



Worst-Case Optimal BGPs on Temporal Graphs

Diego Arroyuelo
Pontificia Universidad
Católica de Chile &
IMFD, Chile
diego.arroyuelo@uc.cl

Aidan Hogan
DCC, Universidad de Chile
&
IMFD, Chile
ahogan@dcc.uchile.cl

Gonzalo Navarro
DCC, Universidad de Chile
&
IMFD, Chile
gnavarro@dcc.uchile.cl

Juan Reutter
Pontificia Universidad
Católica de Chile &
IMFD, Chile
jreutter@ing.puc.cl

ABSTRACT

We study how to evaluate basic graph patterns (BGPs) in a worst-case-optimal (wco) manner over *temporal* labeled graphs, where edges have an interval of temporal validity. We adopt a flexible query language in which users specify m quads of the form (subject, property, object, time), using constants or variables. The time component denotes the instant at which a particular edge is valid, and users may also include order relations between temporal constants or variables. The answer is the set of all valid variable assignments, including time. We describe an index structure that, for a temporal graph with N edges, requires $O(N)$ space and can evaluate extended BGPs in wco time $O(Q^* m \log N)$, where Q^* represents the maximum number of solutions for query Q over any temporal graph with the same number of instants of edge validity. We use our index to adapt Leapfrog Triejoin to the temporal graph setting under any variable evaluation ordering. Our index further yields wco guarantees for related query types, including snapshot evaluation, version queries, and other temporal variants. Experiments on real-world datasets show that our approach answers realistic queries in milliseconds with low space overhead.

PVLDB Reference Format:

Diego Arroyuelo, Aidan Hogan, Gonzalo Navarro, and Juan Reutter.
Worst-Case Optimal BGPs on Temporal Graphs. PVLDB, 19(11): 3174 - 3187, 2026.
doi:10.14778/3836663.3836681

PVLDB Artifact Availability:

The source code, data, extended version and/or other artifacts have been made available at <https://github.com/darroyue/bgps-temporal-graphs>. The extended version is available at <https://arxiv.org/abs/2607.20356>.

1 INTRODUCTION

A significant advance for solving Basic Graph Patterns (BGPs) was the discovery of the AGM bound [8] and worst-case-optimal (wco) joins. A graph database, in the simplest RDF-like format, is a labeled graph, where the edges $s \xrightarrow{p} o$ are seen as a set of triples (s, p, o) for subject, predicate, and object. A BGP is a set of triple patterns (s, p, o) , where each component can be a constant or a variable. Solving a BGP returns all assignments of values to variables so that the triple patterns, when instantiated to triples, appear in the graph.

A wco algorithm, such as Leapfrog TrieJoin (LTJ) [49], takes time proportional to the size of the output of the query on some graph.

Herein we consider *temporal graphs* with quads $(s, p, o, [ti, tf])$ denoting that the graph edge $s \xrightarrow{p} o$ was valid from time ti (inclusive) to time tf (exclusive). Temporal graphs give relations a temporal range of validity (e.g., when a researcher was affiliated with a university), or a temporal instant (e.g., when a conference took place), and support querying complex relations (i.e., subgraphs, paths) with temporal ranges and instants (e.g., which researchers were affiliated with a university when a given conference was held). Temporal graph querying arises naturally in real-world systems where relations evolve over time. Examples include event modeling [23], cybersecurity systems [36], and AI applications [11, 12].

We aim at solving BGPs on temporal graphs. The simplest form of such BGP queries is the *return all time points* query, which looks to find all answers for the BGP that were valid at some point in time, together with the time where each of these answers exists in the database. A folklore way to solve this problem is a so-called *join-first* strategy: we first compute the answers of the BGP disregarding the time information, and then, for each solution to the BGP, we use the time information of the concrete edges to determine the points in time this solution is valid, if any. An advantage of join-first is that one is free to solve the BGP with any wco strategy [41–43, 49], beyond-wco ones [1, 3, 30] and even combinations of wco and non-wco algorithms [21, 38, 51, 52]. Alternatively, Hu et al. [28] introduce a so-called *time-first* strategy, which sweeps over the time component and recomputes the BGP answer each time an event changes the set of edges present in the graph. Though they show good theoretical results for restricted kinds of BGPs, in the general case the whole BGP answer is recomputed for every event.

We argue that these solutions suffer from two main limitations. First, on the usability side, the *return all time points* query may lack expressiveness in several applications. For example, the latest SQL 2011 standard recognizes the importance of returning answers that mix temporal constraints (e.g., joining two triples on the condition that one precedes a certain timestamp and the other comes after it) [34]. Second, while join-first strategies can handle such queries, they generate all solutions without considering time and only then filter by time, which is inefficient if the time filters are selective. Regarding LTJ, in particular, both time- and join-first algorithms translate to fixing particular variable binding orderings, where it is well known that choosing a good variable binding ordering is crucial for efficient query resolution, even if all orderings are theoretically optimal [49]. Hence, systems processing temporal graph queries would benefit from an approach that can effectively integrate time constraints within the core of a wco (or beyond-wco!) algorithm.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 19, No. 11 ISSN 2150-8097.
doi:10.14778/3836663.3836681




Figure 1: A toy temporal graph about trips on buses. For simplicity the timestamps show only hours during some day.

Example 1. As a motivating example, consider a graph describing the bus trips over a year in a big city, with quads $(a, \text{trip}, b, [ti, tf])$ indicating that person a was on bus b during time interval $[ti, tf]$. Figure 1 shows a toy example (all the edges are labeled trip). A query looking for pairs of people that were together in a specific bus “B12” could ask to return all time points for the BGP $\{(x, \text{trip}, B12), (y, \text{trip}, B12)\}$. A join-first strategy would find all the pairs of people (x, y) that had ever traveled in that bus, and only then filter who did so at the same time. If m people take that bus at some time, each taking it v times during the year, join-first involves a time complexity of $O(m^2v)$. A time-first approach would, instead, consider each of the t time instants of the year (at a granularity of, say, minutes) and, for each of those, get the pairs of b people on that bus at that time, for a total time complexity of $O(tb^2)$. A better strategy would be to take each of the m persons x that take that bus, for each of them the v times they took it (this should be once or twice a day), and then find the other b persons on that same bus at that time, for a total time complexity of $O(mvb)$, which should be close to the size of the final output and thus optimal. It is further expected that $b \ll m$ and $mv \ll t$ (since mv is the total number of times bus B12 was taken, a fraction of the universe).

As an example of the first limitation mentioned above, the time-first strategy would have trouble finding pairs of people that were together in the bus for over an hour, or people that took the bus B12 before noon and then the bus B6 in the afternoon (those are beyond a simple *return all time points* query). $\square$

We introduce a linear-space temporal graph representation that overcomes both limitations. Our index structure stores temporal graphs in terms of intervals defined by their updates, using an improvement of versioned binary trees that avoids a space blowup by taking advantage of compressed representations. Using these structures, we obtain the following contributions.

(1) Our structure enables retrieving all time points of any given BGP (via LTJ), allowing any possible variable binding order, with (at least) the same complexity guarantees that time-first or join-first algorithms offer. This already improves the current state of the art with respect to worst-case optimal joins over temporal graphs.

(2) We show that we can answer much more general queries. More precisely, we introduce *temporal* BGPs (tBGPs), which are a temporal analog to BGPs, and show how our data structure can process any tBGP query. This language allows time-aware pattern matching, retrieving different time intervals for different parts of the query, and applying filters to restrict these intervals.

(3) We can answer point-in-time or interval queries that return all answers valid in a given timestamp or time period with stronger guarantees: we do it in wco time with respect to the graph restricted to that point or interval rather than the complete graph. This provides a further speedup for point-in-time or interval queries, which are key primitives of the versioning functionalities of the SQL 2011 standard and several graph database proposals [17, 25, 39, 45].

(4) We provide a prototype implementation and show that our improvements translate to practice, solving realistic temporal queries within milliseconds while using reasonable index space. It also clearly outperforms both join-first and time-first strategies, as well as two prominent solutions from previous work [28, 55].

2 RELATED WORK

Temporal query languages have been extensively studied in relational and graph databases, leading to temporal extensions of SQL:2011 [34] and to proposals for RDF and graph query languages such as T-SPARQL [25], stSPARQL [32], and SPARQL-ST [46]. Systems such as GLENDa [45] and RDF archives [17, 24] further support querying historical graph snapshots and tracking data evolution over time. Our work is complementary, focusing on efficient query evaluation rather than models or languages.

Despite advances in temporal languages and data models, the body of work in algorithms supporting these query extensions is much thinner. Berberich et al. focus on temporal extensions to document databases (see e.g. [4, 9]), but these algorithms cannot be directly applied for graph patterns. Work on temporal constructs or temporal algebras, for relations or graphs, can lead directly to algorithms [10, 18, 39], and there is also body of work on extending traditional relational indexes to support temporal constructs such as duration or range (see e.g. [13, 33]), but all of these are based on traditional techniques and not wco joins. Relatedly, Khamis et al. [2] study the complexity of Boolean conjunctive queries with *intersection joins*, which conceptually align with temporal constraints that intersect time domains. Their techniques can be adapted to temporal joins, typically incurring an additional logarithmic indexing factor, while our work targets general temporal graph patterns in linear space. Finally, Khurana and Deshpande [31] and Hou et al. [27] propose fully-functional systems with support for evolving databases by storing over-time differences. However, their query engines are designed to exploit other components of relational systems such as data partitioning or cache optimization (under the assumption that temporal queries in practice often focus on recent data), which cannot be directly applied to wco strategies.

Recent work on algorithms for temporal graph query processing focuses on evaluation under specific data-time constraints, such as retrieving BGPs that occur at the same time within an interval [55], answers that exist for a long period of time [35, 44], for the longest period [47], or answers where edges respect time constraints, in the sense that every match for (s, p, o) must satisfy that s exists before o [22, 29]. Yet, none of these algorithms can be applied directly in our case because they are designed for specific time constraints. In contrast, our work focuses on answering arbitrary temporal graph patterns, which is a more general problem: all of these specific time constraints can be expressed as temporal graph patterns, whose expressive power goes well beyond these constraints.

