillenniumDB. In particular, if we only consider the queries that run successfully on Virtuoso (i.e., excluding the 59 queries that timed-out or gave an error), we get an average time of 0.85 seconds and a median time of 0.086 seconds on MillenniumDB: less than half the times of Virtuoso with these queries excluded. The boxplots further show the stability of MillenniumDB: the medians of other engines are above the third quartile of MillenniumDB. Their third quartile is 5–10 times higher than MillenniumDB's, and higher than its topmost whisker.

To further test robustness, we also ran all of the queries *without limiting the output size* on MillenniumDB. In this test, the engine timed out in only 15 queries, each returning between 800 thousand and 44 million results before timing out. When running queries to completion, MillenniumDB BFS averaged 13.4 seconds per query (8 seconds excluding timeouts), with a median of 0.1 seconds (both with and without timeouts).

6.3 Wikidata Complete

To show the scalability of MillenniumDB, and to properly leverage its domain graph, we ran experiments with a full version of Wikidata. We call this dataset *Wikidata Complete*, and base it off the Wikidata JSON dump[®] version 20201102-all.json, which is preprocessed and mapped to our data model. In *Wikidata Complete*, we model qualifiers (i.e. edges on edges), put labels on objects, and assign them properties with values. We use properties to store the language value of each string in Wikidata, and also to model elements of complex data values (e.g., for coordinates we would have objects with properties latitude and longitude, and similarly for amounts, date/time, limits, etc.). Each object representing a complex data value also has a label specifying its data type (e.g. coord for geographical coordinates). All qualifiers were loaded. The only elements excluded from the full Wikidata data dump were sitelinks and references. This full version of Wikidata resulted in a knowledge graph with roughly 300 million objects, participating in 4.3 billion edges. The total size on disk of this data was 827GB in MillenniumDB, i.e., more than four times larger than *Wikidata Truthy*. More details about this dataset can be found in the online material accompanying this paper [26].

We ran the same queries from the benchmarks (*Single*, *Multiple* and *Complex* BGPs, as well as *Paths*). The number of outputs on the two versions of the data, while not the same, was within the same order of magnitude averaged over all the queries. The results are presented in Table 6. As we can observe, MillenniumDB shows no deterioration in performance when a larger database is considered for similar queries. This is mostly due to the fact that the buffer only loads the necessary pages into the main memory, and will probably require a rather similar effort in both cases. We also note that, again, no queries resulted in a timeout over the larger dataset.

[®] It is important to note that JSON and RDF dumps of Wikidata do not result in precisely the same knowledge graph due to some restrictions of the particular reification used in RDF; however, they do result in very similar knowledge graphs.

Table 6. Average and median runtimes, in seconds, for MillenniumDB on the complete version of Wikidata.

	Single	Multiple	Complex	Paths
Average	0.08	4.04	0.35	1.04
Median	0.07	0.28	0.095	0.10

6.4 Discussion

Regarding (Q1)—i.e., which join and path algorithms provide the best performance in this setting—regarding join algorithms, we can conclude that the worst-case optimal join algorithm consistently outperforms the greedy and Selinger variants, being particularly notable in the case of more complex graph patterns with many joins (wherein Jena LF—also worst-case optimal—was the next best competitor). Worst-case optimal joins use more space for indexing, but provide superior query runtimes. Regarding path algorithms, we see less difference between BFS and DFS: DFS is slightly faster for returning pairs of nodes connected by paths, while BFS is faster for returning paths.

Regarding (Q2)—i.e., how existing graph database systems compare with MillenniumDB—we found that MillenniumDB, when equipped with the best join and path algorithms, consistently outperforms other competitors in all query sets tested.

7. CONCLUSIONS AND LOOKING AHEAD

This paper presents MillenniumDB, an open-source graph database system with persistent storage implementing the novel (property) domain graph model.

