Fast parallel algorithms for finding elementary circuits of a directed graph: a GPU-based approach

Amira Benachour[1] · Saïd Yahiaoui[2] · Didier El Baz[3] ·
Nadia Nouali-Taboudjemat[2] · Hamamache Kheddouci[4]

Abstract

Circuits in a graph are interesting structures and identifying them is of an important relevance for many applications. However, enumerating circuits is known to be a difficult problem, since their number can grow exponentially. In this paper, we propose fast parallel approaches for enumerating elementary circuits of directed graphs based on graphics processing unit (GPU). Our algorithms are based on a massive exploration of the graph in a breadth-first search strategy. Algorithm V-FEC explores the graph starting from different vertices simultaneously. To further reduce the search space, we present T-FEC, another algorithm that uses triplets as an initial set to start exploring. To the best of our knowledge, those are the first parallel GPU-based algorithms for finding all circuits of a given graph. In addition, they find circuits of a given length and circuits with a specific vertex or edge. The evaluation results show

✉ Amira Benachour
  abenachour@usthb.dz

  Saïd Yahiaoui
  syahiaoui@cerist.dz

  Didier El Baz
  elbaz@laas.fr

  Nadia Nouali-Taboudjemat
  nnouali@cerist.dz

  Hamamache Kheddouci
  hamamache.kheddouci@univ-lyon1.fr

1 Department of Computer Science, University of Sciences and Technology Houari Boumediene, 16111 Bab Ezzouar, Algiers, Algeria

2 CERIST Research Center on Scientific and Technical Information, 16028 Ben Aknoun, Algiers, Algeria

3 LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

4 Université de Lyon, CNRS, Université Lyon 1, LIRIS, UMR5205, F-69622 Lyon, France

⚛ Springer

that the proposed approaches achieve up to 190x speed-up over Johnson's algorithm, one of the most efficient sequential algorithms for finding circuits.

1 Introduction

The problems of enumerating structures such as circuits, paths, trees and cliques in graphs and networks are of fundamental importance. This kind of tasks is usually hard to deal with, due to its challenging time/space complexity. Indeed, even a small graph may contain a huge number of such structures. In the enumeration process, it is important to point that counting and finding are two different processes: Finding is the construction of elements, and counting is determining their number. Generally, knowing the count is rarely useful in finding the objects and even if we are able to count the elements in polynomial time, we might not be able to find all of them due to an exponential number of elements and huge memory requirements. In the following, we will use enumeration and finding interchangeably.

Our interest is in finding the elementary circuits of a directed graph, which is a fundamental problem of graph theory. A circuit is a loop through which information propagates and it is widely used in real-world applications. For example, circuits can reveal money laundering in the financial system [1]. They represent fragile ecosystem dependencies in food webs [2]. Finding circuits is also interesting in fraud detection [3]. Circuits serve as a method for social network analysis and modeling [4] and are used to evaluate networks balance [5]. Finding circuits is also important in biology [6–8], and Internet systems [9].

The first algorithms proposed for finding elementary circuits of a directed graph require either a lot of memory or have exponential running time [10, 11]. As a result, they become impractical as the graph size increases. The algorithm by Johnson [12] is one of the most efficient algorithms for finding elementary circuits. It is still the best state of art algorithm, yet it may fail to handle some large graphs. To deal with larger graphs, Lu et al. [13] introduce a parallel algorithm based on Johnson's algorithm aiming to reduce the execution time. Their algorithm is based on first splitting the job into small ones and then applying Johnson's algorithm on each one in parallel. Giscard et al. [14] proposed an algorithm that counts the circuits of a graph without enumerating them. It gives the number of circuits for a length l, but it is tested for a limit of $l = 10$. In [15] the authors present an algorithm that deals with large graphs but only finds circuits of length less than or equal to k. The limit in their tests is $k = 6$.

In this paper, we present parallel approaches for finding all circuits of a directed graph. We first detect whether the graph contains circuits by searching for the strongly connected components, then we start exploring the graph to find all the circuits. Our algorithms proceed by a massive exploration of paths. We present two approaches based on the initial set used for exploration: The first uses vertices of the graph to which we assign a search area, while the second uses triplets to reduce the

search space of exploration. Our algorithms permit to find circuits of a given length l and circuits with a specific vertex or edge. We develop a parallel implementation of the proposed approaches in a graphics processing unit (GPU) environment that is better suited for massive exploration. Our algorithms can find millions of circuits in a reduced time. We provide, to the best of our knowledge, the first GPU-based algorithms for finding elementary circuits of directed graphs.

The remaining of this paper is organized as follows. In Sect. 2 we present the background and related work. In Sect. 3, we introduce some preliminary definitions. The parallel algorithms are explained in Sect. 4. Details of the GPU implementation and examples are given in Sect. 5. Section 6 describes the experimental tests and results. Conclusion and future work are discussed in Sect. 7.

2 Related work

Finding cyclic structures in a graph has been largely studied in the literature because of its many potential application scenarios. Some existing works have been interested in detecting circuits of directed or undirected graphs [16–19]. Other studies focused on finding the shortest circuit [20–22] or the longest circuits of a graph [23–26]. In this paper, we are interested in finding all the elementary circuits of a directed graph. It is a classical problem whose most efficient solutions date back to the 1970s.