Our work is closest to that of Hu et al. [28], addressing join-first and time-first strategies for wco joins over temporal graphs. In comparison, our approach handles any variable binding ordering, more general queries, and interval or point queries with better guarantees. Their algorithm sweeps over the time domain, which takes $\Theta(T)$ time over T time instants. This corresponds, in our work, to a time-first strategy. Though we support all variable binding plans (including time-first), since Hu et al. first sweep over time and then process the query in two independent steps, they can use more sophisticated algorithms for the (non-temporal) query part, such as EmptyHeaded [1], which uses the Generalized Hyper-tree Decomposition (GHD) of queries. Moreover, as they show in their paper, their idea of only obtaining the solutions where the disappearing edge participates yields better results on hierarchical queries. Finally, their paper also proposes a cross between time-first and join-first combined with GHD-based algorithms, which again may work better or worse, depending on the instance, than our query plans. We discuss in the Conclusions how our algorithms can be adapted to work under GHD-like time guarantees as well.

3 MODEL

Here we describe temporal graphs, queries, and their AGM bound.

3.1 Temporal graphs and queries

As is usual in the literature, we regard a temporal graph G as a set of labeled edges $s \xrightarrow{p} o$, or triples (s, p, o) , that exist during a given time interval $[ti, tf)$ (see, e.g., [25, 46]).

DEFINITION 1. Let $\mathcal{U}$ be any set of values and $(\mathcal{T}, \leq)$ a totally ordered time domain, disjoint from $\mathcal{U}$. A temporal graph G is a set of $N := |G|$ tuples $(s, p, o, [ti, tf))$, where s, p, o are elements from $\mathcal{U}$, and ti, tf are elements from $\mathcal{T}$, with $ti < tf$, so that $[ti, tf)$ represents an interval in the time domain. We also define, respectively $\mathcal{U}_G := \{s, p, o \mid \exists ti, tf, (s, p, o, [ti, tf)) \in G\}$ and $\mathcal{T}_G := \{[ti, tf) \mid \exists s, p, o, (s, p, o, [ti, tf)) \in G\}$ as the set of elements and time instants mentioned in tuples of G .

A tuple $(s, p, o, [ti, tf))$ represents that the triple (s, p, o) is valid in any time t such that $ti \leq t < tf$. We assume that temporal graphs do not have redundant tuples: if G has tuples $(s, p, o, [ti, tf))$ and $(s, p, o, [ti', tf'))$, then $[ti, tf)$ and $[ti', tf')$ must be disjoint. This enforces a unique way to represent temporal graphs. Furthermore, since $\mathcal{U} \cap \mathcal{T} = \emptyset$, node identifiers and labels cannot be time instants.

We focus on graph patterns, which are the foundation of most query languages used in industry, like SPARQL, CYPHER and GQL [5, 19, 20]. Graph patterns are sets of tuples with constants and variables, which must be found in G . To support temporal operations, we add temporal variables to patterns, and temporal filters.

DEFINITION 2. Let $\mathcal{V}_u, \mathcal{V}_t$ be two disjoint sets of variable symbols, additionally disjoint from $\mathcal{U}$ and $\mathcal{T}$. A temporal basic graph pattern (tBGP) is a set of tuple patterns (x, y, z, w) , where each x, y , and z are in $\mathcal{U}_G \cup \mathcal{V}_u$, and w is in $\mathcal{T}_G \cup \mathcal{V}_t$, plus optional comparison clauses of the form $w_1 \leq w_2$, where $w_1, w_2 \in \mathcal{T}_G \cup \mathcal{V}_t$.

We note, in particular, that tBGPs can only mention time instants that exist in $\mathcal{T}_G$. This is not really a restriction since the edges of G at a time $t \in \mathcal{T}$ are exactly the same as those at time t' , where t'

is the predecessor of t in $\mathcal{T}_G$. Queries can then be translated into valid tBGPs: clauses $w_1 \leq w_2$ stay valid after mapping values of $\mathcal{T}$ to their predecessors in $\mathcal{T}_G$. We also assume that those clauses are satisfiable, and do not form cycles (like $w_1 \leq w_2 \leq w_1$, as those are expressed by collapsing the variables in the cycle into one).

Example 2. In Figure 1, $\mathcal{U}_G = \{\text{Anna, Bob, Carl, Dick, Emma, B6, B12, trip}\}$ and $\mathcal{T}_G = \{06:45, 07:22, 09:51, 10:32, 10:40, 10:46, 11:05, 11:25, 16:11, 17:03, 18:30, 19:45, 18:56, 19:10\}$. We cannot write a pattern $(x, \text{trip}, \text{B12}, 10:50)$ because 10:50 is not in $\mathcal{T}_G$, but the answer is the same as for $(x, \text{trip}, \text{B12}, 10:46)$, where x can be Anna or Carl. Only at 11:05 the answer changes: Anna is no longer a solution. $\square$

The semantics of tBGPs is given by assignments, called *solutions*: values given to the tBGP variables so that all tuples occur in G .

DEFINITION 3. A solution to a tBGP with variables $\mathcal{V}_u$ and $\mathcal{V}_t$ over a temporal graph G is an assignment $f : \mathcal{V}_u \rightarrow \mathcal{U}_G$ and $f : \mathcal{V}_t \rightarrow \mathcal{T}_G$. Furthermore, assuming $f(c) = c$ for $c \in \mathcal{U}_G \cup \mathcal{T}_G$, for each tuple pattern (x, y, z, w) in the tBGP, there must exist a tuple $(f(x), f(y), f(z), [ti, tf)) \in G$ where $ti \leq f(w) < tf$. Finally, for each clause $w_1 \leq w_2$ it must hold $f(w_1) \leq f(w_2)$.

Note that the time component of our solutions must also belong to $\mathcal{T}_G$. This is a form of compacting the (possibly infinite) set of all the solutions in $\mathcal{T}$: a solution with time $t \in \mathcal{T}_G$ stays valid for all the times $t' \in \mathcal{T}$ such that $t \leq t' < t'$, where t' follows t in $\mathcal{T}_G$.

Example 3. A tBGP for the first query of Example 1 is $Q = \{(x, \text{trip}, \text{B12}, t), (y, \text{trip}, \text{B12}, t)\}$, with $\mathcal{V}_u = \{x, y\}$ and $\mathcal{V}_t = \{t\}$. Two solutions are $\langle f(x) = \text{Anna}, f(y) = \text{Bob}, f(t) = 10:32 \rangle$ and $\langle f(x) = \text{Anna}, f(y) = \text{Carl}, f(t) = 10:46 \rangle$. The first is valid until 10:40 and the second until 11:05 (not inclusive). A tBGP for the last query is $Q = \{(x, \text{trip}, \text{B12}, t_1), (x, \text{trip}, \text{B6}, t_2), t_1 \leq 12:00, 12:00 \leq t_2\}$. A solution is $\langle f(x) = \text{Anna}, f(t_1) = 10:12, f(t_2) = 18:30 \rangle$. $\square$

DEFINITION 4. Given a temporal graph G , the problem of solving a tBGP Q is that of computing the set $Q(G)$ of all the solutions to Q .

The selection of tBGPs as our language is justified as the key primitive of temporal relational calculus [16], the algebraic building block of most temporal database languages (see e.g. [25, 46, 48]). Hence, tBGPs are a fundamental querying functionality of temporal database languages per BGPs for (non-temporal) graph query languages. We compare tBGPs to other query languages in Section 3.4.

3.2 Temporal worst-case optimal algorithms

The AGM bound refers to cardinalities of relations. To extend the concept to temporal graphs, we use the following notion.

DEFINITION 5. The point-based representation of G is

$$\hat{G} = \{(s, p, o, t) \mid \exists ti, tf, (s, p, o, [ti, tf)) \in G, t \in [ti, tf) \cap \mathcal{T}_G\}.$$

The times t present in $\hat{G}$ are only those constants $ti, tf \in \mathcal{T}$ that appear as interval extremes in G (possibly on another triple); see Figure 2. As explained, we need not consider other times in the domain where G does not change. As a result, we have the bound $|\hat{G}| \leq |G| \cdot |\mathcal{T}_G| \in O(N^2)$, which is tight when the edges are valid for long time intervals and many different time instants exist in G .

We define the AGM bound of query Q over temporal graph G as the maximum number of solutions of Q over any temporal graph G' whose point-wise representation has as many tuples as $\hat{G}$.




Figure 2: Point-based representation of the graph of Figure 1.

DEFINITION 6. The AGM bound Q^* of a tBGP Q for a temporal graph G is $Q^* := \max\{|Q(G')|, |\hat{G}'| \leq |\hat{G}|\}$, over temporal graphs G' . An algorithm computing $Q(G)$ is worst-case optimal (wco) if it takes time $O(Q^*)$, possibly with data-agnostic and polylog factors of $|G|$.

The following theorem then states our main result.

THEOREM 4. Let G be a temporal graph with N tuples. Then, there exists a data structure using $O(N)$ space that can compute $Q(G)$ for every tBGP Q with m tuples in wco time, $O(Q^* m \log N)$.

The wco algorithm in Theorem 4 can be easily obtained by running a wco algorithm directly on $\hat{G}$, interpreting the tuples of the tBGP Q as quad-patterns (let us disregard the time clauses for simplicity in this discussion). The disadvantage is, of course, that $\hat{G}$ is a bloated representation of G , which can require quadratic space. We could, alternatively, build $\hat{G}$ from G and index it on the fly at query time, but this translates the $O(N^2)$ factor to the query time and working space. This is unacceptable in most practical cases.

Our representation of Theorem 4 actually simulates the use of the LTJ algorithm [49] on $\hat{G}$, but uses G in native form within $O(N)$ space, without converting it to $\hat{G}$. Our algorithm is not only wco, but can also simulate LTJ with any desired variable elimination order (VEO), exactly as if run on $\hat{G}$. Choosing specific VEOs is key for practical performance even if all VEOs are wco in theory [49].

3.3 Point & interval queries: better guarantees

We obtain even stronger bounds for point-in-time and point-in-interval queries, which ask for all answers of a BGP valid at a given time point t or at some point in the interval $[ti, tf]$. We can also query for answers valid along a whole time interval. To state these results, we define some (non-temporal) labeled graphs, derived from slicing a temporal graph across the time domain. We use $|G|$ to denote the number of triples in a non-temporal graph.

