

Reachability Labeling for Distributed Graphs

Junhua Zhang[‡], Wentao Li[‡], Lu Qin[‡], Ying Zhang[‡], Dong Wen[§], Lizhen Cui[‡], and Xuemin Lin[§]

[‡]AAIL, FEIT, University of Technology Sydney; [§]The University of New South Wales; [‡]Shandong University

junhua.zhang@student.uts.edu.au; {wentao.li, lu.qin, ying.zhang}@uts.edu.au
dong.wen@unsw.edu.au; lxue@cse.unsw.edu.au; clz@sdu.edu.cn

Abstract—Real-world graphs are typically distributed across multiple data centers. When performing reachability queries on these distributed graphs, reachability labeling methods ensure fast query processing by using lightweight indexes. One of the best-known labeling methods is TOL; however, TOL is a *serial* algorithm and cannot handle distributed graphs. The main goal of this paper is to design new labeling methods that can work in parallel while producing the *same* index as TOL. To this end, we investigate the limitation of TOL and thus propose a filtering-and-refinement framework for index creation. This framework first obtains a super-set of each vertex’s label sets and then eliminates the invalid elements. Based on this framework, we design distributed labeling algorithms and then use batch processing to improve efficiency. Experimental results on real-world graphs show that the proposed algorithms can index distributed graphs efficiently.

I. INTRODUCTION

Graphs are used widely to represent real-world entities and their relationships [1]. As a common graph operation, reachability query $q(s, t)$ asks whether there is a path from vertex s to vertex t . Reachability query plays an important role in applications such as geographic navigation, Internet routing, and ontology reasoning [2]. Moreover, it is a building block for fields such as social sciences, computational biology, and software engineering [3].

Due to the importance of reachability query processing, many approaches have been proposed. These approaches are usually classified into three categories [2]: 1) index-free approaches that use the online search [4] (e.g., breadth-first search or depth-first search) on graphs to determine whether one vertex can reach another [5]; 2) index-assisted approaches that accelerate the online search using auxiliary information [6]–[9]; 3) index-only approaches that *avoid* the online search using an offline index [10]–[14].

Most current approaches are centralized: they assume that graphs reside in the main memory [15]. However, as the graph size increases, real-world graphs are typically distributed in multiple data centers [16]. When performing reachability queries on distributed graphs, the pioneer work [15], [16] implements the online search in a distributed manner for query processing. However, the query latency can be high due to the need to access distributed graphs during the query. This makes methods that require the online search (i.e., index-free or index-assisted approaches) undesirable, especially when a large number of queries need to be processed.

To speed up query processing on distributed graphs, an alternative idea is to use index-only approaches — the index created offline eliminates the dependence on the original graph during querying [14]. Specifically, consider a graph $G(V, E)$ with vertex set V and edge set E , index-only approaches create an index that assigns an out-label set $L_{\text{out}}(v)$ and an in-label set $L_{\text{in}}(v)$ for each vertex $v \in V$. The out-label set $L_{\text{out}}(v)$ of v contains vertices that

v can reach, while the in-label set $L_{\text{in}}(v)$ contains vertices that can reach v . To answer the query $q(s, t)$ between vertices s and t , we *only* need to check whether there exists a common vertex w between $L_{\text{out}}(s)$ and $L_{\text{in}}(t)$: the existence of w means that s can reach t via w .

The state-of-the-art index-only method is Total Order Labeling (TOL) [14]. TOL assigns an order value to each vertex in a graph, and then selects vertices for labeling in a decreasing sequence of order values. When labeling a vertex v , TOL tries to add v to the in-label/out-label sets of other vertices, using a **pruning operation**. When all vertices have finished labeling, the in-label/out-label sets generated by TOL for each vertex are used as an index.

The pruning operation of TOL is a double-edged sword: TOL reduces the index size by eliminating redundancy through the pruning operation; however, each vertex v needs to wait for vertices whose order values are higher than v to finish labeling, thus making the pruning operation of v feasible. This means that the execution of TOL is inherently serial. In other words, TOL *does not work properly on distributed graphs*.

On the other hand, for a distributed graph, the index created by TOL can efficiently support queries. This is because the index of TOL is small enough that we can put it on a single machine to achieve fast in-memory queries. For example, the index size for graph SK (see Table V for graph’s details) with billions of edges is bounded by 1 GB. The main purpose of our paper is to design new labeling methods to handle a distributed graph while obtaining the same index as TOL.

For this purpose, we delve into the labeling process of TOL. We find that when labeling a vertex v , v joins the label sets of some specific vertices. These vertices are defined as the **backward label set** of v . The working process of TOL can be equated to finding the backward label set for each vertex v . To find v ’s backward label set, we use a filtering-and-refinement framework, thereby avoiding the pruning operation used by TOL. Specifically, we first generate a super-set of the backward label set of v as candidates, and then remove invalid elements from candidates to obtain the actual backward label set. This novel framework gets the same index as TOL while letting all vertices run in parallel. Based on this framework, we design efficient labeling algorithms and provide distributed implementations. In addition, we split vertices into batches for labeling to further improve efficiency.

The contributions of this paper are summarized as follows.

- *Analysis of TOL’s limitation (Section II)*. We investigate TOL’s limitation, i.e., the pruning operation of TOL makes parallel work challenging. This motivates the design of new labeling methods.
- *Novel labeling algorithms (Section III)*. We find that each vertex’s labeling process can be replaced by finding the backward label set for that vertex. We use a filtering-and-refinement

Wentao Li is the corresponding author.

framework to find the backward label set of each vertex in parallel. Using this framework, we propose new labeling algorithms and provide implementations in a distributed system.

- *Batch labeling optimization (Section IV)*. To further improve the labeling algorithms' efficiency, we split vertices into batches and then construct the index in batches.
- *Extensive empirical studies (Section VI)*. We conduct numerous experiments to validate the efficiency of the proposed labeling algorithms. On medium-sized graphs, our algorithms can outperform TOL by nearly an order of magnitude. Furthermore, on billion-sized graphs, we can create indexes in half an hour while TOL cannot.

II. PRELIMINARY

We first introduce some notations in Section II-A, then give the labeling algorithm TOL in Section II-B, followed by the problem statement in Section II-C. For ease of understanding, Table I lists frequently used notations.

A. Notations

Given a directed graph $G = (V, E)$ with $n = |V|$ vertices and $m = |E|$ edges, we define the in-neighbor set $N_{in}^G(v)$ of a vertex $v \in V$ as $N_{in}^G(v) = \{u | (u, v) \in E\}$; the out-neighbor set $N_{out}^G(v)$ as $N_{out}^G(v) = \{u | (v, u) \in E\}$. The **in-degree** $d_{in}^G(v)$ (resp. the **out-degree** $d_{out}^G(v)$) of v is the size of its in-neighbor set (resp. out-neighbor set), i.e., $d_{in}^G(v) = |N_{in}^G(v)|$ (resp. $d_{out}^G(v) = |N_{out}^G(v)|$). The path between a vertex pair $s, t \in V$ is defined as $p^G(s, t) = (v_1 = s, v_2, \dots, v_l = t)$, where $(v_i, v_{i+1}) \in E$, for $\forall i \in [1, l-1]$. If there exists a path between s and t , then s can reach t (denoted as $s \rightarrow t$).

Definition 1. The **ancestors** $ANC(v)$ (resp. **descendants** $DES(v)$) of a vertex $v \in V$ contain all the vertices that can reach v (resp. that v can reach). Vertex v is contained both in $ANC(v)$ and $DES(v)$.

If the context is obvious, we drop G from notations. The **inverse graph** $\bar{G} = (V, \bar{E})$ of a graph $G(V, E)$ contains the same vertices but with all edges reversed in direction (i.e., $\bar{E} = \{(v, u) | (u, v) \in E(G)\}$).

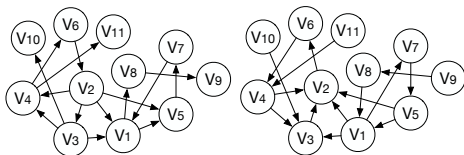


Fig. 1: Graph G

Fig. 2: Inverse Graph \bar{G}

Example 1. Consider the graph G in Fig. 1, which has 11 vertices and 15 edges. For vertex v_2 , $N_{in}(v_2) = \{v_6\}$ and $d_{in}(v_2) = 1$; $N_{out}(v_2) = \{v_1, v_3, v_4, v_5\}$ and $d_{out}(v_2) = 4$. The vertex v_2 can reach vertex v_7 since there exists a path from v_2 to v_7 . For v_2 , $ANC(v_2) = \{v_2, v_3, v_4, v_6\}$ and $DES(v_2) = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8, v_9, v_{10}, v_{11}\}$. The inverse graph \bar{G} of G is shown in Fig. 2.

The in-label/out-label sets over all vertices form index L , which can be used to answer reachability queries in graph G .

Definition 2. The **in-label set** $L_{in}(v)$ of v contains vertices that can reach v , i.e., $L_{in}(v) \subseteq ANC(v)$; the **out-label set** $L_{out}(v)$ of v contains vertices that v can reach, i.e., $L_{out}(v) \subseteq DES(v)$.

TABLE I: Notations

Notation	Meaning
$G(V, E)$	graph
$\bar{G}(V, \bar{E})$	inverse graph
$d_{in}(v), d_{out}(v)$	in-degree, out-degree of v
$ANC(v), DES(v)$	ancestors, descendants of v
$L_{in}(v), L_{out}(v)$	in-label, out-label sets of v
$L_{in}^-(v), L_{out}^-(v)$	backward in-label, out-label sets of v
$DES_{high}(v)$	high-order descendants of v
$BFS_{low}(v), BFS_{high}(v)$	low-order, high-order vertices in trimmed BFS
$IBFS_{low}(v)$	inverted list of v
$[V_1, V_2, \dots, V_g]$	batch sequence of G
$L_{in}^{V_i}, L_{out}^{V_i}$	batch in-label, out-label sets regarding V_i

TABLE II: The Index L

Vertex	L_{in}	L_{out}
v_1	$\{v_1\}$	$\{v_1\}$
v_2	$\{v_2\}$	$\{v_1, v_2\}$
v_3	$\{v_2\}$	$\{v_1, v_2\}$
v_4	$\{v_2\}$	$\{v_1, v_2\}$
v_5	$\{v_1\}$	$\{v_1\}$
v_6	$\{v_2\}$	$\{v_1, v_2\}$
v_7	$\{v_1\}$	$\{v_1\}$
v_8	$\{v_1, v_8\}$	$\{v_8\}$
v_9	$\{v_1, v_8, v_9\}$	$\{v_9\}$
v_{10}	$\{v_2, v_{10}\}$	$\{v_{10}\}$
v_{11}	$\{v_2, v_{11}\}$	$\{v_{11}\}$

Given an index L , the largest label size, denoted as Δ , is defined as $\Delta = \max_{v \in V} (\max(|L_{in}(v)|, |L_{out}(v)|))$. To answer the reachability query $q(s, t)$ between s, t , we check whether there are overlapping vertices between $L_{out}(s)$ and $L_{in}(t)$.

$$q(s, t) = \begin{cases} true, & \text{if } L_{out}(s) \cap L_{in}(t) \neq \emptyset; \\ false, & \text{otherwise.} \end{cases}$$

If the vertices in $L_{out}(s)$ and $L_{in}(t)$ are sorted by IDs, answering $q(s, t)$ takes $O(|L_{out}(s)| + |L_{in}(t)|)$ time [14]. To ensure that index L correctly answers all reachability queries on G , L needs to satisfy the cover constraint.

Definition 3 (Cover Constraint). For $\forall s, t \in V$, $L_{out}(s) \cap L_{in}(t) \neq \emptyset$ if and only if $(\Leftrightarrow) s \rightarrow t$.

Example 2. Consider the graph G in Fig. 1, where Table II lists an index L for G . For v_2 , $L_{out}(v_2) = \{v_1, v_2\} \subseteq DES(v_2)$; for v_3 , $L_{in}(v_3) = \{v_2\} \subseteq ANC(v_3)$. Since $L_{out}(v_2) \cap L_{in}(v_3) = \{v_2\} \neq \emptyset$, the query $q(v_2, v_3)$ returns "true".

