

Fast Parallel Algorithms for Enumeration of Simple, Temporal, and Hop-constrained Cycles

JOVAN BLANUŠA and KUBILAY ATASU, IBM Research Europe-Zurich, Switzerland PAOLO IENNE, Ecole Polytechnique Fédérale de Lausanne, Switzerland

Cycles are one of the fundamental subgraph patterns and being able to enumerate them in graphs enables important applications in a wide variety of fields, including finance, biology, chemistry, and network science. However, to enable cycle enumeration in real-world applications, efficient parallel algorithms are required. In this work, we propose scalable parallelisation of state-of-the-art sequential algorithms for enumerating simple, temporal, and hop-constrained cycles. First, we focus on the simple cycle enumeration problem and parallelise the algorithms by Johnson and by Read and Tarjan in a fine-grained manner. We theoretically show that our resulting fine-grained parallel algorithms are scalable, with the fine-grained parallel Read-Tarjan algorithm being strongly scalable. In contrast, we show that straightforward coarse-grained parallel versions of these simple cycle enumeration algorithms that exploit edge- or vertex-level parallelism are not scalable. Next, we adapt our fine-grained approach to enable the enumeration of cycles under time-window, temporal, and hop constraints. Our evaluation on a cluster with 256 CPU cores that can execute up to 1,024 simultaneous threads demonstrates a near-linear scalability of our fine-grained parallel algorithms when enumerating cycles under the aforementioned constraints. On the same cluster, our fine-grained parallel algorithms achieve, on average, one order of magnitude speedup compared to the respective coarse-grained parallel versions of the state-of-the-art algorithms for cycle enumeration. The performance gap between the fine-grained and the coarse-grained parallel algorithms increases as we use more CPU cores.

CCS Concepts: • Theory of computation \rightarrow Parallel algorithms; Graph algorithms analysis;

Additional Key Words and Phrases: Cycle enumeration, parallel graph algorithms, graph pattern mining

ACM Reference format:

Jovan Blanuša, Kubilay Atasu, and Paolo Ienne. 2023. Fast Parallel Algorithms for Enumeration of Simple, Temporal, and Hop-constrained Cycles. *ACM Trans. Parallel Comput.* 10, 3, Article 15 (September 2023), 35 pages. https://doi.org/10.1145/3611642

1 INTRODUCTION

Graphs are widely adopted for usage as a data representation tool across many domains [47, 50, 70, 71]. A method of analysing graph-based data is to enumerate subgraph patterns, such as cycles, cliques, and motifs, in graphs [3]. However, enumerating subgraph patterns often leads to long

J. Blanuša is also with Ecole Polytechnique Fédérale de Lausanne, School of Computer and Communication Sciences. The support of Swiss National Science Foundation (Project No. 172610) for this work is gratefully acknowledged. Authors' addresses: J. Blanuša and K. Atasu, IBM Research Europe-Zurich, Säumerstrasse 4, CH-8803 Rüschlikon, Switzerland; emails: {jov, kat}@zurich.ibm.com; P. Ienne, Ecole Polytechnique Fédérale de Lausanne, School of Computer and Communication Sciences, Route Cantonale, CH-1015 Lausanne, Switzerland; email: paolo.ienne@epfl.ch.



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

© 2023 Copyright held by the owner/author(s). 2329-4949/2023/09-ART15 \$15.00 https://doi.org/10.1145/3611642



Fig. 1. Per-thread execution time of (a) the coarse-grained Johnson algorithm vs. (b) our fine-grained Johnson algorithm using the WT graph and a 12 h time window. Thanks to a perfect load balancing, our fine-grained method is $3 \times$ faster on a 64-core CPU executing 256 threads.

execution times, because the number of subgraph patterns that can exist in a graph [4, 20] can be orders of magnitude greater than the number of graph vertices. As a result, fast subgraph enumeration algorithms are required that can exploit the parallelisim in modern multi-core processors. In this article, we focus on enumerating simple cycles in directed graphs and introduce scalable parallel algorithms for that problem. A simple cycle is a sequence of edges that starts and ends with the same vertex and visits other vertices at most once. Enumerating simple cycles has important applications in several domains. For example, in electronic design automation, combinatorial loops in circuits are typically forbidden [28, 56], and such loops can be detected by enumerating simple cycles. In a software bug tracking system, a dependency between two software bugs requires one bug to be addressed before the other [63]. Circular bug dependencies are undesirable and can be detected by finding simple cycles. Other applications include detecting feedback loops in biological networks [37, 41] and detecting unstable relationships in social networks [25, 74].

Various types of constraints are often imposed on the simple cycles, because the search for simple cycles may otherwise be computationally impossible [39, 54, 58]. For instance, temporal ordering constraints can be imposed when searching for simple cycles in temporal graphs that have edges annotated with timestamps. Simple cycles enumerated under this constraint contain edges ordered in time; such cycles are referred to as temporal cycles [39]. Enumerating temporal cycles has applications in the financial domain, where a temporal cycle can be an indicator of money laundering [30, 43, 65], credit card fraud [58], or circular trading used for manipulating stock prices [31, 34, 52]. Other types of constraints include hop constraints [54, 58], which limit the length of paths explored during the search for cycles, and time-window constraints [39], which restrict the search to cycles that occur within a time window of a given size. Hop-constrained cycles can be used to detect fraudulent behaviour in e-commerce networks [58]. In addition, long cycles in financial transaction networks are less likely to be associated with money laundering, because they increase the risk for fraudsters of being caught [43], and imposing hop-constraints can filter out such cycles. Furthermore, searching for cycles under temporal ordering, hop, and time-window constraints reduces the number of paths explored during the search, making the cycle enumeration problem more tractable. Therefore, we focus on searching for cycles under these constraints.

Parallelisation challenges. We focus on parallelising the algorithms by Johnson [35] and by Read and Tarjan [60] for enumerating simple cycles, because these algorithms achieve the lowest time complexity bounds reported for directed graphs [27, 46]. Both algorithms are recursively formulated and construct a recursion tree in a depth-first fashion. However, these algorithms employ different pruning techniques to limit the amount of work they perform. In practice, the Johnson algorithm is faster than the Read-Tarjan algorithm due to more aggressive pruning techniques [27, 46]. Furthermore, the state-of-the-art algorithms for temporal and hop-constrained cycle enumeration are extensions of the Johnson algorithm [39, 54]. Thus, parallelising the Johnson algorithms.

Parallel algorithm	Work-efficient	Scalable
Coarse-grained parallel algorithms	\checkmark	
Our fine-grained parallel Johnson		1
Our fine-grained parallel Read-Tarjan	✓	\checkmark

Table 1. Summary of Our Theoretical Analysis

Our fine-grained parallel Read-Tarjan algorithm is the only solution that is both work-efficient and scalable.

The naïve way of parallelising the Johnson and the Read-Tarjan algorithms involves searching for cycles starting from different vertices or edges in parallel, which we refer to as the *coarse-grained parallel* methods. Such coarse-grained parallel approaches are straightforward to implement using the popular vertex-centric [45, 48] and edge-centric [62] graph processing frameworks. However, real-world graphs often exhibit a power-law or a log-normal distribution of vertex degrees [8, 16]. In such graphs, the execution time of coarse-grained parallel approaches is dominated by searches that start from a small set of vertices or edges as illustrated in Figure 1(a). This behaviour leads to a workload imbalance and limits scalability of parallel implementations.

The shortcomings of coarse-grained parallel approaches can be addressed by decomposing the search for cycles starting from a given edge or vertex into finer-grained tasks [1, 11, 20]. However, parallelising the Johnson algorithm using the fine-grained approach is challenging, because the pruning efficiency of this algorithm depends on a strictly sequential depth-first-search-based recursion tree exploration. We demonstrate that the lesser-known Read-Tarjan algorithm does not have such a requirement, and, thus, it is easier to decompose into fine-grained tasks.

Contributions. This article presents an extension of the work by Blanuša et al. [10], which introduces the following contributions:

(*i*) Scalable fine-grained parallelisation of the Johnson and the Read-Tarjan algorithms. To our knowledge, we are the first ones to parallelise these asymptotically optimal cycle enumeration algorithms in a fine-grained manner and achieve an almost linear performance scaling on a system that can execute up to a thousand concurrent software threads. Such a scalability is enabled by our decomposition of long sequential searches into fine-grained tasks, which are then dynamically scheduled across CPU cores. To decompose the Johnson algorithm into fine-grained tasks, we have relaxed its strictly depth-first-search-based exploration, which enables this algorithm to perform multiple independent depth-first searches in parallel. As a result, our fine-grained parallel Johnson algorithm is able to achieve an ideal load balancing as shown in Figure 1(b).

(ii) Theoretical analysis of the coarse- and fine-grained parallel algorithm. We theoretically show that both of our fine-grained parallel algorithms are scalable, which is not the case for the Johnson and the Read-Tarjan algorithms parallelised in a coarse-grained manner. Moreover, we show that our fine-grained parallel Read-Tarjan algorithm performs asymptotically the same amount of work as its serial version, whereas our fine-grained parallel Johnson algorithm does not. Therefore, our fine-grained parallel Read-Tarjan algorithm is the only parallel algorithm based on an asymptotically optimal cycle enumeration algorithm that is both work-efficient and scalable, as shown in Table 1. Interestingly, despite not being work-efficient, our fine-grained Johnson algorithm outperforms our fine-grained parallel Read-Tarjan algorithm in most of our experiments.

In this article, we extend our prior work [10] with the following contributions:

(*iii*) General framework for parallelising temporal and hop-constrained cycle enumeration. We show that our method for parallelising the Johnson algorithm in a fine-grained manner can be adapted to parallelise the state-of-the-art algorithms for temporal and hop-constrained cycle enumeration. This adaptation is possible because these state-of-the-art algorithms, such as the 2SCENT algorithm for temporal cycle enumeration [39] and the BC-DFS algorithm [54] for hop-

Related work	[39]	[58]	[54]	[57]	[29]	Ours
Fine-grained parallelism				~		~
Asymptotic optimality	1		\checkmark		\checkmark	\checkmark
Temporal cycles	1					\checkmark
Time-window constraints	1	\checkmark				\checkmark
Hop constraints		\checkmark	\checkmark	\checkmark	\checkmark	\checkmark

Table 2. Capabilities of the Related Work Versus Our Own

Competing algorithms either fail to exploit fine-grained parallelism or do it on top of asymptotically inferior algorithms.

constrained cycle enumeration, are extensions of the Johnson algorithm. By parallelising these algorithms using our fine-grained method, we were able to achieve speedups of up to $40 \times$ and $61 \times$ compared to the coarse-grained parallel versions of 2SCENT and BC-DFS, respectively.

(iv) Improvements to the pruning efficiency of the Read-Tarjan algorithm. To make this algorithm competitive with the Johnson algorithm, we have introduced several optimisations that enhance the pruning efficiency of the Read-Tarjan algorithm. The optimisations reduce the amount of unnecessary vertex visits that this algorithm performs. As a result, our improved version of the Read-Tarjan algorithm is up to 6.8× faster than the original version of this algorithm.

Article structure. The remainder of this article is organised as follows. The related work and background are presented in Sections 2 and 3, respectively. Coarse-grained parallel versions of the Johnson and the Read-Tarjan algorithms are covered in Section 4. Sections 5 and 6 introduce our fine-grained parallel versions of the Johnson and the Read-Tarjan algorithms, respectively. Section 6 also includes our optimisations for improving the pruning efficiency of the Read-Tarjan algorithm. Our general framework for parallelising temporal and hop-constrained cycle enumeration algorithms is presented in Section 7. In Section 8, we provide an experimental evaluation of our fine-grained parallel algorithms. Finally, we conclude our work in Section 9.

2 RELATED WORK

Simple cycle enumeration algorithms. Enumeration of simple cycles of graphs is a classical computer science problem [2, 9, 27, 35, 44, 46, 60, 66–68, 72]. The backtracking-based algorithms by Johnson [35], Read and Tarjan [60], and Szwarcfiter and Lauer [66] achieve the lowest time complexity bounds for enumerating simple cycles in directed graphs. These algorithms implement advanced recursion tree pruning techniques to improve on the brute-force Tiernan algorithm [68]. Section 3.4 covers such pruning techniques in further detail. A cycle enumeration algorithm that is asymptotically faster than the aforementioned algorithms [35, 60, 66] has been proposed in Birmelé et al. [9]; however, it is applicable only to undirected graphs. Simple cycles can also be enumerated by computing the powers of the adjacency matrix [19, 36, 55] or by using circuit vector space algorithms [24, 46, 73], but the complexity of such approaches grows exponentially with the size of the cycles or the size of the input graphs.

Time-window, temporal ordering, and hop constraints. It is common to search for cycles under some additional constraints. For instance, in temporal graphs, it is common to search for cycles within a sliding time window, such as in Kumar and Calders [39] and Qiu et al [58]. In addition, temporal ordering constraints can be imposed when searching for cycles in temporal graphs, such as in Kumar and Calders [39]. Furthermore, the maximum number of hops in cycles or paths can be constrained, such as in Gupta and Suzumura [29] and Peng et al. [54]. Note that hop-constrained simple cycles can also be enumerated using incremental algorithms, such as in Qiu et al. [58]. However, this algorithm is based on the brute-force Tiernan algorithm [68], which makes it slower than nonincremental algorithms that use recursion tree pruning techniques [54].

Symbol	Description	Symbol	Description
$\overline{\mathcal{G}(\mathcal{V},\mathcal{E})}$	Graph with vertices ${\mathcal V}$ and edges ${\mathcal E}$	П	Current simple path explored by an algorithm
$\mathcal{N}(v)$	The set of neighbours of the vertex v	Blk	Set of blocked vertices
$u \rightarrow v$	A directed edge from u to v	Blist	Unblock list of the Johnson algorithm
n, e	No. vertices and edges in a graph	Ε	Path extension of the Read-Tarjan algorithm
δ	Size of a time window	X_{T_i}	Data structure <i>X</i> is maintained by the thread T_i
L	Hop constraint	p	No. threads used by a parallel algorithm
С	No. simple cycles in a graph	$T_p(n)$	Execution time of a parallel algorithm
S	No. maximal simple paths in a graph	$\hat{W}_p(n)$	Amount of work a parallel algorithm performs

Table 3. Summary of the Notation Used in the Article

Additionally, because incremental algorithms maintain auxiliary data structures, such as paths, to be able to construct cycles incrementally, they are not as memory-efficient as nonincremental algorithms [54]. Table 2 offers comparisons between the capabilities of these methods and ours.