DEFINITION 7. For any $t \in \mathcal{T}_G$, G_t is the set of triples (s, p, o) such that there exists a tuple $(s, p, o, [ti, tf])$ in G with $ti \leq t < tf$. For any $t_1 < t_2$ in $\mathcal{T}_G$, $G_{[t_1, t_2]}$ is the set of triples (s, p, o) such that there exists a tuple (s, p, o) in G_t for every $t \in \mathcal{T}_G$ with $t_1 \leq t < t_2$, and $G_{t_1}^{t_2}$ is the set of triples (s, p, o) such that there exists a tuple (s, p, o) in G_t for some $t \in \mathcal{T}_G$ with $t_1 \leq t < t_2$.

Example 5. Let G be the graph of Figure 1. Then $G_{10:32}$ contains the edges (Anna, trip, B12) and (Bob, trip, B12), $G_{[10:32, 10:46]}$ contains only the edge (Anna, trip, B12), and $G_{10:32}^{10:46}$ contains the edges (Anna, trip, B12), (Bob, trip, B12), and (Carl, trip, B12). $\square$

A point-in-time query amounts to solving classic BGP pattern matching on graph G_t , and a point-in-interval query requires matching BGPs on $G_{t_1}^{t_2}$, with t, t_1, t_2 given at query time. Our data structure can evaluate these queries in wco time with respect to the sizes of G_t or $G_{t_1}^{t_2}$, i.e., the best one can hope for: the algorithm is wco on the labeled graph that has exactly the triples we need to consider. This improves the time of Theorem 4 since all such graphs are smaller than $\hat{G}$. Concretely, we show the following in Section 7.

THEOREM 6. Let G be a temporal graph with N tuples. Then, there is a data structure using $O(N)$ space that allows the following:

- Answer BGPs Q with m tuples on the labeled graph G_t for any time instant t given with Q , in $O(Q_t^* m \log N)$ time, where Q_t^* is the maximum number of solutions for Q in some labeled graph with at most $|G_t|$ triples, and
- Answer BGPs Q with m tuples on the labeled graph $G_{t_1}^{t_2}$ for any time interval $[t_1, t_2]$ given with Q , in $O(Q_{t_1}^{t_2} m \log N)$ time, where $Q_{t_1}^{t_2}$ is the maximum number of solutions for Q in some labeled graph with at most $|G_{t_1}^{t_2}|$ triples.

In Section 7 we also show that our data structure leads to instance-optimal algorithms when matching a single triple pattern, and can also be used to list all time intervals where these matches existed.

Next are duration queries, with two flavors. First, we can specifically input a BGP and time points t_1, t_2 at query time, and require all tuples that existed during the whole interval, or, in other words, to query graph $G_{[t_1, t_2]}$. We can also input a duration δ instead of the interval, in which case we want to look for answers that are valid for at least δ time. We show the following in Section 7.2.

THEOREM 7. Let G be a temporal graph with N tuples. Then, there is a data structure using $O(N)$ space that allows the following:

- Solve BGPs Q with m triples on the labeled graph $G_{[t_1, t_2]}$ for any time interval $[t_1, t_2]$ given with Q , in $O(Q_{[t_1, t_2]}^* m \log N)$ time, where $Q_{[t_1, t_2]}^* = \min_{t_1 \leq t < t_2} Q_t^*$.
- Solve BGPs Q with m triples, reporting all the time intervals of length at least δ where Q holds in G for any duration δ given with Q , in time $O(Q_\delta^* m \log N)$, where Q_δ^* is the maximum number of solutions for this query in the subgraph of G formed by all tuples with duration δ or more.

3.4 On the choice of temporal BGPs

Temporal BGPs play a role for temporal graphs analogous to conjunctive queries (CQs) for relational data: they form the pattern-matching core underlying richer temporal query languages [16].

The study of tBPGs is further justified by the fact that algorithms for standard conjunctive query evaluation cannot be directly translated to the temporal setting, because the language of tBPGs can be more expressive than CQs over a standard relational encoding of temporal graphs (for more details see the extended version) The appendix also shows that, if G is expanded into its point-based representation $\hat{G}$, then every tBGP corresponds to a CQ with inequalities evaluated over $\hat{G}$. This further justifies tBPGs as the CQ analog in our setting. Since the size of $\hat{G}$, for a graph G with N tuples, can be of order $N \cdot |\mathcal{T}_G|$, materializing $\hat{G}$ to process tBPGs as relational CQs is infeasible. Notice that our index structures can correctly evaluate tBPGs over $\hat{G}$ using only $O(N)$ space.




Figure 3: The LTJ trie of order pos for the graph of Figure 1 devoid of temporal annotations.

4 OVERVIEW OF OUR SOLUTION

Our solution is to build a data structure for temporal graphs in which we can run the Leapfrog TrieJoin algorithm with any variable ordering to compute joins in wco time, and where the total space is linear with respect to the number of tuples of the temporal graph.

For a temporal graph G , we replace the ti and tf components in G by their rank in $\mathcal{T}_G$, turning them into integers in $[0, T)$ where $T = |\mathcal{T}_G|$ (we also maintain a dictionary to enable switching between actual time constants and integers in $[0, T)$). We further map the set of node identifiers and labels to the integer interval $[0, U)$, where $U = |\mathcal{U}_G|$. We handle those numeric identifiers using binary tries.

We now define the Leapfrog TrieJoin algorithm and (versioned) binary tries: the two building blocks of our solution. In Section 5, we construct a data structure from these two building blocks that achieves the desired wco optimality, but with $O(N \log N)$ space. Section 6 explains how to further reduce this to a structure using linear space, and how to use this structure for querying.

4.1 Leapfrog Triejoin

We describe the Leapfrog Triejoin (LTJ) algorithm [49], in the version that is adapted for solving BGPs on labeled graphs [26].

LTJ requires that the triples (s, p, o) are represented as tries, which we will call *LTJ tries*, in the 6 possible orders of the components s , p , and o . Tries are labeled trees where no two children of a node have the same label, and represent all the strings that can be read by concatenating the labels of their root-to-leaf paths.

Each LTJ trie stores the components s, p, o in some order, as strings of length 3. We call these orders s, p, o , s, o, p , p, s, o , p, o, s , o, s, p , and o, p, s . For example, the trie for the order pos has one root-to-leaf path with consecutive labels p, o, s for each triple (s, p, o) in the graph (see Figure 3). Each LTJ trie has height 3 and exactly N leaves.

Let $Q = \{t_1, \dots, t_m\}$ be a BGP (i.e., a set of triple patterns) and $\{x_1, \dots, x_v\}$ its set of variables. LTJ carries out v iterations, “eliminating” one variable at a time. The order LTJ chooses to eliminate the variables is known as the *variable elimination order* (VEO).

Each triple pattern t_i is associated with one of six tries, say τ_i . A level ℓ of τ_i corresponds to a constant c (or variable x) if t_i contains c (or x) at the position corresponding to level ℓ in τ_i . To be a valid trie for t_i , the first levels of τ_i (i.e., levels closest to the root) must correspond to constants of t_i , and the subsequent levels must correspond to the variables of t_i , consistently with the VEO.

Example 8. For illustration, consider the graph of Figure 1 devoid of temporal annotations (and thus with only one edge from Carl to B12). If $t_i = (x, \text{trip}, z)$ and the VEO eliminates z and then x , then

τ_i must be the trie of the order pos shown in Figure 3 (this is why LTJ needs the tries in various orders). Its first level corresponds to $p = \text{trip}$, the second to $o = z$, and the third to $s = x$. $\square$

The first step of LTJ is to descend, from the root of each trie τ_i , by the constants of t_i . It then starts the variable elimination step. Say that the chosen VEO is $x_1, \dots, x_v$. LTJ then finds each value c that appears as a child of the current node in every trie τ_i whose next level corresponds to variable x_1 . For each such c , LTJ *binds* $x_1 := c$, descends by c in all the corresponding tries τ_i , and goes on recursively with x_2 . Once we have bound all variables in this way, the set of bindings for $x_1, \dots, x_v$ is a new solution for Q .

The values c are found by intersecting the children of the current nodes v_i in all the suitable tries τ_i ; see the precise algorithm in the extended version. LTJ cycles over all the involved tries looking for the smallest $x' \geq x$, where x is its next candidate, within the children of the current node v_i in each trie τ_i . When all tries find the same x , this is the next output of the intersection. The key primitive used, $\text{leap}(v_i, x)$, finds that smallest $x' \geq x$ within the children of v_i . If leap is executed in $O(\log N)$ time (e.g., using binary search) then LTJ obtains the wco time complexity $O(Q^* m \log N)$.

4.2 Binary Tries

Assume that we aim to represent each G_t as a labeled graph, for the consecutive values $t \in [0, T)$. We now describe a technique inspired by versioned data structures, beginning with binary tries.

A *binary trie* (BT) is a binary tree that represents a set of binary strings, each string being a root-to-leaf path in the trie. If the sequence of left-right turns from root to the leaf are interpreted as 0 and 1, respectively, we obtain the string represented by the leaf. We assume no binary string is a prefix of another.

We use BTs to encode the sets of children of nodes of the LTJ tries. Those children are numbers in $[0, U)$ interpreted as binary strings of length $\ell = \lceil \log_2 U \rceil$, reading from most to least significant digit. Figure 4 (left) illustrates the BT for a toy LTJ trie node.

BTs can simulate the search for the first child of v with value $\geq x$ in $O(\ell) = O(\log N)$ time, as follows. We start at the BT root, with height $h := \ell$. If the h -th highest bit of x (an ℓ -bit number) is 1, we continue recursively by the right child of the BT, decrementing h . Otherwise, we proceed as follows. First, we try to find the answer on the left branch, also decrementing h . If that search returns an answer, we return it too. If not, we return the leftmost leaf of the right child of the node. When we arrive at a leaf, its string represents the first value $\geq x$ in the trie. This takes $O(\ell)$ time because it can descend by both children of a node only once in the whole process. For pseudocode see the extended version.

