

ThunderRW: An In-Memory Graph Random Walk Engine

Shixuan Sun
National University of Singapore
sunsx@comp.nus.edu.sg

Yuhang Chen
National University of Singapore
yuhangc@comp.nus.edu.sg

Shengliang Lu
National University of Singapore
lusl@comp.nus.edu.sg

Bingsheng He
National University of Singapore
hebs@comp.nus.edu.sg

Yuchen Li
Singapore Management University
yuchenli@smu.edu.sg

ABSTRACT

As random walk is a powerful tool in many graph processing, mining and learning applications, this paper proposes an efficient in-memory random walk engine named **ThunderRW**. Compared with existing parallel systems on improving the performance of a single graph operation, ThunderRW supports massive parallel random walks. The core design of ThunderRW is motivated by our profiling results: common RW algorithms have as high as 73.1% CPU pipeline slots stalled due to irregular memory access, which suffers significantly more memory stalls than the conventional graph workloads such as BFS and SSSP. To improve the memory efficiency, we first design a generic step-centric programming model named Gather-Move-Update to abstract different RW algorithms. Based on the programming model, we develop the step interleaving technique to hide memory access latency by switching the executions of different random walk queries. In our experiments, we use four representative RW algorithms including PPR, DeepWalk, Node2Vec and MetaPath to demonstrate the efficiency and programming flexibility of ThunderRW. Experimental results show that ThunderRW outperforms state-of-the-art approaches by an order of magnitude, and the step interleaving technique significantly reduces the CPU pipeline stall from 73.1% to 15.0%.

PVLDB Reference Format:

Shixuan Sun, Yuhang Chen, Shengliang Lu, Bingsheng He, and Yuchen Li. ThunderRW: An In-Memory Graph Random Walk Engine. PVLDB, 14(11): 1992-2005, 2021.
doi:10.14778/3476249.3476257

PVLDB Availability Tag:

The source code of this research paper has been made publicly available at <https://github.com/Xtra-Computing/ThunderRW>.

1 INTRODUCTION

Random walk (RW) is an effective tool for extracting relationships between entities in a graph, and is widely used in many applications such as *Personalized PageRank* (PPR) [45], *SimRank* [20], *Random Walk Domination* [31], *Graphlet Concentration* (GC) [48], *Network Community Profiling* (NCP) [12], *DeepWalk* [47] and *Node2Vec* [16]. For graph analysis tasks such as GC and NCP, RW queries generally

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 11 ISSN 2150-8097.
doi:10.14778/3476249.3476257

Algorithm 1: Execution Paradigm of RW algorithms

Input: a graph G and a set Q of random walk queries;
Output: the walk sequences of each query in Q ;

```
1 foreach  $Q \in \mathbb{Q}$  do
2   do
3     Select a neighbor of the current residing vertex  $Q.cur$  at random;
4     Add the selected vertex to  $Q$ ;
5   while  $Terminate(Q)$  is false;
6 return  $Q$ ;
```

dominate the cost [12, 48]. Even for graph representation learning, the cost of sampling RW is non-trivial, for example, a naive implementation of Node2Vec takes more than eight hours on the *twitter* graph in our experiments. Moreover, increasing the number of RW queries can improve the effectiveness of RW algorithms [16, 48]. Therefore, accelerating RW queries is an important problem.

RW algorithms generally follow the execution paradigm illustrated in Algorithm 1, which consists of massive RW queries. Each query Q starts from a given source vertex. At each step, Q moves to a neighbour of the current residing vertex at random, and repeats this process until satisfying a specific termination condition, e.g., a target length is reached (Lines 2-5). Despite that RW algorithms follow a similar execution paradigm, there are quite some variants of RW algorithms, which can differ significantly in neighbor selections (see Section 2.2). Encouraged by the success of in-memory graph processing engines [44, 55, 58, 68], there have been some recent systems designed specifically for RW algorithms, including C-SAW [46], GraphWalker [60] and KnightKing [66]. They focus on accelerators, disk-based or distributed settings, without specially optimizing in-memory execution of RW queries. However, with the rapid development of hardwares, modern servers equip with hundred gigabytes, even several terabytes memory, which empowers in-memory processing of graphs with hundred billions of edges. This covers many real-world graphs in applications [9]. As such, this paper studies the design and implementation of an in-memory graph engine for RW algorithms.

To crystallize the performance factors for in-memory RW executions, we conduct profiling studies on RW algorithms in comparison with conventional workloads of a single graph operation like BFS and SSSP (see Section 3). Our profiling results show that common RW algorithms have as high as 73.1% CPU pipeline slots stalled due to irregular memory access, which suffers significantly more memory stalls than the conventional workloads. Consequently, the CPUs frequently wait on the high-latency access to the main memory, which becomes the major performance bottleneck. Besides, we observe that the sampling methods such as *inverse transformation sampling* [40], *alias sampling* [59] and *rejection sampling* [50]

have significant varying performance on different RW algorithms (with the difference as much as 6 times). Thus, it requires non-trivial and significant engineering efforts to develop any efficient RW algorithms considering the cache stall bottleneck, as well as parallelization and the choice of sampling methods.

In this paper, we propose **ThunderRW**, a generic and efficient in-memory RW framework. We employ a *step-centric* programming model abstracting the computation from the local view of moving one step of a walker. Users implement their RW algorithms by "thinking like a walker" in user-defined functions (UDF). The framework applies UDFs to each query and parallelizes the execution by regarding a step of a query as a task unit. Furthermore, ThunderRW provides variant sampling methods so that users can select an appropriate one based on the characteristics of workloads. Built upon the step-centric programming model, we propose the *step interleaving* technique to resolve the cache stalls caused by irregular memory access with *software prefetching* [29]. As modern CPUs can process multiple memory access requests simultaneously [64], the core idea of step interleaving is to hide memory access latency by issuing multiple outstanding memory accesses, which exploits *memory level parallelism* [2] among different RW queries.

We demonstrate the generality and programming flexibility of ThunderRW by showcasing four representative algorithms including PPR [45], DeepWalk [47], Node2Vec [16] and MetaPath [57]. We conduct extensive experiments with twelve real-world graphs. The results show that (1) ThunderRW runs 8.6-3333.1X faster than the naive implementation in popular open-source packages; (2) ThunderRW provides speedups of 1.7-14.6X over the state-of-the-art frameworks including GraphWalker [60] and KnightKing [66] running on the same machine; and (3) the step interleaving technique significantly reduces the memory stalls from 73.1% to 15.0%.

2 BACKGROUND AND RELATED WORK

2.1 Preliminary

This paper focuses on the directed graph $G = (V, E)$ where V is a set of vertices and E is a set of edges. An undirected graph can be supported by representing each undirected edge with two directed edges with the same two vertexes in our system. Given a vertex $v \in V$, N_v denotes the neighbors of v , i.e., $\{v' | e(v, v') \in E\}$ where $e(v, v')$ represents the edge between v and v' . The degree d_v denotes the number of neighbors of v . E_v is the set of edges adjacent to v , i.e., $\{e(v, v') | v' \in N_v\}$. Given $e \in E$ (resp. $v \in V$), w_e and l_e (resp. w_v and l_v) represent its weight and label, respectively. Given G , a RW Q is a stochastic process on G , which consists of a sequence of adjacent vertices. $Q[i]$ is the i th vertex in the sequence where i starts from 0. $Q.cur$ is the current residing vertex of Q . $|Q|$ is the number of vertices in Q . Suppose that $Q.cur$ is v . Given $e \in E_v$, we call the probability of e being selected the *transition probability*, which is represented by $p(e)$. Then, the neighbor selection is equivalent to sampling from the discrete probability distribution $P = \{p(e)\}$ where $e \in E_v$. Specifically, it is to pick an element h from E_v based on the distribution of P , i.e., $P[h = e] = p(e)$. For example, if the relative chance of e being selected is proportional to the edge weight w_e , then $p(e) = \hat{w}_e$ is the normalized probability where $\hat{w}_e = \frac{w_e}{\sum_{e' \in E_v} w_{e'}}$.

2.2 Random Walk based Algorithms

RW algorithms generally follow the execution paradigm in Algorithm 1. They mainly differ in the neighbor selection step. We first categorize them into *unbiased* and *biased* based on the transition probability properties. Unbiased RW selects each edge $e \in E_v$ with the same probability where $v = Q.cur$, while the transition probability is nonuniform for biased RWs, e.g., depending on the edge weight. We further classify the biased RWs into *static* and *dynamic*. If the transition probability is determined before execution, then RW is static. Otherwise, it is dynamic, which is affected by states of RW queries. In the following, we introduce four representative RW algorithms that have been used in many applications.

PPR (Personalized PageRank) [45] assigns a score to each vertex v' in the graph from the personalized view of a given source v , which describes how much v is interested in (or similar to) v' . A common solution for this problem is to start a number of RW queries from v , which have a fixed termination probability at each step, and approximately calculates the scores based on the distribution of the end vertices of random walk queries [11, 33]. The algorithms generally set RW queries as unbiased [34].

DeepWalk [47] is a graph embedding technique widely used in machine learning. It is developed based on the SkipGram model [41]. For each vertex, it starts a specified number of RW queries with a target length to generate embeddings. The original DeepWalk is unbiased, while the recent work [6] extends it to consider the edge weight, which becomes biased (static) random walk.

Node2Vec [16] is a popular graph embedding technique based on the second-order random walk. Different from DeepWalk, its transition probability depends on the last vertex visited. Suppose that $Q.cur$ is v . Equation 1 describes the transition probability of selecting the edge $e(v, v')$ where u is the last vertex visited, $dist(v', u)$ is the distance between v' and u , and a and b are two hyperparameters controlling the random walk behaviour. Node2Vec is dynamic because the transition probability relies on the states of queries. Moreover, It can take the edge weight into the consideration by multiplying $p(e)$ with w_e .

$$p(e(v, v')) = \begin{cases} \frac{1}{a} & \text{if } dist(v', u) = 0, \\ 1 & \text{if } dist(v', u) = 1, \\ \frac{1}{b} & \text{if } dist(v', u) = 2. \end{cases} \quad (1)$$

MetaPath [57] is a powerful tool to extract semantics information from heterogeneous information networks, and is widely used in machine learning tasks such as natural language processing [28, 38]. The RW queries are associated with a *meta-path schema* H , which defines the pattern of the walk paths based on the edge type, e.g., "write->publish->mention". Let $H[i]$ be the i th label in H . At each step, the RW query only considers the edges $e \in E_v$ where $v = Q.cur$ such that l_e is equal to $H[|Q|]$. In other words, if $l_e \neq H[|Q|]$, then $p(e) = 0$. Thus, the transition probability depends on the states of the RW, and MetaPath is dynamic.