Parallel and distributed algorithms for cycle enumeration. Cui et al. [18] proposed a multithreaded algorithm for detecting and removing simple cycles of a directed graph. The algorithm divides the graph into its strongly connected components and each thread performs a depth-first search on a different component to find cycles. However, sizes of the strongly connected components in real-world graphs can vary significantly [49], which leads to a workload imbalance. Rocha and Thatte [61] proposed a distributed algorithm for simple cycle enumeration based on the bulksynchronous parallel model [69], but it searches for cycles in a brute-force manner. Qing et al. [57] introduced a parallel algorithm for finding length-constrained simple cycles. It is the only other fine-grained parallel algorithm we are aware of in the sense that it can search for cycles starting from the same vertex in parallel. However, the way this algorithm searches for cycles is similar to the way the brute-force Tiernan algorithm [68] works. To our knowledge, we are the first ones to introduce fine-grained parallel versions of asymptotically optimal simple cycle enumeration algorithms, which do not rely on a brute-force search, as we show in Table 2. Distributed algorithms for detecting the presence of cycles in graphs readily exist [6, 23, 51]. However, our focus is on discovering all simple cycles of a graph rather than detecting whether a graph has a cycle or not.

3 BACKGROUND

This section introduces the main theoretical concepts used in this article and provides an overview of the most prominent simple cycle enumeration algorithms. The notation used is given in Table 3.

3.1 Preliminaries

We consider a directed graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$ having a set of vertices \mathcal{V} and a set of directed edges $\mathcal{E} = \{u \to v \mid u, v \in \mathcal{V}\}$. The set of neighbours of a given vertex v is defined as $\mathcal{N}(v) = \{w \mid v \to w \in \mathcal{E}\}$. We refer to the vertex v of an edge $v \to u$ as its source vertex and to the vertex u as its destination vertex. An outgoing edge of a given vertex v is defined as $v \to w$ and an incoming edge is defined as $u \to v$, where $v \to w, u \to v \in \mathcal{E}$. A *path* between the vertices v_0 and v_k , denoted as $v_0 \to v_1 \dots \to v_k$, is a sequence of vertices such that there exists an edge between every two consecutive vertices of the sequence. A *simple path* is a path with no repeated vertices. A simple path is *maximal* if the last vertex of the path has no neighbours or all of its neighbours are already in the path [21]. A **cycle** is a path of non-zero length from a vertex v to the same vertex v. A **simple cycle** is a cycle with no repeated vertices except for the first and last vertex. The number of maximal simple paths and the number of simple cycles in a graph are denoted as s and c, respectively (see Table 3). Note that s can be exponentially larger than c [67]. A path or a cycle is said to satisfy a **hop-constraint** L if the number of edges in that path or cycle is less than

J. Blanuša et al.



Fig. 2. Two snapshots of a temporal graph associated with two different time windows of size $\delta = 5$. The solid arrows indicate the edges that belong to the respective time windows.

or equal to *L*. The goal of **simple cycle enumeration** is to compute all simple cycles of a directed graph G, ideally without computing all maximal simple paths of it.

A **temporal graph** is a graph that has its edges annotated with timestamps [53]. Such a graph might contain *parallel edges*, which are edges with the same source and destination vertices [7]. An example of a temporal graph with parallel edges is given in Figure 2. In temporal graphs, a **temporal cycle** is a simple cycle, in which the edges appear in the increasing order of their timestamps. A simple cycle or a temporal cycle of a temporal graph occurs within a **time window** $[t_{w1} : t_{w2}]$ if every edge of that cycle has a timestamp t_s such that $t_{w1} \le t_s \le t_{w2}$. Figure 2 shows the simple cycles of a temporal graph that occur within two different time windows of size $\delta = 5$. This graph contains two simple cycles in the time window [2 : 7] (Figure 2(a)), which are also temporal cycles, and two simple cycles in the time window [10 : 15] (Figure 2(b)), neither being a temporal cycle. Note that the existence of parallel edges in temporal graphs makes it possible to have several simple cycles that contain the same sequence of vertices, as shown in Figure 2(a). The union of several cycles that contain the same sequence of vertices is called a *cycle bundle* [39].

3.2 Task-level Parallelism

The parallel algorithms described in this article can be implemented using shared-memory parallel processing frameworks, such as TBB [38], Cilk [13], and OpenMP [59]. These frameworks enable the decomposition of a program into tasks that can be independently executed by different software threads. In our setup, tasks are created and scheduled dynamically. A *parent* task can *spawn* several *child* tasks. The *depth* of a task is the number of its direct ancestors. A dynamic task management system assigns the tasks created to the work queues of the available threads. Furthermore, a work-stealing scheduler [13, 14, 38] enables a thread that is not executing a task to *steal* a task from the work queue of another thread. Stealing tasks enables dynamic load balancing and ensures full utilisation of the threads when there are sufficiently many tasks.

3.3 Work Efficiency and Scalability

We use the notions of *work efficiency* and *scalability* to analyse parallel algorithms [12]. We refer to the time to execute a parallel algorithm on a problem of size n using p threads as $T_p(n)$. The size of a graph is determined by the number of vertices n as well as the number of edges e, but we will refer only to n for simplicity. The *depth* of an algorithm is the length of the longest sequence of dependent operations in the algorithm. The time required to execute such a sequence is equal to the execution time of the parallel algorithm using an infinite number of threads, denoted by T_{∞} . Furthermore, *work* performed by a parallel algorithm on a problem of size n using p threads, denoted as $W_p(n)$, is the sum of the execution times of the individual threads. The *work efficiency* and the *scalability* are formally defined as follows.

Definition 1 (Work Efficiency). A parallel algorithm is work-efficient if and only if $W_p(n) \in O(T_1(n))$.

Definition 2 (Scalability). A parallel algorithm is scalable if and only if $\lim_{n \to \infty} \left(\lim_{p \to \infty} \frac{T_p(n)}{T_1(n)} \right) = 0.$

Informally, a work-efficient parallel algorithm performs the same amount of work as its serial version, within a constant factor. Scalability implies that, for sufficiently large inputs, increasing the number of threads increases the speedup of the parallel algorithm with respect to its serial version.

We also define the notion of strong scalability as follows [32].

Definition 3 (Strong Scalability). A parallel algorithm is strongly scalable if and only if $\frac{T_1(n)}{T_p(n)} = \Theta(p)$ for large enough *n*.

Whereas Definition 2 implies that the speedup $T_1(n)/T_p(n)$ achieved by a parallel algorithm with respect to its serial execution is infinite when the number of threads p is infinite, Definition 3 implies that the speedup is always in the order of p. Another related concept is weak scalability, which requires the speedup to be in the order of p when the input size per thread is constant. Note that both strong scalability and weak scalability imply scalability.

3.4 Simple Cycle Enumeration Algorithms

The following algorithms for simple cycle enumeration perform recursive searches to incrementally update simple paths that can lead to cycles. Each algorithm iterates the vertices or edges of the graph and independently constructs a recursion tree to enumerate all the cycles starting from that vertex or edge. The difference between these algorithms is to what extent they reduce the redundant work performed during the recursive search, which we discuss next.

The Tiernan algorithm [68] enumerates simple cycles using a brute-force search. It recursively extends a simple path Π by appending a neighbour u of the last vertex v of Π provided that u is not already in Π . A clear downside of this algorithm is that it can repeatedly visit vertices that can never lead to a cycle. When searching for cycles in the graph shown in Figure 3(a) starting from the vertex v_0 , this algorithm would explore the path containing b_1, \ldots, b_k 2m times. From each vertex w_i and u_i , with $i \in \{1, \ldots, m\}$, the Tiernan algorithm would explore this path only to discover that it cannot lead to a simple cycle. As noted by Tarjan [67], the Tiernan algorithm explores every simple path and, consequently, all maximal simple paths of a graph. Exploring a maximal simple path takes O(e) time, because it requires visiting each edge of the graph in the worst case. Given a graph with s maximal simple paths (see Table 3), the worst-case time complexity of the Tiernan algorithm is O(se).

The Johnson algorithm [35] improves upon the Tiernan algorithm by avoiding the vertices that cannot lead to simple cycles when appended to the current simple path Π . For this purpose, the Johnson algorithm maintains a set of blocked vertices *Blk* that are avoided during the search. In addition, a list of vertices *Blist*[*w*] is stored for each blocked vertex *w*. Whenever a vertex *w* is unblocked (i.e., removed from *Blk*) by the Johnson algorithm, the vertices in *Blist*[*w*] are also unblocked. This unblocking process is performed recursively until no more vertices can be unblocked, which we refer to as the *recursive unblocking* procedure.

A vertex v is blocked (i.e., added to *Blk*) when visited by the algorithm. If a cycle is found after recursively exploring every neighbour of v that is not blocked, then the vertex v is unblocked. However, v is not immediately unblocked if no cycles are found after exploring its neighbours. Instead, the *Blist* data structure is updated to enable unblocking of v in a later step by adding v to the list *Blist*[w] of every neighbour w of v. This delayed unblocking of the vertices enables the Johnson algorithm to discover each cycle in O(e) time in the worst case. Because this algorithm requires O(n + e) time to determine that there are no cycles, its worst-case time complexity



Fig. 3. (a) An example graph and (b) the recursion tree constructed when searching for cycles that start from v_0 . The nodes of the recursion tree represent the recursive calls of the depth-first search. The dotted path of the right subtree is explored only by the Read-Tarjan algorithm.

is O(n + e + ec) [66]. Note that because *s* can be exponentially larger than *c* [67], the Johnson algorithm is asymptotically faster than the Tiernan algorithm.

In the example shown in Figure 3(a), every simple path Π that starts from v_0 and contains vertices b_1, \ldots, b_k is a maximal simple path, and, thus, it cannot lead to a simple cycle. The Johnson algorithm would block b_1, \ldots, b_k immediately after visiting this sequence once and then keep these vertices blocked until it finishes exploring the neighbours of v_2 . As a result, the Johnson algorithm visits vertices b_1, \ldots, b_k only once, rather than 2m times the Tiernan algorithm would visit them. Note that because these vertices get blocked during the exploration of the left subtree of the recursion tree, they are not going to be visited again during the exploration of the right subtree. Effectively, a portion of the right subtree is pruned (see the dotted path in Figure 3(b)) based on the updates made on *Blk* and *Blist* during the exploration of the left subtree. This strictly sequential depth-first exploration of the recursion tree is critically important for achieving a high pruning efficiency, but it also makes scalable parallelisation of the Johnson algorithm extremely challenging, which we are going to cover in Section 5.

The Read-Tarjan algorithm [60] also has a worst-case time complexity of O(n + e + ec). This algorithm maintains a current path Π between a starting vertex and a frontier vertex. A recursive call of this algorithm iterates the neighbours of the current frontier vertex and performs a **depthfirst search (DFS)**. Assume that v_0 is the starting vertex and v_1 is the frontier vertex of Π (see Figure 3(a)). From each neighbour $y \in \{v_0, v_2\}$ of v_1 , a DFS tries to find a path extension *E* back to v_0 that would form a simple cycle when appended to Π . In the example shown in Figure 3(a), the algorithm finds two path extensions, one indicated as *E* and one that consists of the edge $v_1 \rightarrow v_0$. The algorithm then explores each path extension by iteratively appending the vertices from it to the path Π . For each vertex x added to Π , the algorithm also searches for an alternate path extension from that vertex x to v_0 using a DFS. In the example given in Figure 3(a), the algorithm iterates through the vertices of the path extension E and finds an alternate path extension E' from the neighbour u_1 of v_2 . If an alternate path extension is found, then a child recursive call is invoked with the updated current path Π , which is $v_0 \rightarrow v_1 \rightarrow v_2$ in our example. Otherwise, if all the vertices in *E* have already been added to the current path Π , then Π is reported as a simple cycle. In our example, the Read-Tarjan algorithm explores both E and E' path extensions, and each leads to the discovery of a cycle.

The Read-Tarjan algorithm also maintains a set of blocked vertices *Blk* for recursion-tree pruning. However, differently from the Johnson algorithm, *Blk* only keeps track of the vertices that

ACM Transactions on Parallel Computing, Vol. 10, No. 3, Article 15. Publication date: September 2023.

cannot lead to new cycles when exploring the current path extension within the same recursive call. The vertices in *Blk* are avoided while searching for additional path extensions that branch from the current path extension. For instance, the left subtree of the recursion tree shown in Figure 3(b) demonstrates the exploration of the path extension *E* shown in Figure 3(a). During the exploration of *E*, the vertices b_1, \ldots, b_k are added to *Blk* immediately after visiting w_1 , and they are not visited again while exploring *E*. However, when exploring another path extension *E'* in the right subtree, the vertices b_1, \ldots, b_k are visited once again (see the dotted path of the right subtree). As a result, the Read-Tarjan algorithm visits b_1, \ldots, b_k twice instead of just once. As we are going to show in Section 6, this drawback becomes an advantage when parallelising the Read-Tarjan algorithm, because it enables independent exploration of different subtrees of the recursion tree.

4 COARSE-GRAINED PARALLEL METHODS

The most straightforward way of parallelising the Johnson and the Read-Tarjan algorithms is to search for cycles that start from different vertices in parallel. Each such search can then be executed by a different thread that explores its own recursion tree. This approach is beneficial, because it is work-efficient and can be implemented using one of the existing graph processing frameworks, such as Pregel [45], in a manner similar to the method used by Rocha and Thatte [61]. We refer to this parallelisation approach as the coarse-grained parallel approach.

The coarse-grained approach can express more parallelism if each thread performs a search for cycles that start from a different edge rather than a different vertex. This assumption is supported by the fact that graphs typically have more edges than vertices. Nevertheless, the coarse-grained approach is not scalable, which we prove here.