BTs then support the LTJ leap operation within the same $O(\log N)$ time factor penalty of binary searches. As a consequence, we can use a BT as a local structure to represent the set of children of every LTJ trie node (instead of just a plain array of increasing identifiers), and retain the same wco time $O(Q^* \log N)$ on labeled graphs.

4.3 Versioned Binary Tries

We aim to represent all the graphs G_t with the same tries. When representing G_t , we will have LTJ trie nodes v that will be very similar to the node v for G_{t-1} : the node v in G_t may have a few inserted or deleted children with respect to the node v in G_{t-1} . To




Figure 4: At the top, three consecutive versions of an LTJ trie node v . The leftmost one has three children: values 2, 3, and 6. It is represented (below) with a BT using $\ell = 3$ bits, which stores the binary strings $2 = 010_2$, $3 = 011_2$, and $6 = 110_2$. The root of the BT is v and its leaves are the corresponding children of v in the LTJ trie. The center and right trie are successive versions of the leftmost one: the value 4 appears in the second version, whereas 3 disappears in the third. Both are represented (below) as VBTs: the pointers to preceding versions are shown with dashed lines; note in particular the pointer from the third to the first version.

efficiently represent the children of those LTJ trie nodes v , we use *versioned binary tries* (VBTs). A VBT represents a BT as a set of updates with respect to a *reference* BT. Instead of explicitly storing all the nodes of the BT it represents, the VBT records only the root-to-leaf paths that change with respect to the reference BT. Changes refer to newly inserted or deleted binary strings. The BT subtrees that do not change, instead of being duplicated, are pointed to from the corresponding nodes in the newly created paths to the reference BT. Figure 4 illustrates successive versions of the leftmost subtree.

Note that a top-down traversal on a VBT is identical to that on a standard BT, so we can run the LTJ intersections on the VBTs as well, in $O(\log N)$ time per leap. Note how, in Figure 4, we can traverse the VBT of the third version exactly as if it were a BT.

We use VBTs to represent a sequence of versions of a node’s children. The first version in time is represented as a standard BT. Each new version is a VBT encoded with respect to the preceding BT. From the third VBT onwards, the preceding BT is also represented as a VBT, so pointing to a subtree of the preceding BT may actually correspond to pointing to an earlier BT pointed to by the preceding VBT. This is shown on the third version in Figure 4.

Since the paths are of length $O(\log N)$, it follows that VBTs require $O(\log N)$ space per update they record with respect to a previous BT. VBTs will be part of our solution, as described next.

5 OUR REPRESENTATION

In abstract terms, we regard each tuple $(s, p, o, [ti, tf])$ of G as if it were $tf - ti$ standard quads (s, p, o, t) , one per integer $t \in [ti, tf]$; recall that we have mapped $\mathcal{T}$ to $[0, T]$. This is $\hat{G}$, which as explained may have $\Theta(N^2)$ quads, but we will manage to represent them all within $O(N \log N)$ space (later, we will reduce the space to $O(N)$).

With the quads model, the tBGPs can be solved directly by using the LTJ algorithm on the LTJ tries storing the $4!$ permutations of $\{s, p, o, t\}$. The algorithm is then wco with respect to this model. In

particular, consider a tBGP where the time component is a single constant t for all the triple patterns. Then, using an LTJ trie with an order that starts with t , we descend by the child t of the root and can solve the tBGP in time $O(Q_t^* \log N)$, where Q_t^* is the AGM bound for the corresponding BGP on the graph G_t . As another example, consider the same tBGP where now t is a variable. By solving it on an LTJ trie that starts with the t component, we can mimic the so-called “time-first” approach in previous work [28]; by using a trie that ends with the t component, we mimic the so-called “join-first” strategy. We can use more general strategies, however, by putting the time component elsewhere in the order.

In this scenario, the clauses $w_1 \leq w_2$ are handled as follows. If w_1 is bound before w_2 , then it will have assigned a value $f(w_1) = t_1$ each time w_2 is bound. At this moment we restrict the LTJ intersections for w_2 so that the values stay within $[t_1, T)$. If, instead, w_2 is bound to $f(w_2) = t_2$ before w_1 , we enforce that the values of w_1 be within $[0, t_2]$. This is still wco, as it is equivalent to materializing a table with all the pairs (w_1, w_2) with $w_1 \leq w_2$ and including it in the LTJ algorithm. Since all the time instants appear in both columns, LTJ will not restrict the first bound variable; the second one will be restricted exactly as described.

To do all this within $O(N \log N)$ space, we introduce a special representation for the LTJ tries. Their abstract form will still be a trie of height 4, with one level per tuple component. The difference is in the way we represent the children of the nodes. Given a particular permutation, say (s, t, p, o) , we distinguish the levels before the time component (in our example, the first level, s), the time level, and the levels after the time component (levels 3 and 4, p and o , in our example). Note that which levels are before and after the time level depends on the permutation we are storing.

The levels before the time level can be implemented in traditional form, e.g., as an array of increasing child values. We now describe our implementation of the time level and its subsequent levels.

5.1 The time level

The time level will be represented in a form closer to the tuples of G (with the times already converted to integers in $[0, T)$) than to the set of quads of $\hat{G}$. This level will be an ordered sequence of disjoint time intervals, stored in classical form (e.g., an array). Let v be an LTJ trie node whose children belong to the time level, and $[ti_1, tf_1), [ti_2, tf_2), \dots$ be the time intervals of all the tuples stored in the LTJ subtree of v (this is a subset of all the time intervals in G and they can overlap). Then consider the list $\langle t_1, t_2, \dots \rangle$ containing all the values ti_k and tf_k , sorted in increasing order and with equal values removed. The intervals forming the children of v are then in principle $[t_1, t_2), [t_2, t_3), \dots$, which covers the whole universe $[0, T)$. We remove, however, intervals with no descending tuples. The surviving intervals are then of the form $[ts_1, te_1), [ts_2, te_2), \dots$, where each $[ts_i, te_i)$ is equal to some $[t_k, t_{k+1})$.

Example 9. Consider the oTPs trie for Figure 1. The intervals $[ti_i, tf_i)$ for the node B12 are $\{[9:51, 10:40), [10:32, 11:05), [10:46, 11:25), [16:11, 17:03)\}$. The corresponding sorted list of times t_i is then $\langle 9:51, 10:32, 10:40, 10:46, 11:05, 11:25, 16:11, 17:03 \rangle$. A consecutive pair forms an interval $[ts_i, te_i)$, except for $[11:25, 16:11)$, in which there are no tuples. Figure 5 illustrates the different levels in this trie. $\square$




Figure 5: The OTPS trie corresponding to the graph of Figure 1.

5.2 Levels after the time component

We will use BTs and VBTs to represent the children of all LTJ trie nodes below the time level. Consider an LTJ trie node v whose children, at the time level, are $[ts_1, te_1], [ts_2, te_2], \dots$. We will *concatenate* all the binary strings corresponding to all the levels that follow the time level. In the OTPS trie of Figure 5, the descendants of each node $[ts_i, te_i]$ will be arranged in a single BT (or VBT) holding binary strings of length 2ℓ , formed by concatenating the p and the s values (i.e., $p : s := 2^\ell p + s$) of all the tuples that must be stored below v and exist at time ts_i . Per our construction, this set of tuples does not vary within the time interval $[ts_i, te_i]$.

Figure 6 shows hypothetical values $p : s$ descending from some node v . Note that concatenating the components is almost immaterial with respect to using different BTs for the children of every LTJ trie node. The BT node representing each prefix p (at depth ℓ) becomes the root of the subtree representing each component s of the concatenations $p : s$. We can then interpret the node of p as the LTJ trie node obtained by descending by p from v , and the subtree with the s components as the BT of its children. Our arrangement, however, is more convenient for the versioning we require next.

Let us consider how we store the elements formed by concatenating the remaining attributes of the tuples below v . We first build the basic BT for the node $[ts_1, te_1]$, containing all the elements whose tuples exist in time ts_1 . This set of tuples does not vary until time te_1 . This BT is the child of the node $[ts_1, te_1]$ of the time level.

Now consider the next interval, $[ts_2, te_2]$. If $ts_2 > te_1 + 1$, then the intervals spanning $[te_1, ts_2]$ were empty, and we simply build a new BT for $[ts_2, te_2]$ containing only the new elements that appear at time ts_2 . Otherwise, there exists a nonempty BT that precedes the BT for $[ts_2, te_2]$, so we represent this BT as a VBT, whose reference is the BT for $[ts_1, te_1]$. This VBT must represent two events:

- (1) Consider a tuple $(s, p, o, [ti, tf])$ of G where (s, p, o) must be inserted below the LTJ trie node v and such that $ti = ts_2$. This tuple must be added to the VBT.
- (2) Similarly, each tuple $(s, p, o, [ti, tf])$ belonging to the subtree of v and such that $tf = ts_2$ must be removed from the VBT.

Example 10. In Figure 5, below node B6, the interval $[07:22, 18:30]$ disappears as it contains no tuples, and thus we start a new BT with root $[18:30, 18:45]$ containing Anna. The next interval, $[18:45, 18:56]$, modifies the previous one by adding Dick, represented as a VBT with respect to the previous one. The last interval, $[18:56, 19:10]$, removes Dick again, and is again represented by a VBT. $\square$

We create new paths in the VBT of $[ts_2, te_2]$ as described in Section 4.2, to account for those events. Each event requires creating $O(\log N)$ nodes in the BTs or the VBTs. Note that each graph tuple $(s, p, o, [ti, tf])$ generates two events: the insertion of (s, p, o) at ti and its deletion at tf . Therefore, the trie represents at most $2N$ events, each of which induces $O(\log N)$ new nodes in the VBTs. The total space is thus $O(N \log N)$.