2.3 Sampling Methods

Sampling from a discrete probability distribution $P = \{p_0, p_1, \dots, p_{n-1}\}$ is to select an element h from $\{0, 1, \dots, n-1\}$ based on P (i.e., $P[h = i] = p_i$). In this paper, we focus on five sampling techniques, including *naive sampling*, *inverse transformation sampling*

[40], *alias sampling* [59], *rejection sampling* [50] and a special case of *rejection sampling* [66] because they are efficient and widely used [46, 52, 53, 60, 66]. Naive sampling only works on the uniform discrete distribution, while the other four can handle non-uniform and select the element h in two phases: *initialization*, which preprocesses the distribution P , and *generation*, which picks an element on the basis of the initialization result. Please refer to [52] for the details. In the following, we briefly introduce the sampling methods in the context of this paper, i.e., selecting an edge from E_v based on the transition probability distribution P where $v = Q.cur$.

Naive sampling (NAIVE). This method generates a uniform random integer number x in the range $[0, d_v)$ and picks $E_v[x]$, which is the x th element in E_v . It only works on the uniform discrete distribution. The time and space complexities are both $O(1)$.

Inverse transformation sampling (ITS). The initialization phase of ITS computes the *cumulative distribution function* of P as follows: $P' = \{p'_i = \sum_{j=0}^i p_j\}$ where $0 \leq i < d_v$. After that, the generation phase first generates a uniform real number x in $[0, p'_{d_v-1})$, then uses a binary search to find the smallest index i such that $x < p'_i$, and finally selects $E_v[i]$. The time complexity of the initialization is $O(d_v)$, and that of the generation is $O(\log d_v)$. As ITS needs to store P' , the space complexity is $O(d_v)$.

Alias sampling (ALIAS). The initialization phase builds two tables: the *probability table* H , and the *alias table* A . Both of them have d_v values. $H[i]$ and $A[i]$ represent the i th value of H and A , respectively. Given $0 \leq i < d_v$, $A[i]$ is a bucket containing one or two elements from $\{0, 1, \dots, d_v - 1\}$, which are denoted by $A[i].first$ and $A[i].second$, respectively. $H[i]$ is the probability selecting $A[i].first$. If $A[i]$ has only one element, then $A[i].second$ is *null* and $H[i]$ is equal to 1. The generation phase first generates a uniform integer number x in $[0, d_v)$ and then retrieves $H[x]$ and $A[x]$. Next, it generates a uniform real number y in $[0, 1)$. If $y < H[x]$, then picks $e(v, A[x].first)$. Otherwise, the edge selected is $e(v, A[x].second)$. The time complexity of initialization is $O(d_v)$ and that of generation is $O(1)$. The space complexity is $O(d_v)$.

Rejection sampling (REJ). The initialization phase of REJ gets $p^* = \max_{p \in P} p$. The generation phase can be viewed as throwing darts on a rectangle dartboard until hitting the target area. Specifically, it has two steps: (1) generate a uniform integer number x in $[0, d_v)$ and a uniform real number y in $[0, p^*)$ (i.e., the dart is thrown at the position (x, y)); and (2) if $y < p_x$, then select $E_v[x]$ (i.e., hit the target area); otherwise, repeat Step (1). The time complexity of initialization is $O(d_v)$, and that of generation is $O(\mathbb{E})$ where $\mathbb{E} = \frac{d_v \times p^*}{\sum_{p \in P} p}$ (i.e., the area of the rectangle board divides the target area). Based on the computation method of \mathbb{E} , we can get that $1 \leq \mathbb{E} \leq d_v$. The space complexity is $O(1)$.

A special case of REJ (O-REJ). A special case of REJ is that we can set a value $p^* \geq \max_{p \in P} p$ without the initialization phase, but of keeping $\mathbb{E} = \frac{d_v \times p^*}{\sum_{p \in P} p}$ is close to $\frac{d_v \times \max_{p \in P} p}{\sum_{p \in P} p}$. For example, set p^* to $\max\{1, \frac{1}{a}, \frac{1}{b}\}$ for Node2Vec [66]. The generation phase is the same as REJ. Therefore, the time complexity is $O(\mathbb{E})$ where $\mathbb{E} = \frac{d_v \times p^*}{\sum_{p \in P} p}$ and p^* is specified by users. The space complexity is $O(1)$.

In existing works, unbiased random walks (e.g., PPR [45] and unweighted DeepWalk [47]) adopt NAIVE sampling. In contrast, biased

random walks (e.g., weighted DeepWalk [8, 67], Node2Vec [16] and MetaPath [13, 19]) use ALIAS sampling because the time complexity of the generation phase is $O(1)$. C-SAW [46] adopts ITS to utilize the parallel computation capability of GPUs to calculate the prefix sum. KnightKing [66] uses O-REJ to avoid scanning neighbors of $Q.cur$ to reduce the network communication cost.

2.4 Related Work

Graph computing frameworks. There are a number of generic graph computing frameworks working on different computation environments, for example, (1) Single Machine (CPUs): GraphChi [27], Ligra [55], Graphene [32], and GraphSoft [23]; (2) GPUs: Medusa [69], CuSha [24] and Gunrock [62]; and (3) Distributed Environment: Pregel [39], GraphLab [36], PowerGraph [14], GraphX [15], Blogel [65], Gemini [71], and Grapes [10]. They usually adopt vertex- or edge-centric model, and are highly optimized for a single graph operation. In contrast, ForkGraph [37] targets at graph algorithms consisting of concurrent graph queries, for example, betweenness centrality. However, all of them focus on traditional graph query operations such as BFS and SSSP without considering RW workloads. That motivates the development of engines specially optimized for RW [46, 60, 66].

Random walk frameworks. In contrast to graph computing frameworks abstracting the computation from the view of the graph data, existing RW frameworks adopt the *walker-centric* model, which regards each query as the parallel task. KnightKing [66] is a distributed framework. It adopts the BSP model that moves a step for all queries at each iteration until all queries complete. To reduce data transfers in network, it utilizes O-REJ sampling to avoid scanning E_v where $v = Q.cur$. It exposes an API for users to set a suitable upper bound for the edge transition probability for each edge adjacent to $Q.cur$. Unfortunately, we find that this design introduces an implicit constraint on RW algorithms: a suitable upper bound must be determined without looping over E_v . This works well for Node2Vec by setting the upper bound as $\max\{1.0/a, 1.0, 1.0/b\}$ according to Equation 1. However, it cannot handle MetaPath because the transition probability of each $e \in E_v$ can be zero because of the label filter. Another limitation is that KnightKing can suffer the tail problem since it moves a step for all queries at an iteration, whereas queries can have variant lengths.

C-SAW [46] is a framework on GPUs. It adopts the BSP model as well. To utilize the parallel computation capability in the many-core architecture, C-SAW uses ITS sampling in computation. Particularly, for all random walk types including unbiased, static and dynamic, C-SAW first conducts a prefix sum on the transition probability of edges adjacent to $Q.cur$, and then selects an edge. Consequently, it incurs high overhead for unbiased and static random walks. Moreover, C-SAW cannot support random walks with variant lengths (e.g., PPR) since such RW queries can degrade the utilization of GPUs. Additionally, Node2Vec is not supported by C-SAW, because C-SAW does not support the distance verification on GPUs.

GraphWalker [60] is an I/O efficient framework on a single machine. For a graph that cannot reside in memory, GraphWalker divides it into a set of partitions, and focuses on optimizing the scheduling of loading partitions into memory to reduce the number of I/Os. Specifically, for each partition, GraphWalker records the

number of queries residing in it, and the scheduler prioritizes partitions with more queries. Given a partition G' in memory, GraphWalker adopts the ASP model to execute queries in it. It assigns a query Q to each worker (i.e., a thread), and executes it independently until Q completes or jumps out G' . Once all queries in G' complete or leave G' , GraphWalker swaps it out, and reads the partition with most queries in disk. It repeats this process till all queries complete. GraphWalker supports unbiased RW only.

This paper focuses on accelerating the in-memory execution of RW queries. ThunderRW abstracts the computation of RW algorithms from the perspective of queries as well to exploit the parallelism in RW algorithms, but takes the *step-centric* model, which regards one step of a query as the task unit and factors one step into the gather-move-update operations to empower the step interleaving technique. Moreover, ThunderRW supports all the five sampling methods in Section 2.3 so that users can adopt an appropriate sampling method given a specific workload. ThunderRW supports all the four RW-algorithms in Section 2.2, which demonstrates its programming flexibility over other RW frameworks.

RW algorithm optimization. Due to the importance of the RW-based applications, a variety of algorithm-specific optimizations have been proposed for different RW applications, e.g., PPR [17, 35, 54, 61, 63], Node2Vec [70] and second-order random walks [53]. In contrast, we aim to design a generic and efficient random walk framework on which users can easily implement different kinds of random walk applications. Thus, the algorithm-specific optimizations are beyond the scope of this paper.

Prefetching in databases. Our step-interleaving techniques are inspired by the prefetching techniques in query processing of databases. As the performance gap between main memory and CPU widens, prefetching has been an effective means to improve database performance. There have been studies applying prefetching to B-tree index [5] and hash joins [1, 4, 21, 25]. Hash joins are probably the most widely studied operator for prefetching. The group prefetching (GP) and software pipeline prefetching (SPP) [4] are the classic prefetching technique for hash joins, which rearrange a sequence of operations in a loop to several stages and execute all queries stage by stage in batch. However, GP and SPP cannot efficiently handle queries with irregular access patterns, for example a binary search performs three searches to find the target value, while the other one needs four times. To resolve the problem, AMAC [26] proposes to execute the stages of each query asynchronously by explicitly maintaining the states of each stage. However, AMAC incurs more overhead than GP and SPP, especially when there are a number of stages because it needs to maintain the states of each stage. As in the context of random walk, there is a lack of a model to abstract stages from a sequence of operations and model their dependency relationships to guide the implementation.

3 MOTIVATIONS

In this section, we study the profiling results to assess the performance bottlenecks of in-memory computation of RW algorithms. We execute RW queries with different sampling methods and examine the hardware utilization with the *top-down microarchitecture analysis method* (TMAM). In the following, we first introduce TMAM and then present the profiling results.