Domain graphs adopt the natural idea of adding edge ids to directed labeled edges in order to concisely model higher-arity relations in graphs, as needed in Wikidata, without the need for reserved vocabulary or reification. They can naturally represent popular graph models, such as RDF and property graphs, and allow for combining the features of both models in a novel way. While the idea of using edge ids as a hook for modeling higher-arity relations in graphs is far from new (see, e.g., [23, 32, 33]), it is an idea that is garnering increased attention as a more flexible and concise alternative to reification. Our work proposes a formal data model that incorporates edge ids, a query language that can take advantage of them, and a fully-fledged graph database engine that supports them by design. We also propose to optionally allow (external) annotations on top of the graph structure, thus facilitating better compatibility with property graphs, whereby labels and property-values can be added to graph objects without adding new nodes and edges to the graph itself.

We have also proposed a new query language with a syntax inspired by Cypher, but that additionally enables users to take full advantage of the domain graph model by (optionally) referencing edge ids in their queries, and performing joins on any element of the domain graph. We further combine useful features

present in both Cypher and SPARQL, in order to provide additional expressivity, such as returning the shortest path witnessing a result for a path query (as captured by a 2RPQ expression).

In the implementation of MillenniumDB, we combine both tried-and-trusted techniques that have been successfully used in relational database pipelines for decades [42] (e.g., B+ trees, buffer managers, etc.), with promising state-of-the-art algorithms for computing worst case optimal joins (leapfrog [38]) and evaluating path queries (guided by an automaton [49, 50]). Our experiments over Wikidata, considering real-world queries and data at large-scale, show that this combination outperforms other persistent graph database engines that are commonly found in practice.

Limitations. Many of the current limitations of MillenniumDB relate to the fact that it is still under development. For example, at the moment, MillenniumDB only supports a bulk load of data, where support for (incremental) updates is currently under investigation and development. Currently only the core features of query languages such as Cypher and SPARQL are supported, where we are working on adding support for other features, including negation, value assignment, functions on datatypes, etc. MillenniumDB lacks some of the advanced features supported by other graph database systems, such as geographic, temporal and federated queries, keyword search features, etc. Finally, MillenniumDB does not yet support partitioning the graph over multiple machines in order to achieve horizontal scaling. We do not see such limitations as fundamental, but rather as features that can be added to the engine over time.

Future work. Looking to the future, we foresee extensions such as: returning entire graphs, supporting more complex path constraints, returning sets of paths, path algebra, just to name a few. Regarding more practical features, we aim to add support for full transactions, keyword search, a graph update language, existing graph query languages, and more besides. More importantly, given that MillenniumDB is published as an open source engine, we hope that the research community can view the MillenniumDB code base as a sandbox for incorporating their novel algorithms and ideas into a modern graph database, without the need to remake storage, indexing, access methods, or query parsers. Along these lines, we are currently working on adding an in-memory storage option to MillenniumDB using the ring [47]: a data structure based on the Burrows—Wheeler transform that supports worse-case optimal joins (over triples) in space similar to representing the graph itself. Initial tests show that the ring can store *Wikidata Truthy* in 50GB of space and improve median query times by a factor of 3, with average query times remaining similar. We are working on extending the ring to support edge ids and thus work with domain graphs. We also wish to explore the deployment of MillenniumDB for key use-cases; for example, we plan to provide and host an alternative query service for Wikidata, which may help to prioritize the addition of novel features and optimizations as needed in practice.

AUTHOR CONTRIBUTIONS

Vrgoč acted as the project leader, designed the architecture of MillenniumDB, and adapted the theoretical algorithms used in this work for implementation purposes. Rojas was the lead engineer, working in designing the architecture of MillenniumDB, and implementing the majority of the system. Angles participated in

writing the paper, designing the data model and designing the query language. Arenas participated in writing the paper, designing the data model and designing the query language. Arroyuelo participated in designing and analyzing the experiments, and writing the paper. Buil-Aranda participated in writing the paper, designing the data model, and designing experiments. Hogan participated in writing the paper, designing the data model, and designing experiments. Navarro participated in writing the paper and analyzing experiments. Riveros participated in writing the paper and designing the architecture of MillenniumDB. Romero helped in implementing path queries.