B. Total Order Labeling

Many labeling methods have been proposed to create reachability indexes for graphs [2], and one of the best-known is Total Order Labeling (TOL). TOL works for a total of n rounds, where one vertex is selected for labeling in each round. TOL gives each vertex v an order $ord(v)$ and selects the vertex with the i -th largest order in round i . TOL uses degree to determine the order and vertex IDs to break the tie: we can define $ord(v) = (d_{in}(v) + 1) \cdot (d_{out}(v) + 1) + \frac{ID(v)}{n+1}$ for a vertex v , where $ID(v)$ is the ID of v . There are other ways to define $ord(v)$, but this way is cheap to calculate and works well in practice [3], [17].

Example 3. Consider the graph G in Fig. 1, which has $n = 11$ vertices. For v_1 , $ord(v_1) = (d_{in}(v_1) + 1) \cdot (d_{out}(v_1) + 1) + \frac{1}{12} = 12.08$; for v_{10} , $ord(v_{10}) = (d_{in}(v_{10}) + 1) \cdot (d_{out}(v_{10}) + 1) + \frac{10}{12} = 2.83$. Thus, $ord(v_1) > ord(v_{10})$, which means that the order of v_1 is higher than v_{10} .

Algorithm 1: TOL

Input: Graph $G(V, E)$
Output: Index L

- 1 $L_{in}^1(v) \leftarrow \emptyset, L_{out}^1(v) \leftarrow \emptyset$, for each vertex $v \in V$;
- 2 $G_1 \leftarrow G$;
- 3 **foreach** $i \in [1, n]$ **do**
- 4 $v_i \leftarrow$ the vertex whose order is the i -th largest;
- 5 $DES^{G_i}(v_i) \leftarrow$ a v_i -sourced BFS on G_i ;
- 6 $ANC^{G_i}(v_i) \leftarrow$ a v_i -sourced BFS on $\overline{G_i}$;
- 7 **foreach** $w \in DES^{G_i}(v_i)$ **do**
- 8 // pruning operation
- 9 **if** $L_{out}^i(v_i) \cap L_{in}^i(w) = \emptyset$ **then**
- 10 $L_{in}^{i+1}(w) \leftarrow L_{in}^i(w) \cup \{v_i\}$;
- 11 **foreach** $w \in ANC^{G_i}(v_i)$ **do**
- 12 // pruning operation
- 13 **if** $L_{in}^i(v_i) \cap L_{out}^i(w) = \emptyset$ **then**
- 14 $L_{out}^{i+1}(w) \leftarrow L_{out}^i(w) \cup \{v_i\}$;
- 15 $G_{i+1} \leftarrow G_i \setminus \{v_i\}$;
- 16 **return** $L = \{L_{in}^{n+1}(v) \cup L_{out}^{n+1}(v) | v \in V\}$;

TOL Algorithm. Algorithm 1 shows how TOL creates an index L for graph G . We first initialize the in-label set $L_{in}^1(v)$ and out-label set $L_{out}^1(v)$ of all vertices v to an empty set (Line 1), and then copy G to G_1 (Line 2). Here, $L_{in}^i(v)$ (resp. $L_{out}^i(v)$) refers to the label sets created by vertices $[v_1, v_2, \dots, v_{i-1}]$ of order higher than v_i . The labeling process works in n rounds (Line 3). In round i , the vertex v_i with the i -th largest order starts labeling (Line 4).

Labeling v_i (Line 5-11). First, the descendants $DES^{G_i}(v_i)$ and ancestors $ANC^{G_i}(v_i)$ of v_i are obtained using the v_i -sourced BFS in G_i and $\overline{G_i}$, respectively (Line 5-6). Then, v_i is added to in-label/out-label sets of other vertices when it passes a **pruning operation**: for each vertex $w \in DES^{G_i}(v_i)$, when $L_{out}^i(v_i)$ and $L_{in}^i(w)$ have no overlapping, v_i is appended to $L_{in}^i(w)$ to form $L_{in}^{i+1}(w)$ (Line 8); For each vertex $w \in ANC^{G_i}(v_i)$, when $L_{in}^i(v_i)$ and $L_{out}^i(w)$ have no overlapping, v_i is appended to $L_{out}^i(w)$ to form $L_{out}^{i+1}(w)$ (Line 10). Then, v_i and the incident edges are removed from G_i (denoted as $G_i \setminus v_i$) to form G_{i+1} for the next round (Line 11). After n rounds, the label sets of all vertices are returned as index L (Line 12).

Example 4. Consider the graph G in Fig. 1. We show how to create in-label sets for G , and out-label sets can be created similarly. In round 1, G is copied to G_1 , and v_1 (with the highest order) is inserted into the in-label sets of its descendants $DES^{G_1}(v_1) = \{v_1, v_5, v_7, v_8, v_9\}$ in G_1 , as no pruning occurs. Then, v_1 and its adjacent edges are removed from G_1 to form G_2 . In round 2, v_2 (with the second highest order) finds its descendants $DES^{G_2}(v_2) = \{v_2, v_3, v_4, v_5, v_6, v_7, v_{10}, v_{11}\}$ in G_2 . Then, v_2 performs a pruning operation (Line 8 of Algorithm 1) to test whether v_2 is added to the in-label sets of vertices in $DES^{G_2}(v_2)$. For example, since $L_{out}^2(v_2) \cap L_{in}^2(v_5) = \{v_1\}$, v_2 is not inserted into $L_{in}^3(v_5)$ — pruning occurs. After pruning, v_2 is inserted into the in-label sets of $\{v_2, v_3, v_4, v_6, v_{10}, v_{11}\}$. After processing v_{11} , TOL ends.

Limitation. TOL is a centralized algorithm [14], which means that the graph needs to be stored on one machine for processing. To make matters worse, TOL is non-trivial to be parallelized. To illustrate why, we focus on the process of labeling v_i . When labeling v_i , each vertex w has an in-label set $L_{in}^i(w)$ and an out-label set $L_{out}^i(w)$ created by vertices $[v_1, v_2, \dots, v_{i-1}]$ of order

higher than v_i . We collect the label sets of all vertices w and get the index $L^i = \bigcup_{w \in V} \{L_{in}^i(w) \cup L_{out}^i(w)\}$ generated by vertices of order higher than v_i . The index L^i is necessary when labeling v_i : L^i determines whether or not v_i is added to the label set of another vertex $w \in V$.

Lemma 1. For $\forall w \in V$, whether or not v_i is in the label set of w depends on L^i , specifically,

- $v_i \in L_{in}(w) \Leftrightarrow v_i \rightarrow w, L_{out}^i(v_i) \cap L_{in}^i(w) = \emptyset$;
- $v_i \in L_{out}(w) \Leftrightarrow w \rightarrow v_i, L_{out}^i(w) \cap L_{in}^i(v_i) = \emptyset$.

Proof. See Appendix. \square

Lemma 1 shows that L^i is essential for labeling v_i . However, L^i is generated only after vertices $[v_1, v_2, \dots, v_{i-1}]$ of order higher than v_i have completed labeling. This suggests that labeling v_i cannot begin until those vertices with higher orders have finished labeling. Such a strong *order dependency* prevents TOL from being parallelized. Motivated by this, we aim to design novel labeling methods that can work in parallel while obtaining the same indexes as TOL.

Remark. [14] maintains TOL’s index for dynamic graphs, but we try to generate the same indexes as TOL for distributed graphs. We consider maintaining indexes on distributed dynamic graphs as future work.

C. Problem Statement

We plan to use a **vertex-centric system** [18] to implement the proposed labeling methods. The vertex-centric system performs the tasks in a super-step fashion [19]. In each super-step, each active vertex v calls a user-defined function, compute(), to: 1) compute based on v ’s current state and the messages it received in the previous super-step; 2) update v ’s state; 3) send messages to other vertices (for the next super-step); and 4) (optionally) vote v to make it inactive. The whole computation terminates when there are no messages in the system, or all vertices become inactive. To avoid ambiguity, we use the term “vertex” to denote v as $v \in V$, and the term “node” to denote a computation unit in a cluster.

The problem addressed in this paper is:

Given a distributed graph G , design reachability labeling methods and implement them using a vertex-centric system to create the index as TOL.

Throughout this paper, we do not assume that G is acyclic. This treatment is also used in [16], [20]. We treat it this way for two reasons: 1) our methods are general enough to handle both acyclic and non-acyclic graphs; 2) it is non-trivial to obtain and merge strongly connected components to make graphs acyclic in a distributed environment.

III. DISTRIBUTED REACHABILITY LABELING

We present the concept of backward label sets in Section III-A, and then propose a filtering-and-refinement framework to find backward label sets in Section III-B. New labeling algorithms based on this framework are designed in Section III-C, followed by the algorithm implementation in a distributed system in Section III-D.

A. TOL Revisited

As shown in Algorithm 1, TOL works in n rounds, where a vertex v is selected for labeling in each round. The process of labeling v is to use the pruning operation to determine some

TABLE III: The Backward Label Sets

Vertex	L_{in}^-	L_{out}^-
v_1	$\{v_1, v_5, v_7, v_8, v_9\}$	$\{v_1, v_2, v_3, v_4, v_5, v_6, v_7\}$
v_2	$\{v_2, v_3, v_4, v_6, v_{10}, v_{11}\}$	$\{v_2, v_3, v_4, v_6\}$
v_3	\emptyset	\emptyset
v_4	\emptyset	\emptyset
v_5	\emptyset	\emptyset
v_6	\emptyset	\emptyset
v_7	\emptyset	\emptyset
v_8	$\{v_8, v_9\}$	$\{v_8\}$
v_9	$\{v_9\}$	$\{v_9\}$
v_{10}	$\{v_{10}\}$	$\{v_{10}\}$
v_{11}	$\{v_{11}\}$	$\{v_{11}\}$

vertices (in v 's descendants/ancestors) such that v is added to their label sets. We define these determined vertices as the backward label set of v .

Definition 4. Given $v \in V$ and index L of G , the **backward in-label set** of v is $L_{in}^-(v) = \{w | v \in L_{in}(w)\}$; the **backward out-label set** of v is $L_{out}^-(v) = \{w | v \in L_{out}(w)\}$.

Example 5. Consider the graph G in Fig. 1, where Table II shows the index L . In Table III we list the backward in-label/out-label sets for all vertices. For v_2 , $L_{in}^-(v_2) = \{v_2, v_3, v_4, v_6, v_{10}, v_{11}\}$ since vertices $\{v_2, v_3, v_4, v_6, v_{10}, v_{11}\}$ contain v_2 in their in-label sets; For v_3 , $L_{in}^-(v_3) = \emptyset$ since no vertex contains v_3 in its in-label set. Also, labeling v_2 is to add v_2 to the in-label sets of vertices in $L_{in}^-(v_2)$ and add v_2 to the out-label sets of vertices in $L_{out}^-(v_2)$.

We re-describe the working process of TOL from the perspective of backward label sets: TOL chooses a vertex v to label in each round. The process of labeling v is to determine the vertices in backward in-label/out-label sets where v joins their labels. Recall that TOL applies the pruning operation to determine its backward label sets. According to the analysis in Section II, the pruning operation causes TOL not to work in parallel. Therefore, the main contribution of our paper is to replace the pruning operation of TOL but still get the backward label sets of each vertex $v \in V$. This allows all vertices to work in parallel and get the same index as TOL.

Remark. Label sets and backward label sets are symmetric concepts — v is in $L_{in}(w)$ (resp. $L_{out}(w)$) implies that w is in $L_{in}^-(v)$ (resp. $L_{out}^-(v)$). For this reason, we aim to find the backward label sets $L_{in}^-(v)$ and $L_{out}^-(v)$ of each vertex v to create the index L . Also, since finding $L_{out}^-(v)$ on G is similar to finding $L_{in}^-(v)$ on G , we only discuss how to obtain $L_{in}^-(v)$ in the sequel. The discussions for $L_{in}^-(v)$ can be naturally extended to $L_{out}^-(v)$.

B. Filtering-and-refinement Framework

To determine the backward label set $L_{in}^-(v)$ for each vertex v without relying on the pruning operation of TOL, we give the condition for a certain vertex w to lie in $L_{in}^-(v)$.

Theorem 1. For $\forall v, w \in V$, $w \in L_{in}^-(v) \Leftrightarrow w \in DES(v)$ and w is the highest-order vertex on all paths from v to w .

Proof. We prove $w \in L_{in}^-(v)$, or equivalently $v \in L_{in}(w)$.