PROPOSITION 1. The coarse-grained parallel Johnson and Read-Tarjan algorithms are work-efficient.

The proof of Proposition 1 is trivial, and we omit it for brevity.

THEOREM 1. The coarse-grained parallel Johnson and Read-Tarjan algorithms are not scalable.

PROOF. In this case, the depth $T_{\infty}(n)$ represents the worst-case execution time of a search for cycles that starts from a single vertex or edge, and it depends on the number of cycles found during this search. In the worst case, a single recursive search can discover all cycles of a graph. An example of such graph is given in Figure 4(a), where each vertex v_i , with $i \in \{1, \ldots, n-1\}$, is connected to v_0 and to every vertex v_j such that j > i. In that graph, any subset of vertices v_2, \ldots, v_{n-1} defines a different cycle. Therefore, the total number of cycles in this graph is equal to the number of all such subsets $c = 2^{n-2}$. Before the search for cycles, both the Johnson and the Read-Tarjan algorithm find all vertices that start a cycle, which is only v_0 in this case. Therefore, the search for cycles will be performed only by one thread. Because both the Johnson and the Read-Tarjan algorithms require O(e) time to find each cycle, the depth of the coarse-grained algorithms is $T_{\infty}(n) \in O(ec)$. Because $\lim_{n\to\infty} T_{\infty}(n)/T_1(n) \neq 0$, the coarse-grained algorithms are not scalable based on Definition 2.

Theorem 1 shows that the main drawback of the coarse-grained parallel algorithms is their limited scalability. This limitation is apparent for the graph shown in Figure 4(a), which has an exponential number of cycles in n. When using a coarse-grained parallel algorithm on this graph, all the cycles will be discovered by a single thread, and, thus, the depth of this algorithm grows linearly with c, as shown in Table 4. Because only one thread can be effectively utilised, increasing the number of threads will not result in a reduction of the overall execution time of the coarse-grained parallel algorithm. Figure 1(a) shows the workload imbalance exhibited by the

J. Blanuša et al.



Fig. 4. (a) A graph with an exponential number of simple cycles. (b) The recursion tree of the Johnson algorithm for n = 6 constructed when the algorithm starts from v_0 . Whereas a coarse-grained parallel algorithm explores the complete recursion tree using a single thread, our fine-grained parallel algorithms can explore different regions of the recursion tree in parallel using several threads.

Table 4. Work and Depth of the Coarse- and Fine-grained Parallel Algorithms

Parallel algorithm	Work	Depth
Coarse-grained algorithms	O(n + e + ec)	<i>O</i> (<i>ec</i>)
Fine-grained Johnson algorithm	$O(n + e + \min\{pce, se\})$	<i>O</i> (<i>e</i>)
Fine-grained Read-Tarjan algorithm	O(n + e + ec)	0 (ne)

coarse-grained parallel algorithms in practice. Section 8 demonstrates the limited scalability of coarse-grained parallel algorithms in further detail.

5 FINE-GRAINED PARALLEL JOHNSON

To address the load imbalance issues that manifest themselves in the coarse-grained parallel Johnson algorithm, we introduce the *fine-grained parallel Johnson* algorithm. The main goal of our finegrained algorithm is to enable several threads to explore a recursion tree concurrently, as shown in Figure 4(b), where each thread executes a subset of the recursive calls of this tree. However, enabling concurrent exploration of a recursion tree is in conflict with the sequential depth-first exploration, required by the Johnson algorithm to achieve a high pruning efficiency.

In this section, we first discuss the challenges that arise when parallelising the exploration of a recursion tree of the Johnson algorithm. Then, we introduce the *copy-on-steal* mechanism used to address these challenges and present our fine-grained parallel Johnson algorithm. Finally, we theoretically analyse our algorithm and show that it is scalable.

5.1 Fine-grained Parallelisation Challenges

The requirement of the sequential depth-first exploration of the Johnson algorithm makes it challenging to efficiently parallelise this algorithm in a fine-grained manner. This requirement is enforced by maintaining a set of blocked vertices *Blk* throughout the exploration of a recursion tree. If threads exploring the same recursion tree simply share the same set of blocked vertices *Blk*, then the parallel algorithm could produce incorrect results. For example, considering the graph given in Figure 5(a), a thread exploring the path $\Pi = v_0 \rightarrow v_1 \rightarrow u_1 \rightarrow v_2$ visit and block the vertex u_4 in this case, because u_4 cannot participate in a simple cycle that begins with Π . Because the threads exploring this graph share the blocked vertices, another thread attempting to discover the cycle $v_0 \rightarrow v_1 \rightarrow u_1 \rightarrow v_2 \rightarrow v_0$ would fail to do so, because u_4 is blocked. Therefore, this approach might not discover all cycles in a graph.



Fig. 5. (a) An example graph and (b) the recursion tree of our fine-grained Johnson algorithm when enumerating cycles that start from v_0 . Each thread of our fine-grained Johnson algorithm explores the vertices b_1, \ldots, b_m at most once.

To enable several threads to correctly find all cycles while exploring the same recursion tree, the algorithm could forward a new copy of the *Blk* and *Blist* data structures when invoking each child recursive call. However, this approach would redundantly explore many paths in a graph. The reason is that a recursive call would be unaware of the vertices visited and blocked by other calls that precede it in the depth-first order except for its direct ancestors in the recursion tree. When enumerating the simple cycles of the graph shown in Figure 5(a) starting from v_0 , this approach explores all $4 \times 2^{m-1} + 3$ maximal simple paths instead of just seven, that the Johnson algorithm would explore. Hence, this approach exhaustively explores all maximal simple paths in the graph and is identical to the brute-force solution of Tiernan (see Section 3.4). Next, we propose a fine-grained parallel algorithm that addresses the aforementioned parallelisation challenges.

5.2 Copy-on-steal

To enable different threads to concurrently explore the recursion tree in a depth-first fashion while also taking advantage of the powerful pruning capabilities of the Johnson algorithm, each thread executing our fine-grained parallel Johnson algorithm maintains its own copy of the Π , *Blk*, and *Blist* data structures. These data structures are copied between threads only when these threads attempt to explore the same recursion tree. To achieve this behaviour, our fine-grained parallel Johnson algorithm implements each recursive call of the Johnson algorithm as a separate task. The pseudocode of this task is given in Algorithm 1, where a data structure X, maintained by the thread T_i , is denoted as X_{T_i} (see Table 3). If a child task and its parent task are executed by the same thread T_i , then the child task reuses the Π_{T_i} , Blk_{T_i} , and $Blist_{T_i}$ data structures of the parent task. However, if a child task has been stolen—i.e., it is executed by a thread other than the thread that created it, then the child task will allocate a new copy of these data structures (line 2 of Algorithm 1). We refer to this mechanism as *copy-on-steal*.

The problem with copying data structures between different threads upon task stealing is that the thread that has created the stolen task (i.e., the *victim thread*) can modify its data structures before this task is stolen by another thread (i.e., the *stealing thread*). This problem can be observed in the example shown in Figure 6. There, the victim thread T_1 and the stealing thread T_2 explore the same recursion tree given in Figure 6(b) while searching for cycles that start with $P_1 = v_0 \rightarrow v_1 \rightarrow$ v_2 and $P_2 = v_0 \rightarrow v_1 \rightarrow v_7$, respectively. In this case, T_2 steals a task created by T_1 that explores v_7 , as indicated in Figure 6(b), and receives a copy of the blocked vertices $Blk_{T_1} = \{v_4, v_5, v_6\}$ ALGORITHM 1: FGJ_task (v, v₀, d, T₁) **Input**: v - the current vertex, v₀ - the starting vertex d - the depth of this task **InOut**: T₁ - the thread that created this task \triangleright T₁ maintains Π_{T_1} , Blk_{T1}, Blist_{T1}, and Mutex_{T1} Output: true if a cycle was found \blacktriangleright T_2 maintains $\Pi_{T_2}\text{, }Blk_{T_2}\text{, }Blist_{T_2}\text{, and }Mutex_{T_2}$ $_{1}$ T₂ = the thread executing this task; ² **if** $T_1 \neq T_2$ **then** FGJ_copyOnSteal(d, T_1, T_2); ▶ Check if this task is stolen 3 Mutex_{T2}.lock(); 4 $\Pi_{T_2}.push(v); Blk_{T_2} = Blk_{T_2} \cup \{v\};$ 5 Mutex_{T2}.unlock(); 6 found = false; 7 foreach $u : \mathcal{N}(v)$ s.t. u.id > v₀.id do \triangleright Recursively explore the neighbours of v if $u = v_0$ then 8 report cycle Π_{T_2} ; 9 found = true; 10 else if u ∉ Blk_{T2} then 11 f =**spawn** FGJ_task(u, v₀, d + 1, T₂); ▶ Create a child task 12 found = found \lor f; 13 14 wait for the spawned tasks; 15 Mutex_{T2}.lock(); 16 Π_{T2}.pop(); 17 if found then ▶ Unblock vertices if a cycle was found RecursiveUnblock(v, Blk_{T_2} , $Blist_{T_2}$); 18 19 else **foreach** $u : \mathcal{N}(v)$ **do** $Blist_{T_2}[u] = Blist_{T_2}[u] \cup \{v\};$ 20 21 Mutex_{T2}.unlock();

22 return found;



Fig. 6. (a) An example graph and (b) the recursion tree of our fine-grained Johnson algorithm when enumerating simple cycles that start from v_0 . Here, X_{T_i} denotes a data structure X of the thread T_i . The thread T_2 can prune the dotted part of the tree by avoiding v_5 and v_6 that the thread T_1 has blocked after creating the task stolen by T_2 .

ALGORITHM 2: FGJ_copyOnSteal (d, T ₁ , T ₂)	
Input: d - the depth of the task executing this functi	on
InOut : T_1 - the victim thread	
T_2 - the stealing thread	
1 Mutex _{T1} .lock();	
2 { Π_{T_2} , Bl k_{T_2} , Blist $_{T_2}$ } = copy ({ Π_{T_1} , Bl k_{T_1} , Blist $_{T_1}$ });	\blacktriangleright The data structures of T_1 are copied to T_2
3 Mutex _{T1} .unlock();	
4 while $ \Pi_{T_2} \ge d \mathbf{do}$	▶ Copy-on-steal with recursive unblocking
5 $u = \Pi_{T_2}.pop();$	
6 RecursiveUnblock(u, Blk_{T_2} , $Blist_{T_2}$);	

discovered by T_1 . The thread T_1 blocked these vertices, because they cannot participate in any simple cycle that begins with P_1 . If T_2 simply uses a copy of these blocked vertices Blk_{T_1} without modifications, then T_2 will be unable to find the cycle $v_0 \rightarrow v_1 \rightarrow v_7 \rightarrow v_4 \rightarrow v_2 \rightarrow v_3 \rightarrow v_0$, because v_4 is blocked. Therefore, a method for unblocking vertices after copy-on-steal is required to correctly find all cycles.

We explore two solutions for this problem:

(i) Copy-on-steal with complete unblocking. To enable the threads of our algorithm to find cycles after performing copy-on-steal, the stealing thread could unblock all vertices that the victim thread had blocked after creating the stolen task. In our example given in Figure 6, the stealing thread T_2 unblocks all vertices $Blk_{T_1} = \{v_4, v_5, v_6\}$ it received from the victim thread T_1 . Although this approach enables T_2 to correctly find cycles, it also fails to take advantage of the information collected by T_1 to reduce the redundant work of T_2 . For instance, in Figure 6, T_2 visits v_5 and v_6 , even though T_1 already concluded that these vertices cannot participate in any simple cycle that begins with $P = v_0 \rightarrow v_1$, where P is the largest common prefix of all the paths explored by T_1 and T_2 . As a result, T_2 redundantly visits the dotted part of the recursion tree given in Figure 6(b).

(ii) Copy-on-steal with recursive unblocking. In this approach, the stealing thread capitalises on the information already discovered by the victim thread. The stealing thread T_2 can reuse a subset $B \subset Blk_{T_1}$ of the blocked vertices discovered by T_1 if the vertices in B cannot participate in simple cycles that begin with P, where P is the largest common prefix of all the paths explored by T_1 and T_2 . Because any path discovered by T_2 begins with P, T_2 can avoid visiting vertices from B. Thus, to correctly find simple cycles, it is sufficient for T_2 to unblock the vertices from $Blk_{T_1} \setminus B$. To achieve this behaviour, T_2 invokes a recursive unblocking procedure of the Johnson algorithm for every vertex $v \in \Pi_{T_1} \setminus P$, as shown in Algorithm 2, where Π_{T_1} is the path T_1 is exploring during task stealing. The vertices in B can only be unblocked by a recursive unblocking invoked for $v \in P$; hence, the vertices in B remain blocked. In the example given in Figure 6, T_2 invokes a recursive unblocking procedure for $\Pi_{T_1} \setminus P = \{v_2\}$, which results in unblocking of v_4 . Thus, T_2 is able to discover a cycle that contains v_4 . The vertices $B = \{v_5, v_6\}$ will not be unblocked, because they cannot take part in any simple cycle that begins with $P = v_0 \rightarrow v_1$. Therefore, thread T_2 avoids visiting the dotted part of the recursion tree given in Figure 6(b).

Without countermeasures, our algorithm can suffer from race conditions, because its data structures can be accessed concurrently by different threads. For instance, a stealing thread T_2 can copy the data structures of a victim thread T_1 while T_1 performs a recursive unblocking, in which case T_2 could receive the vertex set Blk_{T_1} that is partially unblocked. When using copy-on-steal with recursive unblocking, T_2 may not be able to continue the interrupted unblocking of Blk_{T_1} , causing the algorithm to miss certain cycles. To avoid this problem, we define critical sections in lines 15– 21 of Algorithm 1 and in lines 1–3 of Algorithm 2 using coarse-grained locking by maintaining a

ALGORITHM 3: FGJ ($\mathcal{G}(\mathcal{V}, \mathcal{E})$) Input: \mathcal{G} - the input graph with vertices \mathcal{V} and edges \mathcal{E} 1 parallel foreach $v_0 \rightarrow v : \mathcal{E}$ do 2 | T_0 = the thread executing this loop iteration; $\triangleright T_0$ maintains Π_{T_0} , Blk_{T_0} , Blist_{T_0} , and Mutex_{T_0}

▶ Create a task

6 **wait** for all spawned tasks;

 $\Pi_{T_0} = v_0; \quad \text{Blk}_{T_0} = \emptyset;$

foreach $u : \mathcal{V}$ **do** Blist_{T0} $[u] = \emptyset$;

spawn FGJ_task(v, v_0 , 1, T₀);

mutex per thread. However, such a locking mechanism is not required when using copy-on-steal with complete unblocking, because T_2 can correctly unblock vertices in Blk_{T_1} simply by removing all vertices from Blk_{T_1} inserted after the stolen task was created. Thus, it is sufficient to enable thread-safe operations on Π , Blk, and Blist using fine-grained locking. As a result, the critical sections are shorter when the copy-on-steal with complete unblocking approach is used.