Tiernan [10] was the first to propose the idea of blocking visited vertices. The algorithm is based on a backtracking procedure. It generates all elementary paths $p = (v_1, v_2, \dots, v_k)$ with $v_1 < v_i$ for all $i \in \{2, \dots, k\}$. Starting from some vertex v_1, it chooses an edge to traverse to some vertex v_2 such that $v_2 > v_1$ and continue this way. Whenever no vertex can be reached, the procedure backs up one vertex and chooses a different edge to traverse. If v_1 is adjacent to v_k the algorithm lists an elementary circuit $(v_1, v_2, \dots, v_k, v_1)$. This approach examines more simple paths than necessary, making the worst-case time bound exponential in the number of elementary circuits and in the graph size.

Tarjan's algorithm [11] is based on Tiernan's depth-first method but it is faster. It lists all the cycles in $O((|V| + |E|)|C|)$ time. It uses backtracking procedure to avoid unnecessary work. Two stacks are needed, the *point stack* for storing the path currently being examined, and a Boolean vector called *marked stack* to indicate whether a vertex is used in a path or not. Whenever a new circuit is found, all vertices in the current point stack will eventually be unmarked when popped from this stack. If no circuit is found involving a vertex, it will be deleted from the point stack, but continue to be marked. Some of the unnecessary work in Tiernan's algorithm is avoided by the condition that a vertex can be added to the *point stack* only if it is unmarked.

The algorithm by Weinblatt [27] begins processing the graph by eliminating vertices which cannot belong to any circuits. Next, it selects a starting vertex and chooses an edge to follow. Circuits are found when a path is cyclic or by combining parts of previously discovered circuits with a part of the path that is being processed. However, this requires more storage and the time saved by efficient search is lost in searching the constructed circuits and paths.

Of all the algorithms analyzed, the algorithm of Johnson [12] is the fastest algorithm having an upper time bound of $O((|V| + |E|)|C|)$. The success of this backtrack algorithm is due to an effective pruning technique which avoids much of the fruitless search present in earlier algorithms. For each start vertex s, a recursive backtracking procedure is invoked and its computation is similar to that of Tarjan's algorithm, except for the marking system, which was considerably enhanced. A vertex is blocked each time it enters the stack. It remains blocked until we backtrack from a vertex v that belongs to the last enumerated circuit. The unblock process recursively unblocks the vertices that have a path to v.

In 2006, an algorithm by Liu and Wang [28] present a way to enumerate circuits in directed and undirected graphs. It uses $k - 1$ paths to generate k circuits by exploring adjacent vertices of the tail (the last vertex in the path), this way it can enumerate circuits by length. Compared to Johnson's and Tiernan's algorithms, their algorithm is slower and not efficient.

In [29], the authors propose a parallelizable algorithm into $|V - 1|$ parts. They use a backtracking strategy using pointers allowing to backtrack more than one vertex at a time, jumping directly to the vertex from which further path extensions are possible.

The algorithm by Lu et al. [13] is capable of handling large-scale graphs. The idea is to split the job into subroutines depending on the number of edges. Then apply Johnson's algorithm to find all elementary circuits including an edge in parallel for each subroutine. Every subroutine has a search area that consists of the graph after deleting all the edges in inputs for previous subroutines. For this, they grade the edges, by estimating the number of all elementary paths from a vertex v to a vertex u. The edge with higher *Id* will have smaller search area. This approach is implemented on Hadoop using the MapReduce technique. Compared to previous approaches, this algorithm reduces the running time for finding elementary circuits of a graph but it uses depth-first search (DFS) sequential algorithm of Johnson. Theoretically, DFS has been proven by Reif [30] to be inherently sequential. It is considered to be a challenging problem for parallelization. Recently, algorithms that find length-constrained circuits have been proposed. In [15], the authors find circuits bounded by length $k = 6$. Another work by [31] finds bounded circuits from a source vertex s in undirected graphs in parallel, with a maximum length of $k = 8$.

The algorithms we propose in this paper exploit the massive parallelism of the multi-cores GPU by running multiple searches simultaneously which makes it capable of finding millions of circuits in a short time. The exploration process is based on the breadth-first search (BFS) algorithm that is better suited to parallel implementation. Details are given in the following.

3 Preliminaries

In this section, we present formal definitions [32] that support our approaches to enumerate all elementary circuits of a graph. Let $G = (V, E)$ be a finite directed simple graph with vertex set V and edge set E. Let $Id(v)$ denote the number associated with the vertex v.

Definition 1 *(Elementary path)* A simple path p in a graph $G = (V, E)$ is a finite sequence of vertices $p = \langle v_1, v_2, \dots, v_k \rangle$, such that $(v_i, v_{i+1}) \in E$ for $i \in \{1, 2, \dots, k - 1\}$. A path is *elementary* if no vertex appears twice.

Definition 2 *(Elementary circuit)* A *circuit* is a path in which the first and the last vertices are identical. An *elementary circuit* is a circuit where no vertex appears twice except the first and the last vertices.

Two elementary circuits are distinct if one is not a cyclic permutation of the other. Note that, in this paper we represent a circuit as a sequence of vertices $c = \langle v_1, v_2, \dots, v_k \rangle$ without writing the redundant vertex twice.

Definition 3 *(Strongly Connected Component)* A strongly connected component (*SCC*) in a directed graph is defined as a subgraph in which every vertex is reachable from every other vertex. That is, for any two vertices u and v in a strongly connected component, there exists a path from u to v, and a path from v to u.

We define *In-edges(v)*, the set of incoming edges of a vertex v such that:

$$In{-}edges(v) = \{(u, v) | (u, v) \in E\}$$

Similarly, we define *Out-edges(v)* the set of outgoing edges of v:

$$Out{-}edges(v) = \{(v, u) | (v, u) \in E\}$$

We name the destination vertex of the outgoing edge of v the *out-neighbor* of v. Likewise, the origin vertex of the incoming edge of v is called the *in-neighbor*.

