

Distributed Subgraph Matching on Timely Dataflow

Longbin Lai[§], Zhu Qing[‡], Zhengyi Yang^{§‡}, Xin Jin[‡], Zhengmin Lai[‡], Ran Wang[‡], Kongzhang Hao[§], Xuemin Lin^{§‡}, Lu Qin[‡], Wenjie Zhang[§], Ying Zhang[‡], Zhengping Qian[‡] and Jingren Zhou[‡]

[§] The University of New South Wales, Sydney, Australia

[‡] East China Normal University, China

[‡] Centre for Artificial Intelligence, University of Technology, Sydney, Australia

[‡] Zhejiang Lab, China

[‡] Alibaba Group, China

[§]{llai,zyang,zhangw,lxue,khao}@cse.unsw.edu.au; [‡]{zhuqing,xinjin,zmlai,rwang}@stu.ecnu.edu.cn;
[‡]{lu.qin,ying.zhang}@uts.edu.au; [‡]{zhengping.qzp, jingren.zhou}@alibaba-inc.com

ABSTRACT

Recently there emerge many distributed algorithms that aim at solving subgraph matching at scale. Existing algorithm-level comparisons failed to provide a systematic view of distributed subgraph matching mainly due to the intertwining of strategy and optimization. In this paper, we identify four strategies and three general-purpose optimizations from representative state-of-the-art algorithms. We implement the four strategies with the optimizations based on the common **Timely** dataflow system for systematic strategy-level comparison. Our implementation covers all representative algorithms. We conduct extensive experiments for both unlabelled matching and labelled matching to analyze the performance of distributed subgraph matching under various settings, which is finally summarized as a practical guide.

PVLDB Reference Format:

Longbin Lai, Zhu Qing, Zhengyi Yang, Xin Jin, Zhengmin Lai, Ran Wang, Kongzhang Hao, Xuemin Lin, Lu Qin, Wenjie Zhang, Ying Zhang, Zhengping Qian and Jingren Zhou. Distributed Subgraph Matching on Timely Dataflow. *PVLDB*, 12(10): 1099-1112, 2019.
DOI: <https://doi.org/10.14778/3339490.3339494>

1. INTRODUCTION

Given a query graph Q and a data graph G , subgraph matching is defined as finding all subgraph instances of G that are isomorphic to Q . In this paper, we assume that the query graph and data graph are undirected¹ simple graphs, and may be unlabelled or labelled. We mainly focus on unlabelled case given that most distributed algorithms are developed under this setting. We also demonstrate some results of labelled matching due to its practical usefulness. Subgraph matching is one of the most fundamental operations in graph analysis, and has been used in a wide spectrum

¹Our implementation can seamlessly handle directed case.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. 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. 12, No. 10

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/3339490.3339494>

of applications [15, 25, 35, 42, 50]. As subgraph matching problem is in general computationally intractable [44], and data graph nowadays is growing beyond the capacity of one single machine, people are seeking efficient and scalable algorithms in the distributed context. Unless otherwise specified, in this paper we consider a simple **hash partition** of the graph data, that is the graph is randomly partitioned by the vertices, and the neighbors will be placed in the same partition.

By treating query vertices as attributes and the matched results as relational tables, we can express subgraph matching via natural joins. The problem is accordingly transformed into seeking optimal join plan, where the optimization goal is typically to minimize the communication cost.

State-of-the-arts In order to solve subgraph matching using join, existing works studied several join strategies, which can be categorized into three classes, namely “Binary-join-based subgraph-growing algorithms” (BINJOIN), “Worst-case optimal vertex-growing algorithms” (WOPTJOIN) and “Shares of Hypercube” (SHRCUBE). We also include OTHERS for algorithms that do not clearly belong to the above categories.

BINJOIN. The strategy computes subgraph matching by solving a series of binary joins. It first decomposes the original query graph into a set of *join units* whose matches can serve the base relations of the join. The strategy then joins the base relations based on a predefined *join order*. The BINJOIN algorithms differ in the use of join unit and join order. Typical choices of join unit are star (a tree of depth 1) in StarJoin [48], TwinTwig (an edge or intersection of two edges) in TwinTwigJoin [32], and clique (a graph whose vertices are mutually connected) in CliqueJoin [34]. Most existing algorithms adopt the easier-solving left-deep join order [29] except CliqueJoin, which explores the optimality-guaranteed bushy join [29].

WOPTJOIN. Given $\{v_0, v_1, \dots, v_n\}$ as the query vertices, WOPTJOIN strategy first computes all matches of $\{v_0\}$ that can present in the results, then matches of $\{v_0, v_1\}$, and so forth until constructing the results. Ngo et al. proposed the worst-case optimal join algorithm GenericJoin [39], based on which Ammar et al. implemented BiGJoin in Timely dataflow system [38] and showed its worst-case optimality [14]. In this paper, we also find out that the BINJOIN algo-

rithm `CliqueJoin` (with “overlapped decomposition”²) is also a variant of `GenericJoin`, and is hence worst-case optimal.

SHRCUBE. `SHRCUBE` strategy treats the computation of the query with n vertices as an n -dimensional *hypercube*. It partitions the hypercube across w workers in the cluster, and then each worker can compute its own share locally with no need of exchanging data. As a result, it typically renders much less communication cost than that of `BINJOIN` and `WOPTJOIN` algorithms. `MultiwayJoin` adopts the idea of `SHRCUBE` for subgraph matching. In order to properly partition the computation without missing results, `MultiwayJoin` needs to duplicate each edge in multiple workers. As a result, `MultiwayJoin` can almost carry the whole graph in each worker for certain queries [32, 14] and thus scale out poorly.

OTHERS. Shao et al. proposed `PSgL` [47] that processes subgraph matching via breadth-first traversal. Starting from an initial query vertex, `PSgL` iteratively expands the partial results by merging the matches of “gray” (visited yet not done) vertex’s neighbors. It has been pointed out in [32] that `PSgL` is actually a variant of `StarJoin`. Very recently, Qiao et al. proposed `CrystalJoin` [41] that aims at resolving the “output crisis” by compressing the (intermediate) results. The idea is to first compute the matches of the vertex cover of the query graph, then the remaining vertices’ matches can be compressed as intersection of the vertex cover’s neighbors to avoid the costly output of cartesian product.

Optimizations. Apart from join strategies, existing algorithms also explored a variety of optimizations, some of which are query- or algorithm-specific, while we spotlight three general-purpose optimizations, `Batching`, `TrIndexing` and `Compression`. `Batching` aims to divide the whole computation into sub-tasks that can be evaluated independently in order to save resource (memory) allocation. `TrIndexing` precomputes and indices the triangles (3-cycles) of the graph to facilitate pruning. `Compression` attempts to maintain the (intermediate) results in a compressed form to reduce maintaining and communication cost.

Motivations In this paper, we survey seven representative algorithms to solve distributed subgraph matching: `StarJoin` [48], `MultiwayJoin` [13], `PSgL` [47], `TwinTwigJoin` [32], `CliqueJoin` [34], `CrystalJoin` [41] and `BiGJoin` [14]. While all these algorithms embody some good merits in theory, existing **algorithm-level** comparisons failed to provide a systematic view of distributed subgraph matching due to several reasons. Firstly, the prior experiments did not take into consideration the differences of languages and the cost of the systems on which each implementation is based (Table 1). Secondly, some implementations hardcode query-specific optimizations for each query, which makes it hard to judge whether the observed performance is from the algorithmic advancement or hardcoded optimization. Thirdly, all `BINJOIN` and `WOPTJOIN` algorithms (more precisely, their implementations) intertwined join strategy with some optimizations of `Batching`, `TrIndexing` and `Compression`. We show in Table 1 how each optimization has been applied in current implementation. For example, `CliqueJoin` only adopted `TrIndexing` and some query-specific `Compression`, while `BiGJoin` considered `Batching` in general, but `TrIndexing` only for one query (`Compression` was discussed in paper, but not implemented). People naturally wonder that “*A strategy may perform better than B strategy if it applies C optimization*”, but unfortunately none

²Decompose the query graph into join units that are allowed to overlap edges

of existing implementation covers that combination. Last but not least, there misses an important benchmarking of the `FULLREP` strategy, that is to maintain the whole graph in each partition and parallelize embarrassingly [27].

Table 1 summarizes the surveyed algorithms regarding the category of strategy, the optimality guarantee, and the status of current implementations including the based platform and how the three optimizations are adopted.

Our Contributions. To address the above issues, we target a systematic, **strategy-level** benchmarking of distributed subgraph matching in this paper. To achieve this goal, we implement all strategies, together with the three general-purpose optimizations for subgraph matching based on the `Timely` dataflow system [38]. Note that our implementation covers all seven representative algorithms. Here, we use `Timely` as the base system as it incurs less cost [37] than other popular systems like Giraph [5], Spark [52] and GraphLab [36], so that the system’s impact can be reduced to the minimum.

We implement the benchmarking platform using our best effort based on the papers of each algorithm and email communications with the authors. Our implementation is (1) **generic** to handle arbitrary query, and does not include any hardcoded optimizations; (2) **flexible** that can configure `Batching`, `TrIndexing` and `Compression` optimizations in any combination for `BINJOIN` and `WOPTJOIN`; and (3) **efficient** that are comparable to and sometimes even faster than the original hardcoded implementation. Note that the three general-purpose optimizations are mainly used to reduce communication cost, and is not useful to the `SHRCUBE` and `FULLREP` strategies. Aware that their performance heavily depends on the local algorithm, we implement and compare the state-of-the-art local subgraph matching algorithms proposed in [12, 31] for unlabelled matching and [15] for labelled matching, and adopt the algorithm of best-possible performance. For `SHRCUBE`, we also refer to [19] to implement “Hypercube Optimization” for better hypercube sharing. We make the following contributions in the paper.

(1) **A benchmarking platform based on Timely dataflow system for distributed subgraph matching.** We implement four distributed subgraph matching strategies (and the general optimizations) that covers seven state-of-the-art algorithms: `StarJoin` [48], `MultiwayJoin` [13], `PSgL` [47], `TwinTwigJoin` [32], `CliqueJoin` [34], `CrystalJoin` [41] and `BiGJoin` [14]. Our implementation is generic to handle arbitrary queries, including the labelled and directed queries, and thus can guide practical use.