Nevertheless, we opt to use the copy-on-steal with recursive unblocking approach in our finegrained parallel Johnson algorithm, because this approach leads to less redundant work and rarely suffers from synchronisation overheads. The pseudocode of our fine-grained parallel Johnson algorithm is given in Algorithm 3.

5.3 Theoretical Analysis

We now show that the fine-grained parallel Johnson algorithm is scalable but not work-efficient.

THEOREM 2. The fine-grained parallel Johnson algorithm is not work-efficient.

PROOF. According to Lemma 3 presented by Johnson [35], a vertex cannot be unblocked more than once unless a cycle is found, and once a vertex is visited, it can be visited again only after being unblocked. Thus, the Johnson algorithm visits each vertex and edge at most *c* times. In the fine-grained parallel Johnson algorithm executed using *p* threads, each thread maintains a separate set of data structures used for managing blocked vertices. Because the threads are unaware of each other's blocked vertices, each edge is visited at most *pc* times, *c* times by each thread. Additionally, an edge cannot be visited more than *s* times, because each maximal simple path of a graph is explored by a different thread in the worst case, and during each simple path exploration, an edge is visited at most once. Therefore, the maximum number of times an edge can be visited by the finegrained parallel Johnson algorithm is min {*pc*, *s*}. Because the algorithm executes in O(n + e) time if there does not exist a cycle or a path in the input graph, the work performed by the fine-grained parallel Johnson algorithm is

$$W_p(n) \in O\left(n + e + \min\{pce, se\}\right). \tag{1}$$

When c > 0, p > 1, and s > c, the work performed by the fine-grained parallel Johnson algorithm $W_p(n)$ is greater than the execution time $T_1(n)$ of the sequential Johnson algorithm. Thus, this algorithm is not work-efficient.

The work inefficiency of our fine-grained parallel Johnson algorithm occurs if more than one thread performs the work the sequential Johnson algorithm would perform between the discovery of two cycles. This behaviour can be illustrated using the graph from Figure 5(a), which contains c = 4 cycles and $s = c \times 2^{m-1} + 3$ maximal simple paths, each starting from vertex v_0 . When discovering each cycle, our fine-grained algorithm explores an infeasible region of the recursion tree, as shown in Figure 5(b), in which the vertices b_1, \ldots, b_m are visited. If this infeasible region is explored using a single thread, then each vertex b_i , with $i \in \{1, \ldots, m\}$, will be visited exactly once. However, if p threads are exploring the same infeasible region of the recursion tree, then

3

4

5

15:15

vertices b_1, \ldots, b_m will be visited up to p times, because the threads are unaware of each other's blocked vertices. In this case, the fine-grained parallel Johnson algorithm performs more work than necessary, and, thus, it is not work-efficient. Additionally, each infeasible region of the recursion tree that visits vertices b_1, \ldots, b_k can be executed by at most $s/c = 2^{m-1}$ threads, because there are 2^{m-1} maximal simple paths that can be explored in each infeasible region. In this case, each vertex b_i , with $i \in \{1, \ldots, m\}$, is visited up to s times, and, thus, the fine-grained parallel Johnson algorithm behaves as the Tiernan algorithm (see Section 3.4).

LEMMA 1. The depth $T_{\infty}(n)$ of the fine-grained parallel Johnson algorithm is in O(e).

PROOF. The worst-case depth of this algorithm occurs when a thread performs copy-on-steal and explores a maximal simple path. A thread explores such a path in O(e) time, because it visits at most *e* edges. As a result, Π and *Blk* contain at most *e* vertices, and *Blist* contains at most *e* pairs of vertices. Therefore, copy-on-steal requires O(e) time to copy Π , *Blk*, and *Blist*, and to unblock vertices in *Blk*. As a result, the depth of this algorithm is $T_{\infty}(n) \in O(e)$.

THEOREM 3. The fine-grained parallel Johnson algorithm is scalable when $\lim_{n\to\infty} c = \infty$.

PROOF. For this algorithm, $T_1(n) \in O(n + e + ec)$ and $T_{\infty}(n) \in O(e)$ (see Lemma 1). Given $e < n^2$ and our assumption that $\lim_{n\to\infty} c = \infty$, we have $\lim_{n\to\infty} \frac{T_{\infty}(n)}{T_1(n)} = \lim_{n\to\infty} \frac{e}{n+e+ec} = 0$. Thus, this algorithm is scalable based on Definition 2.

For the fine-grained parallel Johnson algorithm to be scalable, it is sufficient for c to increase sublinearly with n. Even though this algorithm is scalable, a strong or weak scalability is not guaranteed due to the work inefficiency of this algorithm. Nevertheless, our experiments show that this algorithm is strongly scalable in practice (see Figure 18).

5.4 Summary

Our relaxation of the strictly depth-first-search-based recursion-tree exploration reduces the pruning efficiency of the Johnson algorithm. In the worst case, the fine-grained parallel Johnson algorithm could perform as much work as the brute-force Tiernan algorithm does—i.e., O(se). However, in practice, this worst-case scenario does not happen (see Section 8). In addition, our fine-grained parallel Johnson algorithm can suffer from synchronisation issues in some rare cases (see Section 8), because our copy-on-steal mechanism can lead to long critical sections. In the next section, we introduce a fine-grained parallel algorithm that is scalable, work-efficient, and less prone to synchronisation issues.

6 FINE-GRAINED PARALLEL READ-TARJAN

In this section, we first introduce several optimisations that reduce the number of unnecessary vertex visits performed by the sequential Read-Tarjan algorithm. Then, we present our fine-grained parallel Read-Tarjan algorithm that includes these optimisations. Finally, we show that our parallel algorithm is work-efficient and strongly scalable.

6.1 Improvements to the Pruning Efficiency

To improve the pruning efficiency of the sequential Read-Tarjan algorithm, we include the following optimisations:

(i) Blocked vertex set forwarding enables a recursive call of the Read-Tarjan algorithm to reuse vertices blocked by its parent call, resulting in fewer vertex visits. The original Read-Tarjan algorithm discards blocked vertices after each recursive call [60], even though this information could be reused later. In this optimisation, the algorithm forwards the blocked vertices *Blk* of a



Fig. 7. (a) An example graph and (b) the recursion tree of our fine-grained parallel Read-Tarjan algorithm when enumerating cycles that start from v_0 . The nodes of the recursion tree represent the recursive calls of the depth-first search. Tasks shown in (b) can be executed independently of each other.

recursive call to its child recursive calls, preventing those child calls from unnecessarily visiting the vertices in *Blk* again. For example, in Figure 7, the vertex v_8 is blocked the first time the algorithm visits v_8 while exploring the path extension E_1 . This optimisation prevents the algorithm from visiting v_8 again when exploring the same extension E_1 or another extension E_3 that branches from E_1 . As a result of this optimisation, the algorithm can avoid the dotted part of the recursion tree.

(ii) Path extension forwarding prevents recomputation of the path extension *E* found by a parent recursive call by forwarding this path extension to its child recursive call. In this way, each child recursive call performs one fewer DFS invocation than the original Read-Tarjan algorithm [60].

(iii) Blocking on a successful DFS is another mechanism for discovering vertices to be blocked. As a reminder, the Read-Tarjan algorithm searches for path extensions using a DFS. In the original algorithm, a vertex is blocked only if it is visited during an unsuccessful DFS invocation, which fails to discover a path extension. However, successful DFS invocations could also visit some vertices that have all their neighbours blocked. Such vertices cannot lead to the discovery of new cycles and, thus, can also be blocked. The pseudocode of the DFS function that includes this optimisation is given in Algorithm 4. In our example given in Figure 7, a successful DFS invoked from v_3 finds a path extension E_3 and discovers that the only neighbour v_8 of v_7 is blocked. The algorithm then blocks v_7 , which enables it to avoid visiting v_7 again when exploring E_3 . Therefore, fewer vertices are visited during the execution of the algorithm.

6.2 Fine-grained Parallelisation

Although the optimisations presented in Section 6.1 eliminate some of the redundant work performed by the Read-Tarjan algorithm, this algorithm typically performs more work than the Johnson algorithm (see Section 3.4). However, this redundancy makes it possible to parallelise the Read-Tarjan algorithm in a scalable and work-efficient manner.

Because the Read-Tarjan algorithm allocates a new Blk set for each path extension exploration, a recursive call can explore different path extensions in an arbitrary order. In addition, discovery of a new path extension E results in the invocation of a single recursive call, and these calls can be executed in an arbitrary order. As a result, several threads can concurrently explore different paths of the same recursion tree constructed by the Read-Tarjan algorithm for a given starting

ALGORITHM 4: FGRT DES(11, vo. Blk. Vis)

I	nput : u - the current vertex, v ₀ - the starting vertex	
I	nOut:Blk - blocked vertices	
	Vis - vertices visited during the DFS	
C	Dutput : E - the resulting path extension from u to v_0	
1 if	$\mathbf{f} \mathbf{u} = \mathbf{v}_0$ then return \mathbf{u} ;	
2 V	V is = V is $\cup \{u\}$;	
3 b	lock = true;	
4 f	foreach $w : \mathcal{N}(u)$ s.t. w.id > v ₀ .id do	
5	if $w = v_0$ then	
6	return $u \rightarrow w$	
7	else if $w \notin Blk \land w \notin Vis$ then	
8	$E = FGRT_DFS(w, v_0 Blk, Vis);$	\blacktriangleright Recursively search for the path extension E
9	if $E \neq \emptyset$ then	
10	return E.push_front(u);	
11	if w ∉ Blk then	
12	block = false	
13 i f	f block then $Blk = Blk \cup \{u\};$	▹ Blocking on a successful DFS
14 r	eturn Ø;	

edge. There are neither data dependencies nor ordering requirements between different calls, apart from those that exist between a parent and a child. To exploit the parallelism available during the recursion tree exploration, we execute each path extension exploration in each recursive call as a separate task, all of which can be independently executed. Examples of such tasks are shown in Figure 7. We refer to the resulting algorithm as the *fine-grained parallel Read-Tarjan* algorithm.

Our implementation shown in Algorithm 5 performs only a single path extension exploration in a recursive call and uses all the optimisations we introduced in Section 6.1. We execute each such recursive call as a separate task using a dynamic thread scheduling framework (see Section 3.2). To find all cycles of a graph, we execute a *parallel for* loop iteration for each edge $v_0 \rightarrow v$ that uses Algorithm 4 to search for a path extension E from v to v_0 , as shown in Algorithm 6. If such E exists, then a task is created using v, v_0 , and E as its input parameters. This task then recursively creates new tasks, as shown in lines 14 and 19 of Algorithm 5, until all cycles that start with the edge $v_0 \rightarrow v$ have been discovered.

To prevent different threads from concurrently modifying Π and *Blk*, each task allocates and maintains its own Π and *Blk* sets. A task can receive a copy of Π and *Blk* directly from its parent task at the time of task creation. However, it is possible to minimise the copy overheads by copying these sets only when a task is stolen. For this purpose, we use the copy-on-steal with complete unblocking approach described in Section 5.2, which has shorter critical sections than the copyon-steal with recursive unblocking approach used by our fine-grained parallel Johnson algorithm.

6.3 Theoretical Analysis

We now show that the fine-grained parallel Read-Tarjan algorithm is both work-efficient and strongly scalable.

THEOREM 4. The fine-grained parallel Read-Tarjan algorithm is work-efficient.

PROOF. Because each task of our fine-grained parallel Read-Tarjan algorithm either discovers a cycle or creates at least two child tasks, our algorithm is executed using O(c) tasks. Each task performs several unsuccessful DFS invocations and one successful DFS per each child task it creates. All unsuccessful DFS invocations explore at most *e* edges in total, because they share the same set

ALGORITHM 5: FGRT_task(v, v ₀ , E, d, T ₁)	
Input : v - the current vertex, v_0 - the starting vertex E - the path extension from v to v_0 d - the depth of this task	
$InOut: T_1$ - the thread that created this task	$\blacktriangleright\ T_1$ maintains Π_{T_1} and Blk_{T_1}
$1 T_2$ = the thread executing this task;	\triangleright T ₂ maintains Π_{T_2} and Blk_{T_2}
² if $T_1 \neq T_2$ then	▷ Check if this task is stolen
$\{\Pi_{T_2}, Blk_{T_2}\} = copy(\{\Pi_{T_1}, Blk_{T_1}\});$	\blacktriangleright Operations on Π and Blk are thread-safe
4 while Π_{T_2} .back() \neq v do Π_{T_2} .pop();	
⁵ Remove vertices from Blk_{T_2} inserted at depth $d' \ge d$;	
6 found = false;	
7 while $E \neq \emptyset$ do	\blacktriangleright Exploration of the path extension E
<pre>8 v = E.pop_front();</pre>	
9 $\Pi_{T_2} = \Pi_{T_2}.push(v); Blk_{T_2} = Blk_{T_2} \cup \{v\};$	
10 foreach $u : \mathcal{N}(v)$ s.t. u.id > v_0 .id do	
11 if $u \neq E.front() \land u \notin Blk_{T_2}$ then	
12 $E' = FGRT_DFS(u, v_0, Blk_{T_2}, Vis = \emptyset);$	\triangleright Find an alternate path extension E'
13 if $E' \neq \emptyset$ then	
14 spawn FGRT_task(v, v_0 , E', d + 1, T ₂);	⊳ Create a child task
15 found = $true;$	
16 else $Blk_{T_2} = Blk_{T_2} \cup Vis;$	
17 if found then break ;	
18 if $E = \emptyset$ then report cycle Π_{T_2} ;	
19 else spawn FGRT_task(v, v_0 , E, d + 1, T ₂);	⊳ Create a child task

$\blacktriangleright\ T_0$ maintains Π_{T_0} and Blk_{T_0}
\blacktriangleright Operations on Π and Blk are thread-safe
⊳ Create a task

of blocked vertices. In the worst case, each edge is visited twice per task, once by a successful DFS and once by one of the unsuccessful DFS invocations. Thus, this algorithm performs O(e) work per task. Because this algorithm performs O(n + e) work if there are no cycles in the graph, the total amount of work this algorithm performs is $W_p(n) = O(n + e + ec)$. Hence, this algorithm is work-efficient based on Definition 1.