Table 1: Comparison of pipeline slot breakdown and memory bandwidth (the total value of read and write) between traditional graph algorithms and RW algorithms.

Method	Front End	Bad Spec	Core	Memory	Retiring	Memory Bandwidth
BFS	11.6%	9.1%	20.8%	40.6%	18.0%	51.7 GB/s
SSSP	9.1%	12.5%	24.9%	36.9%	16.6%	38.2 GB/s
PPR	0.6%	0.7%	15.8%	73.1%	9.7%	1.4 GB/s
DeepWalk	1.0%	3.9%	16.7%	69.7%	8.7%	5.6 GB/s
Node2Vec	11.5%	22.1%	24.3%	28.1%	14.1%	17.1 GB/s
MetaPath	6.2%	7.5%	29.7%	33.9%	22.7%	9.9 GB/s

Top-down analysis method (TMAM) [7]. TMAM is a simplified and intuitive model for identifying the performance bottlenecks in out-of-order CPUs. It uses the *pipeline slot* to represent the hardware resources required to process the micro-operations (uOps). In a cycle, a pipeline slot is either empty (*stalled*) or filled with a uOp. The execution stall is caused by the *front-end* or the *back-end* part of the pipeline. The back-end cannot accept new operations due to the lack of required resources. It can be further split into *memory bound*, which represents the stall caused by the memory subsystem, and *core bound*, which reflects the stall incurred by the unavailable execution units. When the slot is filled with a uOp, it will be classified as *retiring* if the uOp eventually retires (Otherwise, the slot is categorized as *bad speculation*). We use Intel Vtune Profiler to measure the percentage of pipeline slots in each category.

3.1 Observations

Varying random walk workloads. We first evaluate the four RW algorithms in Section 2.2. Specifically, we set PPR as unbiased, and configure the termination probability as 0.2. For DeepWalk and Node2Vec, we set the target length as 80. The transition probability of DeepWalk is the edge weight, and that of Node2Vec is calculated based on Equation 1 where $a = 2$ and $b = 0.5$. The schema length of MetaPath is 5, and we generate it by randomly choosing five labels from the edge label set. PPR starts $|V|$ queries from a given vertex, and the others start a query from each vertex in V . Following existing studies [6, 16, 45, 57] (as well as popular open-source packages^{1 2}), we use NAIVE sampling for PPR, while ALIAS sampling for the others. Moreover, we build alias tables for DeepWalk in a preprocessing phase to accelerate the execution of queries. However, this method is prohibitively expensive for high order RW due to the exponential memory consumption [53, 66]. For example, the space complexity of such an index for Node2Vec, which is second-order, is $O(\sum_{v \in V} d_v^2)$, and it can consume more than 1000 TB space for *twitter*. As such, we compute the transition probability and perform the initialization of ALIAS in run time. To compare the performance characteristics with RW algorithms, we evaluate BFS and SSSP, which are two conventional graph algorithms. We develop RW algorithms without any frameworks, whereas implementing BFS and SSSP with Ligra [55].

Table 1 presents the experiment results on *livejournal*, the details of which are listed in Table 5. RW queries randomly visit nodes on the graph that leads to a massive number of random memory accesses. Consequently, as high as 73.1% pipeline slots of PPR and DeepWalk are stalled due to memory access. In contrast, the memory bound of BFS and SSSP is less than 45%, which demonstrates

¹<https://github.com/aditya-grover/node2vec>, Last accessed on 2021/03/20

²<https://github.com/GraphSAINT/GraphSAINT>, Last accessed on 2021/03/20.

much better cache locality than that of PPR and DeepWalk. Due to the large proportional of memory stalls, the retiring of PPR and DeepWalk is less than 10%. Furthermore, we measure the memory bandwidth utilization of these algorithms. Our benchmark shows that the max memory bandwidth of our test bed is 60 GB/s. As shown in the table, the bandwidth utilization of BFS and SSSP are rather high (86.2% and 63.6%, respectively), while that of PPR and DeepWalk is very low (2.3% and 9.3%, respectively).

Compared with PPR and DeepWalk, Node2Vec and MetaPath exhibit different characteristics. The memory bound is lower than PPR and DeepWalk, whereas the retirement and bandwidth are much higher. To achieve more insights, we first examine the execution time breakdown on computing the transition probability (denoted by *compute p(e)*), and the initialization and generation phases of sampling an edge (denoted by *Init* and *Gen*, respectively), and then analyze the complexity of these operations at a step.

Table 2: Comparison of execution time breakdown and the time complexity per step among RW algorithms where $v = Q.cur$ and u is the last vertex of Q .

Method	Time Breakdown			Complexity per Step		
	Compute $p(e)$	Sampling		Compute $p(e)$	Sampling	
		Init	Gen		Init	Gen
PPR	N/A	N/A	100%	N/A	N/A	$O(1)$
DeepWalk	N/A	N/A	100%	N/A	N/A	$O(1)$
Node2Vec	89.9%	9.9%	0.2%	$O(d_v \times \log d_u)$	$O(d_v)$	$O(1)$
MetaPath	29.0%	69.9%	1.1%	$O(d_v)$	$O(d_v)$	$O(1)$

Table 2 lists the results. PPR and DeepWalk are static, and they only need to sample an edge and move Q along it in run time. In contrast, Node2Vec and MetaPath are dynamic, and they first compute the transition probability for each $e \in E_v$ where $v = Q.cur$, and then sample an edge. Consequently, the cost on *Gen* is neglected as shown in Table 2. Moreover, the memory bound is much lower than static RWs in Table 1 since the computation scans E_v in a continuous manner. Given $e \in E_v$ and u is the last vertex of Q , the complexity of computing $p(e)$ in Node2Vec is $O(\log d_u)$ because the distance check in Equation 1 is implemented by a binary search. However, MetaPath computes $p(e)$ with a simple label filter. As a result, computing $p(e)$ accounts for around 90% of the execution time in Node2Vec, whereas *Init* dominates the cost in MetaPath.

Observation 1. *The in-memory computation of common RW algorithms suffers severe performance issues due to memory stalls caused by cache misses and under-utilizes the memory bandwidth. For high order RW algorithms, computing $p(e)$ and initializing the auxiliary data structure for sampling dominate the in-memory computation cost, and their complexities are determined by the RW algorithm and the selected sampling method, respectively.*

Varying sampling methods and RW types. We further examine the performance of sampling methods. We continue to develop a micro benchmark that executes 10^7 RW queries each of which starts from a vertex randomly selected from the graph. The target length is 80. We evaluate three types of RW queries as discussed in Section 2.2: **unbiased**, **static** and **dynamic**. For unbiased RW, we first perform the initialization phase of sampling methods for the neighbor set of each vertex in a *preprocessing* step. We then use the generation phase of a sampling method to select a neighbor of $Q.cur$ in execution. For static RW, we evaluate queries with the

same process as that of unbiased. The only difference is that the edge weight is used to set the transition probability for static RW whereas the transition probability in unbiased RW is the default uniform. For dynamic RW, we set the edge weight as the transition probability, while performing the initialization phase for the neighbor set of $Q.cur$ in execution because the transition probability of dynamic RW varies during the computation.




Figure 1: Effectiveness of sampling methods.

Figure 1 presents the experiment results of the sequential execution with variant sampling methods on different RW types. We have the following findings. First, the NAIVE sampling method performs the best on unbiased RW as it has no initialization phase. Second, among static methods, the ALIAS sampling method outperforms others because its generation phase has lower time complexity. However, ALIAS runs much slower than other methods on dynamic RW since its initialization cost is high in practice. Third, O-REJ performs well on dynamic RW since it does not have the initialization phase. Fourth, we can observe that the cost of evaluating dynamic RW is significantly expensive than that of unbiased and static RW because of the initialization phase (if exists) at each step.

Observation 2. *Sampling methods have an important impact on the performance and no sampling method can dominate on all cases. Generally, dynamic RW is expensive than unbiased and static RW.*

3.2 System Implications

Based on the profiling results, we can categorize the cost of evaluating RW queries into two classes, that of computing $p(e)$ and that of sampling an edge. As the former is determined by the RW algorithms (i.e., algorithm-specific), our framework targets at accelerating the latter operation. Moreover, we have the following implications for the design and implementation of ThunderRW. First, we need to develop mechanisms to reduce the cache stalls. Our profiling results show that in-memory computation of common RW algorithms suffer severe performance issues due to the irregular memory accesses. None of previous random walk frameworks [46, 60, 66] address the problem. On the other hand, there are massive queries in random walk workloads, but the memory bandwidth is under-utilized. Inspired by previous work on accelerating multiple index lookups in database systems with prefetching [4, 26], there are opportunities for prefetching and interleaving executions among different queries.

Second, there is a need to support multiple sampling methods. However, existing frameworks support one sampling method only and generally regard all RW as dynamic (e.g., C-SAW), while (1) the sampling method has an important impact on the performance and none of them can dominate on all cases; and (2) the cost of evaluating dynamic RW is generally much more expensive than that of unbiased and static RW.

4 THUNDERRW ABSTRACTION

In this section, we present the abstraction of the computation in ThunderRW.

4.1 Step-centric Model

To abstract the computation of RW algorithms, we propose the *step-centric* model in this paper. We observe that RW algorithms are built upon a number of RW queries rather than a single query. In spite of limited intra-query parallelism, there is abundant inter-query parallelism in RW-algorithms as each RW query can be executed independently. Therefore, our step-centric model abstracts the computation of RW algorithms from the perspective of queries to exploit the inter-query parallelism.

Specifically, we model the computation from the local view of moving one step of a query Q . Then, we abstract a step of Q into the Gather-Move-Update (GMU) operations to characterize the common structure of RW algorithms. With the step-centric model, users develop RW algorithms by "thinking like a walker". They focus on defining functions setting the transition probability of $e \in E_v$ and updating states of Q at each step, while the framework facilitates applying user-defined step-oriented functions to RW queries.