Example 11. In the trie OTPS of Example 9, the tuple (Anna, trip, B12, $[10:32, 11:05]$) will generate, below the node B12, a time $ti = 10:32$ where we will insert trip:Anna, and a time $tf = 11:05$ where we will delete trip:Anna. Because we use VBTs for intermediate intervals (in this case, for the intervals $[10:40, 10:46]$ and $[10:46, 11:05]$), trip:Anna will not be explicitly represented in those. $\square$

5.3 The intersections

Intersections work exactly as in LTJ over BTs, except at the time level. At this level, each node represents an interval, which must be handled during intersection: each interval $[ts_i, te_i]$ stands for all the time instants $ts_i \leq t < te_i$. We modify the list intersection algorithm of LTJ so as to assume that all those time instants t are explicitly represented and that copies of the BT of $[ts_i, te_i]$ descend from all those implicit time instants t . To find the first time $t^* \geq t_0$, we look for t_0 and, (1) if we find some $ts_i \leq t_0 < te_i$, we answer $t^* := t_0$; (2) if we find some $te_i \leq t_0 < ts_{i+1}$, we answer $t^* := ts_{i+1}$.

We can avoid recomputing the same answers for all the time instants t of an interval: Rather than returning each time instant, we return the range $[t_0, te_i]$ in case (1) above, and $[ts_{i+1}, te_{i+1}]$ in case (2). We then modify the intersection algorithm to record the maximum interval starting at t^* that is included in all the intersected lists. The instantiation of the corresponding time variable is then a whole interval where the instantiated BGP will not change. See pseudocode in the extended version.

Thus we can return solutions to BGPs with ranges of values for the time-bound variables. This is a form of compacting the output, but we cannot ensure we return it in the optimally compacted form.

6 LINEAR SPACE

The base version we have described implements leap in $O(\log N)$ time and $O(N \log N)$ space. This space is worrisome for large graphs. We now describe a more sophisticated storage mechanism that achieves linear space with no penalty in time complexity.




Figure 6: On the left, the two final levels of an LTJ trie node rooted at v . In the middle, the BT representation of the corresponding children. Note that we can also regard all the BTs as a single BT on the concatenation of the path labels, $2 : 5$, $3 : 4$, $3 : 5$, $6 : 4$, $6 : 6$, and $6 : 7$. On the right, a VBT representing two edits: the removal of $2 : 5$ and the insertion of $6 : 1$.

6.1 Data Structure

Consider the sequence of updates that occur in the time-level children $[ts_1, te_1), [ts_2, te_2), \dots$ of a particular LTJ trie node. Instead of creating a sequence of VBTs $v_1, v_2, \dots$, we create a *single* BT V where we insert all the paths corresponding to those updates; each VBT (or BT) v_i corresponds to a sequence of insertions/deletions of tuples (just insertions in case of BTs). If there are b levels after the time level, for $1 \leq b \leq 3$, each such path is of length bl . Updates consisting of insertions correspond to a path of length bl . Updates consisting of deletions are also paths of length bl , plus a deletion mark. Note this differs from the way we represented deletions in VBTs; see Figure 6. If a given tuple is inserted and deleted several times, its leaf will correspond to several updates.

We will record the timestamps of those updates. If V represents L updates, they are numbered 1 to L . Each node $[ts_i, te_i)$ stores the value $p_i \in [1, L]$ for its last update. This means that its BT corresponds to executing the updates $[1, p_i]$ on an empty trie. We store the timestamps, however, in a way that will enable fast navigation.

Each node v of V will conceptually store the subsequence $v.T$ of $[1, L]$ corresponding to the updates that occur below v . The subsequence is $v.T = \langle 1, 2, \dots, L \rangle$ if v is the root of V . If v is a leaf, then $v.T$ contains the timestamps where its particular tuple was created or removed. Figure 7 shows an example (see only the top and bottom-left parts for now).

Instead of storing $v.T[1, L_v]$ explicitly, we will store only a bitvector $v.B[1, L_v]$, where $v.B[i] = 0$ iff $v.T[i]$ appears below the left child of v , and $v.B[i] = 1$ if it appears below the right child. To efficiently support operation leap, we also store a bitvector $v.E[1, L_v]$, where $v.E[i] = 1$ iff at timestamp $v.T[i]$ there exists some tuple below v whose last update was an insertion (i.e., the tuple exists at timestamp i). The leaves of V store only their bitvector $v.E$, not $v.B$.

Example 12. At the bottom of Figure 7, the root's left child has $E[4] = 0$ since the node does not exist at timestamp $T[4] = 6$. Its right child has $B = 0101$ since timestamps $T[1] = 1$ and $T[3] = 4$ go to its left child, whereas $T[2] = 2$ and $T[4] = 6$ go to the right. $\square$

Note that every update induces $2bl + 1$ bits in V : one in the bitvectors $v.B$ of the proper ancestors v of its leaf and one in the bitvectors $v.E$ of those nodes and the leaf itself. We can actually avoid storing $v.E$ at leaves v because it is always an alternating sequence of 1s and 0s. Therefore, the total number of bits in $v.B$ and $v.E$ bitvectors of V is $L \cdot 2bl = O(L \log N)$. The machine word

must hold $\Omega(\log N)$ bits if it can address $\Theta(N)$ tuples in constant time; therefore this number of bits amounts to $O(L)$ words of space. Since the sum of the lengths L of all the nodes below the time level amounts to the $2N$ updates in the graph, those $O(L)$ words amount to $O(N)$ over the whole LTJ trie.

We still have $O(L \log N)$ space to store the topology of V , however, as it stores L paths of length $bl = O(\log N)$. We avoid the need to store those pointers by concatenating all the bitvectors $v.B$ (and, similarly, $v.E$) levelwise, left to right. Note that the total length of the bitvectors $v.B$ or $v.E$ in a given level is always L , so we have bl bitvectors $B_d[1, L]$, for $d \in [0, bl - 1]$ that can be stored as a large concatenated bitvector of length Lbl (similarly, E_d for $d \in [0, bl]$, where the level $d = bl$ can be made implicit, as explained). Those bitvectors are shown at the bottom-right of Figure 7.

6.2 Navigation

To navigate V without pointers, we must know the ranges $B_d[s_v, e_v]$ and $E_d[s_v, e_v]$ where the bitvectors $B_v[1, L_v]$ and $E_v[1, L_v]$ are stored, for any node v of depth d . For the root, these are just $B_0[1, L]$ and $E_0[1, L]$. For the rest, we make use of the function $\text{rank}(B, i, j)$, which counts the number of 1s in $B[i, j]$. This operation can be computed in constant time by storing just $o(|B|)$ additional bits [14, 40]. Assume $B_d[s_v, e_v]$ is the area of B_d corresponding to $v.B$. Then, letting $r := \text{rank}(B_d, s_v, e_v)$, the area for the left child of $v.B$ is $B_{d+1}[s_v, e_v - r]$, and the area for its right child is $B_{d+1}[e_v - r + 1, e_v]$. We know we have reached a null pointer v in V when $s_v > e_v$.

Example 13. Let us start at the root in Figure 7, with $B_0[1, 6] = 001010$. To go to its left child, we compute $\text{rank}(B_0, 1, 6) = 2$; thus its left child corresponds to $B_1[1, 4] = 1111$. If we want to go left again, we compute $\text{rank}(B_1, 1, 4) = 4$, and find that its left child is null because it corresponds to $B_2[1, 0]$. $\square$

Until now, our structure is a variant of the wavelet matrix [15], but the bitvectors E_d extend it so that we can traverse V as if it were the VBT at some timestamp $p_l \in [1, L]$. We start at the root of V (i.e., $d := 0$ and $[s_v, e_v] := [1, L]$) with the local offset $p := p_l - s_v$ for p_l . If $E_d[s_v + p] = 0$, then the VBT node was null at this time instant. Otherwise, we can go left or right. The timestamp p becomes $p - \text{rank}(B_d, s_v, s_v + p)$ on the left child and $\text{rank}(B_d, s_v, s_v + p) - 1$ on the right child. Note that p cannot become negative, since that would mean that there have been no updates




Figure 7: On the top, the BTs of six consecutive timestamps, corresponding to inserting v_1 with value $010_2 = 2$, then v_2 with value $011_2 = 3$, and v_3 with value $110_2 = 6$. The leaf v_1 is then removed, v_4 is inserted with value $100_2 = 4$, and finally v_2 is removed. On the bottom left, our conceptual trie V representing all the events, with the $v.T$ sequences of timestamps besides each node v ; the deletions are marked with a * at the leaves. On the right, the actual compact representation of V , first still with the tree topology and then just like the bitvectors B_d and E_d . Bitvector E_3 can be made virtual, as explained.

on the subtree of V since the beginning, and thus the node would be null at timestamp p (which we detect before entering the node).

Example 14. Assume we want to go to the left child of the root at timestamp $p_l = 6$ in Figure 7. Though V has such a node, which is $B_1[1, 4]$, we see that it did not exist at timestamp 6 since $E_1[1+3] = 0$. This position 3 is where we map the relative value $p = p_l - 1 = 5$ at the root to relative value $p - \text{rank}(B_0, 1, 6) = 3$ at its left child. $\square$

With those operations, we can traverse any VBT top-down in constant time per movement to children, while the total space used by V is $O(L)$ and the total space is $O(N)$ per LTJ trie.

To support leap on this representation, recall that leap is applied on an LTJ trie node $[ts_l, te_l]$ of the time level, and aims to traverse its VBT (or BT) from its root v . Previously we describe this on a VBT of height ℓ , as we always perform leap on a single attribute, but now we have concatenated b attributes in VBTs of height $b\ell$. However, since we can start at any node whose depth is a multiple of ℓ , the procedure remains unchanged over blocks of ℓ elements. When run on V , we are only interested in the timestamps $[1, p_l]$, which represent the VBT of $[ts_l, te_l]$. We use $v.E$ as before to determine if a node is null at timestamp p_l . See pseudocode in the extended version. This completes the proof of Theorem 4.

7 QUERYING FOR POINTS AND INTERVALS

Definition 7 introduces several labeled graphs derived from a temporal graph G , in which it is of interest to answer standard BGPs. We now show how our linear-space data structure on G can be used to answer BGPs over these graphs. We start with point-in-time and point-in-interval queries, corresponding to querying graphs G_t and $G_{t_1}^{t_2}$, where t , t_1 , and t_2 are given at query time (Theorem 6). We finish with $G_{[t_1, t_2]}$ and queries with duration. (Theorem 7).