- \Leftarrow : If v is the highest-order vertex on all paths from v to w , then there is no vertex u such that $v \rightarrow u \rightarrow w$ and $\text{ord}(u) > \text{ord}(v)$. Because TOL processes vertices in a non-increasing sequence of vertex order, this means at the moment v starts labeling: 1) $w \in DES^{G_i}(v)$ since no vertex on a path from v to w can be

removed before labeling v ; 2) $L_{out}(v) \cap L_{in}(w) = \emptyset$ since no vertex on a path from v to w finishes labeling. Therefore, by Line 8 of Algorithm 1, $v \in L_{in}(w)$. With similar logic, we can prove that $v \in L_{out}(w)$ if v is the highest-order vertex on all paths from w to v .

- \Rightarrow : If $v \in L_{in}(w)$, let $u \neq v$ be the highest-order vertex on all paths from v to w . This means that u 's order is the highest on all sub-paths from v to u and from u to w . We reuse the proof in \Leftarrow : 1) u 's order is the highest on all u - w paths, then $u \in L_{in}(w)$; 2) u 's order is the highest on all v - u paths, then $u \in L_{out}(v)$. Thus, $u \in L_{out}(v) \cap L_{in}(w) \neq \emptyset$ when labeling v . By Line 8 of Algorithm 1, $v \notin L_{in}(w)$, contradiction. \square

Example 6. Consider the graph G in Fig. 1. For vertex v_2 , $v_3 \in L_{in}^-(v_2)$ since v_2 has the highest order on all paths from v_2 to v_3 ; $v_5 \notin L_{in}^-(v_2)$ because v_1 , with order higher than v_2 , lies on a path from v_2 to v_5 .

Theorem 1 paves the way for obtaining $L_{in}^-(v)$ of each vertex $v \in V$ in parallel. Specifically, by Theorem 1, $w \in DES(v)$ is a necessary condition for $w \in L_{in}^-(v)$. In other words, $DES(v)$ is a super-set of $L_{in}^-(v)$: $L_{in}^-(v) \subseteq DES(v)$. To obtain $L_{in}^-(v)$, we need to remove invalid elements from $DES(v)$. Thus, we define the higher-order descendants of v .

Definition 5. The **higher-order descendants** of v , denoted as $DES_{\text{hig}}(v)$, are vertices $u \in DES(v)$ whose order is higher than v , that is, $DES_{\text{hig}}(v) = \{u | u \in DES(v), \text{ord}(u) > \text{ord}(v)\}$.

Theorem 1 states that only when v is the highest-order vertex on all paths from v to w , then w is in $L_{in}^-(v)$. In other words, if there is a high-order vertex $u \in DES_{\text{hig}}(v)$ to reach w , w must not be in $L_{in}^-(v)$. Hence, we can use $DES_{\text{hig}}(v)$ to refine the super-set $DES(v)$ by removing invalid elements.

Based on this idea, we propose the filtering-and-refinement framework to obtain $L_{in}^-(v)$: the filtering phase generates the super-set $DES(v)$, and then invalid vertices that can be reached by vertices in $DES_{\text{hig}}(v)$ are removed for refinement. The correctness of this framework is given in Theorem 2.

Theorem 2. $L_{in}^-(v) = DES(v) - \bigcup_{u \in DES_{\text{hig}}(v)} DES(u)$.

Proof. We denote the right-hand side of the equation by RHS and verify that $L_{in}^-(v) = \text{RHS}$.

- $L_{in}^-(v) \subseteq \text{RHS}$: If $w \in L_{in}^-(v) \setminus \text{RHS}$, then there are two possibilities: 1) $w \notin DES(v)$, which contradicts $w \in L_{in}^-(v) \subseteq DES(v)$; 2) $w \in \bigcup_{u \in DES_{\text{hig}}(v)} DES(u)$, but by Theorem 1, $w \notin L_{in}^-(v)$, contradiction.
- $\text{RHS} \subseteq L_{in}^-(v)$: If $w \in \text{RHS}$, then there is no vertex u on any path from v to w for which $\text{ord}(u) > \text{ord}(v)$. By Theorem 1, $w \in L_{in}^-(v)$. \square

Example 7. Consider the graph G in Fig. 1. We show how to obtain $L_{in}^-(v_3)$ of v_3 . We first find $DES(v_3) = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8, v_9, v_{10}, v_{11}\}$; next we get $DES_{\text{hig}}(v_3) = \{v_1, v_2\}$, and $\bigcup_{u \in DES_{\text{hig}}(v_3)} DES(u) = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8, v_9, v_{10}, v_{11}\}$. Thus, $L_{in}^-(v_3) = DES(v_3) - \bigcup_{u \in DES_{\text{hig}}(v_3)} DES(u) = \emptyset$.

C. Two Labeling Methods

C-1. Basic Labeling Method

If we apply Theorem 2 to obtain $L_{in}^-(v)$ for each vertex $v \in V$, we need to perform a v -sourced breadth-first search (BFS) to

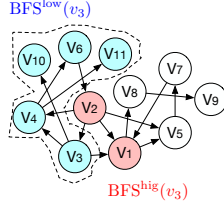


Fig. 3: The v_3 -sourced Trimmed BFS

obtain $\text{DES}(v)$ and $\text{DES}_{\text{high}}(v)$ in the filtering phase, and then perform $|\text{DES}_{\text{high}}(v)|$ BFSs, one BFS for one vertex in $\text{DES}_{\text{high}}(v)$ in the refinement phase. The refinement phase requires a large number of BFSs, rendering this solution inefficient. To improve efficiency, we find that not all vertices in $\text{DES}_{\text{high}}(v)$ are useful for refinement. For example, for vertices $a, b \in \text{DES}_{\text{high}}(v)$, b is unnecessary when $\text{DES}(b)$ is a subset of $\text{DES}(a)$ — vertex a can reach all descendants of b .

So, how to identify unnecessary vertices in $\text{DES}_{\text{high}}(v)$? A simple rule is that it is safe to delete vertex b if vertex $a \in \text{DES}_{\text{high}}(v)$ can reach b : a can reach all descendants of b . Based on this rule, we propose to use a v -sourced BFS but block the expansion branch upon meeting a vertex $a \in \text{DES}_{\text{high}}(v)$, thus implicitly deleting vertices $b \in \text{DES}_{\text{high}}(v)$ that can be reached by a . We denote this BFS as a trimmed BFS.

Algorithm 2: Trimmed BFS

Input: Graph $G(V, E)$, v
Output: $\text{BFS}_{\text{low}}(v)$, $\text{BFS}_{\text{high}}(v)$

```

1 queue  $Q \leftarrow \emptyset$ ;
2  $\text{status}(u) \leftarrow \textcircled{?}$ , for all vertices  $u \in V$ ;
3 push  $v \rightarrow Q$  and  $\text{BFS}_{\text{low}}(v)$ ;
4  $\text{status}(v) \leftarrow \textcircled{\ominus}$ ;
5 while  $Q$  is not empty do
6    $u \leftarrow \text{pop from } Q$ ;
7   foreach  $w \in N_{\text{out}}(u)$  do
8     if  $\text{status}(w) \neq \textcircled{?}$  then continue;
9     if  $\text{ord}(w) < \text{ord}(v)$  then
10       $\text{status}(w) \leftarrow \textcircled{\ominus}$ , push  $w \rightarrow Q$  and  $\text{BFS}_{\text{low}}(v)$ ;
11    else
12      // block the expansion via  $w$ 
13      push  $w \rightarrow \text{BFS}_{\text{high}}(v)$ ;
13 return  $\text{BFS}_{\text{low}}(v)$ ,  $\text{BFS}_{\text{high}}(v)$ ;

```

Trimmed BFS. Algorithm 2 describes the v -sourced trimmed BFS. We initialize an empty queue Q and set the status of all vertices to unvisited (denoted as $\textcircled{?}$) (Line 1-2). Then, v is inserted into Q and $\text{BFS}_{\text{low}}(v)$, and the status of v is set to visited (denoted as $\textcircled{\ominus}$) (Line 3-4). Afterward, we pop a vertex u from Q and check each neighbor w of u (Line 6-7). If w is visited before, we do nothing (Line 8). Otherwise, depending on the order of w , we have two cases: 1) if w is lower in order than v , we continue expanding via w by setting the status of w as visited and inserting w both in Q and $\text{BFS}_{\text{low}}(v)$ (Line 9-10); 2) otherwise, we block the expansion via w and insert w into $\text{BFS}_{\text{high}}(v)$ (Line 12). When the queue Q is empty, the BFS terminates, and $\text{BFS}_{\text{low}}(v)$ and $\text{BFS}_{\text{high}}(v)$ are returned.

Lemma 2. *The time cost of Algorithm 2 is $O(|E| + |V|)$.*

Example 8. Fig. 3 shows the v_3 -sourced trimmed BFS. First, v_3 is inserted into both Q and $\text{BFS}_{\text{low}}(v_3)$. Then, v_3 is popped from Q ,

and for v_3 's out-neighbors $\{v_1, v_4, v_{10}\}$: v_4 and v_{10} are inserted into both $\text{BFS}_{\text{low}}(v_3)$ and Q because they are of lower order than v_3 ; the expansion via v_1 is pruned since $\text{ord}(v_1) > \text{ord}(v_3)$, and v_1 is inserted into $\text{BFS}_{\text{high}}(v_3)$. Then, v_4 is popped from Q , and v_4 's out-neighbors $\{v_6, v_{11}\}$ are examined. The BFS terminates when Q is empty, and we get $\text{BFS}_{\text{low}}(v_3) = \{v_3, v_4, v_{10}, v_6, v_{11}\}$, $\text{BFS}_{\text{high}}(v_3) = \{v_1, v_2\}$.

Modified Framework. During the trimmed BFS sourced from v , we obtain $\text{BFS}_{\text{low}}(v)$ (vertices visited by BFS and of order lower than v) and $\text{BFS}_{\text{high}}(v)$ (vertices with higher order that block the expansion). Using $\text{BFS}_{\text{low}}(v)$ and $\text{BFS}_{\text{high}}(v)$, we optimize the original filtering-and-refinement framework.

Refinement. We first show that in the refinement phase, $\text{BFS}_{\text{high}}(v)$ can replace $\text{DES}_{\text{high}}(v)$ since vertices in $\text{BFS}_{\text{high}}(v)$ reach all the descendants of vertices in $\text{DES}_{\text{high}}(v)$.

Lemma 3. $\bigcup_{u \in \text{BFS}_{\text{high}}(v)} \text{DES}(u) = \bigcup_{u \in \text{DES}_{\text{high}}(v)} \text{DES}(u)$.

Proof. Let LHS be $\bigcup_{u \in \text{BFS}_{\text{high}}(v)} \text{DES}(u)$ and RHS be $\bigcup_{u \in \text{DES}_{\text{high}}(v)} \text{DES}(u)$.

- LHS \subseteq RHS follows from the fact $\text{BFS}_{\text{high}}(v) \subseteq \text{DES}_{\text{high}}(v)$.
- RHS \subseteq LHS: If $\exists s \in \text{RHS} \setminus \text{LHS}$, $s \in \text{RHS}$ implies there is a path from v to s containing some vertex $w \in \text{DES}_{\text{high}}(v)$. On all paths from v to s , we collect the inner vertices in $\text{DES}_{\text{high}}(v)$ and insert them in set S (S is not empty as $w \in \text{DES}_{\text{high}}(v)$ is such a vertex). We choose the vertex $u \in S$ with the smallest distance to v : there is no other higher-order vertex on the v - u path. By Algorithm 2, $u \in \text{BFS}_{\text{high}}(v)$. Thus, $s \in \bigcup_{u \in \text{BFS}_{\text{high}}(v)} \text{DES}(u) = \text{LHS}$, contradiction. \square

Example 9. Consider the graph G in Fig. 1. For v_3 , $\text{BFS}_{\text{high}}(v_3)$ can replace $\text{DES}_{\text{high}}(v_3)$ for refinement since $\text{BFS}_{\text{high}}(v_3) = \{v_1, v_2\}$ reaches all descendants of vertices in $\text{DES}_{\text{high}}(v_3)$: $\bigcup_{u \in \text{BFS}_{\text{high}}(v_3)} \text{DES}(u) = \bigcup_{u \in \text{DES}_{\text{high}}(v_3)} \text{DES}(u) = \{v_1, v_2, v_3, v_4, v_5, v_6, v_7, v_8, v_9, v_{10}, v_{11}\}$.

Filtering. Next, we show $\text{BFS}_{\text{low}}(v)$ is a super-set of $L_{\text{in}}^-(v)$, meaning that $\text{BFS}_{\text{low}}(v)$ can replace $\text{DES}(v)$ for filtering.