(2) **Three general-purpose optimizations - `Batching`, `TrIndexing` and `Compression`.** We investigate the literature on the optimization strategies, and spotlight the three general-purpose optimizations. We propose heuristics to incorporate the three optimizations into `BINJOIN` and `WOPTJOIN` strategies, with no need of query-specific adjustments from human experts. The three optimizations can be flexibly configured in any combination.

(3) **In-depth experimental studies.** In order to extensively evaluate the performance of each strategy and the effectiveness of the optimizations, we use data graphs of different sizes and densities, including sparse road network, dense ego network, and web-scale graph that is larger than each machine’s configured memory. We select query graphs of various characteristics that are either from existing works or suitable for benchmarking purpose. In addition to running time, we measure the communication cost, memory

Table 1: Summarization of the surveyed algorithms.

Algorithm	Category	Worst-case Optimality	Platform	Optimizations
StarJoin [48]	BINJOIN	No	Trinity [46]	None
MultiwayJoin [13]	SHRCUBE	N/A	Hadoop [32], Myria [19]	N/A
PSgL [47]	OTHERS	No	Giraph [5]	None
TwinTwigJoin [32]	BINJOIN	No	Hadoop	Compression [33]
CliqueJoin [34]	BINJOIN	Yes (Section 3.5)	Hadoop	TrIndexing, some Compression
CrystalJoin [41]	OTHERS	N/A	Hadoop	TrIndexing, Compression
BiGJoin [14]	WOPTJOIN	Yes [14]	Timely Dataflow [38]	Batching, specific TrIndexing

usage and other metrics to help reason the performance.

(4) A practical guide of distributed subgraph matching. Through empirical analysis covering the perspectives of join strategies, optimizations and join plans, we propose a practical guide for distributed subgraph matching. We also inspire interesting future work.

2. PRELIMINARIES

2.1 Problem Definition

Graph Notations. A graph g is defined as a 3-tuple, namely $g = (V_g, E_g, L_g)$, where V_g is the vertex set and $E_g \subseteq V_g \times V_g$ is the edge set of g , and L_g is a label function that maps each vertex $\mu \in V_g$ and/or each edge $e \in E_g$ to a label. Note that for unlabelled graph, L_g simply maps all vertices and edges to \emptyset . For a vertex $\mu \in V_g$, denote $\mathcal{N}_g(\mu)$ as the set of neighbors, $d_g(\mu) = |\mathcal{N}_g(\mu)|$ as the degree of μ , $\bar{d}_g = \frac{2|E_g|}{|V_g|}$ and $D_g = \max_{\mu \in V(g)} d_g(\mu)$ as the average and maximum degree, respectively. A *subgraph* g' of g , denoted $g' \subseteq g$, is a graph that satisfies $V_{g'} \subseteq V_g$ and $E_{g'} \subseteq E_g$.

Given $V' \subseteq V_g$, we define induced subgraph $g(V')$ as the subgraph induced by V' , that is $g(V') = (V', E(V'), L_g)$, where $E(V') = \{e = (\mu, \mu') \mid e \in E_g, \mu \in V' \wedge \mu' \in V'\}$. We say $V' \subseteq V_g$ is a *vertex cover* of g , if $\forall e = (\mu, \mu') \in E_g, \mu \in V'$ or $\mu' \in V'$. A minimum vertex cover V_g^c is a vertex cover of g that contains minimum number of vertices. A connected vertex cover is a vertex cover whose induced subgraph is connected, among which a minimum connected vertex cover is denoted as V_g^{cc} .

Data and Query Graph. We denote the data graph as G , and let $N = |V_G|$, $M = |E_G|$. Denote a data vertex of id i as u_i where $1 \leq i \leq N$. Note that the data vertex has been reordered such that if $d_G(u) < d_G(u')$, then $id(u) < id(u')$. We denote the query graph as Q , and let $n = |V_Q|$, $m = |E_Q|$, and $V_Q = \{v_1, v_2, \dots, v_n\}$.


Subgraph Matching. Given a data graph G and a query graph Q , we define *subgraph isomorphism*:

Definition 1. (Subgraph Isomorphism.) Subgraph isomorphism is defined as a bijective mapping $f : V_Q \rightarrow V_G$ such that: (1) $\forall v \in V_Q, L_Q(v) = L_G(f(v))$; (2) $\forall (v, v') \in E_Q, (f(v), f(v')) \in E_G$, and $L_Q((v, v')) = L_G((f(v), f(v')))$. A subgraph isomorphism is called a *Match* in this paper.

With the query vertices listed as $\{v_1, v_2, \dots, v_n\}$, we can simply represent a match f as $\{u_{k_1}, u_{k_2}, \dots, u_{k_n}\}$, where $f(v_i) = u_{k_i}$ for $1 \leq i \leq n$. The *Subgraph Matching* problem aims at finding all matches of Q in G . Denote $R_G(Q)$, or $R(Q)$ when the context is clear, as the result set of Q in G . As prior works [32, 34, 47], we apply *symmetry breaking* for unlabelled matching to avoid duplicate enumeration

caused by automorphism. Specifically, we first assign partial order O_Q to the query graph according to [24]. Here, $O_Q \subseteq V_Q \times V_Q$, and $(v_i, v_j) \in O_Q$ means $v_i < v_j$. In unlabelled matching, a match f must satisfy the *order constraint*: $\forall (v, v') \in O_Q$, it holds $f(v) < f(v')$. Note that we **do not** consider order constraint in labelled matching.

Example 1. In Figure 1, we present a query graph Q and a data graph G . For unlabelled matching, we give the partial order O_Q under the query graph. There are three matches: $\{u_1, u_2, u_6, u_5\}$, $\{u_2, u_5, u_3, u_6\}$ and $\{u_4, u_3, u_6, u_5\}$. It is easy to check that these matches satisfy the order constraint. Without the order constraint, there are actually four automorphic³ matches corresponding to each above match [13]. For labelled matching, we use different fillings to represent the labels. There are two matches accordingly - $\{u_1, u_2, u_6, u_5\}$ and $\{u_4, u_3, u_6, u_5\}$.

Figure 1: Query Graph Q (Left) and Data Graph G (Right).

By treating the query vertices as attributes, data edges and matched results as relational tables, we can write subgraph matching query as a multiway join of the edge relations. For example, regardless of label and order constraints, the query of Example 1 can be written as the following join

$$R(Q) = E(v_1, v_2) \bowtie E(v_2, v_3) \bowtie E(v_3, v_4) \bowtie E(v_1, v_4) \bowtie E(v_2, v_4). \quad (1)$$

This motivates researchers to leverage join operation for large-scale subgraph matching, given that join can be easily distributed, and it is natively supported in many distributed data engines like Spark [52] and Flink [16].

2.2 Timely Dataflow System

Timely is a distributed data-parallel dataflow system [38]. The minimum processing unit of **Timely** is a *worker*, which is simply a thread of execution, and one physical machine can schedule multiple workers. **Timely** follows the *shared-nothing dataflow* computation model [21] that abstracts the computation as a dataflow graph. In the dataflow graph, the vertex (a.k.a. *operator*) defines the computing logics and the edges in between the operators represent the data flow. One operator can accept multiple input streams, feed them to the computing, and produce (typically) one output stream. After the dataflow graph for certain computing task

³Automorphism is an isomorphism from one graph to itself.

is defined, it is distributed to each worker in the cluster, and further translated into a physical execution plan. Based on the physical plan, each worker can accordingly process the task in parallel while accepting a portion of the input.

3. ALGORITHM SURVEY

We survey the distributed subgraph matching algorithms following the categories of BINJOIN, WOPTJOIN, SHRCUBE, and OTHERS. We also show that CliqueJoin is a variant of GenericJoin [39], and is thus worst-case optimal.

3.1 BinJoin

BINJOIN processes subgraph matching by solving a series of binary joins. To improve the performance of BINJOIN, people devoted their efforts into: (1) using more complex base relations other than edge; (2) devising better join plan P . The base relations $B_{[q]}$ represent the matches of a set of sub-structures $[q]$ of the query graph Q . Each $p \in [q]$ is called a join unit, and it must satisfy $V_Q = \bigcup_{p \in [q]} V_p$ and $E_Q = \bigcup_{p \in [q]} E_p$. With the data graph partitioned across the cluster, [34] constrains the join unit to be the structure whose results can be independently computed within each partition (i.e. embarrassingly parallel [27]). It is not hard to see that when each vertex has full access to the neighbors in the partition, we can compute the matches of a k -star (a star of k leaves) rooted on the vertex u by enumerating all k -combinations within $\mathcal{N}_G(u)$. Therefore, star is a qualified and indeed widely used join unit.

Given the base relations, the join plan P determines an order of processing binary joins. A join plan is *left-deep*⁴ if there is at least a base relation involved in each join, otherwise it is *bushy*.

StarJoin. As the name suggests, StarJoin uses star as the join unit, and it follows the left-deep join order. To decompose the query graph, it first locates the vertex cover of the query graph, and each vertex in the cover and its unused neighbors naturally form a star [48]. A StarJoin plan for Equation 1 is

$$(J_1) R(Q) = \text{Star}(v_2; \{v_1, v_3, v_4\}) \bowtie \text{Star}(v_4; \{v_2, v_3\}),$$

where $\text{Star}(r; L)$ denotes a **Star** relation (the matches of the star) with r as the root, and L as the set of leaves.

TwinTwigJoin. Enumerating a k -star on a vertex of degree d will render $O(d^k)$ cost. We refer *star explosion* to the case while enumerating stars on a large-degree vertex. Lai et al. proposed TwinTwigJoin [32] to address the issue of StarJoin by forcing the join plan to use TwinTwig (a star of at most two edges) instead of a general star as the join unit. Intuitively, this would help ameliorate the star explosion by constraining the cost of each join unit from d^k of arbitrary k to at most d^2 . TwinTwigJoin follows StarJoin to use left-deep join order. The authors proved that given any general StarJoin plan in the left-deep join order, we can rewrite it as an alternative TwinTwigJoin plan that draws no more cost (evaluated using random graph model [22]) than the original StarJoin [32]. A TwinTwigJoin plan for Equation 1 is

$$(J_1) R_1(v_1, v_2, v_3, v_4) = \text{TwinTwig}(v_1; \{v_2, v_4\}) \bowtie \text{TwinTwig}(v_2; \{v_3, v_4\}); \quad (2)$$

$$(J_2) R(Q) = R_1(v_1, v_2, v_3, v_4) \bowtie \text{TwinTwig}(v_3; \{v_4\}),$$

⁴More precisely it is deep, and can further be left-deep and right-deep. In this paper, we assume that it is left-deep following the prior work [32].

where $\text{TwinTwig}(r; L)$ denotes a **TwinTwig** relation with r as the root, and L as the leaves.