For the sake of introducing the approaches below, we define the set of triplets as follows:

Definition 4 *(Triplets)* We denote by $T(G)$ the set of triplets of $G = (V, E)$ defined as:

$$T(G) = \{\langle x, v, y\rangle | x, y, v \in V \quad with \quad (x, v) \in \textit{In- edges}(v),$$
$$(v, y) \in \textit{Out- edges}(v),$$
$$Id(v) < Id(x),$$
$$Id(v) < Id(y)\}$$

A triplet is a sequence of vertices that initiate a path of length greater than two. We name v the *source-vertex* and x the *in-vertex* for a triplet $\langle x, v, y\rangle$. Using triplets allows to find every circuit only once and reduce the search space, as we shall see in the sequel.

In the following, we present our algorithms for finding all elementary circuits.

4 Parallel algorithms for finding elementary circuits

We present here our parallel approaches for finding elementary circuits of directed graphs. Circuits are constructed starting from a *source-vertex s*. We build elementary paths by iteratively exploring adjacent vertices.

To avoid visiting vertices that do not belong to any circuit, we first find all the *SCCs* of the graph as a preprocessing task. Definition 3 shows that an *SCC* consisting of more than one vertex must contain circuits and that all vertices in any circuit should be in the same *SCC*. Decomposing the graph into *SCCs* reduces the number of initial source vertices and the number of neighbors to be visited. Only neighbors that are in the same *SCC* are considered, hence, reducing the number of paths to be explored. In this work, we use the algorithm by Tarjan [33] for finding *SCCs* of the graph. The algorithm's running time is linear in the number of edges and vertices in G, it finds *SCCs* in $O(|V| + |E|)$ time.

Once the graph has been partitioned into *SCCs*, we move to the exploration phase. A straightforward approach to obtain all circuits is to enumerate all combinations of distinct edges and check if the path may construct a circuit. On this basis, we propose two algorithms: Vertex-based algorithm for finding elementary circuits (V-FEC) and triplets-based algorithm for finding elementary circuits (T-FEC). Our algorithms are based on a massive exploration of paths. For each *SCC* found, we run a parallel breadth-first search (BFS) starting from a set of initial source vertices and build exploration trees. The detailed algorithm V-FEC is presented in Sect. 4.1.

To reduce the search space and filter out some unnecessary combinations, we propose algorithm T-FEC that uses triplets instead of vertices as initial set. As with the precedent algorithm, V-FEC, we first find *SCCs* of the graph. Then, we generate the set of initial triplets and begin exploration. The same process of exploration is used. It is interesting to note that using triplets helps to balance exploration trees. The detailed algorithm is presented in Sect. 4.2.

For the following, C denotes the set of circuits found. Set P is used to contain the explored paths.

4.1 A vertex-based approach V-FEC

As explained above, the first algorithm is vertex-based. We begin by finding the *SCCs* of the graph. Then, we initialize the set of paths P with vertices of the *SCCs*: $P = \{\langle s_1 \rangle, \langle s_2 \rangle,..., \langle s_k \rangle\}$.

We define the search area (*SA*) of a vertex s by the subgraph induced by s and vertices with *Id* greater than the *Id* of s which are in the same *SCC*: $SA(s) = \{u \in V | SCC(u) = SCC(s) \text{ and } Id(u) > Id(s)\}$, $SCC(v)$ denotes the identifier of the *SCC* where vertex v belongs. In other words, for a vertex s we only consider vertices in the same *SCC* with *Id* greater than *Id* of s. Comparably, in Johnson's algorithm, we consider vertex s to be the smallest vertex in the *SCC*. When all circuits from s are found, we remove s, *In-edges(s)* and *Out-edges(s)*, and then, we explore the rest of the *SCC* in a similar manner.

In the exploration phase, for every vertex s_i in P, we only explore vertices in its search area. We start by visiting the *out-neighbors* of s_i. Then, we iteratively explore the visited vertices until no more vertices can be visited. The above-mentioned process is outlined in Algorithm 1.

Algorithm 1 Vertex-based Finding Elementary Circuits V-FEC(G)

Require: Graph G
Ensure: Set C of all elementary circuits of G
  1: $C \Leftarrow \varnothing$
  2: $SCC \Leftarrow \text{findSCC}(G)$
  3: **foreach** $scc \in SCC$ **do**
  4:     **foreach** $s \in scc$ in parallel **do**
  5:         $C \Leftarrow \text{exploreV-FEC}(s)$ ▷ See Procedure 1
  6: **return** C

4.1.1 Exploration

Now we detail the exploration phase.

Our algorithm is based on running multiple BFS graph exploration simultaneously from every initial paths $p_i \in P$. (P has been already initialized with vertices of the *SCCs* found.) We construct exploration trees for every *source-vertex* s_i of p_i by visiting the *out-neighbors* of the last vertex v of path p_i. For any neighbor u of v, one of the following situations occurs:

- Vertex u is the *source-vertex* s_i: p_i is an elementary circuit. Add p_i to C,

(a) Directed graph G

(b) Decomposition of graph G into Strongly Connected Components (SCC)

Fig. 1 Graph decomposition into *SCCs*

- u is not visited and belongs to the search area of s_i: p_i is extended. We add $\langle p_i, u \rangle$ to the set of paths P.
- u is already visited: the path $\langle p_i, u \rangle$ cannot constitute an elementary circuit.

The algorithm proceeds iteratively by visiting the last visited vertex of each resulting path until no path can be found ($P = \varnothing$). At the end of the exploration phase, we would have listed all the circuits ordered by their length. A pseudo-code of the explained process is presented in Procedure 1.