ACKNOWLEDGEMENTS

This work was supported by ANID—Millennium Science Initiative Program—Code ICN17_002.


REFERENCES

- [1] Angles, R., Gutiérrez, C.: Survey of graph database models. *ACM Comput. Surv.* 40(1), 1:1–1:39 (2008)
- [2] Angles, R., Arenas, M., Barceló, P., et al.: Foundations of Modern Query Languages for Graph Databases. *ACM Comput. Surv.* 50(5), 68:1–68:40 (2017)
- [3] Cattell, R.: Scalable SQL and NoSQL data stores. *SIGMOD Rec.* 39(4), 12–27 (2010)
- [4] Heath, T., Bizer, C.: *Linked Data: Evolving the Web into a Global Data Space*. Synthesis Lectures on the Semantic Web. Morgan & Claypool Publishers (2011)
- [5] Hogan, A., Blomqvist, E., Cochez, M., et al.: Knowledge Graphs. CoRR abs/2003.02320 (2020)
- [6] Webber, J.: A programmatic introduction to Neo4j. In: Leavens, G.T. (ed.) *Conference on Systems, Programming, and Applications: Software for Humanity, SPLASH '12*, Tucson, AZ, USA, October 21–25, 2012, pp. 217–218. ACM (2012)
- [7] Francis, N., Green, A., Guagliardo, P., et al.: Cypher: An Evolving Query Language for Property Graphs. In: Das, G., Jermaine, C.M., Bernstein, P.A. (eds.) *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018*, Houston, TX, USA, June 10–15, 2018, pp. 1433–1445. ACM (2018)
- [8] ArangoDB Team. ArangoDB v3.7.11 Documentation (2021). Available at: <https://www.arangodb.com/docs/stable/>
- [9] JanusGraph Team. JanusGraph Documentation, v.0.5 (2022). Available at: <https://docs.janusgraph.org/>
- [10] OrientDB Team. OrientDB Manual – version 3.0.34 (2022). Available at: <https://orientdb.org/docs/3.0.x/>
- [11] TigerGraph Team. TigerGraph Documentation – version 3.1 (2021)
- [12] Rodriguez, M.A.: The Gremlin graph traversal machine and language (invited talk). In: Cheney, J., Neumann, T. (eds.) *Proceedings of the 15th Symposium on Database Programming Languages*, Pittsburgh, PA, USA, October 25–30, 2015, pp. 1–10. ACM (2015)
- [13] Ali, W., Saleem, M., Yao, B., et al.: A survey of RDF stores & SPARQL engines for querying knowledge graphs. *VLDB J.* 31(3), 1–26 (2022)
- [14] AllegroGraph Team. AllegroGraph 7.1.0 Documentation (2021). Available at: <https://franz.com/agraph/support/documentation/current/>
- [15] Amazon Neptune Team. What Is Amazon Neptune? (2021). Available at: <https://docs.aws.amazon.com/neptune/latest/userguide/intro.html>