CliqueJoin. TwinTwigJoin hampers star explosion to some extent, but still suffers from the problems of long execution ($\Omega(\frac{m}{2})$ rounds) and suboptimal left-deep join plan. CliqueJoin resolves the issues by extending StarJoin in two aspects. Firstly, CliqueJoin applies the “triangle partition” strategy (Section 4.2), which enables CliqueJoin to use clique, in addition to star, as the join unit. The use of clique can greatly shorten the execution especially when the query is dense, although it still degenerates to StarJoin when the query contains no clique subgraph. Secondly, CliqueJoin exploits the bushy join plan to approach optimality. A CliqueJoin plan for Equation 1 is:

$$(J_1) R(Q) = \text{Clique}(\{v_1, v_2, v_4\}) \bowtie \text{Clique}(\{v_2, v_3, v_4\}), \quad (3)$$

where $\text{Clique}(V)$ denotes a **Clique** relation of V .

Implementation Details. We implement the BINJOIN strategy based on the join framework proposed in [34] to cover StarJoin, TwinTwigJoin and CliqueJoin.

We use power-law random graph (PR) model [20] to estimate the cost as [34], and implement the dynamic programming algorithm [34] to compute the cost-optimal join plan. Once the join plan is computed, we translate the plan into Timely dataflow that processes each binary join using a Join operator. We implement the Join operator following Timely’s official “pipeline” HashJoin example⁵. We modify it into a “batching-style” join that is analogous to a MapReduce process - the mappers (senders) shuffle the data based on the join key, while the reducers (receivers) maintain the received key-value pairs in a hash table for join processing (until mappers complete). The reasons that we adopt the “batching-style” join are: (1) buffering all inputs before join can leverage the partial order of the query graph, where binary search can be applied to improve the performance; (2) it replays the original implementation in Hadoop; and (3) it favors the **Batching** optimization (Section 4.1).

3.2 WOPTJoin

WOPTJOIN strategy processes subgraph matching by matching vertices in a predefined order. Given the query graph Q and $V_Q = \{v_1, v_2, \dots, v_n\}$ as the matching order, the algorithm starts from an empty set, and computes the matches of the subset $\{v_1, \dots, v_i\}$ in the i^{th} rounds. Denote the partial results after the i^{th} ($i < n$) round as R_i , and $p = \{u_{k_1}, u_{k_2}, \dots, u_{k_i}\} \in R_i$ is one of the tuples. In the $i + 1^{\text{th}}$ round, the algorithm expands the results by matching v_{i+1} with $u_{k_{i+1}}$ for p iff. $\forall_{1 \leq j \leq i} (v_j, v_{i+1}) \in E_Q$, $(u_{k_j}, u_{k_{i+1}}) \in E_G$. It is immediate that the candidate matches of v_{i+1} , denoted $C(v_{i+1})$, can be obtained by intersecting the relevant neighbors of the matched vertices as

$$C(v_{i+1}) = \bigcap_{\forall_{1 \leq j \leq i} (v_j, v_{i+1}) \in E_Q} \mathcal{N}_G(u_{k_j}). \quad (4)$$

BiGJoin. BiGJoin adopts the WOPTJOIN strategy in Timely dataflow system. The main challenge is to implement the intersection efficiently using Timely dataflow. For that purpose, the authors designed the following three operators: (1) **Count**: Checking the number of neighbors of each u_{k_j} in Equation 4 and recording the worker that stores

⁵<https://github.com/TimelyDataflow/timely-dataflow/blob/master/examples/hashjoin.rs>

the one with the smallest neighbor set; (2) **Propose**: Attaching the smallest neighbor set to p as $(p; C(v_{i+1}))$; (3) **Intersect**: Sending $(p; C(v_{i+1}))$ to the worker that maintains each u_{k_j} and update $C(v_{i+1}) = C(v_{i+1}) \cap \mathcal{N}_G(u_{k_j})$. After intersection, we will expand p by pushing into p every vertex of $C(v_{i+1})$.

Implementation Details. We directly use the authors’ implementation [6], but slightly modify the codes to use the common graph data structure. We do not consider the dynamic version of BiGJoin in this paper, as the other strategies currently only support static context. The matching order is determined using a greedy heuristic that starts with the vertex of the largest degree, and consequently selects the next vertex that connects with most already-selected vertices (id as tie breaker).

3.3 ShrCube

SHRCUBE strategy treats the join as a hypercube of n (n is the number of attributes) dimension. It attempts to divide the hypercube evenly across the workers in the cluster, so that each worker can compute its own share without data communication.

MultiwayJoin. MultiwayJoin applies the SHRCUBE strategy to solve subgraph matching in one single round. Consider w workers in the cluster, a query graph Q with $V_Q = \{v_1, v_2, \dots, v_n\}$ vertices and $E_Q = \{e_1, e_2, \dots, e_m\}$, where $e_i = (v_{i_1}, v_{i_2})$. Regarding each query vertex v_i , assign a positive integer as bucket number b_i that satisfies $\prod_{i=1}^n b_i = w$. The algorithm then divides the candidate data vertices for v_i evenly into b_i parts via a hash function $h : u \mapsto z_i$, where $u \in V_G, 1 \leq z_i \leq b_i$. This accordingly divides the whole computation into w shares, each of which can be indexed via an n -ary tuple (z_1, z_2, \dots, z_n) , and is assigned to one worker. Afterwards, regarding each query edge $e_i = (v_{i_1}, v_{i_2})$, MultiwayJoin maps each data edge (u, u') to n -ary tuples as $(z_1, \dots, z_{i_1} = h(u), \dots, z_{i_2} = h(u'), \dots, z_n)$, where other than z_{i_1} and z_{i_2} , each z_i iterates through $\{1, 2, \dots, b_i\}$, and the edge will be sent to all related workers.

Implementation Details. There are two main impact factors of the performance of SHRCUBE. Firstly, the hypercube sharing by assigning proper b_i for v_i . To achieve better performance, we adopt the ‘‘Hypercube Optimization’’ proposed by Chu et al [19]. Secondly, the local algorithm. When the edges arrive at the worker, we collect them into a local graph (duplicate edges are removed), and use local algorithm to compute the matches. For unlabelled matching, we study the state-of-the-art local algorithms from ‘‘EmptyHeaded’’ [12] and ‘‘DualSim’’ [31]. We implement all our benchmarking queries (Figure 2) using ‘‘EmptyHeaded’’ and ‘‘DualSim’’, ‘‘DualSim’’ generally performs better as it does not need to maintain intermediate results as ‘‘EmptyHeaded’’. As a result, we adopt ‘‘DualSim’’ as the local algorithm for MultiwayJoin. For labelled matching, we implement ‘‘CFLMatch’’ proposed in [15], which is the best-so-far labelled matching algorithm to our best knowledge.

Now we let each worker independently compute matches in its local graph. Simply doing so will result in duplicates, so we process deduplication for MultiwayJoin, the details of which can be found in the full paper [4].

3.4 Others

PSgL. PSgL processes subgraph matching iteratively via breadth-first traversal. All query vertices are configured

three status, ‘‘white’’ (initialized), ‘‘gray’’ (visited, but neighbors unmatched) and ‘‘black’’ (finalized). Denote v_i as the vertex to match in the i^{th} round. The algorithm starts from matching initial query vertex v_1 , and coloring the neighbors as ‘‘gray’’. In the i^{th} round, the algorithm applies the workload-aware expanding strategy at runtime to select the v_i to expand among all current ‘‘gray’’ vertices based on a greedy heuristic to minimize the communication cost [46]; the partial results from previous round R_{i-1} (specially, $R_0 = \emptyset$) will be distributed among the workers based on the candidate data vertices that can match v_i ; the algorithm then computes R_i by merging R_{i-1} with the matches of the star formed by v_i and its neighbors $\mathcal{N}_Q(v_i)$, namely $\text{Star}(v_i; \mathcal{N}_Q(v_i))$; after v_i is matched, v_i is colored as ‘‘black’’ and its ‘‘white’’ neighbors will be colored as ‘‘gray’’; essentially, this process is analogous to StarJoin by processing $R_i = R_{i-1} \bowtie \text{Star}(v_i; \mathcal{N}_Q(v_i))$. As prior work [34] have already shown that PSgL’s performance is dominated by CliqueJoin [34], we will not further discuss this algorithm in this paper.

CrystalJoin. CrystalJoin aims to resolve the ‘‘output crisis’’ by compressing the results of subgraph matching [41]. The authors defined a structure called *crystal*, denoted $\mathcal{Q}(x, y)$. A crystal is a subgraph of Q that contains two sets of vertices V_x and V_y ($|V_x| = x$ and $|V_y| = y$). Here, the induced subgraph $Q(V_x)$ is a x -clique, while V_y forms an independent set and every vertex in V_y connects to all vertices of V_x . The algorithm first obtains the minimum vertex cover V_Q^c , and then applies the *Core-Crystal Decomposition* to decompose the query graph into the *core* $Q(V_Q^c)$ and a set of *crystals* $\{\mathcal{Q}_1(x_1, y_1), \dots, \mathcal{Q}_t(x_t, y_t)\}$. Each crystal satisfies that its clique part is a subgraph of the core.

With core-crystal decomposition, the computation has accordingly split into three stages: (1) **Core computation.** Given that $Q(V_Q^c)$ itself is a query graph, the algorithm can be recursively applied to compute $Q(V_Q^c)$; (2) **Crystal computation.** Given a crystal $\mathcal{Q}(x, y)$, we can compute the matches as (f_x, I_y) , where $f_x = \{u_1, u_2, \dots, u_x\}$ is a matched instance of $Q(V_x)$, and $I_y = \bigcap_{i=1}^x \mathcal{N}_G(u_i)$ denotes the set of vertices that can match V_y . It is not hard to see that each y -combination of I_y together with f_x recover a match; (3) **One-time assembly.** This stage joins the core instances with the compressed crystal matches to produce the final results.