4.2 Step-centric Programming

Framework. Algorithm 2 gives an overview of ThunderRW. Lines 1-6 execute each query one-by-one. Lines 3-5 factor one step into three functions based on the step-centric model. Gather collects the transition probabilities of edges adjacent to $Q.cur$. It loops over $E_{Q.cur}$, applies *Weight*, a user-defined function, to each edge e and add the transition probability of e to C (Lines 10-11). Then, Line 12 executes the initialization phase of a given sampling method to update C . Move picks an edge based on C and moves Q along the selected edge (Lines 14-16). As random memory accesses in the *system space* (i.e., the framework excluding user-defined functions) are mainly in Move, we apply step-interleaving techniques to optimize its performance (see Section 5). Finally, Line 5 invokes *Update*, a user-defined function, to update states of Q based on the movement. The return value of *Update* decides whether Q should be terminated.

Algorithm 2: ThunderRW Framework

Input: a graph G and a set Q of random walk queries;
Output: the walk sequences of each query in Q ;

```

1 foreach  $Q \in \mathcal{Q}$  do
2   do
3      $C \leftarrow \text{Gather}(G, Q, \text{Weight})$ ;
4      $e \leftarrow \text{Move}(G, Q, C)$ ;
5      $stop \leftarrow \text{Update}(Q, e)$ ;
6   while  $stop$  is false;
7 return  $Q$ ;
8 Function  $\text{Gather}(G, Q, \text{Weight})$ 
9    $C \leftarrow \{\}$ ;
10  foreach  $e \in E_{Q.cur}$  do
11    Add  $\text{Weight}(Q, e)$  to  $C$ ;
12   $C \leftarrow$  execute initialization phase of a given sampling method on  $C$ ;
13  return  $C$ ;
14 Function  $\text{Move}(G, Q, C)$ 
15   Select an edge  $e(Q.cur, v) \in E_{Q.cur}$  based on  $C$  and add  $v$  to  $Q$ ;
16   return  $e(Q.cur, v)$ ;
```

Algorithm 3: Preprocessing for Static Random Walk

Input: a graph G ;
Output: the transition probabilities C_v on E_v for each vertex v ;

```

1 foreach  $v \in V(G)$  do
2    $C_v \leftarrow \{\}$ ;
3   foreach  $e \in E_v$  do
4     Add  $\text{Weight}(null, e)$  to  $C_v$ ;
5    $C_v \leftarrow$  execute initialization phase of a given sampling method on  $C_v$ ;
6   Store  $C_v$  for the usage in query execution.
```

The framework described in Algorithm 2 can support unbiased, static and dynamic RW with different sampling methods. Furthermore, we optimize the execution flow of ThunderRW based on the RW type and the selected sampling method. The transition probability of static RW is fixed during the execution. In that case, ThunderRW omits the Gather operation but introducing a preprocessing step to reduce the runtime cost, which obtains transition probabilities in the system initialization. Algorithm 3 presents the preprocessing for static RW. Given a vertex v , Lines 3-4 loop over each edge e in E_v and apply the *Weight* function to e to obtain the transition probability. As the probability does not rely on a query, we set Q as *null*. After that, Lines 5-6 perform the initialization phase of a given sampling method on C_v and store C_v for the usage in the query execution. As such, we can load $C_{Q.cur}$ directly without Gather for static RW in Algorithm 2.

Moreover, the NAIVE and 0-REJ sampling methods have no initialization phase as discussed in Section 2.3. Hence, we do not need to collect the transition probability for initialization. As such, ThunderRW skips both the preprocessing step and the Gather operation in the execution if NAIVE or 0-REJ is used.

Application Programming Interfaces (APIs). ThunderRW provides two kinds of APIs, which include hyperparameters and user-defined functions. Users develop their RW algorithms in two steps. Firstly, set the RW type and the sampling method via hyperparameters *walker_type* and *sampling_method*, respectively. Secondly, define the *Weight* and *Update* functions. The *Weight* function specifies the relative chance of an edge being selected. The *Update* function modifies states of Q given the selected edge. If its return value is *true*, then the framework terminates Q . Otherwise, Q continues walking on G . When using 0-REJ, users need to implement the *MaxWeight* function to set the maximum value of the transition probability. We present an example in the following.

```

WalkerType walker_type = WalkerType::Dynamic;
SamplingMethod sampling_method = SamplingMethod::0-REJ;
double Weight(Walker Q, Edge e) {
    if (Q.length == 0) return max(1.0 / a, 1.0, 1.0 / b);
    else if (e.dst == Q.prev) return 1.0 / a;
    else if (IsNeighbor(e.dst, Q.prev)) return 1.0;
    else return 1.0 / b;
}
bool Update(Walker Q, Edge e) {
    return Q.length == target_length;
}
double MaxWeight() {
    return max(1.0 / a, 1.0, 1.0 / b);
}
```

Listing 1: Node2Vec sample code.

Example 4.1. List 1 shows the sample code of Node2Vec, which is dynamic. As the maximum value can be easily determined by

the parameters a and b , we use O-REJ to avoid scanning each edge adjacent to $Q.cur$ at each step. Thus, we set `sampling_method` to O-REJ and implement `MaxWeight`. The `Weight` function is configured based on Equation 1. Once the length of Q meets the target length, we terminate it.

ThunderRW applies user-defined functions to RW queries, and evaluates the queries based on RW type and selected sampling method in parallel. Thus, users can easily implement customized RW algorithms with ThunderRW, which significantly reduces the engineering effort. For example, users write only around ten lines of code to implement `Node2Vec` as shown in Example 4.1.

Parallelization. RW algorithms contain massive random walk queries each of which can be completed independently and rapidly. Therefore, ThunderRW adopts the static scheduling method to keep load balancing among workers. Specifically, we regard each thread as a worker and evenly assign Q to the workers. A worker independently executes the assigned queries with Algorithm 2. Our experiment results show that the simple scheduling method achieves good performance.

4.3 Analysis

In this subsection, we analyze the space and time cost of Algorithm 2 on different RW types with variant sampling methods. As the cost of `Weight` and `Update` is determined by users' implementation, we assume their cost is a constant value for the ease of analysis.

Space. The space for storing the graph is $O(|E| + |V|)$, and that for maintaining the output is $O(\sum_{Q \in \mathcal{Q}} |Q|)$. `Gather` in Algorithm 2 requires $O(d_{max})$ space to store C where d_{max} is the max degree value of G . Suppose that ThunderRW has n threads. Then, the memory cost is $O(n \times d_{max})$. When there is a preprocessing step, the memory cost of ITS and ALIAS is $O(|E|)$, while that of REJ is $O(|V|)$ based on the analysis in Section 2.3.

Time. Given a sampling method, α and β denote the cost of its initialization phase and generation phase, respectively. Let d_{avg} represent the average degree of G . Thus, the cost of `Gather` in Algorithm 2 is $d_{avg} + \alpha$, and that of `Move` is β . For static RW, the preprocessing cost is $\sum_{v \in V} (d_v + \alpha)$, while the cost of processing one step is β as it does not conduct `Gather` during execution. From Section 2.3 we can get the value of α and β for the sampling methods. Support that $T = \sum_{Q \in \mathcal{Q}} |Q|$, which is the total number of steps of all queries. Table 3 summarizes the time complexity on different RW types with variant sampling methods.

As shown in the table, NAIVE supports unbiased RW only. For ITS, ALIAS and REJ, the cost on unbiased and static RW consists of the preprocessing cost and the execution cost. Because RW algorithms can have massive RW queries with a long length, the execution cost is generally much more expensive than the preprocessing cost. As O-REJ has no initialization phase, it neither performs the preprocessing for unbiased and static RW nor executes `Gather` for dynamic RW. Thus, the time complexity is the same for the three RW types.

Recommendation. From the analysis, we have the following guidelines for setting sampling methods for users: (1) NAIVE is the best sampling method for unbiased RW; (2) ALIAS is a good choice for static RW since the execution time is generally longer than the preprocessing time; and (3) if we can set a reasonable max value for

Table 3: The time complexity of ThunderRW on different random walk types with variant sampling methods

Method	Unbiased	Static	Dynamic
NAIVE	$O(T)$	N/A	N/A
ITS	$O(E + T \times \log d_{avg})$	Same as unbiased	$O(T \times (d_{avg} + \log d_{avg}))$
ALIAS	$O(E + T)$	Same as unbiased	$O(T \times (d_{avg} + 1))$
REJ	$O(E + T \times \mathbb{E})$	Same as unbiased	$O(T \times (d_{avg} + \mathbb{E}))$
O-REJ	$O(T \times \mathbb{E})$	Same as unbiased	Same as unbiased

the transition probability, then use O-REJ for dynamic RW. Users can easily tell the RW type based on the properties of transition probability. To further ease the programming efforts, we set the default sampling method of unbiased, static and dynamic RW to NAIVE, ALIAS and ITS, respectively. We use ITS instead of ALIAS for dynamic RW because the initialization cost of ALIAS at each step is much more than that of ITS in practice. If users can set a good max value for the transition probability, then they can select O-REJ for dynamic RW.

5 STEP-INTERLEAVING

In this section, we present the step interleaving technique, which reduces the pipeline stall caused by random memory accesses.

5.1 General Idea

Based on the step-centric model, ThunderRW processes a step of a query Q with the GMU operations. According to the profiling results in Section 3, there can be two main sources for random memory accesses under the model. First, the `Move` operation picks an edge randomly and moves Q along the selected edge. Second, the operations in user-defined functions can introduce cache misses, for example, the distance check operation in `Node2Vec`. As operations in the user space (i.e., user-defined functions) are determined by RW algorithms, and can be very flexible, we target at memory issues incurred by the system (i.e., the `Move` operation). Motivated by the profiling result, we propose to use the software prefetching technique [29] to accelerate in-memory computation of ThunderRW. However, a step of a query Q does not have enough computation workload to hide memory access latency because steps of Q have dependency relationship. Therefore, we propose to hide memory access latency via executing steps of different queries alternately.

Specifically, given a sequence of operations in `Move`, we decompose them into multiple stages such that the computation of a stage consumes the data generated by previous stages and it retrieves the data for the subsequent stages if necessary. We execute a group of queries simultaneously. Once a stage of a query Q completes, we switch to stages of other queries in the group. We resume the execution of Q when stages of other queries complete. In such a way we hide the memory access latency in a single query and keep CPUs busy. We call this approach *step interleaving*.

Example 5.1. Figure 2 presents an example where a step is divided into four stages. If executing a query step-by-step sequentially, then CPUs are frequently stalled due to memory access. Even with prefetching, the computation of a stage cannot hide the memory access latency. In contrast, the step interleaving hides the memory access latency by executing steps of different queries alternately.