The work-efficiency of our fine-grained parallel Read-Tarjan algorithm can be demonstrated using the example given in Figure 5(a). In this example, the threads of this algorithm independently explore four different path extensions $E_i = v_1 \rightarrow u_i \rightarrow v_2 \rightarrow v_0$, with $i \in \{1 \dots 4\}$. A thread exploring a path extension E_i invokes a DFS from v_2 , which explores vertices b_1, \dots, b_m at most once and fails to find any other path extension. Therefore, the amount of work the fine-grained parallel Read-Tarjan algorithm performs does not increase compared to its single-threaded execution.

LEMMA 2. The depth $T_{\infty}(n)$ of the fine-grained parallel Read-Tarjan algorithm is in O(ne).

PROOF. In the worst case, a thread executing this algorithm creates a task for each vertex of its longest simple cycle, which has a length of at most *n*. Before invoking its first child task, a task executes a sequence of unsuccessful DFS invocations in O(e) and a successful DFS invocation also in O(e). Thus, the depth of this algorithm is O(ne).

The worst-case depth of our algorithm can be observed when this algorithm is executed on the graph given in Figure 4(a). This graph has $c = 2^{n-2}$ cycles and the length of its longest cycle $v_0 \rightarrow \ldots v_{n-1} \rightarrow v_0$ is *n*. The algorithm creates a task for each vertex of the cycle and performs a successful DFS in each such call, which leads to $T_{\infty} \in O(ne)$.

THEOREM 5. The fine-grained parallel Read-Tarjan algorithm is strongly scalable when $\lim_{n\to\infty} c/n = \infty$.

PROOF. Because the fine-grained parallel Read-Tarjan algorithm is work-efficient, we can apply Brent's rule [15]:

$$\frac{T_1(n)}{p} \le T_p(n) \le \frac{T_1(n)}{p} + T_{\infty}(n).$$
(2)

Substituting $T_1(n)$ with O(n + e + ec) and $T_{\infty}(n)$ with O(ne) (see Lemma 2), for a positive constant C_0 , it holds that

$$1 \left| \left(\frac{1}{p} + C_0 \frac{n}{c} \right) < 1 \right| \left(\frac{1}{p} + \frac{T_{\infty}(n)}{T_1(n)} \right) \le \frac{T_1(n)}{T_p(n)} \le p.$$
(3)

Given that $\lim_{n\to\infty} c/n = \infty$, there exist $n_0 > 0$, $C_1 > 0$ such that if $n > n_0$, then $c/n > C_1 p$. Thus, for every $n > n_0$, it holds that $kp \le \frac{T_1(n)}{T_p(n)} \le p$, where $k = C_1/(C_0+C_1) < 1$. As a result, $\frac{T_1(n)}{T_p(n)} = \Theta(p)$, which, based on Definition 3, completes the proof.

As shown in Table 4, our fine-grained parallel Read-Tarjan algorithm has a higher depth than our fine-grained parallel Johnson algorithm, introduced in Section 5. Nevertheless, the former algorithm is strongly scalable when *c* grows superlinearly with *n*, whereas strong scalability cannot be guaranteed for the latter algorithm.

6.4 Summary

The work of our fine-grained parallel Read-Tarjan algorithm does not increase after fine-grained parallelisation. This parallel algorithm performs $W_p(n) \in O(n + e + ec)$ work: the same as the work performed by its serial version. Our optimisations introduced in Section 6.1 do not reduce the work $W_p(n)$ performed by our parallel algorithm in the worst case. However, these optimisations significantly improve its performance in practice (see Section 8.4). In addition, the synchronisation overheads of the fine-grained parallel Read-Tarjan algorithm are not as significant as those of the fine-grained Johnson algorithm because of its shorter critical sections. Furthermore, this algorithm is the only asymptotically optimal parallel algorithm for cycle enumeration for which we are able to prove strong scalability.

7 PARALLELISING CONSTRAINED CYCLE SEARCH

This section describes the methods for adapting our parallel algorithms to search for simple cycles under various constraints. Because state-of-the-art algorithms for temporal and hop-constrained cycle enumeration are extensions of the Johnson algorithm [39, 54], our parallelisation approach described in Section 5 is also applicable to these algorithms. In this section, we describe the changes to the fine-grained parallel Johnson algorithm needed for enumeration of temporal and hop-constrained cycles. We also introduce modifications to the cycle enumeration algorithms required for finding time-window-constrained cycles.

7.1 Cycles in a Time Window

Cycle enumeration algorithms require minimal modifications to support time-window constraints. Such constraints restrict the search for simple, temporal, and hop-constrained cycles to those that occur within a time window of a given size δ , as illustrated in Figure 2. To find time-window-constrained cycles that start with an edge that has timestamp t_0 , only the edges with timestamps that belong to the time window $[t_0 : t_0 + \delta]$ are visited. To avoid reporting the same cycle several times, another edge with the same timestamp t_0 is visited only if the source vertex of that edge has an ID that is smaller than the ID of the vertex from which the search for cycles was started. Overall, imposing time-window constraints reduces the number of cycles discovered, which results in a more tractable problem.

Because temporal graphs may contain parallel edges, several simple cycles that contain the same sequence of vertices may exist in a temporal graph (see Figure 2(a)). As a result, such cycles may be discovered simultaneously, which could accelerate the search for cycles in temporal graphs. For that purpose, we use a method similar to Kumar and Calders [39], in which several simple cycles that contain the same sequence of vertices and whose edges belong to the same time window $[t_0 : t_0 + \delta]$ are grouped into a single cycle bundle and explored together. Once a cycle bundle is discovered, the cycles that belong to that bundle can be extracted and reported.

A **strongly connected component (SCC)** can be used to reduce the number of vertices visited during the search for time-window-constrained cycles. The search for cycles that start with the edge ε can be limited to use only the vertices from the SCC that contains ε [35]. In the case of time-window-constrained cycles, we compute an SCC for ε using only the edges with timestamps that belong to $[t_0 : t_0 + \delta]$, where t_0 is the timestamp of ε . Because an SCC can be computed independently for each edge in $O(\varepsilon)$ time [22], our fine-grained parallel algorithms remain scalable.

7.2 Temporal Cycles

To efficiently enumerate temporal cycles, the 2SCENT algorithm [39] replaces the set of blocked vertices *Blk* in the Johnson algorithm with *closing times*. The closing time *ct* of a vertex v indicates that the outgoing temporal edges of v with a timestamp greater than or equal to *ct* cannot participate in a temporal cycle and are therefore blocked. Increasing the closing time of v to a new value *ct'* unblocks the blocked outgoing edges of v that have timestamps smaller than *ct'*. This operation triggers the recursive unblocking procedure that unblocks the incoming edges of v with a timestamp smaller than the maximal timestamp among the unblocked outgoing edges of v. This process is repeated for every vertex with unblocked outgoing edges.

Because the backtracking phase of 2SCENT is based on the Johnson algorithm, it can be parallelised using our fine-grained approach described in Section 5. For this purpose, we use our copy-on-steal mechanism with recursive unblocking, introduced in Section 5.2, which enables a thread to maintain its own set of data structures used for recursion tree pruning. However, this mechanism is not directly applicable in this case, because the recursive unblocking procedure of 2SCENT requires the new closing time for a vertex as a parameter in addition to the vertex itself. For this reason, an additional data structure called *PrevLocks* is used alongside the current path Π that records the closing time that each vertex v had before it was added to Π . Copy-on-steal then performs the recursive unblocking procedure for each vertex v removed from Π using the original closing time of the vertex v obtained from *PrevLocks*, as shown in Algorithm 7. We refer to the resulting algorithm as the *fine-grained parallel temporal Johnson* algorithm.

The aforementioned modification to the copy-on-steal with recursive unblocking approach also enables a thread of our fine-grained parallel algorithm to reuse the edges blocked by another thread. This behaviour can be observed in the example shown in Figure 8, where the thread T_2 steals the task indicated in Figure 8(b) from the thread T_1 . Copy-on-steal executed by T_2 invokes recursive



Fig. 8. (a) An example graph and (b) the recursion tree of our fine-grained temporal Johnson algorithm when enumerating temporal cycles that start from v_0 . The thread T_2 can avoid the dotted part of the tree by reusing the blocked edges $v_6 \rightarrow v_7$ and $v_7 \rightarrow v_3$ discovered by T_1 .

ALGORITHM 7: CFGJ_copyOnSteal(d, T ₁ , T ₂)	
Input: d - the depth of the task executing this functi	on
InOut : T_1 - the victim thread	
T_2 - the stealing thread	
1 Mutex _{T1} .lock();	
² { Π_{T_2} , Blk _{T₂} , PrevLocks _{T₂} } = copy({ Π_{T_1} , Blk _{T₁} , PrevL	$ocks_{T_1}$ }; > Blk contains closing times or
barriers	
$ {Blist_{T_2}} = copy({Blist_{T_1}}); $	▹ Blist is not used for hop-constrained cycles
4 Mutex _{T1} .unlock();	
5 while $ \Pi_{T_2} > d$ do	▹ Copy-on-steal with recursive unblocking
$u = \Pi_{T_2}.pop();$	
$_{7}$ lock = PrevLocks _{T₂} .pop();	
$\label{eq:recursiveUnblock} 8 \qquad \mbox{RecursiveUnblock}(u, \mbox{lock}, \mbox{Blk}_{T_2}, \mbox{Blist}_{T_2});$	

unblocking that restores the closing time of v_3 to its original value of 9 obtained from *PrevLocks*. Note that this original closing time of v_3 was previously set by T_1 while exploring the path $v_0 \rightarrow v_1 \rightarrow v_3$. The recursive unblocking that T_2 invokes for v_3 unblocks only the edge $v_6 \rightarrow v_3$, because it is the only incoming edge of v_3 with a timestamp smaller than the closing time 9 of the vertex v_3 . Without recording the previous closing times, T_2 could instead unblock all incoming edges of v_3 by invoking recursive unblocking for v_3 with a closing time ∞ , which also unblocks the edges $v_6 \rightarrow v_7$ and $v_7 \rightarrow v_3$. However, because there is no temporal cycle that contains these two edges and starts with v_0 , T_2 would unnecessarily visit them in this case. Thus, restoring the closing time of v_3 to its original value 9 prevents T_2 from performing this redundant work.

We also adapt the Read-Tarjan algorithm and its fine-grained and coarse-grained versions to enumerate temporal cycles using closing times. The necessary changes to the algorithm are trivial, and we omit discussing them for brevity.

To reduce the number of vertices visited during the search for temporal cycles, we use a method similar to the SCC-based technique discussed in Section 7.1. Instead of computing an SCC for each edge ε , we compute a *cycle-union* that represents an intersection of temporal ancestors and



Fig. 9. (a) An example graph and (b) the recursion tree of our fine-grained hop-constrained Johnson algorithm when enumerating cycles of length L = 6 that start from v_0 . Barrier values of unmarked vertices are 0. Copy-on-steal enables the thread T_2 to reuse barriers discovered by the thread T_1 and to avoid exploring the dotted part of the tree.

temporal descendants of ε . The temporal descendants and the temporal ancestors of ε are the vertices that belong to the temporal paths in which ε is the first edge and the last edge, respectively. Defined as such, a cycle-union contains only the vertices that participate in temporal cycles that have ε as their starting edge. Thus, the search for temporal cycles that start with ε can be limited to only those vertices.

7.3 Hop-constrained Cycles

An efficient algorithm for enumerating hop-constrained cycles and paths, called BC-DFS [54], replaces the set of blocked vertices *Blk* in the Johnson algorithm with *barriers*. A barrier value *bar* of a vertex v indicates that the starting vertex v_0 of a cycle cannot be reached within *bar* hops from v. As a result, v is blocked if the length of the current path Π when the algorithm attempts to visit v is greater than or equal to L - bar, where L is the hop constraint. BC-DFS modifies the recursive unblocking of the Johnson algorithm to reduce the barrier *bar* of v to a specified value *bar'* < *bar*. This procedure also sets the barrier of any vertex u that can reach v in k hops to *bar'* + k if the previous barrier of u was greater than *bar'* + k. Maintaining barriers in such a way minimises redundant vertex visits when searching for hop-constrained cycles.

To parallelise BC-DFS in a fine-grained manner, we use the same technique as that used for fine-grained parallelisation of the Johnson algorithm (Section 5) and the 2SCENT algorithm (Section 7.2). In this case, threads exploring a recursion tree of BC-DFS maintain separate data structures, such as the current path Π and barrier values for each vertex, and use the copy-on-steal with the recursive unblocking approach to copy these data structures among threads. Similarly to our algorithm from Section 7.2, each thread also maintains a data structure *PrevLocks* that records the original barrier value of each vertex v from Π . When a thread steals a task, it performs a recursive unblocking procedure for each vertex v removed from Π using its original barrier value obtained from *PrevLocks*, as shown in Algorithm 7. This procedure reduces the barrier value of the vertices that can reach v, enabling the stealing thread to visit those vertices. We refer to the resulting algorithm as the *fine-grained parallel hop-constrained Johnson* algorithm.