- [16] Thompson, B.B., Personick, M., Cutcher, M.: The Bigdata® RRDF Graph Database. In: Harth, A., Hose, K., Schenkel, R. (eds.) *Linked Data Management*, pp. 193–237. Chapman and Hall/CRC (2014)
- [17] Bishop, B., Kiryakov, A., Ognyanoff, D., et al.: OWLIM: A family of scalable semantic repositories. *Semantic Web* 2(1), 33–42 (2011)
- [18] Jena Team. TDB Documentation (2022). Available at: <https://jena.apache.org/documentation/tdb/>
- [19] Stardog Team. Stardog 7.6.3 Documentation (2021). Available at: <https://docs.stardog.com/>
- [20] Erling, O.: Virtuoso, a Hybrid RDBMS/Graph Column Store. *IEEE Data Eng. Bull.* 35(1), 3–8 (2012)
- [21] Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. *Commun. ACM* 57(10), 78–85 (2014)
- [22] Malyshev, S., Krötzsch, M., González, L., et al.: Getting the Most Out of Wikidata: Semantic Technology Usage in Wikipedia’s Knowledge Graph. In: *International Semantic Web Conference (ISWC)*, pp. 376–394. Springer (2018)
- [23] Hernández, D., Hogan, A., Krötzsch, M.: Reifying RDF: What Works Well With Wikidata? In: Liebig, T., Fokoue, A. (eds.) *Proceedings of the 11th International Workshop on Scalable Semantic Web Knowledge Base Systems co-located with 14th International Semantic Web Conference (ISWC 2015)*, Bethlehem, PA, USA, October 11, 2015, volume 1457 of *CEUR Workshop Proceedings*, pp. 32–47. CEUR-WS.org (2015)
- [24] Angles, R., Hogan, A., Lassila, O., et al.: Multilayer graphs: a unified data model for graph databases. In: Kalavri, V., Salihoglu, S. (eds.) *GRADES-NDA ’22: Proceedings of the 5th ACM SIGMOD Joint International Workshop on Graph Data Management Experiences & Systems (GRADES) and Network Data Analytics (NDA)*, Philadelphia, Pennsylvania, USA, 12 June 2022, pp. 11:1–11:6. ACM (2022)
- [25] MillenniumDB Team. MillenniumDB Source Code (2021). Available at: <https://github.com/MillenniumDB/MillenniumDB>
- [26] MillenniumDB Team. Wikidata Benchmark (2021). Available at: <https://github.com/MillenniumDB/benchmark>
- [27] Cyganiak, R., Wood, D., Lanthaler, M.: *RDF 1.1 Concepts and Abstract Syntax*. W3C Recommendation (2014)
- [28] Hartig, O.: Foundations of RDF* and SPARQL* (An Alternative Approach to Statement-Level Metadata in RDF). In: Reutter, J.L., Srivastava, D. (eds.) *Proceedings of the 11th Alberto Mendelzon International Workshop on Foundations of Data Management and the Web*, Montevideo, Uruguay, June 7–9, 2017, volume 1912 of *CEUR Workshop Proceedings*. CEUR-WS.org (2017)
- [29] Hartig, O., Champin, P.-A., Kellogg, G., et al.: *RDF-star and SPARQL-star*. W3C Draft Community Group Report (2021)
- [30] Harris, S., Seaborne, A., Prud’hommeaux, E.: *SPARQL 1.1 Query Language*. W3C Recommendation (2013)
- [31] Sahu, S., Mhedhbi, A., Salihoglu, S., et al.: The ubiquity of large graphs and surprising challenges of graph processing: extended survey. *VLDB J.* 29(2–3): 595–618 (2020)
- [32] Ilievski, F., Garijo, D., Chalupsky, H., et al.: KGTK: A Toolkit for Large Knowledge Graph Manipulation and Analysis. In: *International Semantic Web Conference (ISWC)*, pp. 278–293. Springer (2020)
- [33] Lassila, O., Schmidt, M., Bebee, B., et al.: Graph? Yes! Which one? Help! CoRR abs/2110.13348 (2021)
- [34] Bonifati, A., Martens, W., Timm, T.: An analytical study of large SPARQL query logs. *VLDB J.* 29(2–3): 655–679 (2020)
- [35] Pérez, J., Arenas, M., Gutiérrez, C.: Semantics and complexity of SPARQL. *ACM Trans. Database Syst.* 34(3): 16:1–16:45 (2009)
- [36] Angles, R., Arenas, M., Barceló, P., et al.: G-CORE: A core for future graph query languages. In: *Proceedings of the 2018 International Conference on Management of Data, SIGMOD Conference 2018*, Houston, TX, USA, June 10–15, 2018, pp. 1421–1432 (2018)