Let's perform a simple back-of-envelop calculation on the performance gain of interleaving execution. Given a group containing k queries, we assume that `Move` of each query executes the

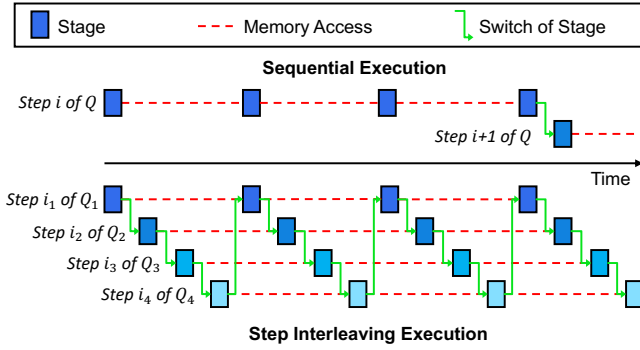


Figure 2: Sequential versus step interleaving.

same number of stages and the cost W_C of each stage is the same for the ease of analysis. Suppose that there are m stages with memory access and \bar{m} without. W_L denotes the latency of memory access. Then, the cost of moving a step for the queries in sequential is equal to $W_0 = k((m + \bar{m})W_C + mW_L)$. Let W_S denote the cost of switching. The cost of Move with step interleaving is $W_1 = k((m + \bar{m})(W_C + W_S) + m(\max(W_L - kW_S - (k - 1)W_C, 0)))$ where the last term calculates whether step interleaving hides memory access latency. Therefore, the gain of step interleaving for a step of k queries can be estimated by Equation 2 where $W_{hide} = \max(W_L - kW_S - (k - 1)W_C, 0)$.

$$W_{gain} = (W_0 - W_1)/k = mW_L - (m + \bar{m})W_S - mW_{hide}. \quad (2)$$

From Equation 2, we can see that step interleaving requires an efficient switch mechanism to reduce the overhead W_S of performing switching, and enough workload to overlap the memory access latency W_{hide} .

5.2 Stage Dependency Graph

We design the *stage dependency graph* (SDG) to model stages of a sequence of operations in a step. Each node in SDG is a stage containing a set of operations and edges represent the dependency relationship among them. Given the sequence of operations, we build SDG in two steps, abstracting stages (nodes) and extracting dependency relationships (edges).

Defining stages: As we hide memory access latency by switching the execution of queries, the constraint on stages is that each stage contains at most one memory access operation and the operations consuming the data are in subsequent stages. Note that we view the operation containing jump operation as a single stage for the ease of the implementation of switching. We present an example in the following.

Example 5.2. The right column of Table 4 illustrates the sequence of operations in the Move function with the ALIAS and REJ sampling methods, respectively, to perform the neighbor selection. The left column lists stages. For example, S_0 of ALIAS loads d_v consumed in O_1 of S_1 . O_5 in REJ has the jump operation. Therefore, we regard it as a separate stage.

Defining edges: Next, we add edges among nodes in SDG based on their dependency relationships. Given stages S and S' , if there is a dependency relationship between S and S' , we add an edge

Table 4: Stages of Move with ALIAS and REJ ($v = Q.cur$).

Stage	ALIAS
S_0	O_0 : Load d_v .
S_1	O_1 : Generate an int random num x in $[0, d_v)$.
	O_2 : Generate a real random num y in $[0, 1)$.
	O_3 : Load $C[x] = (H[x], A[x])$.
S_2	O_4 : If $y < H[x]$, $v' = A[x].first$; Else $v' = A[x].second$.
	O_5 : Add v' to Q and return $e(v, v')$.
Stage	REJ
S_0	O_0 : Load d_v .
S_1	O_1 : Load the maximum value p_v^* .
S_2	O_2 : Generate an int random num x in $[0, d_v)$.
	O_3 : Generate a real random num y in $[0, p_v^*)$.
	O_4 : Load $C[x] = p$.
S_3	O_5 : If $y > C[x]$, jump to O_2 ; Else jump to O_6 .
S_4	O_6 : Load $e(v, v') = E_v[x]$.
S_5	O_7 : Add v' to Q and return $e(v, v')$.




Figure 3: Stage dependency graph.

from S to S' . The edges are categorized into three types, *memory dependency*, *computation dependency* and *control dependency*. We call the first two relationship as *data dependency*. More specifically, if S' consumes the data loaded from memory by S , then the edge type is memory dependency. Otherwise, S' depends on the data computed by S and the edge type is computation dependency. The data leading to the dependency is attached to each edge as properties. Furthermore, if S contains the operation jumping to S' , we add the control dependency from S to S' . SDG allows that there are multiple edges (i.e., dependency relationships) between nodes. If we only consider data dependency, SDG is a directed acyclic graph (DAG), while the control dependency can generate cycles in SDG.

Example 5.3. Continuing with Example 5.2, Figure 3 shows SDGs. In SDG of ALIAS, S_2 relies on x, y , which are random numbers generated by S_1 , while $(H[x], A[x])$ is the data retrieved from memory. As such, S_1 and S_2 have both memory and computation dependency relationships. SDG of ALIAS is a DAG because there is no control dependency. In contrast, there is a cycle containing S_2 and S_3 in SDG of REJ because of the control dependency.

In summary, SDG is a methodology to abstract stages from a sequence of operations in Move and model the dependency relationship among them. Note that the stage design of MOVE does not require user input but it is implemented in the system space.

5.3 State Switch Mechanism

In this subsection, we introduce the implementation of step interleaving under SDG. Based on Equation 2, we need an efficient switch mechanism. For example, using multi-threading is forbidden because the overhead of context switch among threads is in microseconds, whereas the main memory latency is in nanoseconds. As each thread tends to take many RW queries, we switch the execution among stages in a single thread.

We categorize stages of a SDG into two classes based on whether they belong to cycles in SDG. For stages not in cycles (called *non-cycle stages*), a query visits them exactly once to complete Move. Given a group of queries Q' , we execute them in a coupled manner. Particularly, once a query $Q_i \in Q'$ completes a stage S , we switch to the next query $Q_{i+1} \in Q'$ to process S . After all queries complete S , we move to the next stage. In contrast, stages in cycles (called *cycle stages*) can be visited variant times for different queries. To deal with the irregularity, we process them in a decoupled manner. Specifically, each query Q records the stage S to be executed. When switching to Q , we execute S , set the next stage of Q based on SDG, and switch to the next query after completing S . As a result, each query executes asynchronous.

For data communication between different stages in a query, we create two kinds of ring buffers based on SDG, in which the computation dependency edge indicates the information requiring to be stored. In particular, the *task ring* is used for data communication across all stages of a query, while the search ring serves to process cycle stages. As we need to explicitly record states of cycle stages and control the switch of them, processing cycle stages not only causes implementation complexities, but also incurs more overhead. Note that the SDGs of NAIVE and ALIAS have no cycle stages because there are no for loops in their generation phases, whereas that of ITS, REJ and O-REJ have. The implementation details are introduced in the complete version paper [56].

5.4 Ring Size Tuning

The task ring size k and the search ring size k' determine the group size of queries executed simultaneously in a thread, and therefore control memory level parallelism of executing non-cycle stages and cycle stages, respectively. According to Equation 2, we can improve the performance by increasing k to reduce W_{hide} . However, k is limited by hardware. Particularly, modern CPUs can issue a limited number of outstanding memory requests, and the L1 data cache size is only tens of kilobytes. Setting k to a large value can evict data before the usage. In ThunderRW, we tune ring sizes by pre-executing a number of queries. We start a RW query from each vertex with the target length as 10 and set the RW type as static. We first select the NAIVE and ALIAS sampling methods, respectively and vary k from 1, 2, ..., 512, 1024 to pick an optimal value k^* . Next we fix k to k^* and vary k' from 1, 2, ..., k^* to select optimal values for ITS, REJ and O-REJ, respectively.

6 EXPERIMENTS

We conduct experiments to evaluate the performance of ThunderRW in this section.

6.1 Experimental Setup

We conduct experiments on a Linux server equipped with an Intel Xeon W-2155 CPU and 220GB RAM. The CPU has ten physical cores with hyper-threading disabled for consistent measurement. The sizes of L1, L2 and L3 (last level cache, LLC) caches are 32KB, 1MB and 13.75MB, respectively.

Datasets. Table 5 lists the statistics of the twelve real-world graphs in our experiments. *ab* and *ac* are downloaded from [43], *wk* is obtained from [42], *eu*, *uk* and *tw* are obtained from [51],

Table 5: Properties of real-world datasets.

Dataset	Name	$ V $	$ E $	d_{avg}	d_{max}	Memory
amazon	<i>am</i>	0.55M	1.85M	3.38	549	0.01GB
youtube	<i>yt</i>	1.14M	2.99M	5.24	28754	0.03GB
us patents	<i>up</i>	3.78M	16.52M	8.74	793	0.17GB
eu-2005	<i>eu</i>	0.86M	19.24M	44.74	68963	0.15GB
amazon-clothing	<i>ac</i>	15.16M	63.33M	4.18	12845	0.35GB
amazon-book	<i>ab</i>	18.29M	102.12M	5.58	58147	0.52GB
livejournal	<i>lj</i>	4.85M	68.99M	28.45	20333	0.54GB
com-orkut	<i>ot</i>	3.07M	117.19M	76.34	33313	0.89GB
wikidata	<i>wk</i>	40.96M	265.20M	6.47	8085513	1.29GB
uk-2002	<i>uk</i>	18.52M	298.11M	32.19	194955	2.30GB
twitter	<i>tw</i>	41.66M	1.21B	58.08	2997487	9.27GB
friendster	<i>fs</i>	65.61M	1.81B	55.17	5214	13.71GB

and the other graphs are downloaded from [30]. The datasets are from different categories such as web, social and citation, and have different densities. The number of vertices is ranged from hundreds of thousands to tens of millions, and the number of edges scales from millions to billions. Except *am*, all the graphs outsize LLC.

Workloads. We study PPR, DeepWalk, Node2Vec and MetaPath to evaluate the performance and generality of competing methods. The settings of the four algorithms are the same as that in Section 3. *ab* and *ac* are weighted graphs where weights denote review ratings for products. *wk* has 1327 distinct labels, which represents the relationship between entities in a knowledge base. The other graphs are unweighted and unlabeled. Given a graph having no labels or weights, we set the weight and label of edges with the same setting as previous work [66]: (1) We choose a real number from [1, 5) uniformly at random, and assign it to an edge as its weight; and (2) We set the edge label by randomly choosing a label from a set containing five distinct labels.