The modified copy-on-steal with recursive unblocking approach given in Algorithm 7 enables a stealing thread of the aforementioned fine-grained parallel algorithm to reuse barriers discovered by other threads. This behaviour can be observed in the example given in Figure 9. In that example,

platform	Intel KNL [64]	Intel Xeon Skylake [26]
$P \times C/P \times T/C$	$4 \times 64 \times 4$	$5 \times 48 \times 2$
Total no. threads	1,024	480
Frequency	1.3 GHz	2 GHz
Memory per proc.	110 GB	360 GB
L1d/L2/L3 cache	32 KB/512 KB/none	32 KB/1 MB/38.5 MB

Table 5. Hardware Platforms Used in the Experiments

Here, P, C/P, and T/C represent the number of processors, the number of cores per processor, and the number of hardware threads per core, respectively.

the thread T_1 first visits the vertices v_2 , v_6 , v_7 , v_8 and sets the barrier value of each visited vertex to $L - |\Pi| + 1$ (values in red shown in Figure 9(a)), because it was not able to find a cycle of length L = 6 [54]. Here, $|\Pi|$ denotes the length of Π at the moment of exploration of each vertex. When the thread T_2 steals the task indicated in Figure 9(b) from T_1 , the copy-on-steal mechanism executed by T_2 performs a recursive unblocking of the vertex v_1 using the original barrier value 0 of v_1 obtained from *PrevLocks*. This recursive unblocking reduces the barrier value of v_2 from 4 to 1, which enables T_2 to find the cycle that contains v_2 . The barrier values of the vertices v_6 , v_7 , and v_8 are not modified, and, thus, the thread T_2 avoids visiting these vertices unnecessarily.

7.4 Summary

In this section, we described a method to adapt the cycle enumeration algorithms, such as our fine-grained algorithms introduced in Sections 5 and 6, to search for cycles under time window constraints. In addition, we introduced a modified version of our copy-on-steal with recursive unblocking approach, introduced in Section 5, that supports fine-grained parallelisation of temporal and hop-constrained cycle enumeration algorithms [39, 54] derived from the Johnson algorithm. As a result, our fine-grained parallel algorithms can enumerate cycles under time-window, temporal, and hop constraints.

8 EXPERIMENTAL EVALUATION

This section evaluates the performance of our fine-grained parallel algorithms for simple, temporal, and hop-constrained cycle enumeration.¹ As Table 2 shows, we are the only ones to offer fine-grained parallel versions of the asymptotically optimal cycle enumeration algorithms, such as the Johnson and the Read-Tarjan algorithms. However, the methods covered in Table 2 can be parallelised using the coarse-grained approach covered in Section 4. Thus, we use the coarse-grained approach as our main comparison point.

The experiments are performed using two different clusters: Intel² KNL [64] and Intel Xeon Skylake [26]. The details of these two clusters are given in Table 5. We developed our code on the Intel KNL cluster and ran most of the analyses there; yet, for completeness, we also provide the comparisons to competing implementations on the Intel Xeon Skylake cluster available in Google Cloud's Compute Engine [26]. Scalability experiments are conducted on the Intel KNL cluster. In these experiments, the data points that use 64 threads or less are executed on a single Intel KNL processor; two processors are used to execute the data points that use 128 threads; and all four processors are used otherwise. Furthermore, we use more than one thread per core only if the number of threads used is greater than 256.

¹The open-source implementations of our algorithms are maintained here: https://github.com/IBM/parallel-cycle-enumeration

 $^{^{2}}$ Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

graph	abbr.	n	е	Δ_{avg}	Δ_{max}	Pavg	P_{max}	Time span [days]
bitcoinalpha	BA	3.3k	24k	7.4	490	1.0	1.0	1,901
bitcoinotc	BO	4.8k	36k	7.4	763	1.0	1.0	1,903
CollegeMsg	CO	1.3k	60k	44.3	1.1k	2.9	98	193
email-Eu-core	EM	824	332k	403.3	9.8k	13.3	2.8k	803
mathoverflow	МО	16k	390k	23.7	4.5k	1.7	225	2,350
transactions	TR	83k	530k	6.4	1.7k	1.5	290	1,803
higgs-activity	HG	278k	555k	2.0	655	1.2	95	6.0
askubuntu	AU	102k	727k	7.1	7.6k	1.3	154	2,613
superuser	SU	138k	1.1M	8.1	26k	1.3	78	2,773
wiki-talk	WT	140k	6.1M	43.7	233k	1.9	1.1k	2,277
friends2008	FR	481k	12M	25.5	9.1k	1.0	6.0	1,826
wiki-dynamic-nl	NL	1.0M	20M	19.5	30k	1.5	352	3,602
messages	MS	313k	26M	83.4	48k	4.3	10k	1,880
AML-Data	AML	10M	34M	3.4	26k	4.6	96	30
stackoverflow	SO	2.0M	48M	24.2	72k	1.3	594	2,774
		1						

Table 6. Temporal Graphs Used in the Experiments

In this table, Δ_{avg} and Δ_{max} refer to the average and maximum values of Δ , respectively, where Δ is the number of outgoing edges of a vertex, i.e., vertex degree. Similarly, P_{avg} and P_{max} refer to the average and maximum values of *P*, respectively, where *P* represents the number of parallel edges for a given source and destination vertex. Time span refers to the difference between the maximum and minimum timestamps in a graph.

We use the *Threading Building Blocks* (TBB) [38] library to parallelise the algorithms on a single processor. We distribute the execution of the algorithms across multiple processors using the **Message Passing Interface** (MPI) [17]. When using distributed execution, each processor stores a copy of the input graph in its main memory and searches for cycles starting from a different set of graph edges. The starting edges are divided among the processors such that when the edges are ordered in the ascending order of their timestamps, k consecutive edges in that order are assigned to k different processors. Each processor then uses its own dynamic scheduler to balance the workload across its hardware threads. In this setup, workload imbalance across processors may still occur, but its impact is limited in our experiments, because we use at most five processors.

We perform the experiments using the graphs listed in Table 6. The TR, FR, and MS graphs are from *Harvard Dataverse* [33], the NL graph is from *Konect* [40], the AML graph is from the *AML-Data* repository [5], and the rest are from *SNAP* [42]. Except for BA and BO, all of the graphs have parallel edges, as shown in Table 6. To make cycle enumeration problems tractable, we use time-window constraints in all of our experiments. The time window sizes used in our experiments are given in the figures next to the graph names. We stop the execution of an algorithm if it takes more than 24 h on the Intel KNL cluster or more than 6 h on the Intel Xeon Skylake cluster.

8.1 Temporal Cycle Enumeration

The goal of a temporal cycle enumeration problem is to find all simple cycles with edges ordered in time. Here, we evaluate the performance of our fine-grained parallel algorithms for this problem introduced in Section 7.2. Our main comparison points are the coarse-grained parallel versions of the temporal Johnson and temporal Read-Tarjan algorithms. We refer to the backtracking phase of the state-of-the-art 2SCENT algorithm [39] for temporal cycle enumeration as the temporal Johnson algorithm and parallelise it in a coarse-grained manner for the experiments. We do not parallelise the entire 2SCENT algorithm, because the preprocessing phase of 2SCENT is strictly sequential and has a time complexity in the order of the complexity of its backtracking phase. We also provide direct comparisons with the 2SCENT algorithm.



Fig. 10. Performance of parallel algorithms for temporal cycle enumeration on (a) the Intel KNL cluster using 1,024 threads and (b) the Intel Xeon Skylake cluster using 480 threads. The values above the bars show the execution time of each algorithm relative to that of our fine-grained parallel temporal Johnson for the same benchmark. The values that contain the symbol > represent the experiments that did not finish within the given time limit.



Fig. 11. Longer time windows increase the performance gap between the algorithms. The algorithms are executed on the Intel KNL cluster using 1,024 threads. The numbers above the bars show the execution times of the coarse-grained algorithm relative to that of the fine-grained algorithm.



Fig. 12. (a, b) Frequency distribution of temporal cycles for different cycle lengths and (c) the total number of temporal cycles discovered during the experiments shown in Figure 10. The number of temporal cycles discovered is several orders of magnitude greater than the number of vertices or edges of a graph.

Figure 10 shows that our fine-grained parallel algorithms achieve an order of magnitude speedup compared to the coarse-grained algorithms on the Intel KNL cluster. For the NL graph, this speedup reaches up to 40×. Because the Intel Xeon Skylake cluster contains fewer physical cores than the Intel KNL cluster, the speedup between our fine-grained and the coarse-grained parallel Johnson algorithms is smaller on the former cluster. As can be observed in Figure 11, this speedup increases as we increase the time window size used in the algorithms. Note that enumerating cycles in longer time windows is more challenging, because longer time windows contain a larger number of cycles.

Figure 12 shows the number of temporal cycles enumerated in the experiments shown in Figure 10 and their frequency distribution for the given cycle length. The execution time of the cycle enumeration algorithms typically depends on the number of cycles discovered. However, due to the existence of parallel edges, many cycles may consist of the same sequence of vertices and can be explored simultaneously by grouping such cycles into a cycle bundle [39]. For example,



Fig. 13. Scalability evaluation of parallel temporal cycle enumeration algorithms executed on the Intel KNL cluster. The baseline is our fine-grained parallel temporal Johnson algorithm. The relative performance of 2SCENT [39] is shown when it completes in 24 h. Note that the 2SCENT implementation is single-threaded and the single-threaded execution results are not available for all graphs.

in the cases of the CO, TR, and MS graphs, a cycle bundle discovered by our algorithms contains more than 10 M cycles on average. For this reason, despite discovering several orders of magnitude more cycles in the CO, TR, and MS graphs than in the other graphs, the execution time of our fine-grained algorithms on the CO, TR, and MS graphs is comparable to their execution time on the other graphs. In addition, in the cases of BA, BO, FR, and NL, where one temporal cycle per cycle bundle is discovered, our fine-grained algorithms are more time-consuming on the NL graph, because more cycles are discovered in the NL graph than in the BA, BO, and FR graphs. Thus, the execution time of our fine-grained algorithms depends more on the number of cycle bundles explored than the number of cycles.

The scalability evaluation of the parallel temporal cycle enumeration algorithms is given in Figure 13. We also report the performance of the sequential 2SCENT algorithm in the same figure. The performance of our fine-grained parallel algorithms improves linearly until 256 threads, after which it becomes sublinear due to simultaneous multithreading. As a result, our fine-grained versions of the Johnson and the Read-Tarjan algorithms reach 435× and 470× speedups, respectively, compared to their serial versions. Additionally, when using 1,024 threads, our fine-grained Johnson algorithm is on average 260× faster than 2SCENT when 2SCENT completes in 24 h. However,

the coarse-grained Johnson algorithm does not scale as well as the fine-grained algorithms. As a result, the performance gap between the fine-grained and the coarse-grained algorithms increases as we increase the number of threads.

Overall, the fastest algorithm for temporal cycle enumeration that we tested is our fine-grained Johnson algorithm, which is, on average, 60% faster than our fine-grained Read-Tarjan algorithm. When using 1,024 threads, both fine-grained algorithms are an order of magnitude faster than their coarse-grained counterparts. Moreover, our fine-grained parallel algorithms, executed on the Intel KNL cluster using 1,024 threads, are two orders of magnitude faster than the state-of-the-art algorithm 2SCENT [39].

8.2 Hop-constrained Cycle Enumeration

In hop-constrained cycle enumeration, we search for all simple cycles in a graph that are shorter than the specified hop constraint. Here, we compare our fine-grained parallel hop-constrained Johnson algorithm, introduced in Section 7.3, with the state-of-the-art algorithms BC-DFS and JOIN [54] for this problem. For this evaluation, we parallelised BC-DFS and JOIN in the coarse-grained manner. Because adapting the Read-Tarjan algorithm to enumerate hop-constrained cycles is not trivial, we do not report the performance of the fine-grained and coarse-grained versions of this algorithm. We also omit the performance results for the MS graph, because our fine-grained algorithm did not finish under 12 h when using the smallest time window size.

Figure 14 shows that our fine-grained parallel algorithm is, on average, more than $10 \times$ faster than the coarse-grained parallel BC-DFS algorithm for the two largest hop constraints tested. When using the hop-constraint that is less than or equal to ten, the coarse-grained parallelisation approach is able to achieve workload balance across cores, and thus the performance of this approach is similar to that of our fine-grained approach in this case. As we increase the hop constraint, the probability of encountering deeper recursion trees also increases. Exploring such trees using the coarse-grained approach leads to workload imbalance (see Section 4). Our fine-grained algorithm is designed to resolve this problem by exploring a recursion tree using several threads. Therefore, increasing the hop constraint increases the speedup of our fine-grained algorithm with respect to the coarse-grained algorithm.

When the hop constraint is set to 20, our fine-grained parallel algorithm is, on average, $10\times$ faster than the coarse-grained parallel JOIN algorithm, as shown in Figure 14. Although the latter algorithm can be competitive with our fine-grained algorithm, it can also suffer from long execution times, such as in the cases of the AU, NL, and AML graphs. The reason for these long execution times is the fact that the JOIN algorithm might temporarily construct many non-simple cycles while searching for simple cycles. Because this algorithm constructs cycles by combining simple paths, it is not guaranteed that each combination results in a simple cycle. The overhead of combining paths can dominate the execution time of JOIN if this algorithm constructs orders of magnitude more non-simple cycles than simple cycles. For instance, this situation occurs in the case of AU and hop constraint of 20, where JOIN discovers 600× more non-simple cycles than simple cycles. As a result, the speedup of our fine-grained algorithm compared to the coarse-grained JOIN algorithm can reach up to three orders of magnitude.

According to Figure 15, the number of cycles increases exponentially as the hop constraint is increased. Thus, increasing the hop constraint could lead to an exponential increase in the execution time of our fine-grained parallel algorithm for hop-constrained cycle enumeration, which can be observed in Figure 14. Note that Figures 12 and 19 indicate that the frequency distributions of the cycles have a bell shape. As a result, the increase in the number of cycles with increasing hop constraints shown in Figure 15 may not be exponential when the hop constraint is increased beyond 20.



Fig. 14. Performance of parallel algorithms for hop-constrained simple cycle enumeration on (a) the Intel KNL cluster using 1,024 threads and (b) the Intel Xeon Skylake cluster using 480 threads. The values above the bars show the execution time of the coarse-grained parallel algorithms relative to that of our fine-grained parallel algorithm. The values that contain the symbol > represent the experiments that did not finish within the given time limit. Larger hop constraints increase the performance gap between the two algorithms.