We notice two technical obstacles to implement CrystalJoin according to the paper. Firstly, the core $Q(V_Q^c)$ may be disconnected, a case that can produce exponential number of results. The authors originally implemented to maintain a full replica of the graph in each machine to resolve this issue, which is not desirable as we’d rather use FULLREP strategy. Secondly, the authors proposed to precompute the cliques up to certain k , which is often cost-prohibitive in practice. Take UK (Table 2) dataset as an example, the triangles, 4-cliques and 5-cliques are respectively about 20, 600 and 40000 times larger than the graph itself. Therefore, we adapt CrystalJoin in the following. Firstly, we replace the core $Q(V_Q^c)$ with $Q(V_Q^{cc})$. Secondly, instead of implementing CrystalJoin as a new strategy, we use it as an alternative join plan (matching order) for WOPTJOIN. According to CrystalJoin, we first match $Q(V_Q^{cc})$, while the matching order inside and outside V_Q^{cc} still follows WOPTJOIN’s greedy heuristic (Section 3.2). It is worth noting that our adaptation achieves high performance comparable to the original implementation.

FullRep. We implement FULLREP by letting each worker

pick its share of computation via a Round-Robin strategy, that is we settle an initial query vertex v_1 , and let first worker match v_1 with u_1 to continue the remaining process, and second worker match v_1 with u_2 , and so on. This simple strategy already works very well on balancing the load. We use “DualSim” for unlabelled matching and “CFLMatch” for labelled matching as SHRCUBE.

3.5 Worst-case Optimality

Given a query Q and the data graph G , we denote the maximum possible result set as $\overline{R}_G(Q)$. Simply speaking, an algorithm is worst-case optimal if the aggregation of the total output (including intermediate output) is bounded by $\Theta(|\overline{R}_G(Q)|)$. Todd [49] and Ngo et al. [39] almost meanwhile proposed the worst-case join algorithm. Ammar et al. proposed **GenericJoin** for subgraph matching based on Ngo’s algorithm [39]. Following **GenericJoin**, **BigJoin** was developed and shown to be worst-case optimal [14].

As for **CliqueJoin**, the optimality has not been claimed in the paper [34]. In this work, we further contribute to the finding of **CliqueJoin**’s worst-case optimality. We refer interested readers to the full paper [4] for a complete proof.

4. OPTIMIZATIONS

We brief the three general-purpose optimizations, **Batching**, **TrIndexing** and **Compression** in this section, and how we orthogonally apply them to **BINJOIN** and **WOPTJOIN** algorithms. Here we only introduce the basic idea due to short of space, while we refer interested readers to the full paper [4] for details. Note that in the rest of the paper, we will use the strategy **BINJOIN**, **WOPTJOIN**, **SHRCUBE** instead of their corresponding algorithms, as we focus on strategy-level comparison.

4.1 Batching

Let $R(V_i)$ be the partial results that match the given vertices $V_i = \{v_{s_1}, v_{s_2}, \dots, v_{s_i}\}$ (R_i for short if V_i follows a given order), and $R(V_j)$ denote the more complete results with $V_i \subset V_j$. Denote $R_j|R_i$ as the tuples in R_j whose projection on V_i equates R_i . Let’s partition R_i into b **disjoint** parts $\{R_i^1, R_i^2, \dots, R_i^b\}$. We define **Batching** on $R_j|R_i$ as the technique to independently process the following sub-tasks that compute $\{R_j|R_i^1, R_j|R_i^2, \dots, R_j|R_i^b\}$. Obviously, $R_j|R_i = \bigcup_{k=1}^b R_j|R_i^k$.

WOptJoin. Recall from Section 3.2 that **WOPTJOIN** progresses according to a predefined matching order $\{v_1, v_2, \dots, v_n\}$. In the i^{th} round, **WOPTJOIN** will **Propose** on each $p \in R_{i-1}$ to compute R_i . It is not hard to see that we can easily apply **Batching** to the computation of $R_i|R_{i-1}$ by randomly partitioning R_{i-1} . For simplicity, the authors implemented **Batching** on $R(Q)|R_1(v_1)$. Note that $R_1(v_1) = V_G$ in unlabelled matching, which means that we can achieve **Batching** simply by partitioning the data vertices. For short, we also say the strategy batches on v_1 , and call v_1 the batching vertex. We follow the same idea to apply **Batching** to **BINJOIN** algorithms.

BinJoin. While it is natural for **WOPTJOIN** to batch on v_1 , it is non-trivial to pick such a vertex for **BINJOIN**. We propose the heuristic to apply **Batching** to **BINJOIN** on the vertex that presents in the most join units. Note that such vertex can only be in the join key, as otherwise it must at least not present in either side of the join. For complex query, we can still have join unit that does not contain any batching vertex using the above heuristic. In this case, the

sub-query that does not contain the batching vertex will cause huge burden on the **Join** operation. Thus, we devise the *join-level Batching* following the idea of external **MergeSort**. Specifically, we inject a **Buffer-and-Batch** operator for the two data streams before they arrive at the **Join** operator. The **Buffer-and-Batch** operator will buffer the data on disk, and then consume them for join in batches, where each batch is determined by the hash value of the join key. Note that such join-level **Batching** is natively implemented in Hadoop’s “Shuffle” stage, and we incorporate this in **Timely** to improve the scalability of **BINJOIN**.

4.2 Triangle Indexing

As the name suggests, **TrIndexing** precomputes the triangles of the data graph and indices them along with the graph data to prune infeasible results. Based on the default hash partition, Lai et al. proposed “triangle partition” [34] to also incorporate the edges among the neighbors in each partition. “Triangle partition” allows **BINJOIN** to use clique as the join unit [34], which greatly improves the performance of certain queries. “Triangle partition” is in de facto a variant of **TrIndexing**, which instead of explicitly materializing the triangles, maintains them in the local graph structure (e.g. adjacency list). As we will show in the experiment (Section 5), this will save a lot of storage compared to explicit triangle materialization. As a result, we adopt the “triangle partition” for **TrIndexing** optimization in this work.

BinJoin. It is obvious that **BINJOIN** becomes **CliqueJoin** with **TrIndexing**, and **StarJoin** (or **TwinTwinJoin**) otherwise.

WOptJoin. In order to match v_i in the i^{th} round, **WOPTJOIN** utilizes **Count**, **Propose** and **Intersect** to process the intersection of Equation 4. For ease of presentation, suppose v_{i+1} connects to the first s query vertices $\{v_1, v_2, \dots, v_s\}$, and given a partial match, $\{f(v_1), \dots, f(v_s)\}$, we have $C(v_{i+1}) = \bigcap_{j=1}^s \mathcal{N}_G(f(v_j))$. In the original implementation, it is required to send $(p; C(v_{i+1}))$ via network to all machines that contain each $f(v_j)$ ($1 \leq j \leq s$) to process the intersection, which can render massive communication cost. In order to reduce the communication cost, we implement **TrIndexing** for **WOPTJOIN** in the following. We first group $\{v_1, \dots, v_s\}$ such that for each group $U(v_x)$, we have $U(v_x) = \{v_x\} \cup \{v_y \mid (v_x, v_y) \in E_Q\}$. Because of **TrIndexing**, we have $\mathcal{N}_G(f(v_y))$ ($\forall v_y \in U(v_x)$) maintain in $f(v_x)$ ’s partition. Thus, we only need to send the prefix to $f(v_x)$ ’s machine, and the intersection within $U(v_x)$ can be done locally. We process the grouping using a greedy strategy that always constructs the largest group from the remaining vertices.

4.3 Compression

Compression aims to maintain the (intermediate) results in a compressed form to reduce the maintaining and communication cost. In the following, when we say “compress a query vertex”, we mean maintaining its matched data vertices in the form of an array, without unfolding them in line with the one-one mapping of a match (Definiton 1). Qiao et al. proposed **CrystalJoin** to study **Compression** in general [41], which inspires a heuristic for doing **Compression**, that is *to compress the vertices whose matches will not be used in any future computation*. In fact, **Compression** is conventionally studied in relational database as factorization [40], while we follow **CrystalJoin** to call it **Compression** as it literally is.

BinJoin. Obviously we can not compress any vertex that presents in the join key. What we need to do is to locate the

vertices to compress in the join unit, namely star and clique. For star, the root vertex must remain uncompressed, as the leaves’ computation depends on it. For clique, we can only compress one vertex, as otherwise the mutual connection between the compressed vertices will be lost. In a word, we compress two types of vertices: (1) non-key and non-root vertices of a star join unit, (2) one non-key vertex of a clique join unit. Now that the join units are computed with **Compression**, the join results are compressed accordingly.

WOptJoin. Based on a predefined join order $\{v_1, v_2, \dots, v_n\}$, we can compress v_i ($1 \leq i \leq n$), if there does not exist v_j ($i < j$) such that $(v_i, v_j) \in E_Q$. In other words, v_i ’s matches will never be involved in any future intersection (computation). Note that v_n ’s can be trivially compressed. With **Compression**, when v_i is compressed, we will maintain its matches as an array instead of unfolding it into the prefix like a normal vertex.

5. EXPERIMENTS

5.1 Experimental settings

Environments. We deploy two clusters for the experiments: (1) a local cluster of 10 machines connected via one 10GBps switch and one 1GBps switch. Each machine has 64GB memory, 1 TB disk and 1 Intel Xeon CPU E3-1220 V6 3.00GHz with 4 physical cores; (2) an AWS cluster of 40 “r5-2xlarge” instances connected via a 10GBps switch, each with 64GB memory, 8 vCpus and 500GB Amazon EBS storage. By default we use the local cluster of 10 machines with 10GBps switch. We run 3 workers in each machine in the local cluster, and 6 workers in the AWS cluster for **Timely**. The codes are implemented based on the open-sourced **Timely** dataflow system [9] using Rust 1.32. We are still working towards open-sourcing the codes, and the bins together with their usages are temporarily provided⁶ to verify the results.

Metrics. In the experiments, we measure query time T as the slowest worker’s wall clock time from an average of three runs. We allow 3 hours as the maximum running time for each test. We use **OT** and **OOM** to indicate a test case runs out of the time limit and out of memory, respectively. By default we will not show the **OOM** results for clear presentation. We divide T into two parts, the computation time T_{comp} and the communication time T_{comm} . We measure T_{comp} as the time the slowest worker spends on actual computation by timing every computing function. We are aware that the actual communication time is hard to measure as **Timely** overlaps computation and communication to improve throughput. We consider $T - T_{comp}$, which mainly records the time the worker waits data from the network channel, and it belongs solely to communication time. While the other part of communication that overlaps computation is of less interest as it does not affect the query progress. As a result, we simply let $T_{comm} = T - T_{comp}$ in the experiments. We measure the peak memory usage using Linux’s “**time -v**” in each machine. We define the communication cost as the number of integers a worker receives during the process, and measure the maximum communication cost among the workers accordingly.