Lemma 4. $L_{\text{in}}^-(v) \subseteq \text{BFS}_{\text{low}}(v)$.

Proof. Suppose there is $s \in L_{\text{in}}^-(v) \setminus \text{BFS}_{\text{low}}(v)$, then there must be a path from v to s through a high-order vertex $u \in \text{BFS}_{\text{high}}(v)$. By Theorem 1, $s \notin L_{\text{in}}^-(v)$, contradiction. \square

Example 10. Consider the graph G in Fig. 1. For v_3 , $\text{BFS}_{\text{low}}(v_3)$ can replace $\text{DES}(v_3)$ for filtering, because $\text{BFS}_{\text{low}}(v_3) = \{v_3, v_4, v_6, v_{10}, v_{11}\}$ includes all vertices in $L_{\text{in}}^-(v_3) = \emptyset$.

Basic Labeling Method. Lemma 4 shows that $\text{BFS}_{\text{low}}(v)$ is a super-set of $L_{\text{in}}^-(v)$, while Lemma 3 shows that $\text{BFS}_{\text{high}}(v)$ is sufficient to eliminate invalid elements not in $L_{\text{in}}^-(v)$. Thus, we give the basic labeling method for labeling $v \in V$.

- Step 1. In the filtering phase, we use a v -sourced BFS to find $\text{BFS}_{\text{low}}(v)$ and $\text{BFS}_{\text{high}}(v)$;
- Step 2. In the refinement phase, we perform a BFS for each vertex in $\text{BFS}_{\text{high}}(v)$.
- Step 3. Return $\text{BFS}_{\text{low}}(v) - \bigcup_{u \in \text{BFS}_{\text{high}}(v)} \text{DES}(u)$ as $L_{\text{in}}^-(v)$.

Combing Lemma 3, Lemma 4, and Theorem 2, the correctness of this method is given in Theorem 3.

Theorem 3. $L_{\text{in}}^-(v) = \text{BFS}_{\text{low}}(v) - \bigcup_{u \in \text{BFS}_{\text{high}}(v)} \text{DES}(u)$.

C-2. Improved Labeling Method

Compared with the framework given in Theorem 2, the basic method based on Theorem 3 reduces the number of BFSs needed in the refinement phase from $|\text{DES}_{\text{hig}}(v)|$ to $|\text{BFS}_{\text{hig}}(v)|$. But the number $|\text{BFS}_{\text{hig}}(v)|$ may still be very large. Therefore, can we avoid using a large number of BFSs in the refinement phase?

To answer this question, we revisit the refinement phase of the basic method (i.e., Lemma 3): when labeling v , w is eliminated when a vertex in $\text{BFS}_{\text{hig}}(v)$ can reach w . We focus on a specific vertex $u \in \text{BFS}_{\text{hig}}(v)$, which has the highest order on paths from v to w : when performing a u -sourced trimmed BFS in G , u can reach w , so w is in $\text{BFS}_{\text{low}}(u)$ ¹; when performing a u -sourced trimmed BFS in the inverse graph \bar{G} , u can reach v , so v is in $\text{BFS}_{\text{low}}^{\bar{G}}(u)$. Thus, by examining whether there exist vertices u of order higher than v such that $w \in \text{BFS}_{\text{low}}(u)$ and $v \in \text{BFS}_{\text{low}}^{\bar{G}}(u)$, we can eliminate w to complete the refinement without using any BFSs: the existence of higher-order vertices u on the v - w paths implies that w can be eliminated.

Example 11. Consider the graph G in Fig. 1. For v_3 in G , v_4 can be eliminated because $\exists v_2 \in \text{BFS}_{\text{hig}}(v_3)$ s.t., 1) in G (Fig. 1), v_2 -sourced BFS visits v_4 and hence $v_4 \in \text{BFS}_{\text{low}}(v_2)$; 2) in \bar{G} (Fig. 2), v_2 -sourced BFS visits v_3 , and hence $v_3 \in \text{BFS}_{\text{low}}^{\bar{G}}(v_2)$.

Improved Refinement. Based on the above idea, in the refinement phase of labeling v , to check whether some vertex $w \in \text{BFS}_{\text{low}}(v)$ should be removed, we need to check if there exists $u \in \text{BFS}_{\text{hig}}(v)$ such that $w \in \text{BFS}_{\text{low}}(u)$ and $v \in \text{BFS}_{\text{low}}^{\bar{G}}(u)$. Determining $w \in \text{BFS}_{\text{low}}(u)$ can be done intuitively in G , since $\text{BFS}_{\text{low}}(u)$ is known; but determining $v \in \text{BFS}_{\text{low}}^{\bar{G}}(u)$ is not so simple, since the information on \bar{G} needs to be used.

To make it feasible to determine $v \in \text{BFS}_{\text{low}}^{\bar{G}}(u)$, we create an inverted list $\text{IBFS}_{\text{low}}(v)$ for vertex $v \in V$.

Definition 6. If v is visited by the u -sourced trimmed BFS in \bar{G} , vertex u is in the **inverse list** $\text{IBFS}_{\text{low}}(v)$ of v , i.e., $\text{IBFS}_{\text{low}}(v) = \{u | v \in \text{BFS}_{\text{low}}^{\bar{G}}(u)\}$.

With $\text{IBFS}_{\text{low}}(v)$, we can eliminate w by Lemma 5.

Lemma 5. For a vertex $w \in \text{BFS}_{\text{low}}(v)$, $w \notin L_{\text{in}}^-(v)$ if $\exists u \in \text{IBFS}_{\text{low}}(v)$, and $w \in \text{BFS}_{\text{low}}(u)$.

Proof. A vertex $u \in \text{IBFS}_{\text{low}}(v)$ with $w \in \text{BFS}_{\text{low}}(u)$ means there is a higher-order vertex u on the path from v to w . Then, $w \notin L_{\text{in}}^-(v)$ by Theorem 1. \square

Improved Labeling Method. With the refinement given in Lemma 5, we give an improved labeling method for labeling v .

- Step 1. In the filtering phase, we use a v -sourced BFS in G to find $\text{BFS}_{\text{low}}(v)$ and $\text{BFS}_{\text{hig}}(v)$, for $\forall v \in V$;
- Step 2. We use a v -sourced BFS in \bar{G} to find $\text{BFS}_{\text{low}}^{\bar{G}}(v)$, and then get $\text{IBFS}_{\text{low}}(v)$ by Definition 6, for $\forall v \in V$;
- Step 3. In the refinement phase, if $\exists u \in \text{IBFS}_{\text{low}}(v)$, and $w \in \text{BFS}_{\text{low}}(u)$, the vertex w can be eliminated;
- Step 4. Return the non-eliminated vertices as $L_{\text{in}}^-(v)$.

Combing Lemma 5 and Theorem 3, the correctness of this method is given below.

Theorem 4. $L_{\text{in}}^-(v) = \text{BFS}_{\text{low}}(v) - S$, where $S = \{w | w \in \text{BFS}_{\text{low}}(v), \exists u \in \text{IBFS}_{\text{low}}(v), w \in \text{BFS}_{\text{low}}(u)\}$.

¹Without ambiguity, $\text{BFS}_{\text{low}}(u)$ and $\text{BFS}_{\text{low}}^{\bar{G}}(u)$ refers to the same thing.

Note that in Step 2 of the improved labeling method, we need a v -sourced trimmed BFS on \bar{G} to obtain $L_{\text{in}}^-(v)$. This step does not introduce additional costs because $\text{BFS}_{\text{low}}^{\bar{G}}(v)$ is needed to obtain $L_{\text{out}}^-(v)$. In other words, only trimmed BFSs are required to obtain both $L_{\text{in}}^-(v)$ and $L_{\text{out}}^-(v)$.

So far, in addition to finding the backward label sets using the framework given in Theorem 2, we have proposed a basic method based on Theorem 3 and an improved method based on Theorem 4. We give in Table IV the number of BFSs required in the filtering and refinement phases for each method.

TABLE IV: The Comparison Between Labeling Methods

Method	Filtering	Refinement
Theorem 2	1	$ \text{DES}_{\text{hig}}(v) $
Theorem 3 (Basic)	1	$ \text{BFS}_{\text{hig}}(v) \leq \text{DES}_{\text{hig}}(v) $
Theorem 4 (Improved)	1	1

D. Distributed Implementation

To handle distributed graphs, we implement the improved labeling method using a vertex-centric system, which is denoted as DRL. We omit the distributed implementation of the basic labeling method, as this can be implemented in a similar way.

Algorithm 3: Compute() for DRL

```

1 Data:  $in\text{-msgs} \leftarrow$  messages from in-neighbors;
2    $out\text{-msgs} \leftarrow$  messages to out-neighbors;
3 if  $super\text{-step} = 1$  then
4   | //  $w$  is vertex to perform computations
5   |  $w = vertex\_id()$ ;
6   |  $w.status(z) \leftarrow \textcircled{?}$ , for each vertex  $z \in V$ ;
7   |  $w.status(w) \leftarrow \textcircled{\ominus}$ ;
8   | // message format:  $\{ID, order\}$ 
9   | message  $\leftarrow \{w, ord(w)\}$ ;
10  | send message to out-neighbors;
11 foreach message  $\in in\text{-msgs}$  do
12  | //  $v$  is the source to do trimmed BFS
13  |  $v \leftarrow message.ID$ ;
14  |  $ord(v) \leftarrow message.order$ ;
15  | if  $w.status(v) = \textcircled{\ominus}$  then continue;
16  | if  $ord(v) > ord(w)$  then
17  |   | if  $Check(v, w) = true$  then continue;
18  |   |  $w.status(v) \leftarrow \textcircled{\ominus}$ ;
19  |   | message  $\leftarrow \{v, ord(v)\}$ ;
20  |   | send message to out-neighbors;
21  |   | // works on  $\bar{G}$ 
22  |   | insert  $v$  into  $\text{IBFS}_{\text{low}}(w)$ ;
23  | // only run after the final super-step
24 foreach  $v$ , s.t.,  $w.status(v) = \textcircled{\ominus}$  do
25  | if  $Check(v, w) = true$  then  $w.status(v) \leftarrow \textcircled{?}$ ;
26 Procedure  $Check(v, w)$ 
27   | foreach  $u \in \text{IBFS}_{\text{low}}(v)$  do
28   |   | if  $w.status(u) = \textcircled{\ominus}$  then return true
29   |   | return false

```

Algorithm. Algorithm 3 describes DRL, where the compute() function is executed on each vertex $w \in V$ in super-steps. We record the visited status of w using a status array² $w.status$. Specifically, if the value of $w.status(v)$ is $\textcircled{?}$, then w is not visited by the vertex v ; if the value is $\textcircled{\ominus}$, then w is visited by v . By reading the values of status arrays, the backward in-label sets of all vertices can be obtained.

²In the implementation, the hash table can be used to replace the array because of the sparsity of the array.

In the first super-step (Line 3), vertex w initializes its status array by assigning the unvisited status $\textcircled{?}$ to all vertices (Line 5), except for w itself, which is assigned as $\textcircled{\ominus}$ (Line 6). Then, w sends the message containing its vertex ID (w) and vertex order ($\text{ord}(w)$) to out-neighbors (Line 7-8). In subsequent super-steps, once vertex w receives the message from in-neighbors (Line 9), w extracts vertex ID v (Line 10) and order $\text{ord}(v)$ (Line 11) of the message. If $w.\text{status}(v)$ is $\textcircled{\ominus}$, we do nothing as v visited w before (Line 12).

If $w.\text{status}(v)$ is $\textcircled{?}$ and the order v is higher than w , we continue the v -sourced trimmed BFS via w (Line 13). We mark the status of $w.\text{status}(v)$ as $\textcircled{\ominus}$ (Line 15), and we send the message $\{v, \text{ord}(v)\}$ to w 's out-neighbors to continue the v -sourced BFS. Also, on the inverse graph \bar{G} , if v can reach w , then v is inserted in $\text{IBFS}_{\text{low}}(w)$ (Line 18). Note that we will call the procedure $\text{Check}(v, w)$ (Line 21-24) for an expansion pruning (Line 14): if $\text{IBFS}_{\text{low}}(v)$ contains a vertex u that can reach w , it follows from Lemma 5 that w is not in $L_{\text{in}}^-(v)$, and we prune the expansion of v -sourced BFS via w .