Fig. 15. (a, b) Frequency distribution of hop-constrained cycles for different cycle lengths and (c) the total number of hop-constrained cycles discovered during the experiments shown in Figure 14 using the hop-constraint of 20. In most cases, the number of cycles increases exponentially with hop-constraint.

Figure 16 shows that the speedup of our fine-grained parallel Johnson algorithm with respect to the coarse-grained parallel BC-DFS can be increased by using more threads. The performance of our fine-grained parallel algorithm scales linearly with the number of threads, whereas the scaling of the coarse-grained parallel BC-DFS eventually slows down. Thus, in addition to being, on average, an order of magnitude faster than the coarse-grained parallel BC-DFS, our fine-grained algorithm is also more scalable.

8.3 Simple Cycle Enumeration

Here, we evaluate our fine-grained parallel algorithms for simple cycle enumeration. The computational complexity of simple cycle enumeration is higher than the complexity of temporal and hop-constrained cycle enumeration, because simple cycle enumeration does not impose temporal ordering or hop constraints. The only constraint we impose is the time-window constraint. Because the complexity of enumerating simple cycles is higher, we use smaller time windows compared to the cases of temporal and hop-constrained cycle enumeration. We use the coarse-grained parallel versions of the Johnson and the Read-Tarjan algorithms as our main comparison points. We do not report the results for the MS graph, because our algorithms did not finish in 12 h even if we set the time window to one second. We also provide a comparison with the Tiernan algorithm [68] parallelised in a fine-grained manner, which is a more efficient version of the previous algorithm by Qing et al. [57] (see Table 2 and Section 2). We parallelise the Tiernan algorithm in a fine-grained manner by wrapping each recursive call of this algorithm into a task and by using a dynamic task scheduler to balance the workload across the threads. Note that the algorithm by



Fig. 16. Scalability evaluation of parallel hop-constrained cycle enumeration algorithms executed on the Intel KNL cluster using the hop constraint of 15. The speedup values are relative to the single-threaded execution of BC-DFS. Evaluation on other graphs is omitted for brevity.



Fig. 17. Performance of parallel algorithms for simple cycle enumeration on (a) the Intel KNL cluster using 1,024 threads and (b) the Intel Xeon Skylake cluster using 480 threads. The values above the bars show the execution time of each algorithm relative to that of our fine-grained parallel Johnson algorithm for the same benchmark. The values that contain the symbol > represent the experiments that did not finish within the given time limit.

Qing et al. [57] uses a static load balancing mechanism, which makes it less efficient than our fine-grained parallelisation of the Tiernan algorithm.

As we can see in Figure 17, our fine-grained parallel algorithms show an order of magnitude average speedup compared to coarse-grained parallel algorithms on two different platforms. The reason for this speedup is better scalability of our fine-grained algorithms, which we demonstrate in Figure 18. Similarly to the cases of temporal and hop-constrained cycle enumeration (see Figures 13 and 16), our fine-grained parallel algorithms scale linearly with the number of physical cores used whereas the coarse-grained parallel Johnson algorithm does not scale as well. Thus, the speedup between the fine-grained and the coarse-grained algorithms increases by utilising more threads.

Figure 19 shows the number of simple cycles enumerated in the experiments shown in Figure 17 and their frequency distribution for the given cycle length. Similarly to temporal cycle enumeration, the execution time mainly depends on the number of cycle bundles explored rather than on the number of cycles enumerated. For example, each cycle bundle explored in the BA, BO, SU, FR, and NL graphs contains only two or fewer simple cycles on average, and the execution time of our fine-grained Johnson algorithm is the longest for the NL graph, which also has the most reported



Fig. 18. Scalability evaluation of parallel simple cycle enumeration algorithms executed on the Intel KNL cluster. The speedup values are relative to the single-threaded execution of the Johnson algorithm. Evaluation on other graphs is omitted for brevity.



Fig. 19. (a, b) Frequency distribution of simple cycles for different cycle lengths and (c) the total number of simple cycles discovered during the experiments shown in Figure 17. Simple cycles tend to be longer than temporal cycles, despite using a smaller time window for simple cycle enumeration (see Figure 12).

cycles (see Figure 19). Furthermore, despite BA and BO having similar sizes, the execution time of our fine-grained Johnson algorithm is an order of magnitude longer for BO than for BA. The reason for this difference is that more cycles were discovered in the BO graph than in the BA graph for the time window sizes given in Figure 17. As a result, the execution time of the simple cycle enumeration can significantly vary, even for graphs of similar sizes.

Figure 20 presents the comparison of our fine-grained parallel Johnson and Read-Tarjan algorithms with the Tiernan algorithm [68] parallelised in a fine-grained manner. Our fine-grained parallel Johnson algorithm is up to $7\times$ faster than the fine-grained parallel Tiernan algorithm. The main reason for this performance gap is that the Tiernan algorithm performs more redundant work than the Johnson algorithm (see Section 3.4). The fine-grained parallel Read-Tarjan algorithm is slower in the case of NL than the fine-grained parallel Tiernan algorithm, because the redundant work performed by the algorithms is limited, and the Tiernan algorithm performs less work per visited edge than the Read-Tarjan algorithm. However, the fine-grained parallel Read-Tarjan algorithm for other benchmarks. Therefore, our fine-grained parallel Johnson and Read-Tarjan algorithms are preferable to the parallel formulation of the Tiernan algorithm, such as the algorithm by Qing et al. [57].

The synchronisation overheads caused by recursive unblocking of our fine-grained parallel Johnson algorithm (see Section 5.2) are visible only in the case of AML. In this case, the fine-grained



Fig. 20. Comparison of our fine-grained parallel algorithms for simple cycle enumeration with the finegrained parallel Tiernan algorithm [68] on the Intel KNL cluster using 1,024 threads. The values above the bars show the execution time of each algorithm relative to that of our fine-grained parallel Johnson algorithm for the same benchmark. Our fine-grained parallel Johnson algorithm is up to 7× faster than the fine-grained parallel Tiernan algorithm.



Fig. 21. Effect of the pruning improvements to our fine-grained parallel Read-Tarjan algorithm for (a) simple and (b) temporal cycle enumeration. Execution times are normalised to the case that includes all optimisations. Our optimisations accelerate this algorithm by up to $6.8\times$.

parallel Johnson algorithm performs 60% fewer edge visits than the fine-grained parallel Read-Tarjan; however, it is 25% slower. These synchronisation overheads can be explained by a very low cycle-to-vertex ratio. Because a vertex is blocked if it cannot take part in a cycle, the probability of a vertex being blocked is higher when the cycle-to-vertex ratio is lower (see Table 6 and Figure 19). In consequence, more vertices are unblocked during the recursive unblocking of the fine-grained parallel Johnson algorithm, leading to longer critical sections and more contention on the locks. Nevertheless, our fine-grained parallel Johnson algorithm achieves a good trade-off between pruning efficiency and lock contention in most cases.

Overall, our fine-grained parallel Johnson and fine-grained parallel Read-Tarjan algorithms have comparable performance, as shown in Figure 17. Although the former algorithm is slightly faster, it can suffer from synchronisation overheads in some cases. Nevertheless, both parallel algorithms achieve linear scaling with the number of physical cores used and achieve, on average, more than $10 \times$ speedup with respect to coarse-grained parallel versions of the algorithms. These conclusions also hold in the cases of temporal and hop-constrained cycle enumeration.

8.4 Improvements to the Read-Tarjan Algorithm

Figure 21 shows the effect of our pruning improvements, introduced in Section 6.1, on the performance of our fine-grained Read-Tarjan algorithm. The experiments are performed using a single Intel KNL processor using 256 threads. Note that using one processor instead of the entire cluster results in longer execution times, but it enables us to eliminate the effect of workload imbalance across processors in this experiment. The execution time of the fine-grained parallel Read-Tarjan algorithm decreases after activating each optimisation, because fewer redundant vertex and edge visits are performed during the execution of this algorithm. When all optimisations are enabled, the average speedup of our algorithm for simple cycle enumeration compared to its unoptimised version is 2×. In the case of temporal cycle enumeration, the average speedup increases to 3.4×. As a result, our pruning improvements enable the fine-grained parallel Read-Tarjan algorithm to be competitive with the fine-grained parallel Johnson algorithm.

9 CONCLUSIONS

This work has made three contributions to the area of parallel cycle enumeration. First, we have introduced scalable fine-grained parallel versions of the state-of-the-art Johnson and Read-Tarjan algorithms for enumerating simple cycles. In particular, we have shown that the novel fine-grained parallel approach we contributed for parallelising the Johnson algorithm can be adapted to support the enumeration of temporal and hop-constrained cycles as well. Our fine-grained parallel algorithms for enumerating the aforementioned types of cycles achieve a near-linear performance scaling on a compute cluster with a total number of 256 CPU cores that can execute 1,024 simultaneous software threads.

Second, we have shown that our fine-grained parallel cycle enumeration algorithms are scalable both in theory and in practice. In contrast, their coarse-grained parallel versions do not share this property. When using 1,024 software threads, our fine-grained parallel algorithms are on average an order of magnitude faster than their coarse-grained counterparts. In addition, the performance gap between the fine-grained and coarse-grained parallel algorithms widens as we use more physical CPU cores. This performance gap also widens when increasing the time window in the case of temporal cycle enumeration and when increasing the hop constraint in the case of hop-constrained cycle enumeration.

Third, we have shown that, whereas our fine-grained parallel Read-Tarjan algorithm is work efficient, our fine-grained parallel Johnson algorithm is not. In general, the former is competitive against the latter because of the new pruning methods we introduced, yet the latter outperforms the former in most experiments. In some rare cases, our fine-grained parallel Johnson algorithm can suffer from synchronisation overheads. In such cases, our fine-grained parallel Read-Tarjan algorithm offers a more scalable alternative.

REFERENCES

- [1] Ehab Abdelhamid, Ibrahim Abdelaziz, Panos Kalnis, Zuhair Khayyat, and Fuad Jamour. 2016. ScaleMine: Scalable parallel frequent subgraph mining in a single large graph. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'16). IEEE, 716–727. DOI: 10.1109/SC.2016.60
- [2] Udit Agarwal and Vijaya Ramachandran. 2016. Finding k simple shortest paths and cycles. In Proceedings of the 27th International Symposium on Algorithms and Computation (ISAAC'16) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 64), Seok-Hee Hong (Ed.). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 8:1–8:12. https://doi.org/10.4230/LIPIcs.ISAAC.2016.8
- [3] Charu C. Aggarwal and Haixun Wang (Eds.). 2010. Managing and Mining Graph Data. Advances in Database Systems, Vol. 40. Springer US, Boston, MA. DOI: 10.1007/978-1-4419-6045-0
- [4] R. E. L. Aldred and Carsten Thomassen. 2008. On the maximum number of cycles in a planar graph. J. Graph Theory 57, 3 (Mar. 2008), 255–264. DOI: 10.1002/jgt.20290
- [5] Erik Altman. 2021. AML-Data. Retrieved from https://github.com/IBM/AML-Data. Accessed: 2022-05-30.
- [6] David A. Bader. 1999. A Practical Parallel Algorithm for Cycle Detection in Partitioned Digraphs. Retrieved from https://digitalrepository.unm.edu/ece_rpts/45
- [7] V. K. Balakrishnan. 1997. Graph Theory. McGraw-Hill Professional, New York, NY.
- [8] Albert-László Barabási and Márton Pósfai. 2016. Network Science. Cambridge University Press, Cambridge, UK, Chapter The scale-free property, 1–57.
- [9] Etienne Birmelé, Rui Ferreira, Roberto Grossi, Andrea Marino, Nadia Pisanti, Romeo Rizzi, and Gustavo Sacomoto. 2013. Optimal listing of cycles and st-paths in undirected graphs. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, Philadelphia, PA, 1884–1896. DOI: 10.1137/1.9781611973105.134
- [10] Jovan Blanuša, Paolo Ienne, and Kubilay Atasu. 2022. Scalable fine-grained parallel cycle enumeration algorithms. In Proceedings of the ACM Symposium on Parallelism in Algorithms and Architectures (SPAA'22). ACM, Philadelphia, PA, 247–258. https://doi.org/10.1145/3490148.3538585