Dataset Formats. We preprocess each dataset as follows: we treat it as a simple undirected graph by removing self-loop and duplicate edges, and format it using “Compressed

Sparse Row” (CSR) [3]. We relabel the vertex id according to the degree and break the ties arbitrarily.

Compared Strategies. In the experiments, we implement **BINJOIN** and **WOPTJOIN** with all **Batching**, **TrIndexing** and **Compression** optimizations (Section 4). **SHRCUBE** is implemented with “Hypercube Optimization” [19], and “DualSim” (unlabelled) [31] and “CFLMatch” (labelled) [15] as local algorithms. **FULLREP** is implemented with the same local algorithms as **SHRCUBE**.

Auxiliary Experiments. We have also conducted several auxiliary experiments to study the strategies of **BINJOIN**, **WOPTJOIN**, **SHRCUBE** and **FULLREP**: (1) “Scalability” experiment in unlabelled matching. The experiment shows that **FULLREP** scales out the best, followed by **WOPTJOIN**, **BINJOIN** and **SHRCUBE**; (2) “Vary Density” experiment in labelled matching. This experiment shows that all strategies’ performance decreases with the densities; (3) “Vary Labels” experiment in labelled matching. This experiment shows the transition from unlabelled to labelled matching for all strategies. We present these experiments in the appendix of the full paper [4].

5.2 Unlabelled Experiments

Datasets. The datasets used in this experiment are shown in Table 2. All datasets except **SY** are downloaded from public source, which are indicated by the letter in the bracket (**S** [10], **W** [11], **D** [1]). All statistics are measured as G is an undirected graph. Among the datasets, **GO** is a small dataset to study cases of extremely large (intermediate) result set; **LJ**, **UK** and **FS** are three popular datasets used in prior works, featuring statistics of real social network and web graph; **GP** is the google plus ego network, which is exceptionally dense; **US** and **EU**, on the other end, are sparse road networks. These datasets vary in number of vertices and edges, densities and maximum degree, as shown in Table 2. We synthesize the **SY** data according to [17] that generates data with real-graph characteristics. Note that the data occupies roughly 80GB space, and is larger than the configured memory of our machine. We synthesize the data because we do not find public accessible data of this size. Larger dataset like **Clueweb** [2] is available, but it is beyond the processing power of our current cluster.

Each data is hash partitioned (“hash”) across the cluster. We also implement the “triangle partition” (“tri.”) for **TrIndexing** optimization (Section 4.2). To do so, we use **BiGJoin** to compute the triangles and send the triangle edges to corresponding partition. We record the time T_* and average number of edges $|\overline{E_*}|$ of the two partition strategies. The partition statistics are recorded while running in the local cluster, except for **SY** that is processed in the AWS cluster. From Table 2, we can see that $|\overline{E_{tri.}}|$ is noticeably larger, around 1-10 times larger than $|\overline{E_{hash}}|$. Note that in **GP** and **UK**, which either is dense, or must contain a large dense community, the “triangle partition” can maintain a large portion of data in each partition. While compared to complete triangle materialization, “triangle partition” turns out to be much cheaper. For example, the **UK** dataset contains around 27B triangles, which means each partition in our local cluster should by average take 0.9B triangles (three integers); in comparison, **UK**’s “triangle partition” only maintains an average of 0.16B edges (two integers) according to Table 2.

We use **US**, **GO** and **LJ** as default datasets in the experiments “Exp-1”, “Exp-2” and “Exp-3” in order to collect

⁶<https://goo.gl/Xp5BrW>

useful feedbacks from successful queries, while we may not present certain cases when they do not give new findings.

Queries. The queries are presented in Figure 2. We also give the partial order under each query for symmetry breaking. The queries except q_7 and q_8 are selected based on all prior works [14, 32, 34, 41, 47], while varying in number of vertices, densities, and the vertex cover ratio $|V_Q^{cc}|/|V_Q|$, in order to better evaluate the strategies from different perspectives. The three queries q_7 , q_8 and q_9 are relatively challenging given their result scale. For example, the smallest dataset GO contains 2,168B(illion) q_7 , 330B q_8 and 1,883B q_9 , respectively. We present the number of results of each successful query on each dataset in the auxiliary materials (Section 9). Note that q_7 and q_8 are absent from existing works, while we benchmark q_7 considering the importance of path query in practice, and q_8 considering the varieties of the join plans.




Figure 2: The unlabelled queries.

Exp-1: Optimizations. We study the effectiveness of **Batching**, **TrIndexing** and **Compression** for both BINJOIN and WOPTJOIN strategies, by comparing BINJOIN and WOPTJOIN with their respective variants with one optimization off, namely “without Batching”, “without Trindexing” and “without Compression”. In the following, we use the suffix of “(w.o.b.)”, “(w.o.t.)” and “(w.o.c.)” to represent the three variants. We use the queries q_2 and q_5 , and the results of US and LJ are shown in Figure 3. By default, we use the batch size of 1,000,000 for both BINJOIN and WOPTJOIN (according to [14]) in this experiment, and we reduce the batch size when it runs out of memory, as will be specified.

While comparing BINJOIN with BINJOIN_(w.o.b.), we observe that **Batching** barely affects the performance of q_2 , but severely for q_5 on LJ (1800s vs 4000s (w.o.b.)). The reason is that we still apply join-level **Batching** for BINJOIN_(w.o.b.) that dumps the intermediate data to the disk (Section 4.1). While q_5 ’s intermediate data includes the massive results of sub-query $Q(\{v_2, v_3, v_4, v_5\})$, which incurs huge amount of disk I/O (US does not have this problem as it produces very few results). We also run q_5 without the join-level **Batching** on LJ, but it fails with OOM. For BINJOIN, **TrIndexing** is a critical optimization, with the observed performance of BINJOIN better than that of BINJOIN_(w.o.t.), especially so on LJ. This is expected as BINJOIN_(w.o.t.) actually degenerates to **StarJoin**. **Compression**, on the one hand, allows BINJOIN to run much faster than BINJOIN_(w.o.c.) for both queries on LJ, on the other hand, makes it slower on US. The reason is that US is a sparse dataset with few room for **Compression**, while **Compression** itself incurs extra cost. We also compare BINJOIN with BINJOIN_(w.o.c.) on the other sparse graph EU, and the results are the same.

For WOPTJOIN strategy, **Batching** has little impact to the performance. Surprisingly, after using **TrIndexing** to WOPTJOIN, the improvement by average is only around 18%. We do another experiment in the same cluster but using 1GBps switch, which shows WOPTJOIN is over 6 times faster than WOPTJOIN_(w.o.t.) for both queries on LJ. Note that **Timely** uses separate threads to buffer received data from the network. Given the same computing speed, a faster network allows the data to be more fully buffered and hence less wait for the following computation. Similar to BINJOIN, **Compression** greatly improves the performance while querying on LJ, but the opposite on US.

Exp-2 Challenging Queries. We study the challenging queries q_7 , q_8 and q_9 in this experiment. We run this experiment using BINJOIN, WOPTJOIN, SHRCUBE and FULLREP, and show the results of US and GO (LJ failed all cases) in Figure 4. Recall that we split the time into computation time and communication time (Section 5.1), here we plot the communication time as gray filling in each bar of Figure 4.

FULLREP beats all the other strategies, while SHRCUBE fails q_8 and q_9 on GO because of OT. Although SHRCUBE uses the same local algorithm as FULLREP, it spends a lot of time on deduplication (Section 3.3).

We focus on comparing BINJOIN and WOPTJOIN on GO dataset. On the one hand, WOPTJOIN outperforms BINJOIN for q_7 and q_8 . Their join plans of q_7 are nearly the same except that BINJOIN relies on a global shuffling on v_3 to processing join, while WOPTJOIN sends the partial results to the machine that maintains the vertex to grow. It is hence reasonable to observe BINJOIN’s poorer performance for q_7 as shuffling is typically a more costly operation. The case of q_8 is similar, so we do not further discuss. On the other hand, even living with costly shuffling, BINJOIN still performs better for q_9 . Due to the vertex-growing nature, WOPTJOIN’s “optimal plan” will have to process the costly sub-query $Q(\{v_1, v_2, v_3, v_4, v_5\})$. On US dataset, WOPTJOIN consistently outperforms BINJOIN for these queries. This is because that US does not produce massive intermediate results as LJ, thus BINJOIN’s shuffling cost consistently dominates.

While processing complex queries like q_8 and q_9 , we can study varieties of join plans for BINJOIN and WOPTJOIN. First of all, we want the readers to note that BINJOIN’s join plan for q_8 is different from the optimal plan originally given [34]. The original “optimal” plan computes q_8 by joining two tailed triangles (triangle tailed with an edge), while this alternative plan works better by joining the uppers “house-shape” sub-query with the bottom triangle. In theory, the tailed triangle has worse-case bound (AGM bound [39]) of $O(M^2)$, smaller than the house’s $O(M^{2.5})$, and BINJOIN’s actually favors this plan based on cost estimation. However, we find out that the number of tailed triangles is very close to that of the houses on GO, which renders costly process for the original plan to join two tailed triangles. This indicates insufficiency of both cost estimation proposed in [34] and worst-case optimal bound [14] while computing the join plan, which will be further discussed in Section 6.

Secondly, it is worth noting that we actually report the result of WOPTJOIN for q_9 while using the **CrystalJoin** plan, as it works better than WOPTJOIN’s original “optimal” plan. For q_9 , **CrystalJoin** will first compute $Q(V_Q^{cc})$, namely the 2-path $\{v_1, v_3, v_5\}$, thereafter it can compress all remaining vertices v_2, v_4 and v_6 . In comparison, the “optimal” plan can only compress v_2 and v_6 . In this case, **CrystalJoin** performs better because it configures larger compression. In [41], the

Table 2: The unlabelled datasets.