In the final super-step, we check for w the vertices v for which $w.\text{status}(v)$ is $\textcircled{\ominus}$: if the procedure $\text{Check}(v, w)$ returns true, we reset $w.\text{status}(v)$ to $\textcircled{?}$ (Line 19-20). After this check, the vertices v for which $w.\text{status}(v)$ is $\textcircled{\ominus}$ form the backward in-label set $L_{\text{in}}^-(v)$ of v . Finally, we can collect the backward label sets of each vertex on one machine to obtain an index the same as TOL to support reachability queries.

Analysis. We give the correctness analysis of DRL, i.e., the vertices w whose $w.\text{status}(v)$ value is $\textcircled{\ominus}$ form $L_{\text{in}}^-(v)$ of v .

Theorem 5. *Given a graph G and a vertex v , $L_{\text{in}}^-(v) = \{w | w.\text{status}(v) = \textcircled{\ominus}\}$ for Algorithm 3.*

Proof. We denote RHS by $\{w | w.\text{status}(v) = \textcircled{\ominus}\}$ and we prove $L_{\text{in}}^-(v) = \text{RHS}$.

- $\text{RHS} \subseteq L_{\text{in}}^-(v)$: $w.\text{status}(v) = \textcircled{\ominus}$ means that v can reach w and $\text{ord}(w) < \text{ord}(v)$. Then we verify that there are no higher-order vertices on any path from v to w , thus deriving $w \in L_{\text{in}}^-(v)$ by Theorem 1. Suppose there are high-order vertices on paths from v to w , we choose the highest-order vertex u . Since u 's order is the highest on all v - w paths, u can reach v in \bar{G} , thus $u \in \text{IBFS}_{\text{low}}(v)$; u can reach w in G , thus $w.\text{status}(u) = \textcircled{\ominus}$. So the procedure $\text{Check}(v, w)$ will set $w.\text{status}(v) = \textcircled{?}$, contradiction.
- $L_{\text{in}}^-(v) \subseteq \text{RHS}$: Suppose there exists a vertex $w \in L_{\text{in}}^-(v)$ and $w.\text{status}(v) = \textcircled{?}$, then there are two cases: 1) if v cannot reach w , then $w \notin L_{\text{in}}^-(v)$ by Theorem 1, contradiction; 2) there exists a higher-order vertex u in $\text{IBFS}_{\text{low}}(v)$ to block the expansion branch from v to w to set $w.\text{status}(v)$ to $\textcircled{?}$ (Line 14) or to reset $w.\text{status}(v)$ to $\textcircled{?}$ (Line 19-20), which contradicts with Theorem 1. \square

We then analyze the computation and communication costs.

Lemma 6. *The computation cost of labeling a vertex $v \in V$ using Algorithm 3 is $O(|E| + |\text{IBFS}_{\text{low}}(v)| \cdot |V|)$, where $|\text{IBFS}_{\text{low}}(v)|$ is the inverted list size of v .*

Proof. The time required to perform a v -sourced trimmed BFS is $O(|E|)$. Also, Algorithm 3 triggers at most $|V|$ times of the $\text{Check}()$ procedure for v , each requiring $|\text{IBFS}_{\text{low}}(v)|$ time. \square

Lemma 7. *The communication cost of labeling a vertex $v \in V$ using Algorithm 3 is $O(|E| + |\text{IBFS}_{\text{low}}(v)|)$.*

Proof. Each vertex needs to send/read a message to its neighbors at most once for labeling a certain vertex v . In addition, we need to share $\text{IBFS}_{\text{low}}(v)$ to implement the $\text{Check}()$ procedure. Hence, the communication cost is $\sum_{v \in V} d_{\text{out}}(v) + \sum_{v \in V} d_{\text{in}}(v) + |\text{IBFS}_{\text{low}}(v)| = |E| + |\text{IBFS}_{\text{low}}(v)|$. \square

Remark. Although each vertex v needs to share its inverted list $\text{IBFS}_{\text{low}}(v)$, the size of $\text{IBFS}_{\text{low}}(v)$ is pretty small (empirical studies show that the average size of $\text{IBFS}_{\text{low}}(v)$ of each vertex v is less than one), so the communication overhead associated with sharing the inverted list is not significant. The efficiency of DRL is validated in Section VI.

IV. BATCH LABELING OPTIMIZATION

DRL creates backward label sets for all vertices in parallel. However, DRL misses the opportunity provided by the serial execution of TOL — the already processed high-order vertices strongly prune the search space when labeling the current vertex. As a remedy, we further improve the labeling efficiency by splitting vertices into batches to trade-off between pruning power and parallelization.

Batch Sequence. We split the vertices into a batch sequence for *batch labeling*: we label all the vertices within a batch simultaneously, while vertices in different batches perform the labeling process sequentially.

Definition 7. $[V_1, V_2, \dots, V_g]$ is a **batch sequence** when

- $\bigcup_{i \in g} V_i = V$ and $V_i \cap V_j = \emptyset$, for $\forall i \neq j$;
- for vertex $u \in V_i$ and vertex $v \in V_j$ with $i < j$, it must be ensured that $\text{ord}(u) > \text{ord}(v)$.

The batch sequence $[V_1, V_2, \dots, V_g]$ is a graph partition since it disjointly covers all vertices. Also, the vertices with high order are placed before the vertices with low order in the sequence. When the batch size $|V_i|$ ($1 \leq i \leq g$) is fixed to one, we get $|V|$ batches of vertices for labeling. This fully serial execution is how TOL works. When the batch size is fixed to $|V|$, we get 1 batch of vertices for labeling. This fully parallel execution is how DRL works. By setting the batch size flexibly, we make a trade-off between TOL and DRL.

To obtain a valid batch sequence, we need two parameters: an initial batch size variable b (ranging from 1 to $|V|$) and an increment factor k . The specific procedure is given below.

- Step 1. Sort vertices V in a non-increasing order of ord values, and then copy sorted vertices into the set S ;
- Step 2. In iteration i , remove b vertices with the highest order from S to form V_i (i.e., $S \leftarrow S \setminus V_i$), and then multiply b by k for the next iteration (i.e., $b \leftarrow b \cdot k$);
- Step 3. Stop at round $g + 1$ when $S = \emptyset$ and return $[V_1, V_2, \dots, V_g]$; otherwise, increase i by 1 and go to Step 2.

The number of vertices in the last batch V_g may not exceed b . We set the values of both b and k to 2. The effect of b and k on the labeling efficiency is discussed in Section 5.

Example 12. Consider the graph G in Fig. 1. Suppose $b = 2$ and $k = 2$. In the first round ($b = 2$), we get $V_1 = \{v_1, v_2\}$; in the second round ($b = 4$), we get $V_2 = \{v_3, v_4, v_5, v_6\}$; in the third round ($b = 8$), we get $V_3 = \{v_7, v_8, v_9, v_{10}, v_{11}\}$. $[V_1, V_2, V_3]$ is a batch sequence of G .

Algorithm 4: Compute() for DRL_b

```

1 Data:  $in\_msgs \leftarrow$  messages from in-neighbors;
2    $out\_msgs \leftarrow$  messages to out-neighbors;
3 Input:  $V_i$ ;
4 if  $super\_step = 1$  then
5   //  $w$  is vertex to perform computations
6    $w = vertex\_id()$ ;
7   if  $w \notin V_i$  or  $L_{out}^{V_i}(w) \cap L_{in}^{V_i}(w) \neq \emptyset$  then Return;
8   the same as Line 5-8 of Algorithm 3;
9   broadcast  $L_{out}^{V_i}(w)$  and  $L_{in}^{V_i}(w)$  to all computation nodes;
10 for each message  $\in in\_msgs$  do
11   //  $v$  is source to do trimmed BFS
12    $v \leftarrow message.ID$ ;
13    $ord(v) \leftarrow message.order$ ;
14   if  $L_{out}^{V_i}(w) \cap L_{in}^{V_i}(w) \neq \emptyset$  then Continue;
15   the same as Line 12-20 in Algorithm 3;
16   // only run after the final super-step
17  $L_{in}^{V_{i+1}}(w) \leftarrow \{v | w.status(v) = \ominus\}$ ;

```

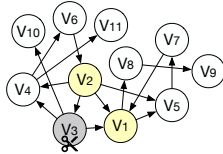


Fig. 4: The Illustration of Batch Labeling

Batch Label Sets. Since we process vertices in batches, the vertices in previous batches $[V_1, V_2, \dots, V_{i-1}]$ completed labeling before the current batch V_i begins. We define label sets generated by vertices in batches $[V_1, V_2, \dots, V_{i-1}]$ as batch label sets regarding V_i .

Definition 8. Given the batch V_i , the **batch in-label set** $L_{in}^{V_i}(w)$ of a vertex $w \in V$ is defined as $L_{in}^{V_i}(w) = \{u | u \in L_{in}(w), ord(u) > ord(V_i)\}$; the **batch out-label set** $L_{out}^{V_i}(w)$ of a vertex $w \in V$ is defined as $L_{out}^{V_i}(w) = \{u | u \in L_{out}(w), ord(u) > ord(V_i)\}$, where $ord(V_i) = \max\{ord(v) | v \in V_i\}$.

Similar to TOL, we can use the batch label sets to perform the pruning operation during the current batch, thereby optimizing the efficiency of DRL.

Example 13. Consider the graph G in Fig. 1. Suppose v_1 and v_2 finished labeling in the previous batch V_1 and only v_3 is in current batch V_2 . Before V_2 starts labeling, the batch in-label set $L_{in}^{V_2}(v_9)$ of v_9 is $\{v_1\}$; v_1 in $L_{in}(v_9) = \{v_1, v_8, v_9\}$ has a higher order than v_3 ; the batch out-label set $L_{out}^{V_2}(v_9)$ of v_9 is \emptyset ; $L_{out}(v_9) = \{v_9\}$ has no vertex of higher order than v_3 .

Algorithm. We incorporate the idea of batch labeling into DRL and implement it on a vertex-centric system to obtain algorithm DRL_b (Algorithm 4). DRL_b resembles DRL (Algorithm 3), and we only list the differences. First, only vertices in the current batch V_i are selected for labeling (Line 6); Also, if there is a higher-order vertex on the path from w to w ($L_{out}^{V_i}(w) \cap L_{in}^{V_i}(w) \neq \emptyset$), w is pruned (note that the graph is unnecessary to be acyclic) (Line 6). The batch label sets are then sent to all computation nodes (Line 8). The batch label sets are also used for pruning in Line 12. Then, at the end of batch V_i , vertices v with $w.status(v) = \ominus$ form batch in-label set of w for the next round (Line 14).

Example 14. Fig. 4 shows how batch labeling works. When labeling v_3 in the current batch V_i , as $L_{in}^{V_i}(v_3) = \{v_2\}$ intersects with $L_{out}^{V_i}(v_3) = \{v_1, v_2\}$, v_3 is pruned immediately — the search

space for labeling v_3 is dramatically reduced.

Analysis. We analyze the correctness of DRL_b.

Theorem 6. Given a graph G and a vertex v , $L_{in}^-(v) = \{w | w.status(v) = \ominus\}$ for Algorithm 4.

Proof. We denote $\{w | w.status(v) = \ominus\}$ by RHS and prove that $L_{in}^-(v) = \text{RHS}$.

- $\text{RHS} \subseteq L_{in}^-(v)$: Suppose there is a vertex w with $w.status(v) = \ominus$ but $w \notin L_{in}^-(v)$, $w \notin L_{in}^-(v)$ implies 1) $v \not\rightarrow w$, which shows that $w.status(v) = \text{?}$, or 2) there is a vertex with order higher than v on a path from v to w , so we select the highest-order vertex s . If s is in the previous batches and the current batch is denoted as V_i , since there are no vertices to prune s , $s \in L_{out}^{V_i}(v)$ and $s \in L_{in}^{V_i}(w)$. Hence, $w.status(v) = \text{?}$ by Line 12 of Algorithm 4; or s is in the current batch, then $w.status(v) = \text{?}$ by the correctness of DRL, contradiction.
- $L_{in}^-(v) \subseteq \text{RHS}$: We prove this by induction on the batch number. When $v \in V_1$, since no pruning occurs, then DRL_b is correct by the correction of DRL. Suppose V_{i-1} finishes labeling and DRL_b is correct, we prove DRL_b is correct for $v \in V_i$. Since $w \in L_{in}^-(v)$, then v can reach w and no pruning occurs at Line 12 of Algorithm 4. Therefore, by the correction of DRL, DRL_b is correct for $v \in V_i$. \square

We then provide its computation and communication costs.