- [37] Nebula Team. Nebula Graph Query Language (nGQL) (2021). Available at: <https://docs.nebula-graph.io/1.2.0/manual-EN/1.overview/1.concepts/2.nGQL-overview/>
- [38] Veldhuizen, T.L.: Triejoin: A simple, worst-case optimal join algorithm. In: Proc. 17th International Conference on Database Theory (ICDT), Athens, Greece, March 24–28, 2014, pp. 96–106 (2014)
- [39] Baier, J.A., Daroch, D., Reutter, J.L., et al.: Evaluating navigational RDF queries over the web. In: Proceedings of the 28th ACM Conference on Hypertext and Social Media, HT 2017, Prague, Czech Republic, July 4–7, 2017, pp. 165–174 (2017)
- [40] Ngo, H.Q., Porat, E., Ré, C., et al.: Worst-case optimal join algorithms. *J. ACM* 65(3): 16:1–16:40 (2018)
- [41] Hogan, A., Arenas, M., Mallea, A., et al.: Everything you always wanted to know about blank nodes. *J. Web Semant.* 27–28: 42–69 (2014)
- [42] Ramakrishnan, R., Gehrke, J.: Database management systems. McGraw-Hill (2000)
- [43] Hogan, A., Riveros, C., Rojas, C., et al.: A worst-case optimal join algorithm for SPARQL. In: Ghidini, C., Hartig, O., Maleshkova, M., et al. (eds.) *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference*, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I, volume 11778 of *Lecture Notes in Computer Science*, pp. 258–275. Springer (2019)
- [44] Sciore, E.: Database Design and Implementation - Second Edition. *Data-Centric Systems and Applications*. Springer (2020)
- [45] Garcia-Molina, H., Ullman, J.D., Widom, J.: Database systems - the complete book (2. ed.). Pearson Education (2009)
- [46] Mhedhbi, A., Kankanamge, C., Salihoglu, S.: Optimizing One-time and Continuous Subgraph Queries using Worst-case Optimal Joins. *ACM Trans. Database Syst.* 46(2): 6:1–6:45 (2021)
- [47] Arroyuelo, D., Hogan, A., Navarro, G., et al.: Worstcase optimal graph joins in almost no space. In: Li, G.L., Li, Z.H., Idreos, S., et al. (eds.) *SIGMOD '21: International Conference on Management of Data*, Virtual Event, China, June 20–25, 2021, pp. 102–114. ACM (2021)
- [48] Yu, C.T., Ozsoyoglu, M.Z.: An algorithm for tree-query membership of a distributed query. In: *The IEEE Computer Society's Third International Computer Software and Applications Conference, COMPSAC 1979*, 6–8 November, 1979, Chicago, Illinois, USA, pp. 306–312 (1979)
- [49] Mendelzon, A.O., Wood, P.T.: Finding regular simple paths in graph databases. In: *Proceedings of the Fifteenth International Conference on Very Large Data Bases*, August 22–25, 1989, Amsterdam, The Netherlands, pp. 185–193 (1989)
- [50] Baeza, P.B.: Querying graph databases. In: *Proceedings of the 32nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, PODS 2013*, New York, NY, USA - June 22–27, 2013, pages 175–188 (2013)
- [51] Wikimedia Foundation. Wikidata: database download (2021). Available at: https://www.wikidata.org/wiki/Wikidata:Database_download

AUTHOR BIOGRAPHY



Domagoj Vrgoč is an Assistant Professor at the Institute for Mathematical and Computational Engineering at the Pontificia Universidad Católica de Chile, and an Associate Researcher at the Millennium Institute for Foundational Research on Data (IMFD). He received a MS at the University of Zagreb in 2007, and a PhD at the University of Edinburgh in 2014. His work mainly focuses on Graph Databases, Semantic Web, and Information Extraction.



Carlos Rojas is a software engineer at the Millennium Institute for Foundational Research on Data (IMFD). He received an MS at the Pontificia Universidad Católica de Chile in 2019 and has been working on MillenniumDB since then.