Datasets	Name	$ V_G /\text{mil}$	$ E_G /\text{mil}$	\bar{d}_G	D_G	T_{hash}/s	$ \overline{E_{\text{hash}}} /\text{mil}$	$T_{\text{tri.}}/\text{s}$	$ \overline{E_{\text{tri.}}} /\text{mil}$
google(S)	GO	0.86	4.32	5.02	6,332	1.53	0.28	2.31	1.23
gplus(S)	GP	0.11	12.23	218.2	20,127	5.57	0.80	46.5	10.68
usa-road(D)	US	23.95	28.85	2.41	9	12.43	1.89	3.69	1.90
livejournal(S)	LJ	4.85	43.37	17.88	20,333	14.25	2.81	20.33	12.49
uk2002(W)	UK	18.50	298.11	32.23	194,955	61.99	17.16	266.60	156.05
eu-road(D)	EU	173.80	342.70	3.94	20	72.96	22.47	16.98	22.98
friendster(S)	FS	65.61	1806.07	55.05	5,214	378.26	118.40	368.95	395.31
Synthetic	SY	372.00	10,000.00	53	613,461	2027	493.75	5604.00	660.61

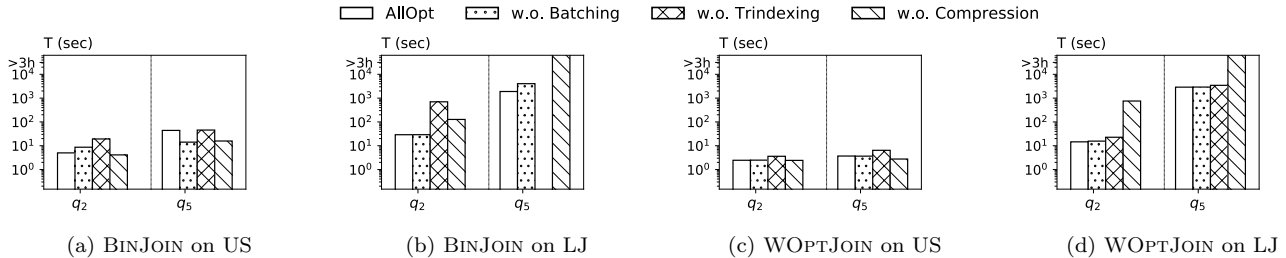


Figure 3: Effectiveness of optimizations.




Figure 4: Challenging queries.

authors proved that it renders maximum compression to use the vertex cover as the uncompressed core. However, this may not necessarily result in the best performance, considering that it can be costly to compute the core part. In our experiments, the unlabelled q_4 , q_8 and labelled q_8 are cases that CrystalJoin plan performs worse than the original BiGJoin plan (with **Compression** optimization), where CrystalJoin plan does not render strictly larger compression while having to process the costly core part. As a result, we only recommend CrystalJoin plan when it leads to strictly larger compression.

The final observation is that the computation time dominates most of the evaluated cases, except BINJOIN’s q_8 , WOPTJOIN and SHRCUBE’s q_9 on US. We will further discuss this in Exp-4.

Exp-3 All-Around Comparisons. In this experiment, we run $q_1 - q_6$ using BINJOIN, WOPTJOIN, SHRCUBE and FULLREP across the datasets GP, LJ, UK, EU and FS. We also run WOPTJOIN with CrystalJoin plan in q_4 as it is the only query that renders different CrystalJoin plan from BiGJoin plan, and the results show that the performance with BiGJoin plan is consistently better. We report the results in Figure 5, where the communication time is plotted as gray filling. As a whole, among all 35 test cases, FULLREP achieves the best 85% completion rate, followed by WOPTJOIN and BINJOIN which complete 71.4% and 68.6% respectively, and SHRCUBE performs the worst with just 8.6% completion rate.

FULLREP typically outperforms the other strategies. Observe that WOPTJOIN’s performance is often very close to FULLREP. The reason is that the WOPTJOIN’s computing plans for these evaluated queries are similar to “DualSim” adopted by FULLREP. The extra communication cost of WOPTJOIN has been reduced to very low while adopting **TrIndexing** optimization. While comparing WOPTJOIN with BINJOIN, BINJOIN is better for q_3 , a clique query (join unit) that requires no join (a case of embarrassingly parallel). BINJOIN performs worse than both FULLREP and WOPTJOIN in most other queries, which, as we mentioned before, is due to the costly shuffling. There is an exception - querying q_1 on GP - where BINJOIN performs better than both FULLREP and WOPTJOIN. We explain this using our best speculation. GP is a very dense graph, where we observe nearly 100 vertices with degree around 10,000. To process q_1 , after computing the sub-query $Q(\{v_1, v_2, v_4\})$, WOPTJOIN (and “DualSim”) processes the intersection of v_1 and v_4 (their matches) for v_3 . Those larger-degree vertices are now frequently pairing, leading to expensive intersection. In comparison, BINJOIN computes q_1 by joining the sub-query $Q(\{v_1, v_2, v_3\})$ with $Q(\{v_1, v_3, v_4\})$. Because both strategies compute $Q(\{v_1, v_2, v_3\})$, we consider how BINJOIN computes $Q(\{v_1, v_3, v_4\})$. BINJOIN first locate the matched vertex of v_3 , then matches v_1 and v_4 among its neighbors, which is generally cheaper than intersecting the neighbors of v_1 and v_4 to compute v_3 . Due to the existence of these high-degree pairs, the cost WOPTJOIN’s intersection can exceed BINJOIN’s shuffling.

We observe that the computation time T_{comp} dominates in most cases as we mentioned in “Exp-3”, with details listed in Table 3 for all queries using BINJOIN and WOPTJOIN on LJ. This is trivially true for SHRCUBE and FULLREP, but it may not be clearly so for WOPTJOIN and BINJOIN given that they all need to transfer a massive amount of intermediate data. We investigate this and find out two potential reasons. The first one attributes to **Timely**’s highly optimized communication component, which allows the computation to overlap communication by using extra threads to receive and buffer the data from the network so that it can be mostly ready for the following computation. The second one is the fast network. We re-run these queries using

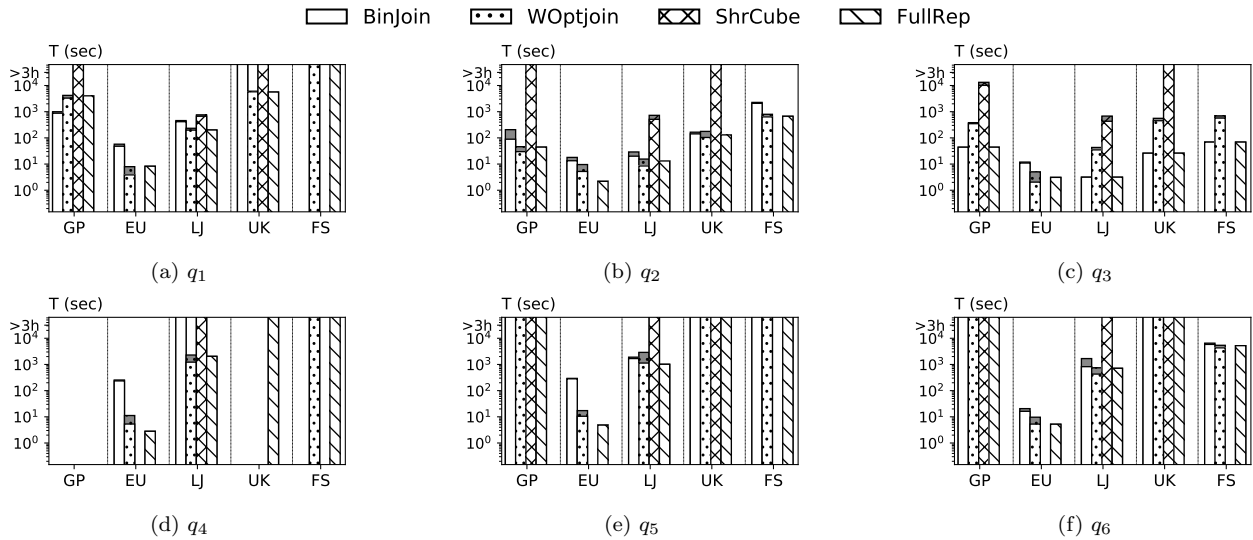


Figure 5: All-around comparisons.

Table 3: Computation time vs. communication time (all queries using BINJOIN and WOPTJOIN on LJ).

Queries	BINJOIN/s		WOPTJOIN/s	
	T_{comp}	T_{comm}	T_{comp}	T_{comm}
q_1	415	43	197	35
q_2	20	9	9	6
q_3	3	0	34	9
q_4	0T	0T	1522	769
q_5	1022	654	428	314
q_6	95	0	1511	332

the 1GBps switch, while the results show the opposite trend that the communication time T_{comm} in turn takes over.

Table 4: The web-scale experiments.

Queries	BINJOIN (T_{comp})/s	WOPTJOIN (T_{comp})/s
q_2	8810 (6893)	1751 (1511)
q_3	76 (75)	518 (443)

Exp-4 Web-Scale. We run the SY datasets in the AWS cluster of 40 instances. Note that FULLREP can not be used as SY is larger than the machine’s memory. We use the queries q_2 and q_3 , and present the results of BINJOIN and WOPTJOIN (SHRCUBE fails all cases due to OOM) in Table 4. The results are consistent with the prior experiments, but observe that the gap between BINJOIN and WOPTJOIN while querying q_1 is larger. This is because that we now deploy 40 AWS instances, and BINJOIN’s shuffling cost increases.

5.3 Labelled Experiments

We use the LDBC social network benchmarking (SNB) [7] for labelled matching experiment due to the lack of labelled big graphs in the public. SNB provides a data generator that generates a synthetic social network of required statistics, and a document [8] that describes the benchmarking tasks, in which the complex tasks are actually subgraph matching. The join plans for BINJOIN and WOPTJOIN for labelled experiments are generated as unlabelled case, but we use the label frequencies to break tie.

Table 5: The labelled datasets.

Name	$ V_G $	$ E_G $	\bar{d}_G	D_G	# Labels
DG10	29.99	176.48	11.77	4,282,812	10
DG60	187.11	1246.66	13.32	26,639,563	10

Datasets. We list the datasets and their statistics in Table 5. These datasets are generated using the “Facebook” mode with a duration of 3 years. The dataset’s name, denoted as DG_x , represents a scale factor of x . The labels are preprocessed into integers.




Figure 6: Labelled queries.