Comparison. In our experiments, we compare the performance of ThunderRW (called *TRW* for short) with the following methods.

- *BL*: Baseline approaches that first load a graph entirely into memory and then execute random walks, the detail of which is presented in Section 3.
- *HG*: Our homegrown implementation optimizing *BL* from two aspects: (1) select a suitable sampling method for each algorithm according to the recommendation in Section 4.3; and (2) regard each query as a parallel task with OpenMP.
- *GW*: GraphWalker [60], the state-of-the-art RW framework in a single machine. For the fair of comparison, we configure GraphWalker to execute in-memory, without any disk I/O.
- *KK*: KnightKing [66], the state-of-the-art distributed RW framework. It supports to execute in a single machine.

We implement all our methods including *BL*, *HG* and *TRW* in C++. *GW*³ and *KK*⁴ are programmed in C++ as well. All the source code is compiled by g++ 8.3.1 with -O3 enabled. *BL* executes in serial, while the other methods are running on all the cores of the single socket, with one thread per core.

We consider C-SAW [46], the state-of-the-art RW framework on GPUs, as well. However, its open source package⁵ supports 4000 queries at most, which cannot handle the workload containing massive queries in the experiment. Previous experiment results [60, 66] show that *KK* and *GW* significantly outperform generic graph computing frameworks such as Gemini [71] on RW algorithms.

³<https://github.com/ustcadsl/GraphWalker>, Last accessed on 2020/12/07.

⁴<https://github.com/KnightKingWalk/KnightKing>, Last accessed on 2020/12/20.

⁵<https://github.com/concept-inversion/C-SAW>, Last accessed on 2020/12/07.

Table 6: Overall performance comparison (seconds).

Dataset	PPR				DeepWalk				Node2vec				MetaPath			
	BL	HG	GW	KK	TRW	BL	HG	KK	TRW	BL	HG	KK	TRW	BL	HG	TRW
<i>am</i>	0.06	0.008	0.42	0.012	0.007	2.16	0.21	0.44	0.07	9.97	0.26	2.08	0.14	0.22	0.018	0.012
<i>yt</i>	0.33	0.04	1.68	0.05	0.015	9.78	0.98	1.93	0.26	853.13	1.30	5.94	1.03	6.18	0.23	0.24
<i>up</i>	1.24	0.13	7.19	0.19	0.07	45.44	4.33	8.41	0.95	369.00	6.20	16.92	4.01	4.88	0.40	0.24
<i>eu</i>	0.16	0.02	0.99	0.03	0.011	8.16	0.82	1.56	0.20	2731.07	1.47	4.43	1.14	90.55	3.18	3.55
<i>ac</i>	4.84	0.51	19.31	0.65	0.19	173.66	17.86	31.88	3.31	6951.12	24.54	87.86	6.26	45.01	2.01	1.69
<i>ab</i>	8.86	0.94	26.74	1.09	0.26	212.80	22.24	40.07	4.01	26231.45	32.04	100.78	7.87	128.35	5.06	4.47
<i>lj</i>	1.69	0.19	7.90	0.23	0.06	55.63	5.44	10.67	1.19	2951.33	9.09	24.95	6.20	18.08	0.94	0.75
<i>ot</i>	1.49	0.16	5.25	0.19	0.04	38.54	3.70	7.97	0.80	5891.28	7.28	15.16	4.82	40.77	1.72	1.57
<i>wk</i>	21.86	2.21	47.05	3.07	0.59	502.27	49.67	95.17	9.26	OOT	68.43	216.24	27.68	5.98	0.54	0.55
<i>uk</i>	6.47	0.69	27.72	0.90	0.24	203.86	20.42	21.40	4.56	12630.01	34.36	94.69	28.68	322.66	12.84	12.56
<i>tw</i>	26.42	2.73	77.12	3.61	1.16	575.43	61.18	115.92	11.13	OOT	130.72	232.41	91.00	OOT	12300.32	9780.20
<i>fs</i>	79.14	8.20	223.81	10.72	4.10	1043.93	108.23	208.45	17.67	OOT	178.15	364.51	120.16	683.05	28.69	25.01

Therefore, our experiment does not involve C-SAW as well as any generic graph computing frameworks.

As for RW algorithms, *GW* only supports unbiased RW. Thus, we execute PPR without considering edge weights, and evaluate *GW* on PPR only. Despite that *KK* studies MetaPath in the original paper [66], its open source package cannot handle labeled graphs. As such, it cannot execute MetaPath. In contrast, *TRW* supports all the four algorithms, which demonstrates its flexibility.

As for sampling methods, *BL* uses NAIVE for PPR, while adopts ALIAS for the other three algorithms. As discussed in Section 3, building alias tables for dynamic RW in an indexing phase can consume a huge amount of memory. Therefore, in the experiments, *BL* dynamically computes the alias table (i.e., perform the initialization of ALIAS) at each step of a query, which is the same as the computation flow of *TRW* for dynamic RW. Different from *BL*, *HG* adopts 0-REJ for Node2Vec, and ITS for MetaPath. This is because (1) the max value of transition probability of Node2Vec can be easily set as $\max(1, 1/a, 1/b)$, and 0-REJ can avoid scanning the neighbors of $Q.cur$ at each step; and (2) the probability distribution of MetaPath is skewed due to filtering based on labels, which increases the generation cost of rejection sampling, and the initialization phase of ITS is much faster than that of ALIAS in practice. *TRW* adopts the same sampling method as *HG* for each algorithm.

Ring Size Setting. We tune the ring size with the method in Section 5.4. Despite that the graphs have variant structures, the optimal setting for them is close. First, the optimal value for the graphs except *am* is $k = 64$ and $k' = 32$ because the optimal ring size is closely related to the instructions available for computation, the switch overhead, the memory access latency, and the maximum number of outstanding memory requests, which are determined by the program and hardwares. Second, the optimal value for *am* is $k = 32$ and $k' = 32$ as *am* fits in LLC and the memory access latency is smaller than that of the other graphs. Additionally, the tuning process is very efficient, which takes less than one minute for most of the graphs. Even for *fs* with more than 1.8 billion edges, the tuning is completed with around four minutes.

Metrics. The *total time* is the elapsed time on evaluating RW algorithms without counting the time on loading data from the disk. For static random walk, the total time consists of the *preprocessing time*, which is the time spent on the preprocessing, and the *execution time*, which is the time spent on executing queries. To complete experiments in a reasonable time, we set the time limit for each algorithm as eight hours. If an algorithm cannot be completed

within the limit, we terminate it and record the execution time as OOT (i.e., out-of-time). We measure the *throughput* (steps per second) by dividing the number of steps of all queries by the execution time. To provide more insights, we adopt *Intel Vtune Profiler* to examine the pipeline slot utilization and use *Linux Perf* to examine the *instructions per step* and *cycles per step*, which are the number of instructions and the number of cycles on one step, respectively.

Supplement experiments. More experiment results including the impact of ring sizes, memory bandwidth utilization, the effectiveness of prefetching data to different cache levels, the impact of the step interleaving on existing systems and the comparison with AMAC [26] are presented in the full version paper [56].

6.2 Overall Comparison

Table 6 gives an overall comparison of competing methods on the four RW algorithms. Although *GW* is parallel, it runs slower than *BL*, the sequential baseline algorithm. *KK* runs faster than *GW* and *BL*, but slower than *HG* because (1) the framework incurs extra overhead compared with *HG*; and (2) *HG* adopts an appropriate sampling method for each algorithm. *TRW* runs 54.6-131.7X and 1.7-14.6X faster than *GW* and *KK*, respectively.

Benefiting from parallelization, *HG* achieves 7.5-10.5X speedup over *BL* on PPR and DeepWalk. Moreover, *HG* runs 38.3-1857.9X and 11.1-28.5X faster than *BL* on Node2Vec and MetaPath, respectively, because *HG* adopts 0-REJ sampling for Node2Vec, which avoids scanning the neighbors of $Q.cur$ at each step, and uses ITS sampling for MetaPath, the initialization phase of which is more efficient than that of ALIAS in practice. *TRW* runs 8.6-3333.1X faster than *BL*. Even compared with *HG*, *TRW* achieves up to 6.1X speedup benefiting from our step-centric model and step interleaving technique. As MetaPath is dynamic and both *TRW* and *HG* use ITS sampling, the gather operation at each step dominates the cost. Still, MetaPath on ThunderRW outperforms that on *HG* for nine out of twelve graphs, and is slightly slower on the other three graphs. *tw* is dense but highly skewed (as shown in Table 5) and the vertices with high degrees are frequently visited. Consequently, the execution time on MetaPath against *tw* is much longer than that on other graphs.

In summary, ThunderRW significantly outperforms state-of-the-art frameworks and homegrown solutions (e.g., *BL* takes more than eight hours for Node2Vec on *tw*, while *TRW* completes the algorithm in two minutes). Furthermore, ThunderRW saves a lot of engineering effort on the implementation and parallelization of RW algorithms compared with *BL* and *HG*.

6.3 Evaluation of Step Interleaving

We evaluate the effectiveness of step interleaving in this subsection. For brevity, we use *lj* as the representative graph by default.

Varying RW algorithms. We first evaluate the effectiveness of step interleaving on different RW algorithms. Figure 4 presents the pipeline slot breakdown and speedup among the RW algorithms. *wo/si* and *w/si* denote ThunderRW without and with the step interleaving technique, respectively. Enabling step interleaving drastically reduces memory bound on PPR and DeepWalk, and improves the instruction retirement. Correspondingly, *w/si* achieves significant speedup over *wo/si* in Figure 4b. The speedup on PPR is lower than that on DeepWalk because PPR issues all queries from a given vertex and the expected length of a query is 5, which by default exhibits better memory locality than DeepWalk. The memory bound on Node2Vec is reduced from around 60% to 40% because the Weight function checks whether two vertices are neighbors with a binary search, which causes a number of random memory access. The speedup on MetaPath is small because MetaPath is dynamic and the gather operation dominates the cost at each step.




Figure 4: Vary RW-algorithms on *lj*.

Varying sampling methods. We next examine the performance of step interleaving on variant sampling methods. As the gather operation dominates the cost on dynamic random walk, we focus on unbiased and static random walk. Particularly, we use DeepWalk as the representative RW algorithm and evaluate it with the five sampling methods in Section 2.3, respectively. When adopting NAIVE, we regard DeepWalk as unbiased random walk (i.e., without considering edge weight). Figure 5 presents the pipeline slot breakdown and speedup on *lj* with variant sampling methods. We can see that the step interleaving technique significantly reduces memory bound on all the five sampling methods and achieves remarkable speedup. The results demonstrate both the generality and effectiveness of the step interleaving technique.