Procedure 1 exploreV-FEC

Require: Vertex s
Ensure: Set of circuits C
1: $P \Leftarrow \{\langle s \rangle\}$
2: $visited(s) \Leftarrow true$
3: **while** $P \neq \varnothing$ **do**
4:     **foreach** $p \in P$ **do**
5:         $v \Leftarrow lastVertex(p)$ ▷ returns the last vertex of p
6:         **foreach** $u \in out\text{-}neighbors(v)$ **do**
7:             **if** $Id(u) = Id(s)$ **then**
8:                 $C \Leftarrow C \cup \{p\}$
9:             **else**
10:                 **if** $\neg\, visited(u)$ and $u \in SA(s)$ **then**
11:                     $P \Leftarrow P \cup \{\langle p, u \rangle\}$
12:                     $visited(u) \Leftarrow true$
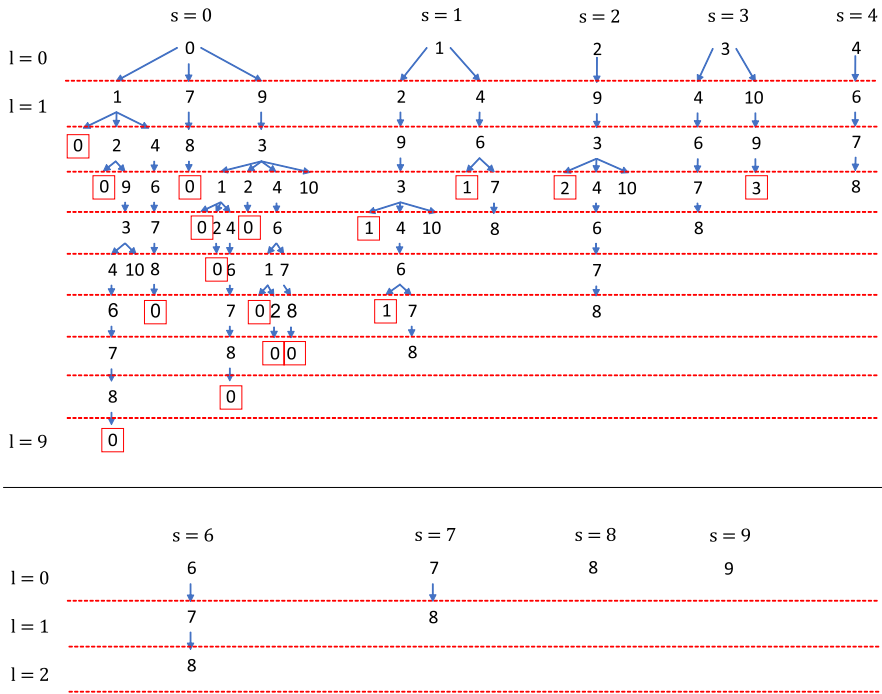13: $P \Leftarrow P - \{p\}$
14: **return** C

Fig. 2 Exploration trees from different source vertices: V-FEC approach

Example 1 To better explain the proposed algorithm we run an example on graph $G = (V, E)$, $|V| = 12$ and $|E| = 22$ shown in Fig. 1a. We begin with a search for the *SCCs* of the graph to eliminate unnecessary vertices and edges. Graph G contains three strongly connected components, a large *SCC*: SCC_0 and two trivial ones: SCC_1 and SCC_2, illustrated in Fig. 1b. We remove all the trivial *SCCs* and the edges going to or from them and begin processing the remaining graph.

We initialize the set P of paths with the vertices of SCC_0 ($l = 0$ in Fig. 2). Then we start constructing exploration trees for each *source-vertex s*. We visit *out-neighbors* of s that are in its search area. We recall that a vertex u is in the search area of vertex s if the two vertices are in the same *SCC* and $Id(u) > Id(s)$. Consider vertex s, $Id(s) = 1$. Vertex s has three *out-neighbors*: vertex 0, 2 and 4. Vertex 0 is not in $SA(s)$ since $0 < 1$. In this case, vertex 0 is not considered. If a vertex has no visited neighbors, the path is removed, e.g., paths: $\langle 8 \rangle$ and $\langle 9 \rangle$ have no visited neighbors, so they are removed. ($l = 0$ in Fig. 2).

4.2 A triplets approach T-FEC

We aim to reduce the search space for the exploration phase by generating triplets using direct neighbors of each vertex. The idea is to construct what can constitute potential circuits by joining incoming and outgoing edges of a given vertex. We divide the process into two stages and define a parallel algorithm for each one of them. As a first step, (1) we generate the set of triplets $T(G)$, and then, (2) we run a parallel breadth-first search (BFS) for each resulting triplet in a likewise strategy as in the V-FEC approach. Algorithm 2 outlines the proposed approach.

Algorithm 2 Triplet-based Finding Elementary Circuits T-FEC(G)

Require: Graph G
Ensure: Set C of all elementary circuits of G
 1: $C \Leftarrow \varnothing$
 2: $T \Leftarrow \varnothing$
 3: $SCC \Leftarrow \mathrm{findSCC}(G)$
 4: **foreach** $scc \in SCC$ **do**
 5:     **foreach** $v \in scc$ in parallel **do**
 6:         $T \Leftarrow \mathrm{generateTriplets}(v)$ ▷ See Procedure 2
 7: **foreach** t $\in T$ in parallel **do**
 8:     $C \Leftarrow \mathrm{exploreT\text{-}FEC}(t)$ ▷ See Procedure 3
 9: **return** C

4.2.1 Triplets generation

The idea of using triplets to begin exploration helps reducing the search space and avoids redundant circuits. Triplets also permit to have more balanced exploration trees, thus, a more balanced work charge. We build the set of triplets by making every possible combination $\langle x, v, y \rangle$ of the *in-neighbors* and *out-neighbors* of vertex v such as:

- x is the *in-neighbor* of v with $x \in SA(v)$, i.e., $Id(v) < Id(x)$ and v and x are in the same *SCC*.
- Similarly, y is the *out-neighbor* of v with $y \in SA(v)$.