Queries. The queries, shown in Figure 6, are selected from the SNB’s complex tasks with some adaptations, and the details are in the full paper.

Exp-5 All-Around Comparisons. We now conduct the experiment using all queries on DG10 and DG60, and present the results in Figure 7. Here we compute the join plans for BINJOIN and WOPTJOIN by using the unlabelled method, but further using the label frequencies to break tie. The gray filling again represents communication time. FULLREP outperforms the other strategies in many cases, except that it performs slightly slower than BINJOIN for q_3 and q_5 . This is because that q_3 and q_5 are join units, and BINJOIN processes them locally in each machine as FULLREP, and it does not build indices as “CFLMatch” used in FULLREP. When comparing to WOPTJOIN, Among all these queries, we only have q_8 that configures different CrystalJoin plan (w.r.t. BiGJoin plan) for WOPTJOIN. The results show that the performance of WOPTJOIN drops

about 10 times while using CrystalJoin plan. Note that the core part of q_8 is a 5-path of “Psn-City-Cty-City-Psn” with enormous intermediate results. As we mentioned in unlabelled experiments, it may not always be wise to first compute the vertex-cover-induced core.

We now focus on comparing BINJOIN and WOPTJOIN. There are three cases that intrigue us. Firstly, observe that BINJOIN performs much better than WOPTJOIN while querying q_4 . The reason is high intersection cost as we discovered on GP dataset in unlabelled matching. Secondly, BINJOIN performs worse than WOPTJOIN in q_7 , which again is because of BINJOIN’s costly shuffling. The third case is q_9 , the most complex query in the experiment. BINJOIN performs much better while querying q_9 . The bad performance of WOPTJOIN comes from the long execution plan together with costly intermediate results. The two algorithms all expand the three “Psn”s, and then grow via one of the “City”s to “Cty”, but BINJOIN approaches this using one join (a triangle \bowtie a TwinTwig), while WOPTJOIN will first expand to “City” then further “Cty”, and the “City” expansion is the culprit of the slower run.

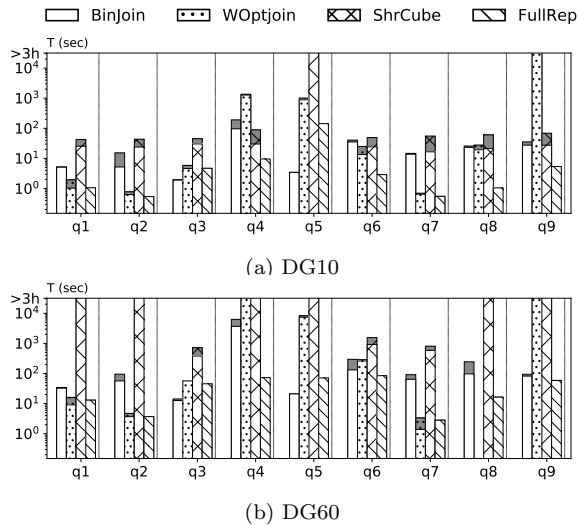


Figure 7: All-around comparisons of labelled matching.

6. DISCUSSIONS AND FUTURE WORK.

We discuss our findings and potential future work based on the experiments in Section 5. Eventually, we summarize the findings into a practical guide.

Strategy Selection. FULLREP is obviously the preferred choice when the machine can hold the graph data, while both WOPTJOIN and BINJOIN are good alternatives when the graph is larger than the capacity of the machine. For BINJOIN and WOPTJOIN, on one side, BINJOIN may perform worse than WOPTJOIN (e.g. unlabelled q_2 , q_4 , q_5) due to the expensive shuffling operation, on the other side, BINJOIN can also outperform WOPTJOIN (e.g. unlabelled and labelled q_9) while avoiding costly sub-queries due to query decomposition. One way to choose between BINJOIN and WOPTJOIN is to compare the cost of their respective join plans, and select the one with less cost. For now, we can either use cost estimation proposed in [34], or summing the worst-case bound, but none of them consistently gives the best solution, as will be discussed in “Optimal Join Plan”. Alternatively, we refer to “EmptyHeaded” [12] to study a potential hybrid strategy of BINJOIN and WOPTJOIN. Note

that “EmptyHeaded” is developed in single-machine setting, and it does not take into consideration the impact of Compression, we hence leave such hybrid strategy in the distributed context as an interesting future work.

Optimizations. Our experimental results suggest always using *Batching*, using *TrIndexing* when each machine has sufficient memory to hold “triangle partition”, and using *Compression* when the data graph is not very sparse (e.g. $\bar{d}_G \geq 5$). *Batching* often does not impact performance, so we recommend always using *Batching* due to the unpredictability of the size of (intermediate) results. *TrIndexing* is critical for BINJOIN, and it can greatly improve WOPTJOIN by reducing communication cost, while it requires extra storage to maintain “triangle partition”. Amongst the evaluated datasets, each “triangle partition” maintains an average of 30% data in our 10-machine cluster. Thus, we suggest a memory threshold of $60\%|E_G|$ (half for graph and half for running algorithm) for *TrIndexing* in a cluster of the same or larger scale. Note that the threshold does not apply to extremely dense graph. Among the three optimizations, *Compression* is the primary performance booster that improves the performance of BINJOIN and WOPTJOIN by 5 times on average in all but the cases on the very sparse road networks. For such very sparse data graphs, *Compression* can render more cost than benefits.

Optimal Join Plan. It is challenging to systematically determine the optimal join plans for both BINJOIN and WOPTJOIN. From the experiments, we identify three impact factors: (1) the worst-case bound; (2) cost estimation based on data statistics; (3) favoring the optimizations, especially *Compression*. All existing works only partially consider these factors, and we have observed sub-optimal join plans in the experiments. For example, BINJOIN bases the “optimal” join plan on minimizing the cost estimation, but the join plan does not render the best performance for unlabelled q_8 ; WOPTJOIN follows the worst-case optimality, while it may encounter costly sub-queries for labelled and unlabelled q_9 ; CrystalJoin focuses on maximizing the compression, while ignoring the facts that the vertex-cover-induced core part itself can be costly to compute. Additionally, there are other impact factors such as the partial orders of query vertices and the label frequencies, which have not been studied in this work due to short of space. It is another very interesting future work to thoroughly study the optimal join plan while considering all above impact factors.

Computation vs. Communication. We argue that distributed subgraph matching nowadays is a computation-intensive task. This claim holds when the cluster configures high-speed network (e.g. ≥ 10 Gbps), and the data processor can efficiently overlap computation with communication. Note that computation cost (either BINJOIN’s join or WOPTJOIN’s intersection) is lower-bounded by the output size that is equal to the communication cost. Therefore, computation becomes the bottleneck if the network condition is good to guarantee the data to be delivered in time. Nowadays, the bandwidth of local cluster commonly exceeds 10Gbps, and the overlapping of computation and communication is widely used in distributed systems (e.g. Spark [52], Flink [16]). As a result, we tend to see distributed subgraph matching as a computation-intensive task, and we advocate future research to devote more efforts into optimizing the computation while considering the following perspectives: (1) the new advancements of hardware, for example the co-processing on GPU in the coupled CPU-GPU architectures

Table 6: All Query’s Results.

Dataset	q_1	q_2	q_3	q_4	q_5	q_6	q_7	q_8	q_9
GO	539.58M	621.18M	39.88M	38.20B	27.80B	9.28B	2,168.86B	330.68B	1.88T
GP	1.42T	1.16T	78.40B	-	-	-	-	-	-
US	1.61M	21,599	90	117,996	2,186	1	160.93M	2,891	89
LJ	51.52B	76.35B	9.93B	53.55T	44.78T	18.84T	-	-	-
UK	2.49T	2.73T	157.19B	-	-	-	-	-	-
EU	905,640	2,223	6	12,790	450	0	342.48M	436	71
FS	-	185.19B	8.96B	-	-	3.18T	-	-	-
SY	-	834.78B	5.47B	-	-	-	-	-	-
DG10*	40.14M	26.76M	28.73M	22.59M	23.08B	1.49M4	47,556	42.56M	10.07M
DG60*	302.41M	169.86M	267.38M	161.69M	203.33B	12.44M	983,370	4.14B	114.19M

[26] and the SIMD programming model on modern CPU [28]; (2) general computing optimizations such as load balancing strategy and cache-aware graph data accessing [51].

A Practical Guide. Based on the experimental findings, we propose a practical guide for distributed subgraph matching in Figure 8. Note that this program guide is based on current progress of the literature, and future work is needed, for examples to study the hybrid strategy and the impact factors of the optimal join plan, before we can arrive at a solid decision-making to choose between BINJOIN and WOPTJOIN.

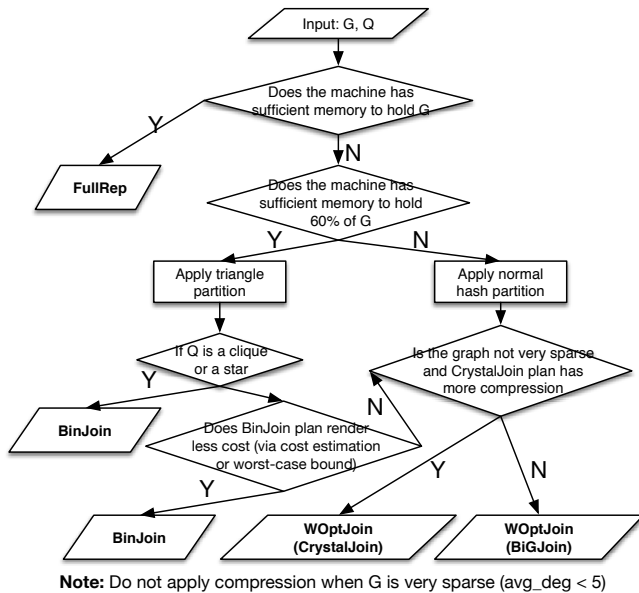


Figure 8: A practical guide of distributed subgraph matching.

7. RELATED WORK

Isomorphism-based Subgraph Matching. In the labelled case, Shang et al. [45] used the spanning tree of the query graph to filter infeasible results. Han et al. [25] observed the importance of matching order. In [43], the authors proposed to utilize the symmetry properties in the data graph to compress the results. Bi et al. [15] proposed CFLMatch based on the “core-forest-leaves” matching order, and obtained performance gain by postponing the notorious cartesian product.