- [11] Jovan Blanuša, Radu Stoica, Paolo Ienne, and Kubilay Atasu. 2020. Manycore clique enumeration with fast set intersections. Proc. Very Large Data Base 13, 12 (Aug. 2020), 2676–2690. DOI: 10.14778/3407790.3407853
- [12] Guy E. Blelloch and Bruce M. Maggs. 2010. Parallel Algorithms. CRC Press, London, England, Chapter 25, 25.1–25.40.
- [13] Robert D. Blumofe, Christopher F. Joerg, Bradley C. Kuszmaul, Charles E. Leiserson, Keith H. Randall, and Yuli Zhou. 1996. Cilk: An efficient multithreaded runtime system. J. Parallel and Distrib. Comput. 37, 1 (Aug. 1996), 55– 69. DOI:10.1006/jpdc.1996.0107
- [14] Robert D. Blumofe and Charles E. Leiserson. 1999. Scheduling multithreaded computations by work stealing. J. ACM 46, 5 (Sept. 1999), 720–748. DOI: 10.1145/324133.324234
- [15] Richard P. Brent. 1974. The parallel evaluation of general arithmetic expressions. J. ACM 21, 2 (Apr. 1974), 201–206. DOI: 10.1145/321812.321815
- [16] Anna D. Broido and Aaron Clauset. 2019. Scale-free networks are rare. Nat. Commun. 10, 1 (Dec. 2019), 1017. DOI:10.1038/s41467-019-08746-5
- [17] CORPORATE The MPI Forum. 1993. MPI: A message passing interface. In Proceedings of the ACM/IEEE Conference on Supercomputing (Supercomputing'93). ACM Press, 878–883. DOI: 10.1145/169627.169855
- [18] Huanqing Cui, Jian Niu, Chuanai Zhou, and Minglei Shu. 2017. A multi-threading algorithm to detect and remove cycles in vertex- and arc-weighted digraph. *Algorithms* 10, 4 (Oct. 2017), 115. DOI:10.3390/a10040115
- [19] G. Danielson. 1968. On finding the simple paths and circuits in a graph. IEEE Trans. Circ. Theory 15, 3 (Sept. 1968), 294–295. DOI: 10.1109/TCT.1968.1082837
- [20] Apurba Das, Seyed-Vahid Sanei-Mehri, and Srikanta Tirthapura. 2020. Shared-memory parallel maximal clique enumeration from static and dynamic graphs. ACM Trans. Parallel Comput. 7, 1 (Apr. 2020), 1–28. https://doi.org/10.1145/3380936
- [21] P. Erdős and T. Gallai. 1959. On maximal paths and circuits of graphs. Acta Mathematica Academiae Scientiarum Hungaricae 10, 3-4 (Sept. 1959), 337–356. DOI: 10.1007/BF02024498
- [22] Lisa K. Fleischer, Bruce Hendrickson, and Ali Pinar. 2000. On identifying strongly connected components in parallel. In Parallel and Distributed Processing. Vol. 1800. Springer, Berlin, 505–511. DOI:10.1007/3-540-45591-4_68
- [23] Pierre Fraigniaud and Dennis Olivetti. 2019. Distributed detection of cycles. ACM Trans. Parallel Comput. 6, 3 (Dec. 2019), 1–20. https://doi.org/10.1145/3322811
- [24] Norman E. Gibbs. 1969. A cycle generation algorithm for finite undirected linear graphs. J. ACM 16, 4 (Oct. 1969), 564–568. DOI: 10.1145/321541.321545
- [25] Pierre-Louis Giscard, Paul Rochet, and Richard C. Wilson. 2017. Evaluating balance on social networks from their simple cycles. J. Complex Netw. 5 (May 2017), 750–775. DOI: 10.1093/comnet/cnx005
- [26] Google Cloud. 2022. General-purpose machine family: N1 machine series. Retrieved from https://cloud.google.com/ compute/docs/general-purpose-machines. Accessed: 2022-11-14.
- [27] Roberto Grossi. 2016. Enumeration of paths, cycles, and spanning trees. In *Encyclopedia of Algorithms*. Springer, New York, NY, 640–645. DOI: 10.1007/978-1-4939-2864-4_728
- [28] A. Gupta and C. Selvidge. 2005. Acyclic modeling of combinational loops. In Proceedings of the IEEE/ACM International Conference on Computer-Aided Design (ICCAD'05). IEEE, San Jose, CA, 343–348. DOI: 10.1109/ICCAD.2005.1560091
- [29] Anshul Gupta and Toyotaro Suzumura. 2021. Finding All Bounded-Length Simple Cycles in a Directed Graph. Retrieved from https://arXiv:2105.10094
- [30] László Hajdu and Miklós Krész. 2020. Temporal network analytics for fraud detection in the banking sector. In Proceedings of the ADBIS, TPDL, and EDA Common Workshops and Doctoral Consortium. Springer, Cham, Switzerland, 145–157. DOI: 10.1007/978-3-030-55814-7_12
- [31] M. d. Nazrul Islam, S. M. Rafizul Haque, Kaji Masudul Alam, and M. d. Tarikuzzaman. 2009. An approach to improve collusion set detection using MCL algorithm. In *Proceedings of the 12th International Conference on Computers and Information Technology*. IEEE, Dhaka, Bangladesh, 237–242. DOI: 10.1109/ICCIT.2009.5407133
- [32] Joseph JaJa. 1992. Introduction to Parallel Algorithms. Addison Wesley, Boston, MA.
- [33] Jaroslaw Jankowski, Radosłlaw Michalski, and Piotr Bródka. 2017. Spreading processes in multilayer complex network within virtual world. https://www.nature.com/articles/sdata2017144
- [34] Zhi-Qiang Jiang, Wen-Jie Xie, Xiong Xiong, Wei Zhang, Yong-Jie Zhang, and Wei-Xing Zhou. 2013. Trading networks, abnormal motifs and stock manipulation. *Quant. Finance Lett.* 1, 1 (Dec. 2013), 1–8. DOI: 10.1080/21649502.2013.802877
- [35] Donald B. Johnson. 1975. Finding all the elementary circuits of a directed graph. SIAM J. Comput. 4, 1 (Mar. 1975), 77-84. DOI: 10.1137/0204007
- [36] T. Kamae. 1967. A systematic method of finding all directed circuits and enumerating all directed paths. IEEE Trans. Circ. Theory 14, 2 (June 1967), 166–171. DOI: 10.1109/TCT.1967.1082699
- [37] Steffen Klamt and Axel von Kamp. 2009. Computing paths and cycles in biological interaction graphs. BMC Bioinform. 10, 1 (Dec. 2009), 181. DOI: 10.1186/1471-2105-10-181

- [38] Alexey Kukanov. 2007. The foundations for scalable multicore software in intel threading building blocks. ITJ 11, 04 (Nov. 2007), 309–322. DOI: 10.1535/itj.1104.05
- [39] Rohit Kumar and Toon Calders. 2018. 2SCENT: An efficient algorithm for enumerating all simple temporal cycles. Proc. Very Large Data Base 11, 11 (July 2018), 1441–1453. DOI: 10.14778/3236187.3236197
- [40] Jérôme Kunegis. 2013. KONECT: The Koblenz network collection. In Proceedings of the 22nd International Conference on World Wide Web (WWW'13). ACM Press, Rio de Janeiro, Brazil, 1343–1350. DOI: 10.1145/2487788.2488173
- [41] Yung-Keun Kwon and Kwang-Hyun Cho. 2007. Analysis of feedback loops and robustness in network evolution based on Boolean models. BMC Bioinform. 8 (2007), 430. DOI: 10.1186/1471-2105-8-430
- [42] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. Retrieved from https://snap.stanford.edu/data. Accessed: 2022-05-30.
- [43] Xiangfeng Li, Shenghua Liu, Zifeng Li, Xiaotian Han, Chuan Shi, Bryan Hooi, He Huang, and Xueqi Cheng. 2020. FlowScope: Spotting money laundering based on graphs. *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (Apr. 2020), 4731–4738. https://doi.org/10.1609/aaai.v34i04.5906
- [44] G. Loizou and P. Thanisch. 1982. Enumerating the cycles of a digraph: A new preprocessing strategy. Info. Sci. 27, 3 (Aug. 1982), 163–182. DOI: 10.1016/0020-0255(82)90023-8
- [45] Grzegorz Malewicz, Matthew H. Austern, Aart J. C. Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. 2010. Pregel: A system for large-scale graph processing. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. ACM, Indianapolis, Indiana, 135–146. DOI: 10.1145/1807167.1807184
- [46] Prabhaker Mateti and Narsingh Deo. 1976. On algorithms for enumerating all circuits of a graph. SIAM J. Comput. 5, 1 (Mar. 1976), 90–99. DOI: 10.1137/0205007
- [47] Nav Mathur. 2017. Graph Technology for Financial Services. Technical Report. Neo4J. 1–14 pages. Retrieved from https: //neo4j.com/use-cases/financial-services. Accessed: 2022-05-30.
- [48] Robert Ryan McCune, Tim Weninger, and Greg Madey. 2015. Thinking like a vertex: A survey of vertex-centric frameworks for large-scale distributed graph processing. ACM Comput. Surv. 48, 2 (Nov. 2015), 1–39. DOI: 10.1145/2818185
- [49] Robert Meusel, Sebastiano Vigna, Oliver Lehmberg, and Christian Bizer. 2014. Graph structure in the web revisited: A trick of the heavy tail. In Proceedings of the 23rd International Conference on World Wide Web (WWW'14). ACM Press, 427–432. DOI: 10.1145/2567948.2576928
- [50] S. Noel, E. Harley, K. H. Tam, M. Limiero, and M. Share. 2016. CyGraph: Graph-based analytics and visualization for cybersecurity. In *Handbook of Statistics*. Vol. 35. Elsevier, Oxford, UK, 117–167. DOI: 10.1016/bs.host.2016.07.001
- [51] Gabriele Oliva, Roberto Setola, Luigi Glielmo, and Christoforos N. Hadjicostis. 2018. Distributed cycle detection and removal. *IEEE Trans. Control Netw. Syst.* 5, 1 (Mar. 2018), 194–204. https://doi.org/10.1109/TCNS.2016.2593264
- [52] Girish Keshav Palshikar and Manoj M. Apte. 2008. Collusion set detection using graph clustering. Data Min. Knowl. Disc. 16, 2 (Apr. 2008), 135–164. DOI: 10.1007/s10618-007-0076-8
- [53] Ashwin Paranjape, Austin R. Benson, and Jure Leskovec. 2017. Motifs in temporal networks. In Proceedings of the 10th ACM International Conference on Web Search and Data Mining. ACM, Cambridge, UK, 601–610. DOI:10.1145/3018661.3018731
- [54] You Peng, Ying Zhang, Xuemin Lin, Wenjie Zhang, Lu Qin, and Jingren Zhou. 2019. Towards bridging theory and practice: hop-constrained s-t simple path enumeration. *Proc. Very Large Data Base* 13, 4 (Dec. 2019), 463–476. DOI:10.14778/3372716.3372720
- [55] J. Ponstein. 1966. Self-avoiding paths and the adjacency matrix of a graph. SIAM J. Appl. Math. 14, 3 (Mar. 1966), 600–609. DOI: 10.1137/0114051
- [56] Sri Harsha Pothukuchi and Amit Dhuria. 2021. Deterministic loop breaking in multi-mode multi-corner static timing analysis of integrated circuits. Patent No. 11003821.
- [57] Zhu Qing, Long Yuan, Zi Chen, Jingjing Lin, and Guojie Ma. 2020. Efficient parallel cycle search in large graphs. In *Database Systems for Advanced Applications*. Vol. 12113. Springer International, Cham, Switzerland, 349–367. DOI:10.1007/978-3-030-59416-9_21
- [58] Xiafei Qiu, Wubin Cen, Zhengping Qian, You Peng, Ying Zhang, Xuemin Lin, and Jingren Zhou. 2018. Real-time constrained cycle detection in large dynamic graphs. Proc. Very Large Data Base 11, 12 (Aug. 2018), 1876–1888. DOI:10.14778/3229863.3229874
- [59] Michael J. Quinn. 2004. Parallel Programming in C with MPI and openMP. McGraw-Hill, Dubuque, Iowa.
- [60] R. C. Read and R. E. Tarjan. 1975. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. Networks 5, 3 (July 1975), 237–252. DOI: 10.1002/net.1975.5.3.237
- [61] Rodrigo Caetano Rocha and Bhalchandra D. Thatte. 2015. Distributed cycle detection in large-scale sparse graphs. In Proceedings of the Simpósio Brasileiro de Pesquisa Operacional (SBPO'15). SOBRAPO, Porto de Galinhas, Pernambuco, Brasil, 1–12. DOI:10.13140/RG.2.1.1233.8640
- [62] Amitabha Roy, Ivo Mihailovic, and Willy Zwaenepoel. 2013. X-Stream: Edge-centric graph processing using streaming partitions. In Proceedings of the 24th ACM Symposium on Operating Systems Principles. ACM, 472–488. DOI:10.1145/2517349.2522740

ACM Transactions on Parallel Computing, Vol. 10, No. 3, Article 15. Publication date: September 2023.

15:34

- [63] SAS. 2021. SAS OPTGRAPH Procedure: Graph Algorithms and Network Analysis. Retrieved from https:// documentation.sas.com/doc/en/pgmsascdc/9.4_3.5/procgralg/procgralg_optgraph_examples.htm. Accessed: 2022-05-30.
- [64] Avinash Sodani. 2015. Knights landing (KNL): 2nd Generation Intel Xeon Phi processor. In Proceedings of the IEEE Hot Chips 27 Symposium (HCS'15). IEEE, 1–24. DOI: 10.1109/HOTCHIPS.2015.7477467
- [65] Toyotaro Suzumura and Hiroki Kanezashi. 2021. Anti-Money Laundering Datasets: InPlusLab Anti-Money Laundering Datasets. Retrieved from https://github.com/IBM/AMLSim. Accessed: 2022-05-30.
- [66] J. Szwarcfiter and P. Lauer. 1976. A search strategy for the elementary cycles of a directed graph. BIT Numer. Math. 16 (1976), 192–204.
- [67] Robert Tarjan. 1973. Enumeration of the elementary circuits of a directed graph. SIAM J. Comput. 2, 3 (Sept. 1973), 211–216. DOI: 10.1137/0202017
- [68] James C. Tiernan. 1970. An efficient search algorithm to find the elementary circuits of a graph. Commun. ACM 13, 12 (Dec. 1970), 722–726. DOI: 10.1145/362814.362819
- [69] Leslie G. Valiant. 1990. A bridging model for parallel computation. Commun. ACM 33, 8 (Aug. 1990), 103–111. DOI:10.1145/79173.79181
- [70] Fei Wang, Peng Cui, Jian Pei, Yangqiu Song, and Chengxi Zang. 2020. Recent advances on graph analytics and its applications in healthcare. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. ACM, 3545–3546. DOI: 10.1145/3394486.3406469
- [71] Jim Webber. 2021. Powering Real-Time Recommendations with Graph Database Technology. Technical Report. Neo4J. 1–7 pages. Retrieved from https://neo4j.com/use-cases/real-time-recommendation-engine. Accessed: 2022-05-30.
- [72] Herbert Weinblatt. 1972. A new search algorithm for finding the simple cycles of a finite directed graph. J. ACM 19, 1 (Jan. 1972), 43–56.
- [73] J. T. Welch. 1965. Numerical applications: Cycle algorithms for undirected linear graphs and some immediate applications. In Proceedings of the 20th National Conference. ACM Press, 296–301. DOI:10.1145/800197.806053
- [74] Xiaoping Zhou, Xun Liang, Jichao Zhao, and Shusen Zhang. 2018. Cycle-based network centrality. *Sci. Rep.* 8, 1 (Dec. 2018), 11749. DOI: 10.1038/s41598-018-30249-4

Received 20 February 2023; revised 7 July 2023; accepted 7 July 2023