If $x = y$, $c = \langle v, y \rangle$ is a 2-length circuit. The generation process is presented in Procedure 2.

Procedure 2 generateTriplets

Require: Vertex v
Ensure: Set T of initial triplets
  1: **foreach** $x \in in\text{-}neighbors(v)$ **do**
  2:     **foreach** $y \in out\text{-}neighbors(v)$ **do**
  3:        **if** $x = y$ **then**
  4:           $C \Leftarrow C \cup \{\langle x, v \rangle\}$
  5:        **else**
  6:           **if** $x \in SA(v)$ and $y \in SA(v)$ **then**
  7:              $T \Leftarrow T \cup \{\langle x, v, y \rangle\}$
  8:              $visited(x) \Leftarrow true$
  9:              $visited(v) \Leftarrow true$
10:              $visited(y) \Leftarrow true$
11: **return** T

4.2.2 Exploration

The exploration phase for algorithm T-FEC is similar to the exploration in algorithm V-FEC. Initially we visit the *out-neighbors* of y for every triplet $\langle x, v, y \rangle \in T$. Only neighbors in *SA(v)* are considered. Then, we iteratively visit the *out-neighbors* of the last visited vertex v_k of every path $p_i \in P$. We recall that for a triplet $\langle x, v, y \rangle$, v denote the *source-vertex* and x is the *in-vertex*. For any neighbor u of vertex v_k, the exploration goes as follows:

- u is the *in-vertex* ($u = x$): p_i is an elementary circuit. Add p_i to C,
- u is not visited and $u \in SA(v)$: p_i is extended, its an elementary path. We add $\langle p_i, u \rangle$ to P.
- u is already visited: $\langle p_i, u \rangle$ is not an elementary path.

The exploration is over when no path can be expanded ($P = \varnothing$). The exploration process is described in Procedure 3.

Procedure 3 exploreT-FEC

Require: Triplet $t\langle x, v, y\rangle$
Ensure: Set of circuits C
1: $P \Leftarrow \{t\}$
2: **while** $P \neq \varnothing$ **do**
3:      **foreach** $p \in P$ **do**
4:          $v_k = lastVertex(p)$
5:          **foreach** $u \in out\text{-}neighbors(v_k)$ **do**
6:              **if** $u = x$ **then**
7:                  $C \Leftarrow C \cup \{p\}$
8:              **else**
9:                  **if** $\neg\ visited(u)$ and $u \in SA(v)$ **then**
10:                      $P \Leftarrow P \cup \{\langle p, u\rangle\}$
11:                      $visited(u) \Leftarrow true$
12:          $P \Leftarrow P - \{p\}$
13: **return** C

4.3 Proof of correctness

In the following, we will prove that Algorithms V-FEC and T-FEC find every elementary circuit of a graph exactly once.

We start by showing that both algorithms find only elementary circuits of a graph.

Lemma 1 *Algorithms V-FEC and T-FEC find elementary circuits only.*

Proof In algorithm V-FEC, every vertex in a path is marked as visited and a new vertex is added to a path if it is not already visited. If a vertex is visited we do not visit it again (line 10 of Procedure 1). Hence, a path cannot contain redundant vertices. Let a path $p = \langle v_1, \ldots, v_i, \ldots, v_k\rangle$ to be expanded, and let vertices v_l and v_i be in *out-neighbors* of vertex v_k. Suppose that vertex v_l is not visited yet. If $v_l \in SA(v_1)$, it will be added to p and marked as visited. However, since vertex v_i is previously visited, it will not be explored again. This confirms that algorithm V-FEC finds only elementary circuits.

Similarly, for algorithm T-FEC, in Procedure 3 *exploreT-FEC*, in line 9, visited vertices are not explored again, thus, the resulting circuits are elementary. □

Theorem 1 *Algorithms V-FEC and T-FEC find all the elementary circuits of a graph exactly once.*

Proof The correctness of this lemma can be verified by the following: (1) Algorithms V-FEC and T-FEC find any elementary circuit exactly once, (2) all the elementary circuits of the graph are found.

To prove the first property, we suppose a circuit is found more than once. Let $\langle v_1, v_i, \ldots, v_l \rangle$ be a circuit, such that $Id(v_i) = i$, $\forall v_i \in V$, and vertex v_1 has the smallest Id. Many permutations of the same circuit can be found by starting from different source vertices. A circuit can be expressed in a permuted form as many times as the number of vertices it has: $\langle v_1, v_i, \ldots, v_l \rangle$, $\langle v_i, \ldots, v_l, v_1 \rangle$, $\cdots$, $\langle v_l, v_1, v_i, \ldots \rangle$.

In algorithm V-FEC, during the exploration phase every vertex of each *SCC* starts building its exploration tree by visiting its *out-neighbors*. A new vertex is added to a path if its *Id* is greater than the *Id* of the *source-vertex* (line 10 in Procedure 1: see the definition of *SA(s)* in Sect. 4.1). As a consequence, the first vertex in a path is always the vertex with the smallest *Id*.