7.1 Point-in queries

To answer a query Q over G_t , we transform it into a temporal BGP by adding the constant t as the fourth component in all the triples. Interestingly, if we force the variable order to descend first by the attribute $\tau = t$ we can show that the resulting algorithm takes exactly the same amount of steps that the normal LTJ algorithm on Q would take, immediately giving us worst-case optimality over G_t . This proves the first part of Thm. 6.

For the case of $G_{t_1}^{t_2}$, we must extend the navigation of G_t using VBTs, described in Section 6.2, to intervals $[t_1, t_2]$. The key idea is that if t_1 and t_2 are represented by p_1 and $p_2 + 1$ at some node $[s_v, e_v]$, then $\text{rank}(E_d, s_v + p_1, s_v + p_2) = 0$ iff the node s_v did not exist along the whole period between the local offsets p_1 and p_2 , that is, it did not exist in $G_{t_1}^{t_2}$. Hence, we can quickly find the time intervals that overlap or are contained in $[t_1, t_2]$, and continue the LTJ algorithm while keeping track of these intervals. This yields the second part of Thm. 6; details in the extended version.

We note that the so-called “join-first” strategy, which solves Q on the labeled graph $G_{\text{all}} = G_1^T$ of all the tuples that ever existed and then filters the results by time, can offer the AGM bound only on $|G_{\text{all}}|$, whereas our result bounded by $|G_{t_1}^{t_2}|$ is the best one can hope for: the algorithm is wco on the labeled graph that has exactly the triples we want to consider.

Given a single triple pattern (s, p, o) , our algorithm lists all matches that existed during the interval $[t_1, t_2]$ in instance-optimal time: $O(\log N)$ per reported triple. We can also list all the maximal intervals where each such match existed during $[t_1, t_2]$, each in $O(\log N)$ time. This also permits tracking the differences between G_{t_1-1} and G_{t_2-1} : an important operation on versioned graphs. Neither “join-first” nor “time-first” strategies can handle these queries near-optimally. Full proof in the extended version.

THEOREM 15. *Let G be a temporal graph with N tuples. Then, there is a data structure using $O(N)$ space that can report all the*

occ occurrences of a single-triple-pattern query $Q = \{(s, p, o)\}$ in the labeled graph $G_{t_1}^{t_2}$, for any time interval $[t_1, t_2]$ given with Q , in time $O((1 + \text{occ}) \log N)$. It can also list each maximal time interval where each occurrence appears in time $O(\log N)$.

7.2 Queries with duration

In some cases we are interested in requiring that the solutions to the queries last for some time. We consider two cases of such queries. The first establishes a time interval and looks for solutions that always hold during this interval. According to our definitions, this amounts to running the query over $G_{[t_1, t_2]}$; recall Definition 7. A second case does not fix the exact times, but sets a minimum duration δ along which the reported solutions must hold [28].

In order to query $G_{[t_1, t_2]}$ in optimal time, we can reuse some of the ideas used to query $G_{t_1}^{t_2}$. Indeed, we have that $\text{rank}(E_d, s_v + p_1, s_v + p_2) = p_2 - p_1 + 1$ iff the node s_v exists throughout the whole period between the local offsets p_1 and p_2 . Therefore, we can use our leap algorithm with a single change to require that all the bits in the area of E_d are 1s, not just one of them. This may not be optimal in the size of $G_{[t_1, t_2]}$, because we may be stuck exploring a BT node that existed all along $[t_1, t_2]$, while none of its children did. However, the algorithm is still wco with respect to the size of G_t for any $t \in [t_1, t_2]$: if the subtree of a node v at depth d does not exist in G_t , then the corresponding bit E_d will be zero and the algorithm will not attempt to explore it. Instead of satisfying the AGM bound on $|G_{[t_1, t_2]}| = |\cap_{t \in [t_1, t_2]} G_t|$, we satisfy it in terms of $\min_{t \in [t_1, t_2]} |G_t|$. This proves the first part of Thm. 7.

For the case where we look for solutions of a query Q that hold over a minimum duration δ , we first solve the query on G , adding the same variable t as the fourth component of all the triple patterns, and use the LTJ tries where the time component τ is at the end, as in join-first approaches. For the time level we use a more sophisticated data structure, based on geometric grids, that allows us to quickly discard time intervals smaller than δ while at the same time storing the children of each LTJ trie node v in an efficient way; see the extended version for more details. This is the second part of Thm. 7.

8 EXPERIMENTS

In this section we evaluate its empirical performance over real temporal graphs, in two scenarios. First, we evaluate our index on a temporal graph and real queries from Wikidata, aiming to assess its practicality and the advantage of being able to bind the time at any point, not just at the beginning (time-first) or at the end (join-first). Second, we compare our index with two other systems that offer comparable functionality [28, 55] over real temporal graphs with synthetic queries.

We ran our experiments on a 16-core Intel Xeon Silver 4110 with Debian 5.10.127-2, clocked at 2.1 GHz, with 768 GB of RAM and 10 MB cache. Our C++ code was compiled with `-Ofast`. Our benchmark, code and datasets are available on Github. For details on the implementation, we refer readers to the extended version.

8.1 Wikidata with real query logs

Benchmark. We perform experiments over the Wikidata knowledge graph [50]. We choose Wikidata as it provides a (1) large-scale, (2) real-world, (3) diverse knowledge graph featuring (4) temporal

Table 1: The datasets we use, including space in bytes per tuple and construction time in microseconds per edge.

Dataset	N	$ \mathcal{U}_G $	$ \mathcal{T}_G $	Space	Time
Wiki	844,172,299	116,145,285	155,956	241	38.0
WikiT	15,943,638	9,384,937	155,956	230	34.7
Divvy	21,243,344	7,158	20,673,270	190	25.0
Yellow	30,003,832	8,552	20,797,780	194	29.1
Caida	15,792,089	110,704	336	141	20.8

annotations, further providing a comprehensive log of (5) real-world queries [37]. As per many real-world temporal graphs, some relations include temporal annotations, while others are asserted without temporal qualification (and considered universally valid).

We extract the temporal graph from Wikidata considering all 844 million triples (see Table 1) consisting of items and properties. We extract temporal qualifiers on Wikidata statements, where available, for start time (ti), end time (tf), and point in time (ti and tf). In the resulting temporal graph, about 16 million triples have non-trivial temporal annotations. The graph features 12,894 distinct predicates.

We extract 1000 BGPs from the real-world query logs published for Wikidata [37]; more details are given in the extended version. We filter queries whose constants do not occur in the graph, resulting in 974 queries. We run queries with a result limit of 1,000 and a timeout of 600 seconds.

Indexing. Table 1 shows that our index (Wiki) uses 241 bytes per edge, or 60 32-bit words. This is 12 times the space needed to store the edges in plain form (as five 32-bit integers). For comparison, a standard index on non-temporal graphs uses 157 bytes per edge [7]. From the total space, 84% is used by the 6 tries starting with the time component, with TSPO , TPOS , and TOSP being the largest tries.

Indexing took 9 hours, growing linearly with the data size (verified by indexing increasing subsets). About 9% of the time was for building the 12 regular tries; the rest is used to build the 6 tries starting with the time component: building the linear-space data structure we designed in Section 6 requires sorting the data $\lceil \log_2 |\mathcal{U}_G| \rceil = 28$ times per level, not just once per the standard tries (those use faster radix-sorts, however). Our construction runs essentially in-place, using no more space than the final index.

Variable elimination. A standard LTJ strategy [26] is to choose as the next variable the one whose trie nodes have the least amount of children among those (if possible) that are connected to a previously bound variable. Lonely variables (i.e., those appearing once in the BGP) are bound at the end. For temporal graphs we modify this strategy in three ways: join-first (**JF**) leaves the temporal variables to the end, time-first (**TF**) binds the temporal variables first, and least-children (**LC**) treats the time as any other variable (the ability to do this distinguishes our index from previous work). For nodes at the level τ , we divide the total temporal span divided by their number of children, to estimate the number of children of its nodes.

Results. Figure 8 shows boxplots of query times, along with averages and number of queries for which one approach is faster than another. Our **LC** strategy is on average twice as fast as **JF** and orders




Figure 8: Left: Boxplots of runtimes in msec. Right: number of queries for which the approach in the row runs faster than that in the column, with average times at the bottom.




Figure 9: Left: boxplots of runtimes in msec for the query with clause $t_1 \leq t_2$. Right: comparison with Hu et al. [28] for point-in-time queries on WikiT.

of magnitude faster than **TF**. The boxplots of **LC** and **JF** seem comparable, yet **JF** is 20% slower in the median. The boxplots and the averages show that **TF** has many bad cases, making it an unstable strategy. By not always binding the temporal variables at the start (per **TF**) nor at the end (per **JF**), the **LC** strategy outperforms **TF** in 64% of the queries and **JF** in 69%. The average of the query-wise minima of the three times is 0.94 msec, just 11% faster than **LC**.

Discussion. The results show that our index is practical on real-life scenarios: using less than twice the space needed for a non-temporal graph, we answer realistic queries within a millisecond. We also conclude that it is often highly beneficial to allow temporal variables to be eliminated at any point: **LC** is the fastest strategy in the majority of the queries, clearly outperforming **JF** and **TF**. This highlights the relevance of a unique capability of our index, and leaves room for designing better heuristics along the lines of **LC**.

A more complex query. Our index solves more than just point-in-time queries. We demonstrate its performance on more complex queries by using now two time variables, t_1 and t_2 , and randomly assigning t_1 or t_2 (but at least once each) to the BGP triples in order to build the tBGP, and complete the query with a clause $t_1 \leq t_2$. We convert the 825 BGPs having more than one triple. Figure 9 (left) shows the distribution of times using **JF** and **LC**; note that the **TF** strategy is very inefficient when there are two (or more) time variables, as it would lead to $\Omega(T^2)$ query time.

In this case, **LC** and **JF** perform similarly, with a distribution close to that of the simpler point-in-time query. For example, the median of **LC** is 65 microseconds, just 30% higher than the median of the point-in-time queries (50 microseconds). This shows that our index can efficiently solve more complex queries as well.