Renzo Angles is an Associate Professor in the Department of Computer Science at the Universidad de Talca (Chile). He received the degree of Bachelor in Systems Engineering from Universidad Católica de Santa María (Arequipa, Perú), and a Ph.D. degree in Computer Science from Universidad de Chile in 2009. During 2013, he carried out postdoctoral research at the Department of Computer Science, VU University Amsterdam, as part of his participation in the Linked Data Benchmark Council Project. Currently he participates as a researcher in the Millennium Institute for Foundational Research on Data (Chile). His research interests lie in the intersection of graph databases and the Semantic Web. Specifically, he works in the theory and design of graph query languages, and the interoperability between RDF and graph databases.




Marcelo Arenas is a Professor at the Department of Computer Science and the Institute for Mathematical and Computational Engineering, at the Pontificia Universidad Católica de Chile. He is a Distinguished Member of the Association for Computing Machinery (ACM), the director of the Millennium Institute for Foundational Research on Data and the former director of the Center for Semantic Web Research. He received a Ph.D. from the University of Toronto in 2005. His research interests are in the areas of data management, applications of logic in computer science and Semantic Web.



Diego Arroyuelo is an Assistant Professor of the Department of Informatics, Universidad Técnica Federico Santa María, Chile, and an Adjunct Researcher of the Millennium Institute for Foundational Research on Data (IMFD). He received a Ph.D. in Computer Science from University of Chile in 2009, and then carried out a postdoctoral stay at the David Cheriton School of Computer Science, University of Waterloo, Canada. His research interests relate to data structures and the design and analysis of algorithms, in particular space-efficient data structures for sets, trees, graphs, strings, and database joins.



Carlos Buil-Aranda is an Assistant Professor at Universidad Técnica Federico Santa María and Junior Researcher at the Millennium Institute for Foundational Research on Data (IMFD). He received a PhD at the Polytechnical University of Madrid in 2012. His work mainly focuses on Graph Databases and the Semantic Web.



Aidan Hogan is an Associate Professor of the Department of Computer Science, University of Chile and an Associate Researcher of the Millennium Institute for Foundational Research on Data (IMFD). He received a B.Eng. and Ph.D. from the National University of Ireland, Galway, in 2006 and 2011, respectively. His main research interests relate to the Semantic Web and Knowledge Graphs.



Gonzalo Navarro completed his PhD in Computer Science in 1998 at the University of Chile, where he is currently Full Professor. His areas of interest include algorithms and data structures, compression, graph databases, and text searching. He has directed the Millennium Nucleus Center for Web Research and projects funded by Yahoo! Research and Google. He currently participates in the Center for Biotechnology and Bioengineering (CeBiB) and the Millennium Institute for Foundational Research on Data (IMFD). He has been PC (co-)chair of over ten international conferences and guest editor of several special issues in relevant journals. He is the Editor in Chief of the ACM Journal of Experimental Algorithmics and a member of the Editorial Board of the ACM Transactions on Algorithms and Information Systems. He has co-authored two books published by Cambridge University Press, about 25 book chapters, 200 papers in international journals, and around 275 in international conferences. He is one of the most prolific and highly cited authors in Latin America. He has received 7 Best Paper Awards in conferences, 4 Google Research awards, a Highest Cited Paper Award from Elsevier, and a Scopus Chile Award. He is an ACM Fellow.



Cristian Riveros is an Associate Professor at the Department of Computer Science at the Pontificia Universidad Católica de Chile, and a Junior Researcher at the Millennium Institute for Foundational Research on Data (IMFD). He received a D. Phil from the University of Oxford in 2013. His research interests are mostly in data management systems, specifically, in data streams management systems, information extraction, and graph data. He also does research in theoretical computer science, like automata theory, logics, and computational complexity.



Juan Romero received an engineering degree at the Pontificia Universidad Católica de Chile in 2023. He completed several internships at the Millennium Institute for Foundational Research on Data (IMFD), where he worked on implementing path queries in MillenniumDB.