Lemma 8. The computation cost of labeling a vertex $v \in V$ using Algorithm 4 is $O(|E'| + (|\text{IBFS}_{\text{low}}(v)| + \Delta) \cdot |V|)$, where $E' \subseteq E$, and Δ is the largest label size.

Proof. For each vertex v , Algorithm 4 needs to explore the reduced search space (denoted as E' , $E' \subseteq E$) due to the pruning operation. Moreover, Algorithm 4 requires at most $|V|$ times of Check() procedure (each costing $O(|\text{IBFS}_{\text{low}}(v)|)$) and $|V|$ label queries (each costing $O(\Delta)$). \square

Lemma 9. The communication cost of labeling a vertex $v \in V$ using Algorithm 4 is $O(|E'| + |\text{IBFS}_{\text{low}}(v)| + \Delta)$.

Proof. The cost comes from: 1) sharing label sets with other computation nodes, which incurs $O(\Delta)$ cost; 2) sending $\text{IBFS}_{\text{low}}(v)$ for refinement; 3) reduced search space E' . \square

Remark. Compared to DRL (Algorithm 3), DRL_b (Algorithm 4) requires additional costs to share and query batch label sets. However, empirical studies in Section VI show that the benefit of reducing the search space from E to $E' \subseteq E$ outweighs the additional overhead.

V. RELATED WORK

Index-free Approaches. The online search [16] such as breadth-first [21] or depth-first search [22] can be used to process reachability queries. However, because the graph needs to be used at the query time, the query latency can be large.

Index-assisted Approaches. Index-assisted methods speed up the online search by using auxiliary structures. The auxiliary structures can be subgraphs [6], multiple intervals [7], independent permutations [8], bloom filters [9], or partial label sets [15]. One of the best-known approaches in this category is BFL [9]. The basic idea of BFL is that if vertex s can reach vertex t , then s can reach all descendants $\text{DES}(t)$ of t , i.e., $\text{DES}(t) \subseteq \text{DES}(s)$. Through a Bloom filter, BFL maps the descendants $\text{DES}(v)$ of

each vertex v to a subset as the out-label set of v [23]. At query time, the labels alone can be used to determine that s cannot reach t : if t 's out-label set is not fully contained in s 's out-label set, then $\text{DES}(t) \not\subseteq \text{DES}(s)$ and thus $s \not\rightarrow t$. However, if s can reach t , then BFL needs to perform a graph search to report the answer. Similarly, BFL uses the Bloom filter to map each vertex's ancestors to generate its in-label set. For BFL, as the index cannot answer all queries, the graph needs to be loaded into memory at query time (whereas the index-only approach does not need to rely on the graph at query time). BFL is undesirable for distributed graphs because (1) its index construction strictly follows the post-order of depth-first search (DFS), and thus requires performing distributed DFS, and (2) it needs to traverse distributed graphs during query processing. These two operations incur high costs in the distributed environment, as will be validated in Section VI.

Index-only Approaches. Reachability relations between all pairs of vertices can be stored as an index using a transit closure (TC) in $O(n^2)$ space. Due to its huge size, TC is compressed by intervals [24], bit vectors [25], or other structures (e.g., a spanning tree and additional intervals [26], multiple chains [27], and a path-tree [28]). The limitation of TC is the huge space required. Another type of index is created by reachability labeling methods. Reachability labeling methods were pioneered by Cohen et al. [10]. After that, [11] provided a divide-and-conquer strategy for labeling; geometric-based and graph partition-based labeling methods were proposed in [12] and [13], respectively. The state-of-the-art index-only approach TOL is introduced in Section II.

Parallelized Distance Labeling. Parallelized distance labeling methods are proposed for shortest distance queries. Li et al. [29] designed a parallel distance labeling algorithm for unweighted graphs; Lakhotia et al. [30] provided a distributed distance labeling algorithm for weighted graphs. For distance labeling, as shown in [29], v is inserted into the label set of w if and only if v is the highest-order vertex on *all shortest paths* from v to w . But according to Theorem 1 in Section III, v is inserted into the label set of w if and only if v is the highest-order vertex on *all paths* from v to w . Current distance labeling methods cannot find a higher-order vertex *not* on the shortest path, thus resulting in much larger indexes than the reachability labeling methods. Therefore, there is still a need to investigate new techniques for parallelizing reachability labeling, which is the focus of this paper.

VI. EXPERIMENTS

A. Settings

Algorithms. We aim to propose distributed labeling algorithms that produce the same index as TOL. Our methods include:

- DRL (Algorithm 3), a distributed labeling algorithm based on Theorem 4.
- DRL^- , a basic labeling algorithm based on Theorem 3. Since its distributed implementation is similar to DRL, we omit its implementation details.
- DRL_b (Algorithm 4), a distributed algorithm obtained by applying batch labeling to DRL.

Datasets. The experiments were conducted on 18 real-world directed graphs that are widely used in recent work related to reachability queries [3], [6], [31]. The properties of the graphs are shown in Table V. The largest graph has more than 3.7 billion edges. All the datasets are from Stanford Large Network Dataset

TABLE V: Datasets

Name	Dataset	$ V $	$ E $	Type
WEBW	Web-wikipedia	1,864,433	4,507,315	Web
DBPE	Dbpedia	3,365,623	7,989,191	Knowledge
CITE	Citeseerx	6,540,401	15,011,260	Citation
CITP	Cit-patent	3,774,768	16,518,947	Citation
TW	Twitter	18,121,168	18,359,487	Social
GO	Go-uniprot	6967956	34,770,235	Biology
SINA	Soc-sinaweibo	58,655,849	261,321,071	Social
LINK	Wikipedia-link	13,593,032	437,217,424	Web
WEBB	Webbase-2001	118,142,155	1,019,903,190	Web
GRPH	Graph500	17,043,780	1,046,934,896	Synthetic
TWIT	Twitter-2010	41,652,230	1,468,365,182	Social
HOST	Host-linkage	57,383,985	1,643,624,227	Web
GSH	Gsh-2015-host	68,660,142	1,802,747,600	Web
SK	Sk-2005	50,636,154	1,949,412,601	Web
TWIM	Twitter-mpi	52,579,682	1,963,263,821	Social
FRIE	Friendster	68,349,466	2,586,147,869	Social
UK	Uk-2006-05	77,741,046	2,965,197,340	Web
WEBS	Webspam-uk	105,896,555	3,738,733,648	Web

Collection³ [32], Koblenz Network Collection⁴ [33], Laboratory for Web Algorithms⁵ [34], [35], Network Data Repository⁶ [36], and the links in [37]. Note that, to verify the generality of our algorithms for processing distributed graphs, we do not transform the graphs into acyclic graphs, but build the indexes directly on the original graphs.

Environment. We implement all algorithms in C++ and compile them using GNU GCC 4.8.5. Our distributed algorithms (DRL^- , DRL , DRL_b) are designed to run on a vertex-centric system. To eliminate the dependence on specific features provided by the vertex-centric system, we implement this system ourselves using MPI. We map graph vertices to different computation nodes via vertex IDs to make it suitable for distributed setups.

Our algorithms are executed on a cluster of 32 computation nodes — each node contains an Intel Xeon 2.7 GHz CPU, 32 GB main memory, and runs Linux (Red Hat Linux 4.8.5, 64 bits). In contrast, centralized algorithms such as TOL [14] and BFL [9] are executed on only one node with the same system settings. If not explicitly stated, we only run one thread on each computation node. We set the cut-off time to 2 hours. If the algorithm runs out of memory or cannot complete the computation within the cut-off time, the execution time is marked as “INF”.

B. Comparison with Competitor Methods

Exp 1: Comparison with TOL. TOL is an index-only algorithm [14]. To illustrate the necessity of the proposed methods, we compare our best method DRL_b with TOL. The results of the comparison are given in Table VI. When a method cannot complete index creation because it exceeds the memory limit, we mark its results with the notation “-” in the table.

On index time. The time of DRL_b includes both computation time and communication time. On a medium-sized graph that can be accommodated by a single computation node, DRL_b 's index time can be at most 9.37 times faster than TOL. Note that TOL is a centralized algorithm and cannot handle distributed graphs. When a single computation node cannot accommodate a graph (e.g., WEBS), TOL fails to work. In contrast, DRL_b can index all graphs within half an hour. This shows that our method can efficiently handle large-scale graphs that are beyond the ability of TOL.

³<http://snap.stanford.edu/data/>

⁴<http://konect.uni-koblenz.de/>

⁵<http://law.di.unimi.it>

⁶<http://networkrepository.com/>

TABLE VI: The Comparison with Competitor Methods

Name	Index Time (sec)					Index Size (MB)					Query Time (sec)				
	BFL ^C	BFL ^D	TOL	DRL _b	DRL _b ^M	BFL ^C	BFL ^D	TOL	DRL _b	DRL _b ^M	BFL ^C	BFL ^D	TOL	DRL _b	DRL _b ^M
WEBW	1.51	59.21	61.84	9.08	7.31	85.35	85.35	432.06	432.06	432.06	8.58E-07	5.39E-05	2.09E-07	2.09E-07	2.09E-07
DBPE	2.10	110.17	2.21	0.92	0.64	154.07	154.07	63.91	63.91	63.91	2.25E-07	4.63E-05	1.51E-07	1.51E-07	1.51E-07
CITE	3.39	195.62	4.95	2.34	1.42	299.40	299.40	138.80	138.80	138.80	1.24E-07	4.52E-05	1.78E-07	1.78E-07	1.78E-07
CITP	5.10	138.30	125.21	13.36	11.17	172.80	172.80	622.04	622.04	622.04	5.68E-07	5.07E-05	3.12E-07	3.12E-07	3.12E-07
TW	3.74	469.34	7.27	1.13	1.40	829.52	829.52	271.60	271.60	271.60	1.95E-07	6.72E-05	1.84E-07	1.84E-07	1.84E-07
GO	3.56	365.38	7.40	1.76	2.03	318.97	318.97	274.43	274.43	274.43	1.11E-07	4.41E-05	2.02E-07	2.02E-07	2.02E-07
SINA	41.35	2822.48	—	136.32	—	2,685.05	2,685.05	—	13,691.20	—	2.82E-06	8.64E-05	—	6.76E-07	—
LINK	16.29	213.43	55.16	15.64	9.38	622.24	622.24	239.76	239.76	239.76	2.29E-07	7.31E-05	1.35E-07	1.35E-07	1.35E-07
WEBB	—	1,181.08	—	103.98	—	—	5,408.12	—	2,578.85	—	—	2.37E-04	—	1.84E-07	—
GRPH	46.30	6.36	76.31	24.00	16.44	780.20	780.20	325.01	325.01	325.01	9.61E-08	7.03E-05	8.71E-08	8.71E-08	8.71E-08
TWIT	57.44	304.55	134.87	62.82	35.79	1,906.69	1,906.69	766.17	766.17	766.17	1.55E-07	1.43E-04	1.06E-07	1.06E-07	1.06E-07
HOST	—	1,655.19	—	66.87	—	—	2,626.83	—	926.77	—	—	5.30E-04	—	2.26E-07	—
GSH	—	512.04	—	77.93	—	—	3,143.01	—	1,266.78	—	—	2.85E-04	—	1.37E-07	—
SK	—	219.47	—	82.85	—	—	2,317.94	—	975.83	—	—	1.80E-04	—	1.01E-07	—
TWIM	—	359.05	—	68.36	—	—	2,406.91	—	958.17	—	—	2.60E-04	—	2.76E-07	—
FRIE	—	688.47	—	112.77	—	—	3,128.79	—	1,240.01	—	—	3.58E-04	—	4.38E-07	—
UK	—	296.52	—	217.94	—	—	3,558.70	—	1,567.59	—	—	2.25E-04	—	2.81E-07	—
WEBS	—	747.92	—	188.58	—	—	4,847.56	—	2,063.97	—	—	4.81E-04	—	4.15E-07	—

On index size. For DRL_b, we aggregate the index distributed on different computation nodes on one node. Hence, our algorithms have the same index size as TOL. The index size of DRL_b (and TOL) on all graphs is very small. For example, the largest graph WEBS has an index size of 2.06 GB. This indicates that although distributed graphs may be large, the generated index can be accommodated on an ordinary machine to support in-memory queries. So, it is feasible to create the index of TOL for a distributed graph because the index size is not so large.