Comparison with previous work. Figure 9 (right) compares the time to solve point-in-time queries using our index and that of Hu

et al. [28]. That index implements queries with duration δ , which for $\delta = 1$ are point-in-time queries. We chose the strategies that performed best: generic join for cyclic queries and their acyclic baseline for the rest. As they do not limit the output size, we run both systems in that mode over a reduced Wikidata graph having only the triples with non-trivial time annotations; see WikiT in Table 1. We only report on the 76 BGPs that are supported by the code of Hu et al., that is, with variable subject and object. It can be seen that **JF** and **LC** outperform Hu et al. by orders of magnitude, which shows that our index has a competitive advantage over state-of-the-art solutions. **TF** is much slower in this case (still with better median but with worse distribution than Hu et al.); we consider next a scenario where **TF** is more competitive.

8.2 Other datasets with synthetic queries

Zhu et al. [55] study BGPs with fixed topologies (stars, chains, cycles, diamonds, cliques of various sizes) on various datasets. Table 1 describes the datasets we use; see the extended version for more details. For each topology, they generate 100 point-in-time queries where all nodes are variables and edge labels are assigned a random constant so that the tBGP occurs at least once in the graph. We evaluate our system and related work on this setting.

Hu et al. Table 2 shows the times of Hu et al. [28] on Divvy, with a 60-second timeout, compared with our system using **LC** and **TF**, and no output size limit. On the acyclic queries (stars and chains), **TF** is 1.7 to 4.5 times faster than Hu et al.; on the cyclic queries it is up to 50 times faster (and more in the shapes where Hu et al. have many timeouts). **LC** outperforms **TF** on the stars, but it is (sometimes much) slower in the other shapes. As choosing good VEOs makes a sharper difference on cyclic queries, this suggests that we take advantage from the freedom to choose good query plans. But it also shows that finding the best plan can be challenging.

Zhu et al. This paper implements point-in-time queries restricted to a time window $[t_s, t_e]$ given with the query; this corresponds in our language to adding the same temporal variable t to all BGPs and adding constraints $t_s \leq t$ and $t \leq t_e$. They restrict the results to random time windows that cover a fraction of the time domain, set an output size limit, and set the timeout to 60 seconds. While they do not offer public code, we follow their descriptions to perform similar experiments on Divvy, which is close to their description in size and other parameters. Their TSRJoin index (the one that performs best) takes 483 bytes per tuple (2.5 times our space) and is built 4 times faster than ours in their machine (which is clocked at 2.9 GHz, but its architecture is older). Table 2 compares query performance with a time window of 10% of the domain and limiting the results to 100,000, which is the only configuration where we can compare all query shapes. We show their best time (TSRJoin). Our approach, using only **TF** strategy this time, is 2–100 times faster than TSRJoin. While the 2–4 speedup factors on star shapes can be attributed to our more modern machine¹, we are 1–2 orders of magnitude faster on the other shapes. We conjecture that this is mainly due to our ability to choose arbitrary VEOs for the other variables: we outperform TSRJoin more sharply on the queries

¹According to Gemini, our hardware should be from 25%–35% faster for pure-CPU calculations, to 3 times faster for computations bounded by RAM transfers.

Table 2: Average time in msec to solve the different shapes on Divvy, with different windows of the time domain (100%, 10%, 1%) and limiting results to 100,000 (except on the 100% windows, which set no limit to match Hu et al. [28]). Times from Zhu et al. [55] are approximated from their plots. (*): Early termination due to many queries exceeding the 60,000 msec timeout.

	System	3-star	4-star	5-star	3-chain	4-chain	5-chain	3-circle	4-circle	5-circle	diamond	4-clique	5-clique
100%	Hu et al. [28]	250	349	454	180	206	235	262	2,382	*	3,614	1,389	*
	Ours (LC)	100	116	138	57	85	1,054	17	213	3,263	150	62	638
	Ours (TF)	146	169	195	46	53	52	47	50	56	96	55	63
10%	Zhu et al. [55]	50	50	50	50	50	70	150	150	70	100	300	700
	Ours (TF)	12.9	23.9	18.4	4.30	6.04	4.52	4.36	4.39	4.98	8.60	4.39	6.58
1%	Ours (TF)	2.22	4.22	2.39	0.97	1.07	0.83	0.80	0.86	1.04	2.06	1.18	1.13

Table 3: Average time in msec on three collections with our system. We use a 10% time window and limit the results to 100,000

Dataset	3-star	4-star	5-star	3-chain	4-chain	5-chain	3-circle	4-circle	5-circle	diamond	4-clique	5-clique
Divvy	12.9	23.9	18.4	4.30	6.04	4.52	4.36	4.39	4.98	8.60	4.39	6.58
Yellow	0.23	0.12	0.47	2.80	20.7	401	17.0	346	13,913	197	204	34.6
Caida	230	240	234	236	242	247	224	1,416	3,032	762	868	4,203

where it is not trivial to find the best VEO. On the star queries, instead, there is only one join variable apart from time.

More on our performance. By looking at the “Ours” rows in the areas 100%, 10%, and 1% of Table 2, we can see that the time window sensitivity has a significant impact on the query time of our index, showing that it filters effectively using the time component.

Table 3 measures our times on other datasets mentioned in Zhu et al.’s paper [55] (but much larger than their versions, so the results are not comparable): Yellow and Caida. We use a 10% time window and limit the results to 100,000. The differences in performance are explained mostly by the distribution of predicates: there are a few thousand in Divvy and Yellow, but those in Yellow are considerably skewed, while those in Divvy distribute uniformly. Caida, instead, has just one predicate. Since queries fix a predicate at random and leave all nodes as variables, the times in Caida are high and relatively uniform, growing as larger structures are reported: essentially we are generating all the ways a shape appears in the graph. Predicates help filter edges on Divvy, which yields low and relatively uniform times. On Yellow, with a skewed distribution, filtering is effective for most predicates, but the high times obtained on the larger structures when using a popular predicate dominate the average.

9 CONCLUSIONS

We have described a data structure that enables processing temporal BGPs in wco time, and provides further guarantees for other key applications such as point-in-time queries or queries with duration. Our experiments show that we solve realistic queries on large graphs within milliseconds, and that our strategy clearly outperforms simple solutions, as well as previous work. Although our index is static, in the extended version we describe how to support adding new events (i.e., edge insertions and deletions) to the graph.

One interesting challenge for future work is to obtain beyond-wco query times. For example, we can easily obtain time related to the generalized hypertree width of the query [1], which is never larger than its fractional hypertree width (the AGM bound). This

can be done by optimally partitioning the query into a tree of cyclic queries, using our wco algorithm to solve each cyclic component, and then applying Yannakakis’ algorithm [54] to solve the resulting acyclic join of intermediate results. A problem is that those intermediate results might be obtained with the bindings of the time variables in non-compacted form (akin to representing $\hat{G}$ instead of G), and then they might require a lot of space. A solution would be to represent the intermediate results optimally using time intervals, yet finding the minimal set of hyperrectangles covering all the instantiations of the temporal variables seems difficult.

A related challenge is to choose the most efficient VEOs. As seen in Figure 8, which approach works best depends on the query. The least-children (LC) strategy—akin to greedy optimization on cardinality estimates—works well in many cases, but there is room for better heuristics, including adaptive VEOs [7].

Another line of future work is to extend the language of the temporal clauses. For example, we may include timestamp arithmetic like $w_1 \leq w_2 + 3$ if $\mathcal{T}$ is numeric, which requires representing the original time domain $\mathcal{T}$ (rather than $\mathcal{T}_G$). In this case, the point graph $\hat{G}$ is only bounded by $N \cdot |\mathcal{T}|$, which can even be infinite, and worst-case optimality must be reconsidered. Still, our structures would work without many changes, yet it will be crucial to always bind temporal variables to ranges, as discussed in Section 5.3.

Finally, wco algorithms can be deployed for standard DBMS [21], which suggests our techniques could be deployed in a broader setting. This may include temporal path languages [6, 53] or even becoming part of temporal DBMS pipelines [27].

ACKNOWLEDGMENTS

This work was supported by ANID – Millennium Science Initiative Program – Code ICN17_002. GN was also supported by Fondecyt Grant 1260080. JR was also supported by Fondecyt Grant 1221799.