Figure 5: Vary sampling methods on *lj*.

Varying datasets. To explore the impact of graph structures on the performance, we evaluate the speedup of enabling step interleaving for DeepWalk on different datasets. Figure 6 presents the experiment results. The speedup on *am* and *yt* is smaller than

that on other graphs because *am* can fit in LLC, and *yt* is only two times larger than LLC. The speedup on *eu* and *uk* is lower than the other graphs that are much larger than LLC since *eu* and *uk* have dense communities (e.g., *uk* has a clique containing around 1000 vertices [3]), and RW queries exhibit good memory locality. In contrast, the speedup on *ac* and *ab* is generally higher than the other graphs because they are bipartite graphs and very sparse, and RW queries have poor memory locality. In summary, the optimization tends to achieve higher speedup on large and sparse graphs than small graphs and graphs with dense community structures because RW queries have poorer memory locality on the former one. Nevertheless, the optimization brings up to 3X speedup even on graphs entirely fitting in LLC (i.e., *am*) since L1 cache is only tens of kilobytes, but around ten times faster than LLC, and the step interleaving directly fetches the data to L1 cache.




Figure 6: Vary datasets for DeepWalk.

6.4 Scalability Evaluation

In this section, we evaluate the scalability of ThunderRW. By default, we execute 10^7 RW queries on *lj* with the target length as 80. Each query starts from a vertex selected from the graph randomly. We first evaluate the throughput in terms of steps per second with the number of queries and the length of queries varying, respectively. In that case, we set the RW as static and use the ALIAS sampling method as the representative. Next, we evaluate the speedup with the number of threads varying. When setting the RW as unbiased, we use the NAIVE sampling method, while we examine the speedup on ITS, ALIAS, REJ and O-REJ, respectively, when setting the RW as static and dynamic.




Figure 7: Throughput on *lj* with number and length of queries varying.

Varying number and length of queries. Figure 7a presents the throughput with the number of queries varying from 10^2 to 10^7 . For $10^2 - 10^4$ queries, the execution time is very short and the start up and shut down time can dominate it. For example, for 10^2 queries, each thread spends less than 0.1 ms on performing random walks, while the execution time is around 2 ms because of the cost on resource (e.g., memory and threads) initialization and release. As a result, the benefit of the optimization is limited, and

the throughput is lower than that with a large number of queries. The throughput is more than 3×10^8 and keeps stable with the number of queries varying from 10^6 to 10^7 . Figure 7b presents the throughput with the length of queries varying from 5 to 160. The throughput is steady. In summary, ThunderRW has good scalability in terms of the number and length of queries.




Figure 8: Speedup on *lj* with number of threads varying.

Varying number of threads. Figure 8 shows the speedup with the number of threads varying from 1 to 10 (i.e., the number of cores in the machine). For all the five sampling methods on unbiased/static RW, ThunderRW achieves nearly linear speedup with the number of threads as shown in Figure 8a. Particularly, when the number of threads is 10, the speedup is from 8.8X to 9.6X. Figure 8b presents the speedup on dynamic RW. The speedup is from 7.8X to 9.0X. Overall, ThunderRW achieves good scalability in terms of the number of threads.

6.5 Generality Evaluation

To evaluate the generality of ThunderRW, we repeat the first experiment in Section 6.4 on a machine equipped with an Intel Xeon Gold 6246R CPU, which has 16 physical cores. The sizes of L1, L2 and LLC caches are 32KB, 1MB and 35.75MB, respectively. The CPU is based on the *Cascade Lake* microarchitecture, while that used in other experiments is based on *Skylake*. As the CPU has 16 physical cores, we set the number of workers as 16. As shown in Figure 9, enabling the optimization significantly improves the throughput. Moreover, using the new CPU increases the throughput, for example, when the length of queries is 160, the throughput grows from 3×10^8 to 4.1×10^8 . The experiment results show that the techniques proposed in this paper are generic to different architectures.




Figure 9: Throughput on *lj* with number and length of queries varying on processors with different architectures.

6.6 Discussions

ThunderRW regards a step of a query as a parallel task unit, which parallelizes the computation from the perspective of queries instead of the graph data to exploit the inter-query parallelism. As RW algorithms consist of massive queries and the cost of moving a

step is extremely small (e.g., around 34 ns for DeepWalk on *lj*), there are a large number of small parallel tasks, which can be easily parallelized. As such, the parallelization of ThunderRW can achieve significant speedup over the sequential despite that graph structures are complex and flexible. Moreover, the sampling method has an important impact on the performance, and therefore providing variant sampling methods is essential.

The step interleaving technique executes different queries alternately to reduce memory bound incurred by random memory accesses. Its effectiveness is closely related to the memory locality of workloads, which is determined by RW algorithms and graph structures. In general, the optimization tends to achieve higher speedup on large and sparse graphs than small graphs and graphs with dense community structures because RW queries have poorer memory locality on the former graphs. Nevertheless, the random memory access is a common issue for RW algorithms since (1) graphs are much larger than cache sizes; and (2) RW queries wander randomly in the graph. Thus, the step interleaving can achieve significant speedup even on graphs entirely fitting LLC.

However, the speedup achieved by the step interleaving on high order RW algorithms can be lower than that on first order algorithms. First, the operations in user-defined functions can introduce random memory accesses, while the step interleaving technique targets at random memory accesses in the Move operation. Despite that, the optimization still brings 1.2-4.3X speedup on Node2Vec. Second, the Gather operation dominates the cost at each step when performing it in run time because it loops over the neighbors of the current residing vertex of a query.

7 CONCLUSION

In this paper, we propose ThunderRW, an efficient in-memory RW engine on which users can easily implement customized RW algorithms. We design a step-centric model to abstract the computation from the local view of moving one step of a query. Based on the model, we propose the step interleaving technique to hide memory access latency by executing multiple queries alternately. We implement four representative RW algorithms including PPR, DeepWalk, Node2Vec and MetaPath with our framework. Experimental results show that ThunderRW outperforms state-of-the-art RW frameworks by up to one order of magnitude and the step interleaving reduces the memory bound from 73.1% to 15.0%. Currently, we implement the step interleaving technique in ThunderRW by explicitly and manually storing and restoring states of each query. An interesting future work is to implement the method with *coroutines*, an efficient technique supporting interleaved execution [18, 22, 49].

ACKNOWLEDGMENTS

This project is supported by the grant “Asian Institute of Digital Finance” awarded by National Research Foundation, Singapore and administered by the Infocomm Media Development Authority under its Smart Systems Strategic Research Programme in 2020. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. This research is also supported by the Ministry of Education, Singapore, under its Academic Research Fund Tier 2 (Award No.: MOE2019-T2-2-065).