Consider a circuit $c = \langle v_1, v_i, \ldots, v_l \rangle$ and two cyclic permutations of c: $c' = \langle v_i, \ldots v_l, v_1 \rangle$ and $c'' = \langle v_l, v_1, v_i, \ldots \rangle$. Following Procedure 1 *exploreV-FEC*, we try to construct circuit c and its permutations starting from different source vertices v_1, v_i,..., and v_l. Each time we visit a neighbor we check if its *Id* is greater than the *Id* of the *source-vertex* of the path. Since $Id(v_1) < Id(v_i)$ and $Id(v_1) < Id(v_l)$, in both cases, $\langle v_i, \ldots v_l, v_1 \rangle$ and $\langle v_l, v_1, v_i, \ldots \rangle$, vertex v_1 is not explored, thus, circuits c' and c'' cannot be constructed. Hence, no circuit can be found more than once.

In a similar way, when generating triplets for a vertex v in algorithm T-FEC, we only consider neighbors that have *Ids* greater than $Id(v)$, and likewise for the exploration phase, this guarantees that circuits are not duplicated. We come to the conclusion that both algorithms find each circuit exactly once.

Now we prove the second property: Algorithms V-FEC and T-FEC find all the elementary circuits of a graph. We know from the definition of *SCC* that every *SCC* with more than one vertex must have circuits, and that all vertices in any circuit belong to the same *SCC*. Finding all elementary circuits of a graph is equivalent to finding all elementary circuits starting from a vertex, for every vertex of the graph.

In algorithm V-FEC, Procedure *exploreV-FEC* finds for a vertex s all the elementary circuits starting from s such that $Id(s)$ is the smallest, by visiting every neighbor of every vertex within *SA(s)* not visited yet in the same *SCC*, making every possible combination of possible paths (follows directly from Lemma 1 and the proof of property(1)). Thus, for every elementary circuit of the graph, with all vertices in *SA(s)*, its cyclic permutation starting by vertex s is returned at the end of the algorithm. Since Procedure *exploreV-FEC*, is called for every vertex of every *SCC* of the graph, every elementary circuit is found.

Similarly for T-FEC, in the triplets generation phase, we consider all neighbors of vertex v having *Id* greater than $Id(v)$. Procedure *exploreT-FEC* visits every neighbor of every vertex within *SA(v)* not visited yet in the same *SCC*. It finds for a triplet $\langle x, v, y \rangle$ all the elementary circuits starting from v and including vertices x and y, such that $Id(v)$ is the smallest. The procedure is called for every generated triplet; thus, every elementary circuit is found. We conclude that the proposed algorithms find all the elementary circuits of a graph exactly once. This completes the proof. $\square$

4.4 Circuits and paths given a source vertex

For some applications, we are interested in the enumeration of circuits that goes through a certain vertex v. The proposed approaches can be easily adapted to only find circuits that contains vertex v.

To achieve this, we make an adjustment in the initialization phase. In the vertex-based approach, we initialize the set of paths to v and then call Procedure 1 *exploreV-FEC*, which only runs for vertex v. For the triplets approach, we generate triplets for vertex v by calling Procedure 2 *generateTriplets* and omit the condition about considering only vertices greater than v, which was mainly put to avoid enumerating circuits that are enumerated in the exploration process of other vertices. Since we just list circuits of v, circuits containing vertices of inferior *Id* are not found if we do not consider all vertices of the graph. That means, the set of triplets is generated using all of v's neighbors and there will be no redundant circuits. For both approaches, the exploration procedure remains the same, except for the condition about vertices *Ids* for the same reason explained above. In this case, the entire graph is considered.

In a similar manner, we can find circuits with a given edge (u, v) by initializing *Paths* to $\langle u, v \rangle$ in the vertex based approach. Or by generating triplets of form $\langle x_i, u, v \rangle$ with $x_i \in$ *in-neighbors* of u, for the triplet based approach. It is also possible to enumerate existing paths from a given source vertex s to a destination vertex d by searching circuits containing the edge (d, s) or Using triplets $\langle d, s, y_i \rangle$ where $y_i \in$ *out-neighbors* of s.

5 GPU implementation

A GPU is a highly parallel architecture that adopts the single instruction multiple data (SIMD) processing model. A GPU program consists of sequential segments and parallel segments of codes. The sequential codes run on the CPU, and the parallel codes, called kernels, run on the GPU. When a kernel is invoked, a large number of threads are launched to exploit data parallelism. Threads are grouped into blocks and blocks can be grouped into a grid. Thus, the kernel configuration has a significant effect on the degree of parallelism and hence influences the computing efficiency.

In this section, we present a GPU implementation of the proposed approaches. We proceed in parallel, multiple BFS explorations. This model fits the multi-threaded GPU architecture.

The *SCC* decomposition algorithm was not parallelized due to the low impact in the processing time of the algorithm. A GPU-kernel is used to implement the exploration phase of the vertex-based version. For T-FEC two kernels are used: The first kernel generates the set of initial triplets, the second one explores the graph.

Developing efficient graph processing algorithms in GPU is a challenging task. A GPU has limitations such as a relatively small memory size. Thus, the choice of data structures and strategies employed to take advantage of the characteristics of the different types of GPU memories is important to overcome these limitations. In our

(a) CSR representation **(b)** CSC representation

Fig. 3 Data structure for graph representation

implementation, we use data structures that are best suited to the GPU. We also take advantage of the GPU global, local and shared memories.

In the following we discuss our choices for data structures and detail the implementation of our algorithms.