On query time. DRL_b creates the same index as TOL, so the query time is the same for TOL and DRL_b. The query time of DRL_b (and TOL) on all graphs is less than one microsecond. This result reinforces our research motivation that efficient reachability queries on a distributed graph can be supported by proposing new labeling methods to obtain the same index as TOL.

Exp 2: Comparison with BFL. We compare our best method DRL_b with BFL, which is an index-assisted method [9]. BFL uses DFS to create the index and may also use the online search (e.g. DFS) at query time. We use the code provided in [9] to implement the *centralized* BFL algorithm, which is denoted as BFL^C, where all parameters are set by default. Also, we implemented a distributed version of DFS (which is a core operation of BFL) to obtain the *distributed* BFL algorithm, which is denoted as BFL^D. BFL^C runs on one computation node, while BFL^D runs on 32 nodes. We compare DRL_b, BFL^C, and BFL^D and record the results in Table VI.

On index time. (1) First, we compare DRL_b with the centralized algorithm BFL^C. On medium-sized graphs, the index time of BFL^C is normally better than that of DRL_b. But the index time of DRL_b is comparable to that of BFL^C: the index time of DRL_b is within seven times that of BFL^C. Moreover, DRL_b can handle large-scale graphs for which BFL^C cannot create indexes. (2) Then, we compare DRL_b with the distributed algorithm BFL^D. BFL^D can process large-scale graphs by partitioning them to different computation nodes. But because BFL^D requires the distributed DFS to create indexes, its index time is very high: BFL^D runs on average 52.54 times slower than DRL_b.

On index size. Since the index sizes of BFL^C and BFL^D are the same on all graphs, we only compare the index sizes of BFL^D and DRL_b. It can be seen that the index size of BFL^D is small on all graphs (no larger than 6 GB), and the index of BFL^D is smaller than that of DRL_b on WEBW, CITP, and SINA. However, the

index size of DRL_b is smaller than that of BFL^D on the other graphs: the index size of DRL_b is on average 2.38 times smaller than that of BFL^D on these graphs.

On query time. (1) We first compare the query time of BFL^C and DRL_b. Both BFL^C and DRL_b can answer queries in microseconds, but the performance of DRL_b is better than BFL^C: on average, the query time of DRL_b is 1.8 times faster than that of BFL^C on graphs that BFL^C can process. (2) Then we compare the query time of BFL^D and DRL_b. Because BFL^D cannot avoid traversing the distributed graph to answer queries, the performance of BFL^D can be very bad: the query time of BFL^D is on average 867.6 times longer than that of DRL_b.

Overall, BFL performs well only when index creation can be done on a single node (see performance of BFL^C); when dealing with large graphs, BFL has high index time and query time due to the high cost of distributed DFS and graph search (see performance of BFL^D). This clarifies why we parallelize the index-only method TOL instead of BFL. Note that the graphs we use are not converted to be acyclic because of the high cost of performing such conversions in a distributed environment using DFS. This partly explains why there are some minor inconsistencies between our conclusions and [9].

Exp 3: Comparison with Multi-core Version. Our algorithm DRL_b (see Algorithm 4) achieves parallelism among multiple computation nodes in a distributed environment. Besides, we can also achieve parallelism among multiple threads (instead of multiple nodes), resulting in a multi-core version of DRL_b, which is denoted as DRL_b^M. For a fair comparison, we test DRL_b^M in a machine configured similarly to the single computation node used by DRL_b, and this machine contains 32 cores and has a memory size of 32 GB. OpenMP [38] is used to implement DRL_b^M. Since DRL_b and DRL_b^M have the same index size and query time, we compare only the index time between them. We record the results in Table VI and have the following findings.

On medium-sized graphs. Because DRL_b^M can use shared memory for data exchange [39], it avoids the communication cost of DRL_b. This leads to a better index time for DRL_b^M than for DRL_b in most cases: DRL_b^M is 1.34 times faster than DRL_b on graphs where DRL_b^M can create indexes. However, this speedup is limited. One possible reason is that the communication cost of DRL_b is relatively small compared to the computation cost (see the communication time and the computation time of DRL_b in Exp

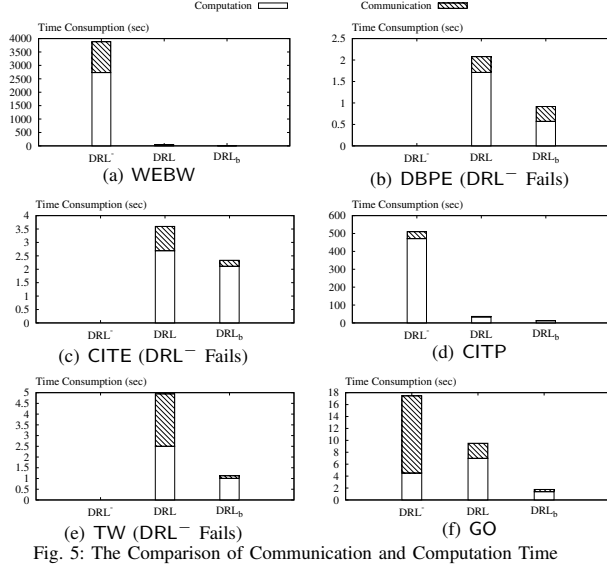


Fig. 5: The Comparison of Communication and Computation Time

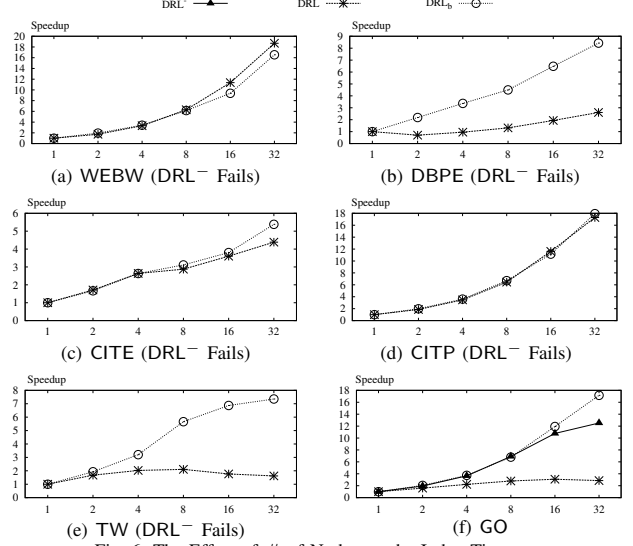


Fig. 6: The Effect of # of Nodes on the Index Time

4 for details), leading to a less prominent advantage of shared memory.

On large-scale graphs. Since DRL_b^M is a centralized algorithm, the usability of DRL_b^M is limited by memory and therefore cannot build indexes for massive graphs. For example, DRL_b^M ran out of memory when building the index for WEBS. On the other hand, DRL_b can allocate graphs to multiple computation nodes and thus is more suitable for large graph processing.

C. Comparison Between Proposed Algorithms

Exp 4: Communication and Computation Time. We compare our proposed labeling algorithms, DRL^- , DRL , and DRL_b . We divide the index time of the proposed algorithms into computation time and communication time. If an algorithm is unable to finish labeling within the cut-off time, we do not report its time and mark the failure at the title of that graph. Due to space constraints, we present in Fig. 5 the results on the first 6 graphs (WEBW, DBPE, CITE, CITP, TW, and GO) with the following findings. *Comparison of DRL^- and DRL .* Compared to DRL^- , DRL uses the inverted list to implement the refinement phase. We find that DRL can index on DBPE, CITE and TW, while DRL^- cannot. In addition, on the other three graphs, DRL has an average of 88.2 times shorter index time than DRL^- , thanks to the new refinement technique used by DRL .

Comparison of DRL and DRL_b . DRL_b uses batch labeling to further optimize DRL . We find that DRL_b has an average of 3.5 times shorter index time over DRL . Moreover, DRL_b reduces the computation time while substantially reducing the communication cost of DRL , which validates the effectiveness of the optimization strategy used by DRL_b .

Exp 5: Effect of Node Number. We used 32 computation nodes by default. To test the impact of the number of computation nodes on the proposed algorithms, we varied the number of nodes from 1, 2, 4, 8, 16 to 32 and recorded the corresponding index time. We define the **speedup** as the ratio of the index time on one node to the index time on x nodes, i.e., $speedup = \frac{\text{the index time on 1 node}}{\text{the index time on } x \text{ node}}$. If an algorithm fails to finish labeling within the cut-off time on 1 node, we do not report its speedup ratio and mark the failure

at the title of that graph. We show the speedup ratios on the first six graphs in Fig. 6 with the following findings.

DRL_b has a satisfying speedup ratio. The maximum speedup ratio for DRL_b with 32 nodes is 17.93 (on CITP) compared to using a single node. Moreover, the speedup ratio of DRL_b shows an increasing trend as the number of nodes increases.

The speedup of DRL^- and DRL has limitations. Although the maximum speedup ratio of DRL^- is 12.54 (on GO), DRL^- cannot finish labeling on other five graphs using a single node within the cut-off time. On the other hand, although the maximum speedup ratio of DRL is 18.69 (on WEBW), on TW, the ratio of DRL is only 2.86 while that of DRL_b is 17.2, which shows that introducing the batch labeling optimization maintains a better speedup ratio.

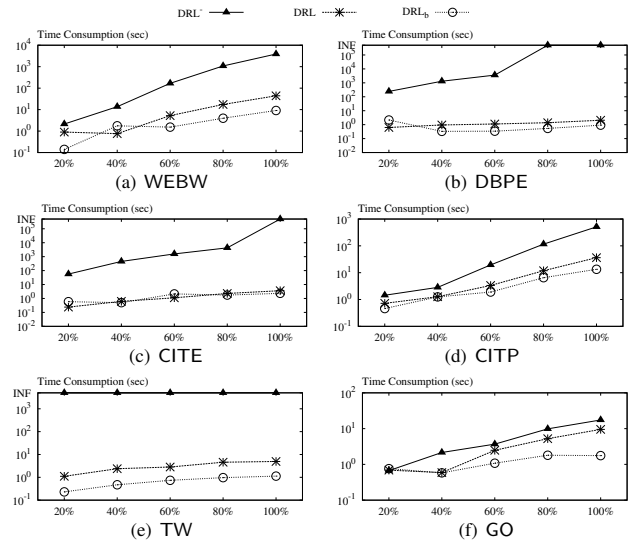


Fig. 7: The Test of the Scalability on the Index Time

Exp 6: Test of Scalability. In testing the scalability of the proposed algorithms, we divide the edges of the graph into five disjoint groups, each group consisting of $\frac{1}{5}$ edges of the

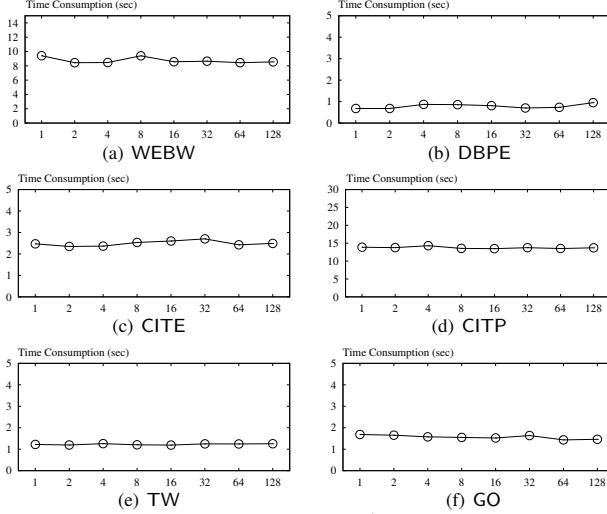


Fig. 8: The Effect of Initial Batch Size b on the Index Time

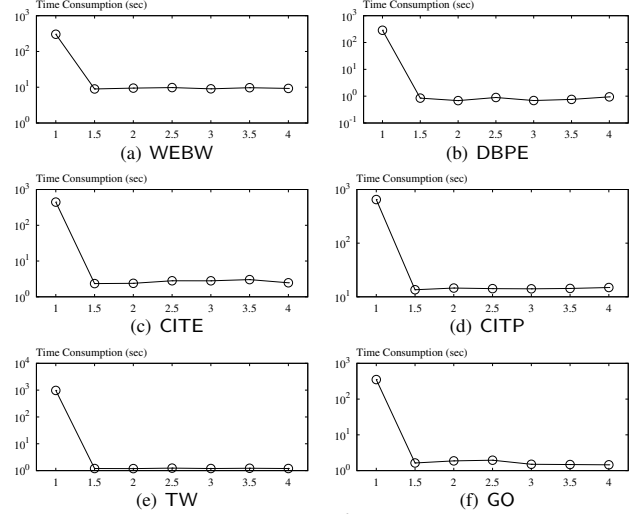


Fig. 9: The Effect of Factor k on the Index Time

original graph. We generate five test graphs, where the i -th test graph contains edges in the first i groups. The experiments are conducted on five test graphs for each dataset. Since the index size and query time are the same for all algorithms, we omit the discussion of them and only provide the effect on the index time in Fig. 7. We have the following finding.