REFERENCES

- [1] C. Balkesen, J. Teubner, G. Alonso, and M. T. Özsu. 2013. Main-memory hash joins on multi-core CPUs: Tuning to the underlying hardware. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, 362–373.
- [2] Scott Beamer, Krste Asanovic, and David Patterson. 2015. Locality exists in graph processing: Workload characterization on an ivy bridge server. In *2015 IEEE International Symposium on Workload Characterization*. IEEE, 56–65.
- [3] Lijun Chang. 2019. Efficient maximum clique computation over large sparse graphs. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 529–538.
- [4] Shimin Chen, Anastasia Ailamaki, Phillip B Gibbons, and Todd C Mowry. 2007. Improving hash join performance through prefetching. *ACM Transactions on Database Systems (TODS)* 32, 3 (2007), 17–es.
- [5] Shimin Chen, Phillip B. Gibbons, and Todd C. Mowry. 2001. Improving Index Performance through Prefetching. *SIGMOD Rec.* 30, 2 (2001), 235–246.
- [6] Michael Cochez, Petar Ristoski, Simone Paolo Ponzetto, and Heiko Paulheim. 2017. Biased graph walks for RDF graph embeddings. In *Proceedings of the 7th International Conference on Web Intelligence, Mining and Semantics*, 1–12.
- [7] Intel Corporation. 2016. Intel 64 and IA-32 architectures optimization reference manual.
- [8] Quanyu Dai, Qiang Li, Jian Tang, and Dan Wang. 2018. Adversarial network embedding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [9] Laxman Dhulipala. [n.d.]. Provably Efficient and Scalable Shared-Memory Graph Processing. ([n. d.]).
- [10] Wenfei Fan, Wenyuan Yu, Jingbo Xu, Jingren Zhou, Xiaojian Luo, Qiang Yin, Ping Lu, Yang Cao, and Ruiqi Xu. 2018. Parallelizing sequential graph computations. *ACM Transactions on Database Systems (TODS)* 43, 4 (2018), 1–39.
- [11] Dániel Fogaras, Balázs Rác, Károly Csalogány, and Tamás Sarlós. 2005. Towards scaling fully personalized pagerank: Algorithms, lower bounds, and experiments. *Internet Mathematics* 2, 3 (2005), 333–358.
- [12] Santo Fortunato and Darko Hric. 2016. Community detection in networks: A user guide. *Physics reports* 659 (2016), 1–44.
- [13] Tao-yang Fu, Wang-Chien Lee, and Zhen Lei. 2017. Hin2vec: Explore meta-paths in heterogeneous information networks for representation learning. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, 1797–1806.
- [14] Joseph E Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. 2012. Powergraph: Distributed graph-parallel computation on natural graphs. In *Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, 17–30.
- [15] Joseph E Gonzalez, Reynold S Xin, Ankur Dave, Daniel Crankshaw, Michael J Franklin, and Ion Stoica. 2014. Graphx: Graph processing in a distributed dataflow framework. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, 599–613.
- [16] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*, 855–864.
- [17] Wentian Guo, Yuchen Li, Mo Sha, and Kian-Lee Tan. 2017. Parallel personalized pagerank on dynamic graphs. *Proceedings of the VLDB Endowment* 11, 1 (2017), 93–106.
- [18] Yongjun He, Jiacheng Lu, and Tianzheng Wang. 2020. CoroBase: coroutine-oriented main-memory database engine. *Proceedings of the VLDB Endowment* 14, 3 (2020), 431–444.
- [19] Binbin Hu, Chuan Shi, Wayne Xin Zhao, and Philip S Yu. 2018. Leveraging meta-path based context for top-n recommendation with a neural co-attention model. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1531–1540.
- [20] Glen Jeh and Jennifer Widom. 2002. SimRank: a measure of structural-context similarity. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 538–543.
- [21] Saurabh Jha, Bingsheng He, Mian Lu, Xuntao Cheng, and Huynh Phung Huynh. 2015. Improving Main Memory Hash Joins on Intel Xeon Phi Processors: An Experimental Approach. *Proc. VLDB Endow.* 8, 6 (2015), 642–653.
- [22] Christopher Jonathan, Umar Farooq Minhas, James Hunter, Justin Levandoski, and Gor Nishanov. 2018. Exploiting coroutines to attack the “killer nanoseconds”. *Proceedings of the VLDB Endowment* 11, 11 (2018), 1702–1714.
- [23] Sang-Woo Jun, Andy Wright, Sizhuo Zhang, Shuotao Xu, et al. 2018. GrafBoost: Using accelerated flash storage for external graph analytics. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 411–424.
- [24] Farzad Khorasani, Keval Vora, Rajiv Gupta, and Laxmi N Bhuyan. 2014. CuSha: vertex-centric graph processing on GPUs. In *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*, 239–252.
- [25] Changkyu Kim, Tim Kaldewey, Victor W. Lee, Eric Sedlar, Anthony D. Nguyen, Nadathur Satish, Jatin Chhugani, Andrea Di Blas, and Pradeep Dubey. 2009. Sort vs. Hash Revisited: Fast Join Implementation on Modern Multi-Core CPUs. 2, 2 (2009), 1378–1389.
- [26] Onur Kocberber, Babak Falsafi, and Boris Grot. 2015. Asynchronous memory access chaining. *Proceedings of the VLDB Endowment* 9, 4 (2015), 252–263.
- [27] Aapo Kyröla, Guy Blleloch, and Carlos Guestrin. 2012. Graphchi: Large-scale graph computation on just a PC. In *Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12)*, 31–46.
- [28] Ni Lao, Tom Mitchell, and William Cohen. 2011. Random walk inference and learning in a large scale knowledge base. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, 529–539.
- [29] Jaekyu Lee, Hyesoon Kim, and Richard Vuduc. 2012. When prefetching works, when it doesn't, and why. *ACM Transactions on Architecture and Code Optimization (TACO)* 9, 1 (2012), 1–29.
- [30] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [31] Rong-Hua Li, Jeffrey Xu Yu, Xin Huang, and Hong Cheng. 2014. Random-walk domination in large graphs. In *2014 IEEE 30th International Conference on Data Engineering*. IEEE, 736–747.
- [32] Hang Liu and H Howie Huang. 2017. Graphene: Fine-grained IO management for graph computing. In *15th USENIX Conference on File and Storage Technologies (FAST 17)*, 285–300.
- [33] Qin Liu, Zhenguo Li, John CS Lui, and Jiefeng Cheng. 2016. Powerwalk: Scalable personalized pagerank via random walks with vertex-centric decomposition. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, 195–204.
- [34] Peter Lofgren. 2015. Efficient algorithms for personalized pagerank. *arXiv preprint arXiv:1512.04633* (2015).
- [35] Peter A Lofgren, Siddhartha Banerjee, Ashish Goel, and C Seshadhri. 2014. FAST-PPR: scaling personalized pagerank estimation for large graphs. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 1436–1445.
- [36] Yucheng Low, Joseph Gonzalez, Aapo Kyröla, Danny Bickson, Carlos Guestrin, and Joseph M Hellerstein. 2012. Distributed graphlab: A framework for machine learning in the cloud. *arXiv preprint arXiv:1204.6078* (2012).
- [37] Shengliang Lu, Shixuan Sun, Johns Paul, Yuchen Li, and Bingsheng He. 2021. Cache-Efficient Fork-Processing Patterns on Large Graphs. In *Proceedings of the 2021 International Conference on Management of Data*, 1208–1221.
- [38] Xin Lv, Yuxian Gu, Xu Han, Lei Hou, Juanzi Li, and Zhiyuan Liu. 2019. Adapting Meta Knowledge Graph Information for Multi-Hop Reasoning over Few-Shot Relations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3367–3372.
- [39] Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, 135–146.
- [40] George Marsaglia. 1963. Generating discrete random variables in a computer. *Commun. ACM* 6, 1 (1963), 37–38.
- [41] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [42] n.d. [n.d.]. Wikimedia Downloads. <https://dumps.wikimedia.org/>.
- [43] n.d. 2018. Amazon Review Data. <https://nijianmo.github.io/amazon/index.html>.
- [44] Donald Nguyen, Andrew Lenharth, and Keshav Pingali. 2013. A lightweight infrastructure for graph analytics. In *Proceedings of the twenty-fourth ACM symposium on operating systems principles*, 456–471.
- [45] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [46] Santosh Pandey, Lingda Li, Adolfo Hoisie, Xiaoye Li, and Hang Liu. [n.d.]. C-SAW: A Framework for Graph Sampling and Random Walk on GPUs. In *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE Computer Society, 780–794.
- [47] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710.
- [48] Nataša Pržulj. 2007. Biological network comparison using graphlet degree distribution. *Bioinformatics* 23, 2 (2007), e177–e183.
- [49] Georgios Psaropoulos, Thomas Legler, Norman May, and Anastasia Ailamaki. 2017. Interleaving with coroutines: a practical approach for robust index joins. *Proceedings of the VLDB Endowment* 11, CONF (2017), 230–242.
- [50] Christian Robert and George Casella. 2013. *Monte Carlo statistical methods*. Springer Science & Business Media.
- [51] Ryan A. Rossi and Nesreen K. Ahmed. 2015. The Network Data Repository with Interactive Graph Analytics and Visualization. In *AAAI*. <http://networkrepository.com>
- [52] Keith Schwarz. 2011. Darts, dice, and coins: Sampling from a discrete distribution. *Retrieved* 3, 28 (2011), 2012.
- [53] Yingxia Shao, Shiyue Huang, Xupeng Miao, Bin Cui, and Lei Chen. 2020. Memory-Aware Framework for Efficient Second-Order Random Walk on Large Graphs. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of*

- Data. 1797–1812.
- [54] Jieming Shi, Renchi Yang, Tianyuan Jin, Xiaokui Xiao, and Yin Yang. 2019. Real-time top-k personalized pagerank over large graphs on gpus. *Proceedings of the VLDB Endowment* 13, 1 (2019), 15–28.
- [55] Julian Shun and Guy E Blelloch. 2013. Ligra: a lightweight graph processing framework for shared memory. In *Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming*. 135–146.
- [56] Shixuan Sun, Yuhang Chen, Shengliang Lu, Bingsheng He, and Yuchen Li. 2021. ThunderRW: An In-Memory Graph Random Walk Engine (Complete Version). <https://arxiv.org/abs/2107.11983>.
- [57] Yizhou Sun and Jiawei Han. 2013. Mining heterogeneous information networks: a structural analysis approach. *Acm Sigkdd Explorations Newsletter* 14, 2 (2013), 20–28.
- [58] Narayanan Sundaram, Nadathur Rajagopalan Satish, Md Mostofa Ali Patwary, Subramanya R Dulloor, Satya Gautam Vadlamudi, Dipankar Das, and Pradeep Dubey. 2015. Graphmat: High performance graph analytics made productive. *arXiv preprint arXiv:1503.07241* (2015).
- [59] Alastair J Walker. 1977. An efficient method for generating discrete random variables with general distributions. *ACM Transactions on Mathematical Software (TOMS)* 3, 3 (1977), 253–256.
- [60] Rui Wang, Yongkun Li, Hong Xie, Yinlong Xu, and John CS Lui. 2020. GraphWalker: An I/O-Efficient and Resource-Friendly Graph Analytic System for Fast and Scalable Random Walks. In *2020 USENIX Annual Technical Conference (USENIX ATC 20)*. 559–571.
- [61] Sibowang, Renchi Yang, Xiaokui Xiao, Zhewei Wei, and Yin Yang. 2017. FORA: simple and effective approximate single-source personalized pagerank. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 505–514.
- [62] Yangzihao Wang, Andrew Davidson, Yuechao Pan, Yuduo Wu, Andy Riffel, and John D Owens. 2016. Gunrock: A high-performance graph processing library on the GPU. In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 1–12.
- [63] Zhewei Wei, Xiaodong He, Xiaokui Xiao, Sibowang, Shuo Shang, and Ji-Rong Wen. 2018. Topppr: top-k personalized pagerank queries with precision guarantees on large graphs. In *Proceedings of the 2018 International Conference on Management of Data*. 441–456.
- [64] Samuel Williams, Andrew Waterman, and David Patterson. 2009. Roofline: an insightful visual performance model for multicore architectures. *Commun. ACM* 52, 4 (2009), 65–76.
- [65] Da Yan, James Cheng, Yi Lu, and Wilfred Ng. 2014. Blogel: A block-centric framework for distributed computation on real-world graphs. *Proceedings of the VLDB Endowment* 7, 14 (2014), 1981–1992.
- [66] Ke Yang, MingXing Zhang, Kang Chen, Xiaosong Ma, Yang Bai, and Yong Jiang. 2019. Knightking: a fast distributed graph random walk engine. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 524–537.
- [67] Zhonglin Ye, Haixing Zhao, Ke Zhang, Yu Zhu, Yuzhi Xiao, and Zhaoyang Wang. 2019. Improved DeepWalk Algorithm Based on Preference Random Walk. In *CCF International Conference on Natural Language Processing and Chinese Computing*. Springer, 265–276.
- [68] Yunming Zhang, Mengjiao Yang, Riyadh Baghdadi, Shoaib Kamil, Julian Shun, and Saman Amarasinghe. 2018. Graphit: A high-performance graph dsl. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 1–30.
- [69] Jianlong Zhong and Bingsheng He. 2013. Medusa: Simplified graph processing on GPUs. *IEEE Transactions on Parallel and Distributed Systems* 25, 6 (2013), 1543–1552.
- [70] Dongyan Zhou, Songjie Niu, and Shimin Chen. 2018. Efficient graph computation for Node2Vec. *arXiv preprint arXiv:1805.00280* (2018).
- [71] Xiaowei Zhu, Wenguang Chen, Weimin Zheng, and Xiaosong Ma. 2016. Gemini: A computation-centric distributed graph processing system. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. 301–316.