5.1 Data structures

5.1.1 Graph representation

Compressed adjacency list is the most common data structure for graph representation in GPU. It provides compact storage for large sparse graphs and regular memory access. It uses two arrays to store the graph $G = (V, E)$ and only requires $O(|V| + |E|)$ memory space. Adjacency lists of all vertices are packed into a single large array that we name the *edge array*. We store the starting position of the adjacency list in the *edge array* for each vertex [34]. When outgoing edges are used in the *edge array*, we name this adjacency list format, the compressed sparse row (CSR). If incoming edges are used in the *edge array*, it is called compressed sparse column (CSC) [35]. To generate triplets both CSR and CSC representations are used, since we need to access both incoming and outgoing edges of each vertex. Once the triplets are generated we no longer need the incoming edges, CSC arrays are deleted. An example of the CSR and CSC representations of the graph in Fig. 1a is represented in Fig. 3.

5.1.2 Intermediate data storage and results

When searching for circuits, a number of paths is generated in every iteration. We keep track of them by using a global integer array. Initially, it contains the initial vertices or generated triplets. Another array stores the circuits. Since the GPU do not allow dynamic memory allocation, every memory that becomes necessary must be previously allocated. We declare two static structures *Circuits* and *Paths* larger enough for the purpose they were intended to.

During exploration, previously visited vertices in a path must not be visited again. This guarantees that circuits found are elementary. A single bit indicating whether or not the vertex is in the path is sufficient. Accordingly, a bitmap (bit array) *Visited* is employed to mark the already visited vertices. This map is defined by a bi-dimensional matrix that contains a row for each path and $|V|$ columns of bits, one for each vertex of the graph. In terms of bytes, the number of columns is $\lceil \frac{|V|}{8} \rceil$, e.g.,
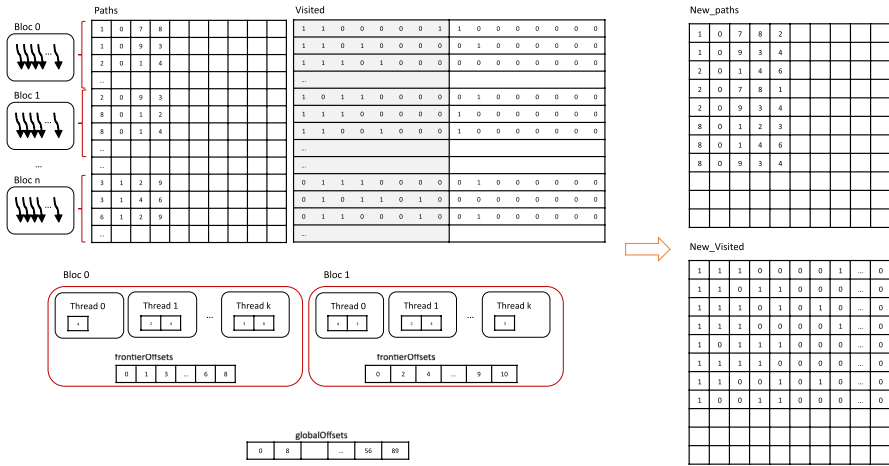
Fig. 4 A schema for paths exploration process

in a graph G with $|V| = 12$ a path storage occupy only 2 bytes. A Vertex v_j belongs to path i if and only if, bit j of row i is set to 1. Figure 4 shows an example of these arrays. The *Paths* matrix stores the vertices of an explored path in the order in which they are discovered. In the *Visited* matrix, each row contains a combination of bits that corresponds to the visited vertices of a path.

In addition to the small occupied space, adding a vertex or checking if it is already in the solution is done using a simple logical operation in the desired position. Because the bit-level operations are among the least computationally expensive, this method of storage contributes to improve the algorithm running time.

5.2 Parallel implementation

Now we detail the parallel implementations of V-FEC and T-FEC algorithms. We first describe the triplets generation process for T-FEC; then, we detail the exploration phase for both V-FEC and T-FEC.

5.2.1 Triplets generation

We begin by creating the set $T(G)$ of initial triplets $\langle x, v, y \rangle$. Such a set is created by selecting vertex $v \in V$ and analyzing all pairs of its adjacent vertices: x in the *in-neighbors* of v and y in *out-neighbors* of v. Procedure 2 *generateTriplets* describes the process of generating triplets. We launch a kernel *generateTriplets* and we assign a thread to every vertex v in the *SCC*. Each thread takes a vertex x from the *edge array E_{in}* and another vertex y from the *edge array E_{out}* and verify that they are in $SA(x)$. When this condition is satisfied the thread adds the combination $\langle x, v, y \rangle$ to the set of triplets. Another pair of neighbors is analyzed until all neighbors are checked.

In the case where vertex x equals vertex y, $\langle x, v \rangle$ is a 2-length circuit and is added to the set C.

At the end of the triplets generation phase, the set of incoming edges is no longer needed. We free memory used by deleting V_{in} and E_{in} arrays. We only keep the set of outgoing edges: V_{out} and E_{out} arrays, for the graph representation.

5.2.2 Exploration

The exploration kernels, exploreV-FEC and exploreT-FEC, are quite similar. Kernel exploreV-FEC launches a number of threads equal to the number of initial vertices. Kernel exploreT-FEC launches a number of threads equal to the number of generated triplets. Each thread is assigned a vertex/triplet and starts exploring. During exploration, it is important to save:

- for V-FEC: the *source-vertex* and the last visited vertex.
- for T-FEC: the *source-vertex*, the *in-vertex* and the last visited vertex of a path.