The proposed algorithms exhibit good scalability. The index time of all algorithms improves as the graph size becomes larger. However, the increase of all methods is smooth. For example, the index time of DRL_b for the test graph with 100% edges is 4.8 times longer than that of the test graph with 20% edges on TW. This indicates the good scalability of the proposed algorithms.

D. Effect of Parameters on Index Time

In Exp-4, we verified that using the batch labeling optimization (forming DRL_b) can speed up the index time of DRL. For DRL_b , two parameters need to be set to generate the batch sequence: the initial batch size b and the incremental factor k . We set both parameters to 2 by default, but we need to further test the effect of these two parameters on the index performance (time) of DRL_b .

Exp 7: Effect of Initial Batch Size b . We first analyze the effect of b . We vary the value of b from 1, 2, 4, 8, 16, 32, 64, to 128. We record the index time of DRL_b for different values of b and present the results of DRL_b on the first 6 graphs in Fig 8. We have the following findings.

b has little effect on the index time. As the value of b varies, the difference between the maximum index time and the minimum index time on all used graphs is no more than 1.5 times. This indicates that DRL_b is not sensitive to the parameter b .

The default value of 2 is a good choice. On some graphs (e.g., WEBW and DBPE), setting b to 2 leads to a local minimum index time. This explains why 2 is used as the default value.

Exp 8: Effect of Factor k . We analyze the effect of another parameter k on the index time. We vary the value of k from 1, 1.5, 2, 2.5, 3, 3.5, to 4. We report the index time of DRL_b for different k in Fig. 9 and obtain the following findings.

When k is not 1. When k is taken other than 1, the index time does not vary much: the difference between the maximum and

minimum index time on all graphs does not exceed 1.4. Also, on some graphs (e.g., DBPE and CITE), the index time reaches a local minimum when k is 2, so 2 is used as the default value.

When k is 1. The index time becomes very slow when k is taken as 1: the index time is up to 812 times slower when k is 1 than when k is taken as other. This further explains why k needs to be set to 2 as the default value.

VII. CONCLUSIONS

We develop novel labeling methods to produce the same indexes as TOL on distributed graphs. To overcome the limitation that TOL cannot be executed in parallel, we resort to finding the backward label set of each vertex. We propose to use a filtering-and-refinement framework to find backward label sets. Using this framework, we design new labeling algorithms and further improve the efficiency by batch labeling optimization. Experimental results show that our algorithms can efficiently handle distributed graphs. The follow-up work is to maintain the indexes on distributed dynamic graphs. Another important research direction is to further explore the features of distributed and multi-core systems to accelerate the construction of indexes.

APPENDIX

Proof of Lemma 1. We verify only the case $v_i \in L_{in}(w)$.

- \Leftarrow : If $L_{out}^i(v_i) \cap L_{in}^i(w) = \emptyset$, then there is no vertex u with $ord(u) > ord(v_i)$ such that $v_i \rightarrow u \rightarrow w$. Then, v_i is the highest-order vertex on all paths from v_i to w . By Theorem 1, $v_i \in L_{in}(w)$.
- \Rightarrow : If $v_i \in L_{in}(w)$ but $L_{out}^i(v_i) \cap L_{in}^i(w) = S \neq \emptyset$, we choose $s \in S$ whose order is the highest in S . $s \neq v_i$ since $v_i \notin L_{in}^i(v_i)$ by definition. Moreover, the fact that $s \in S$ yields $v_i \rightarrow s \rightarrow u$ and $ord(s) > ord(v_i)$. By Theorem 1, $v_i \notin L_{in}(w)$, contradiction.

ACKNOWLEDGMENT

Lu Qin is supported by ARC FT200100787 and DP210101347. Ying Zhang is supported by ARC FT170100128 and ARC DP210101393. Xuemin Lin is supported by ARC DP180103096 and ARC DP170101628.

REFERENCES

- [1] S. Sahu, A. Mhedhbi, S. Salihoglu, J. Lin, and M. T. Özsu, “The ubiquity of large graphs and surprising challenges of graph processing: extended survey,” *The VLDB Journal*, vol. 29, no. 2, pp. 595–618, 2020.
- [2] J. X. Yu and J. Cheng, “Graph reachability queries: A survey,” in *Managing and Mining Graph Data*. Springer, 2010, pp. 181–215.
- [3] R. Jin and G. Wang, “Simple, fast, and scalable reachability oracle,” *Proceedings of the VLDB Endowment*, vol. 6, no. 14, pp. 1978–1989, 2013.
- [4] S. Beamer, K. Asanovic, and D. Patterson, “Direction-optimizing breadth-first search,” in *SC’12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2012, pp. 1–10.
- [5] R. Tarjan, “Depth-first search and linear graph algorithms,” *SIAM journal on computing*, vol. 1, no. 2, pp. 146–160, 1972.
- [6] R. Jin, N. Ruan, S. Dey, and J. Y. Xu, “Scarab: scaling reachability computation on large graphs,” in *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*. ACM, 2012, pp. 169–180.
- [7] H. Yildirim, V. Chaoji, and M. J. Zaki, “Grail: Scalable reachability index for large graphs,” *Proceedings of the VLDB Endowment*, vol. 3, no. 1-2, pp. 276–284, 2010.
- [8] H. Wei, J. X. Yu, C. Lu, and R. Jin, “Reachability querying: An independent permutation labeling approach,” *Proceedings of the VLDB Endowment*, vol. 7, no. 12, pp. 1191–1202, 2014.
- [9] J. Su, Q. Zhu, H. Wei, and J. X. Yu, “Reachability querying: Can it be even faster?” *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 3, pp. 683–697, 2016.
- [10] E. Cohen, E. Halperin, H. Kaplan, and U. Zwick, “Reachability and distance queries via 2-hop labels,” *SIAM Journal on Computing*, vol. 32, no. 5, pp. 1338–1355, 2003.
- [11] R. Schenkel, A. Theobald, and G. Weikum, “Hopi: An efficient connection index for complex xml document collections,” in *International Conference on Extending Database Technology*. Springer, 2004, pp. 237–255.
- [12] J. Cheng, J. X. Yu, X. Lin, H. Wang, and S. Y. Philip, “Fast computation of reachability labeling for large graphs,” in *International Conference on Extending Database Technology*. Springer, 2006, pp. 961–979.
- [13] J. Cheng, J. X. Yu, X. Lin, H. Wang, and P. S. Yu, “Fast computing reachability labelings for large graphs with high compression rate,” in *Proceedings of the 11th international conference on Extending database technology: Advances in database technology*. ACM, 2008, pp. 193–204.
- [14] A. D. Zhu, W. Lin, S. Wang, and X. Xiao, “Reachability queries on large dynamic graphs: a total order approach,” in *Proceedings of the 2014 ACM SIGMOD international conference on Management of data*. ACM, 2014, pp. 1323–1334.
- [15] T. Zhang, Y. Gao, C. Li, C. Ge, W. Guo, and Q. Zhou, “Distributed reachability queries on massive graphs,” in *International Conference on Database Systems for Advanced Applications*. Springer, 2019, pp. 406–410.
- [16] W. Fan, X. Wang, and Y. Wu, “Performance guarantees for distributed reachability queries,” *Proceedings of the VLDB Endowment*, vol. 5, no. 11, pp. 1304–1316, 2012.
- [17] Y. Yano, T. Akiba, Y. Iwata, and Y. Yoshida, “Fast and scalable reachability queries on graphs by pruned labeling with landmarks and paths,” in *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*. ACM, 2013, pp. 1601–1606.
- [18] L. G. Valiant, “A bridging model for parallel computation,” *Communications of the ACM*, vol. 33, no. 8, pp. 103–111, 1990.
- [19] X. Feng, L. Chang, X. Lin, L. Qin, and W. Zhang, “Computing connected components with linear communication cost in pregel-like systems,” in *2016 IEEE 32nd International Conference on Data Engineering (ICDE)*. IEEE, 2016, pp. 85–96.
- [20] T. Zhang, Y. Gao, L. Chen, W. Guo, S. Pu, B. Zheng, and C. S. Jensen, “Efficient distributed reachability querying of massive temporal graphs,” *The VLDB Journal*, vol. 28, no. 6, pp. 871–896, 2019.
- [21] A. Buluç and K. Madduri, “Parallel breadth-first search on distributed memory systems,” in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011, pp. 1–12.
- [22] B. Awerbuch, “A new distributed depth-first-search algorithm,” *Information Processing Letters*, vol. 20, no. 3, pp. 147–150, 1985.
- [23] B. H. Bloom, “Space/time trade-offs in hash coding with allowable errors,” *Communications of the ACM*, vol. 13, no. 7, pp. 422–426, 1970.
- [24] E. Nuutila, “Efficient transitive closure computation in large digraphs,” 1998.
- [25] S. J. van Schaik and O. de Moor, “A memory efficient reachability data structure through bit vector compression,” in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data*. ACM, 2011, pp. 913–924.
- [26] R. Agrawal, A. Borgida, and H. V. Jagadish, “Efficient management of transitive relationships in large data and knowledge bases,” in *ACM SIGMOD Record*, vol. 18, no. 2. ACM, 1989, pp. 253–262.
- [27] Y. Chen and Y. Chen, “An efficient algorithm for answering graph reachability queries,” in *2008 IEEE 24th International Conference on Data Engineering*. IEEE, 2008, pp. 893–902.
- [28] R. Jin, Y. Xiang, N. Ruan, and H. Wang, “Efficiently answering reachability queries on very large directed graphs,” in *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*. ACM, 2008, pp. 595–608.
- [29] W. Li, M. Qiao, L. Qin, Y. Zhang, L. Chang, and X. Lin, “Scaling distance labeling on small-world networks,” in *Proceedings of the 2019 International Conference on Management of Data*, 2019, pp. 1060–1077.
- [30] K. Lakhotia, R. Kannan, Q. Dong, and V. Prasanna, “Planting trees for scalable and efficient canonical hub labeling,” *Proceedings of the VLDB Endowment*, vol. 13, no. 4.
- [31] J. Cheng, S. Huang, H. Wu, and A. W.-C. Fu, “Tf-label: a topological-folding labeling scheme for reachability querying in a large graph,” in *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. ACM, 2013, pp. 193–204.
- [32] J. Leskovec and A. Krevl, “Snap datasets: Stanford large network dataset collection,” 2014.
- [33] J. Kunegis, “Konect: the koblenz network collection,” in *Proceedings of the 22nd International Conference on World Wide Web*. ACM, 2013, pp. 1343–1350.
- [34] P. Boldi and S. Vigna, “The WebGraph framework I: Compression techniques,” in *Proc. of the Thirteenth International World Wide Web Conference (WWW 2004)*. Manhattan, USA: ACM Press, 2004, pp. 595–601.
- [35] P. Boldi, M. Rosa, M. Santini, and S. Vigna, “Layered label propagation: A multiresolution coordinate-free ordering for compressing social networks,” in *Proceedings of the 20th international conference on World Wide Web, S. Srinivasan, K. Ramamritham, A. Kumar, M. P. Ravindra, E. Bertino, and R. Kumar, Eds.* ACM Press, 2011, pp. 587–596.
- [36] R. A. Rossi and N. K. Ahmed, “The network data repository with interactive graph analytics and visualization,” in *AAAI*, 2015. [Online]. Available: <http://networkrepository.com>
- [37] J. Zhou, S. Zhou, J. X. Yu, H. Wei, Z. Chen, and X. Tang, “Dag reduction: Fast answering reachability queries,” in *Proceedings of the 2017 ACM International Conference on Management of Data*. ACM, 2017, pp. 375–390.
- [38] L. Dagum and R. Menon, “Openmp: an industry standard api for shared-memory programming,” *IEEE computational science and engineering*, vol. 5, no. 1, pp. 46–55, 1998.
- [39] R. Chandra, L. Dagum, D. Kohr, R. Menon, D. Maydan, and J. McDonald, *Parallel programming in OpenMP*. Morgan kaufmann, 2001.