REFERENCES

- [1] Christopher R Aberger, Andrew Lamb, Susan Tu, Andres Nötzli, Kunle Olukotun, and Christopher Ré. Emptyheaded: A relational engine for graph processing. *ACM Transactions on Database Systems*, 42(4):1–44, 2017.
- [2] Mahmoud Abo Khamis, George Chichirim, Antonia Korpma, and Dan Olteanu. The complexity of boolean conjunctive queries with intersection joins. In *Proc. 41st ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems (PODS)*, pages 53–65, 2022.
- [3] Mahmoud Abo Khamis, Hung Q Ngo, and Dan Suciu. What do Shannon-type inequalities, submodular width, and disjunctive datalog have to do with one another? In *Proc. 36th ACM Symposium on Principles of Database Systems (PODS)*, pages 429–444, 2017.
- [4] Avishek Anand, Srikanta Bedathur, Klaus Berberich, and Ralf Schenkel. Index maintenance for time-travel text search. In *Proc. 35th International ACM Conference on Research and Development in Information Retrieval (SIGIR)*, pages 235–244, 2012.
- [5] Renzo Angles, Marcelo Arenas, Pablo Barceló, Aidan Hogan, Juan Reutter, and Domagoj Vrgoč. Foundations of modern query languages for graph databases. *ACM Computing Surveys*, 50(5):1–40, 2017.
- [6] Marcelo Arenas, Pedro Bahamondes, Amir Aghasadeghi, and Julia Stoyanovich. Temporal regular path queries. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 2412–2425. IEEE, 2022.
- [7] Diego Arroyuelo, Daniela Campos, Adrián Gómez-Brandón, Yuval Linker, Gonzalo Navarro, Carlos Rojas, and Domagoj Vrgoč. CompactLTJ: Space & time efficient Leapfrog Triejoin on graph databases. *The Very Large Databases Journal*, 34:article 67, 2025.
- [8] Albert Atserias, Martin Grohe, and Daniel Marx. Size bounds and query plans for relational joins. *SIAM Journal on Computing*, 42(4):1737–1767, 2013.
- [9] Klaus Berberich, Srikanta Bedathur, Thomas Neumann, and Gerhard Weikum. A time machine for text search. In *Proc. 30th Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR)*, pages 519–526, 2007.
- [10] Michael H Böhlen, Christian S Jensen, and Richard Thomas Snodgrass. Temporal statement modifiers. *ACM Transactions on Database Systems*, 25(4):407–456, 2000.
- [11] Borui Cai, Yong Xiang, Longxiang Gao, He Zhang, Yunfeng Li, and Jianxin Li. Temporal knowledge graph completion: A survey. *arXiv preprint arXiv:2201.08236*, 2022.
- [12] Li Cai, Xin Mao, Yuhao Zhou, Zhaoguang Long, Changxu Wu, and Man Lan. A survey on temporal knowledge graph: Representation learning and applications. *arXiv preprint arXiv:2403.04782*, 2024.
- [13] Matteo Ceccarello, Anton Dignös, Johann Gamper, and Christina Khnaisser. Indexing temporal relations for range-duration queries. In *Proceedings of the 35th International Conference on Scientific and Statistical Database Management*, pages 1–12, 2023.
- [14] David R. Clark. *Compact PAT Trees*. PhD thesis, University of Waterloo, Canada, 1996.
- [15] Francisco Claude, Gonzalo Navarro, and Alberto Ordóñez. The wavelet matrix: An efficient wavelet tree for large alphabets. *Information Systems*, 47:15–32, 2015.
- [16] Albert Croker and James Clifford. On completeness of historical relational data models. NYU Working Paper No. IS-89-002, 1989.
- [17] Ignacio Cuevas and Aidan Hogan. Versioned queries over RDF archives: All you need is SPARQL? In *MEPDAW@ISWC*, pages 43–52, 2020.
- [18] Anton Dignös, Michael H Böhlen, and Johann Gamper. Temporal alignment. In *Proc. ACM International Conference on Management of Data (SIGMOD)*, pages 433–444, 2012.
- [19] Nadime Francis, Amélie Gheerbrant, Paolo Guagliardo, Leonid Libkin, Victor Marsault, Wim Martens, Filip Murlak, Liat Peterfreund, Alexandra Rogova, and Domagoj Vrgoč. A researcher’s digest of GQL. In *Proc. 26th International Conference on Database Theory (ICDT)*, pages 1:1–1:22, 2023.
- [20] Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaer, Victor Marsault, Stefan Plankikow, Mats Rydberg, Petra Selmer, and André Taylor. Cypher: An evolving query language for property graphs. *SIGMOD Record*, 47(1):143–149, 2018.
- [21] Michael J. Freitag, Maximilian Bandle, Tobias Schmidt, Alfons Kemper, and Thomas Neumann. Adopting worst-case optimal joins in relational database systems. *Proceedings of the VLDB Endowment*, 13(11):1891–1904, 2020.
- [22] Yunjun Gao, Tianming Zhang, Linshan Qiu, Qingyuan Linghu, and Gang Chen. Time-respecting flow graph pattern matching on temporal graphs. *IEEE Transactions on Knowledge and Data Engineering*, 33(10):3453–3467, 2020.
- [23] Simon Gottschalk and Elena Demidova. Eventkg: A multilingual event-centric temporal knowledge graph. In *European semantic web conference*, pages 272–287. Springer, 2018.
- [24] Fabio Grandi et al. Multi-temporal RDF ontology versioning. In *IWOD@ISWC*, 2009.
- [25] Fabio Grandi et al. T-SPARQL: A TSQ2-like temporal query language for RDF. In *ADBIS (local proceedings)*, volume 639, pages 21–30, 2010.
- [26] Aidan Hogan, Cristian Riveros, Carlos Rojas, and Adrián Soto. A worst-case optimal join algorithm for SPARQL. In *Proc. 18th International Semantic Web Conference (ISWC)*, pages 258–275, 2019.
- [27] Jiamin Hou, Zhanhao Zhao, Wei Lu, Shiming Yang, Shuang Liu, Quanqing Xu, Chuanhui Yang, and Xiaoyong Du. An efficient and scalable graph database with built-in temporal support: J. hou et al. *The VLDB Journal*, 34(4):53, 2025.
- [28] Xiao Hu, Stavros Sintos, Junyang Gao, Pankaj K. Agarwal, and Jun Yang. Computing complex temporal join queries efficiently. In *Proc. International Conference on Management of Data (SIGMOD)*, pages 2076–2090, 2022.
- [29] Chengying Huan, Heng Zhang, Yongchao Liu, Likang Chen, Xuran Wang, Yongchun Jiang, Shaonan Ma, and Yanjun Wu. TeMatch: A fast temporal subgraph matching framework with temporal-aware subgraph matching algorithms. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*, pages 1029–1042, 2025.
- [30] Mahmoud A. Khamis, Hung Q. Ngo, Cristopher Ré, and Atri Rudra. Joins via geometric resolutions: Worst case and beyond. *ACM Transactions on Database Systems*, 41(4):22, 2016.
- [31] Udayan Khurana and Amol Deshpande. Storing and analyzing historical graph data at scale. In *International Conference on Extending Database Technology*. OpenProceedings.org, 2016.
- [32] Manolis Koubarakis and Kostis Kyzirakos. Modeling and querying metadata in the semantic sensor web: The model stRDF and the query language stSPARQL. In *Proc. Extended Semantic Web Conference*, pages 425–439, 2010.
- [33] Hans-Peter Kriegel, Marco Pötke, and Thomas Seidl. Managing intervals efficiently in object-relational databases. In *VLDB*, volume 20, page 0, 2000.
- [34] Krishna Kulkarni and Jan-Eike Michels. Temporal features in SQL: 2011. *ACM Sigmod Record*, 41(3):34–43, 2012.
- [35] Faming Li, Zhaonian Zou, and Jianzhong Li. Durable subgraph matching on temporal graphs. *IEEE Transactions on Knowledge and Data Engineering*, 35(5):4713–4726, 2022.
- [36] Min Lu, Qianzhen Zhang, and Xianqiang Zhu. Temporal multi-query subgraph matching in cybersecurity. *Technologies*, 13(8):335, 2025.
- [37] Stanislav Malyshev, Markus Krötzsch, Larry González, Julius Gonsior, and Adrian Bielefeldt. Getting the most out of Wikidata: Semantic technology usage in Wikipedia’s knowledge graph. In *Proc. 17th International Semantic Web Conference (ISWC)*, pages 376–394, 2018.
- [38] Amine Mhedhbi and Semih Salihoglu. Optimizing subgraph queries by combining binary and worst-case optimal joins. *Proceedings of the VLDB Endowment*, 12(11):1692–1704, 2019.
- [39] Vera Zaychik Moffitt and Julia Stoyanovich. Temporal graph algebra. In *Proc. 16th International Symposium on Database Programming Languages*, pages 1–12, 2017.
- [40] J. Ian Munro. Tables. In *Proc. 16th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pages 37–42, 1996.
- [41] Hung Q. Ngo. Worst-case optimal join algorithms: Techniques, results, and open problems. In *Proc. 37th Symposium on Principles of Database Systems (PODS)*, pages 111–124, 2018.
- [42] Hung Q. Ngo, Ely Porat, Cristopher Ré, and Atri Rudra. Worst-case optimal join algorithms. In *Proc. 31st Symposium on Principles of Database Systems (PODS)*, pages 37–48, 2012.
- [43] Dung Nguyen, Molham Aref, Martin Bravenboer, George Kollias, Hung Q Ngo, Christopher Ré, and Atri Rudra. Join processing for graph patterns: An old dog with new tricks. In *Proceedings of the GRADES’15*, pages 1–8, 2015.
- [44] Ashwin Paranjape, Austin R Benson, and Jure Leskovec. Motifs in temporal networks. In *Proc. 10th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 601–610, 2017.
- [45] Olivier Pelgrin, Ruben Taelman, Luis Galárraga, and Katja Hose. GLEND: Querying RDF archives with full SPARQL. In *Proc. European Semantic Web Conference (ESWC)*, pages 75–80, 2023.
- [46] Matthew Perry, Prateek Jain, and Amit P Sheth. SPARQL-st: Extending SPARQL to support spatiotemporal queries. In *Geospatial Semantics and the Semantic Web: Foundations, Algorithms, and Applications*, pages 61–86. Springer, 2011.
- [47] Konstantinos Semertzidis and Evaggelia Pitoura. Top-k durable graph pattern queries on temporal graphs. *IEEE Transactions on Knowledge and Data Engineering*, 31(1):181–194, 2018.
- [48] Richard Snodgrass. The temporal query language TQuel. *ACM Transactions on Database Systems*, 12(2):247–298, 1987.
- [49] Todd L. Veldhuizen. Triejoin: A simple, worst-case optimal join algorithm. In *Proc. International Conference on Database Theory (ICDT)*, pages 96–106, 2014.
- [50] Denny Vrandečić and Markus Krötzsch. Wikidata: a free collaborative knowledgebase. *Commun. ACM*, 57(10):78–85, 2014.
- [51] Jialing Wang, Immanuel Trummer, Ahmet Kara, and Dan Olteanu. ADOPT: Adaptively optimizing attribute orders for worst-case optimal join algorithms via reinforcement learning. *Proceedings of the VLDB Endowment*, 16(11):2805–2817, 2023.
- [52] Yisu Remy Wang, Max Willsey, and Dan Suciu. Free Join: Unifying worst-case optimal and traditional joins. *Proceedings of the ACM on Management of Data (SIGMOD)*, 1(2):150:1–150:23, 2023.
- [53] Huanhuan Wu, James Cheng, Silu Huang, Yiping Ke, Yi Lu, and Yanyan Xu. Path problems in temporal graphs. *Proceedings of the VLDB Endowment*, 7(9):721–732,

- 2014.
- [54] Mihalis Yannakakis. Algorithms for acyclic database schemes. In *Proc. 7th International Conference on Very Large Databases (VLDB)*, pages 82–94, 1981.
- [55] Kaijie Zhu, George Fletcher, and Nikolay Yakovets. Leveraging temporal and topological selectivities in temporal-clique subgraph query processing. In *Proc. 37th IEEE International Conference on Data Engineering (ICDE)*, pages 672–683, 2021.