The unlabelled case is also known as subgraph listing/enumeration, and due to the gigantic (intermediate) results, people have been either seeking scalable algorithms

in parallel, or devising techniques to compress the results. Other than the algorithms studied in this paper (Section 3), Kim et al. proposed the external-memory-based parallel algorithm DUALSIM [31], which maintains the data graph in blocks on the disk, and matches the query graph by swapping in/out blocks of data to improve I/O efficiency.

Incremental Subgraph Matching. Computing subgraph matching in a continuous context has recently drawn a lot of attentions. Fan et al. [23] proposed incremental algorithm that identifies a portion of the data graph affected by the update regarding the query. The authors in [18] used the join scheme as BINJOIN algorithms (Section 3.1). The algorithm maintained a left-deep join tree for the query, with each vertex maintaining a partial query and the corresponding partial results. Then one can compute the incremental answers of each partial query in response to the update, and utilizes the join tree to re-construct the results. Graphflow [30] solved incremental subgraph matching using join, in the sense that the incremental query can be transformed into m independent joins, where m is the number of query edges. Then they used the worst-case-optimal join algorithm to solve these joins in parallel. Most recently, Kim et al. proposed TURBOFLUX that maintains data-centric index for incremental queries, which achieves good tradeoff between performance and storage.

8. CONCLUSIONS

In this paper, we implement four strategies and three general-purpose optimizations for distributed subgraph matching based on Timely dataflow system, aiming for a systematic, strategy-level comparison. Based on thorough empirical analysis, we summarize a practical guide, and we also motivate interesting future work for distributed subgraph matching.

9. AUXILIARY MATERIALS

All Query Results. In Table 6, We show the number of results of every successful query on each dataset evaluated in this work, where M , B , T stand for millions, billions and trillions, respectively. Note that DG10 and DG60 record the labelled queries of $q_1 - q_9$.

10. ACKNOWLEDGEMENTS

Xuemin Lin is supported by 2019DH0ZX01, 2018-YFB1003504, NSFC61232006, ARC DP180103096 and DP170101628. Lu Qin is supported by ARC DP160101513. Wenjie Zhang is supported by ARC DP180103096. Ying Zhang is supported by FT170100128 and DP180103096.

11. REFERENCES

- [1] The challenge9 datasets. <http://www.dis.uniroma1.it/challenge9>.
- [2] The clubweb12 dataset. <https://lemurproject.org/clueweb12>.
- [3] Compressed sparse row. https://en.wikipedia.org/wiki/Sparse_matrix.
- [4] Distributed subgraph matching on timely dataflow - the full paper. <https://goo.gl/zkTkL4>.
- [5] Giraph. <http://giraph.apache.org/>.
- [6] The implementation of bigjoin. <https://github.com/frankmcsherry/dataflow-join/>.
- [7] Ldbc benchmarks. <http://ldbcouncil.org/benchmarks>.
- [8] The ldbc social network benchmark. https://ldbc.github.io/ldbc_snb_docs/ldbc_snb_specification.pdf.
- [9] The open-sourced timely dataflow system. <https://github.com/frankmcsherry/timely-dataflow>.
- [10] The snap datasets. <http://snap.stanford.edu/data/index.html>.
- [11] The webgraph datasets. <http://law.di.unimi.it/datasets.php>.
- [12] C. R. Aberger, S. Tu, K. Olukotun, and C. Ré. Emptyheaded: A relational engine for graph processing. SIGMOD '16, pages 431–446.
- [13] F. N. Afrati, D. Fotakis, and J. D. Ullman. Enumerating subgraph instances using map-reduce. In *Proc. of ICDE'13*, 2013.
- [14] K. Ammar, F. McSherry, S. Salihoglu, and M. Joglekar. Distributed evaluation of subgraph queries using worst-case optimal low-memory dataflows. *PVLDB*, 11(6):691–704, 2018.
- [15] F. Bi, L. Chang, X. Lin, L. Qin, and W. Zhang. Efficient subgraph matching by postponing cartesian products. SIGMOD '16, pages 1199–1214, 2016.
- [16] P. Carbone, A. Katsifodimos, . Kth, S. Sweden, S. Ewen, V. Markl, S. Haridi, and K. Tzoumas. Apache flink: Stream and batch processing in a single engine. 38, 01 2015.
- [17] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-mat: A recursive model for graph mining. In *SDM*, 2004.
- [18] S. Choudhury, L. B. Holder, G. Chin, K. Agarwal, and J. Feo. A selectivity based approach to continuous pattern detection in streaming graphs. In *EDBT*, 2015.
- [19] S. Chu, M. Balazinska, and D. Suciu. From theory to practice: Efficient join query evaluation in a parallel database system. SIGMOD '15, pages 63–78.
- [20] F. R. K. Chung, L. Lu, and V. H. Vu. The spectra of random graphs with given expected degrees. *Internet Mathematics*, 1(3), 2003.
- [21] D. J. DeWitt and J. Gray. Parallel database systems: The future of database processing or a passing fad? *SIGMOD Rec.*, 19(4):104–112.
- [22] P. Erdos and A. Renyi. On the evolution of random graphs. In *Publ. Math. Inst. Hungary. Acad. Sci.*, 1960.
- [23] W. Fan, J. Li, J. Luo, Z. Tan, X. Wang, and Y. Wu. Incremental graph pattern matching. SIGMOD '11, pages 925–936, 2011.
- [24] J. A. Grochow and M. Kellis. Network motif discovery using subgraph enumeration and symmetry-breaking. In *Proc. of RECOMB'07*, 2007.
- [25] W.-S. Han, J. Lee, and J.-H. Lee. Turboiso: Towards ultrafast and robust subgraph isomorphism search in large graph databases. In *Proc. of SIGMOD'13*, 2013.
- [26] J. He, S. Zhang, and B. He. In-cache query co-processing on coupled cpu-gpu architectures. *PVLDB*, 8(4):329–340, 2014.
- [27] M. Herlihy and N. Shavit. *The Art of Multiprocessor Programming*. 2008.
- [28] H. Inoue, M. Ohara, and K. Taura. Faster set intersection with simd instructions by reducing branch mispredictions. *PVLDB*, 8(3):293–304, 2014.
- [29] Y. E. Ioannidis and Y. C. Kang. Left-deep vs. bushy trees: An analysis of strategy spaces and its implications for query optimization. In *SIGMOD'91*, pages 168–177, 1991.
- [30] C. Kankanamge, S. Sahu, A. Mhedbhi, J. Chen, and S. Salihoglu. Graphflow: An active graph database. SIGMOD '17, pages 1695–1698.
- [31] H. Kim, J. Lee, S. S. Bhowmick, W.-S. Han, J. Lee, S. Ko, and M. H. Jarrah. Dualsim: Parallel subgraph enumeration in a massive graph on a single machine. SIGMOD '16, pages 1231–1245, 2016.
- [32] L. Lai, L. Qin, X. Lin, and L. Chang. Scalable subgraph enumeration in mapreduce. *PVLDB*, 8(10):974–985, 2015.
- [33] L. Lai, L. Qin, X. Lin, and L. Chang. Scalable subgraph enumeration in mapreduce: A cost-oriented approach. *The VLDB Journal*, 26(3):421–446, June 2017.
- [34] L. Lai, L. Qin, X. Lin, Y. Zhang, L. Chang, and S. Yang. Scalable distributed subgraph enumeration. *PVLDB*, 10(3):217–228, 2016.
- [35] B. Liu, L. Yuan, X. Lin, L. Qin, W. Zhang, and J. Zhou. Efficient (α, β) -core computation: An index-based approach. WWW, pages 1130–1141, 2019.
- [36] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein. Distributed graphlab: A framework for machine learning and data mining in the cloud. *PVLDB*, 5(8):716–727, 2012.
- [37] F. McSherry, M. Isard, and D. G. Murray. Scalability! but at what cost? HOTOS'15, 2015.
- [38] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: A timely dataflow system. SOSP '13, pages 439–455, 2013.
- [39] H. Q. Ngo, E. Porat, C. Ré, and A. Rudra. Worst-case optimal join algorithms. *J. ACM*, 65(3), 2018.
- [40] D. Olteanu and M. Schleich. Factorized databases. *SIGMOD Rec.*, 45(2):5–16, Sept. 2016.
- [41] M. Qiao, H. Zhang, and H. Cheng. Subgraph matching: On compression and computation. *PVLDB*, 11(2):176–188, 2017.
- [42] X. Qiu, W. Cen, Z. Qian, Y. Peng, Y. Zhang, X. Lin, and J. Zhou. Real-time constrained cycle detection in large dynamic graphs. *PVLDB*, 11(12):1876–1888.
- [43] X. Ren and J. Wang. Exploiting vertex relationships in speeding up subgraph isomorphism over large graphs. *PVLDB*, 8(5):617–628, 2015.
- [44] R. Shamir and D. Tsur. Faster subtree isomorphism. In *Proceedings of the Fifth Israeli Symposium on Theory of Computing and Systems*, pages 126–131, 1997.
- [45] H. Shang, Y. Zhang, X. Lin, and J. X. Yu. Taming verification hardness: An efficient algorithm for testing subgraph isomorphism. *PVLDB*, 1(1):364–375, 2008.

- [46] B. Shao, H. Wang, and Y. Li. Trinity: A distributed graph engine on a memory cloud. SIGMOD '13, pages 505–516, 2013.
- [47] Y. Shao, B. Cui, L. Chen, L. Ma, J. Yao, and N. Xu. Parallel subgraph listing in a large-scale graph. In *SIGMOD'14*, pages 625–636. ACM, 2014.
- [48] Z. Sun, H. Wang, H. Wang, B. Shao, and J. Li. Efficient subgraph matching on billion node graphs. *PVLDB*, 5(9):788–799, 2012.
- [49] T. L. Veldhuizen. Triejoin: A simple, worst-case optimal join algorithm. In *ICDT*, 2014.
- [50] K. Wang, X. Lin, L. Qin, W. Zhang, and Y. Zhang. Vertex priority based butterfly counting for large-scale bipartite networks. *arXiv preprint arXiv:1812.00283*, 2018.
- [51] H. Wei, J. X. Yu, C. Lu, and X. Lin. Speedup graph processing by graph ordering. SIGMOD '16, pages 1813–1828.
- [52] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: Cluster computing with working sets. In *HotCloud'10*, pages 10–10.