These data are essential for exploration. The last visited vertex of a path permits to expand the set of paths with new paths. At each iteration, a thread visits the *out-neighbors* of the last visited vertex of the path it is assigned to. In V-FEC, the *source-vertex* determines when a circuit is found. If the new vertex is the *source-vertex*, it is a circuit. For T-FEC, it is the *in-vertex* that determines a circuit. We use an atomic operation to add the discovered circuits by copying the current path to *Circuits*. If the new vertex do not constitute a circuit, the path can be expanded when the two following conditions are met: we first check if the explored vertex is not visited using a simple bit-wise operation, then, we check if the vertex is in the search area of the *source-vertex*. Considering these data are used by a thread for as many *out-neighbors* there is to explore, it is more appropriate to copy data to shared memory because it is faster. In V-FEC, we use two vectors V_{source} and V_{last} to store the *source-vertex* and the last visited vertex of a path, respectively. Another vectors V_{in}, that stores *in-vertices*, is necessary for the triplet based approach. The number of copied paths equals the number of threads per block. Chunks of the *Visited* array corresponding to copied paths are also transferred to shared memory.

Each thread uses a local array called *frontier* array to store the newly visited vertices to be copied back in global memory as new paths for the next iteration. We use a shared array *frontierOffsets*, set to 0, to compute the number of new paths. Each thread increments the value corresponding to its *Id* when it founds a new path. These values are used to copy the results of each thread in parallel to global memory. We first compute offsets for each thread on a block level by applying a parallel scan [34] on the *frontierOffsets* array. We use another array *blockOffsets* array to store the offsets of the blocks. Since there is no synchronization between blocks of threads, we use atomic operations to calculate the offsets of blocks. These two indices (*frontierOffsets* and *blockOffsets*) are used to compute the global memory address for each thread to copy its result. The schema represented in Fig. 4 encapsulates the process.

Compared to the V-FEC algorithm, the exploration process in T-FEC requires less iterations considering that two vertices are already explored for each path, when

generating triplets. Moreover, the triplets approach exposes more parallelism in early iterations, since there is more initial triplets than initial vertices which results in more thread occupancy.

Example 2 (Using triplets) Consider the same graph $G = (V, E)$ in Fig. 1a. In the T-FEC algorithm we start by initializing *triplets* and its corresponding *Visited* bitmap. A thread is mapped to a vertex, e.g., thread 2 generates triplets corresponding to vertex 2. We generate combinations of *in-neighbors* and *out-neighbors* of vertex 2. The *in-neighbors* of vertex 2 are vertex 1 and vertex 3. Vertex 1 is not considered since $1 < 2$. For vertex 3, $3 > 2$, so vertex 3 is considered. The *out-neighbors* of vertex 2 are vertex 0 and 9. $0 < 2$ vertex 0 is not considered as well. We are left with two valid vertices: vertex 3 as *in-neighbor* and vertex 9 as *out-neighbor*. As a result the triplet $\langle 3, 2, 9 \rangle$ is generated. Triplets from all *SCCs* are generated before starting exploration. The visited vertices of each triplet are marked as visited in the *Visited* array. Array *Visited* is an array of bytes (unsigned integers of length = 8 bits), we only use one bit to mark the visited vertices. Thus, $\lceil \frac{n}{8} \rceil$ columns are needed. In this example $n = 12$, as a consequence, *Visited* is a two columns array. During exploration, if a circuit is found, it is added to the *Circuits* array next to the circuits already found at earlier iterations. We take as an example triplet $\langle 2, 0, 1 \rangle$, vertex 2 is the *in-vertex*, vertex 0 is the *source-vertex* and vertex 1 is the last vertex of the path. We explore *out-neighbors* of vertex 1 that are vertices 0, 2 and 4. Vertex 0 is already visited so it is not added to *Paths*. Vertex 2 is the *in-vertex*, path $\langle 2, 0, 1, 2 \rangle$ is a circuit, we add it to *Circuits*. *Circuits* = *Circuits* ∪ {$\langle 2, 0, 1 \rangle$}. Vertex 4 is not visited and is in *SA*(0): the new path $\langle 2, 0, 1, 4 \rangle$ is added to *Paths*.

6 Experimental results

In this section, we present the results obtained by the CUDA implementation of the two proposed algorithms and the CPU multi-core implementation of algorithm V-FEC using Posix Threads library in C++. We compare the results to those of Johnson's algorithm. We use the Python implementation of Johnson's algorithm in the NetworkX package[1].

6.1 Experiment settings

Experiments were conducted on a GPU station with 32 GB RAM, with an Intel Xeon Silver 4216 2.10GHz CPU. The GPU card is an Nvidia GeForce RTX 2080 Ti with 11GB of GDDR6 VRAM and 4,352 CUDA cores, CUDA Version 11.0.

---

[1] https://networkx.org/.

Table 1 Experimental comparison of V-FEC, T-FEC, mCore-V-FEC (using 8 CPU cores) and the algorithm by Johnson for synthetic graphs

Graph	|V|	|E|	|C|	Execution time (ms)			
				V-FEC	T-FEC	mCore-V-FEC	Johnson
Graph-40-1-1-1	40	113	54121	24.66	**19.72**	244,98	1140.63
Graph-40-1-1-0	40	70	1676	**9.83**	10.54	12,099	85.68
Graph-50-1-1-1	50	149	120519	40.21	**28.71**	404,36	2779.66
Graph-50-1-1-0	50	93	5767	**11.86**	18.31	33,66	176.45
Graph-60-1-1-1	60	186	4942125	1547.52	**1161.07**	21520,82	121862.10
Graph-60-1-1-0	60	120	164783	**76.65**	386.73	791,39	8398.95
Graph-70-1-1-1	70	225	281100	209.36	**129.30**	2349,22	7143.13
Graph-70-1-1-0	70	147	1820687	**764.34**	3712.17	8873,87	96148.63
Graph-80-1-1-1	80	264	329844	124.04	**93.69**	1485,44	7043.80
Graph-80-1-1-0	80	174	56067879	**18200.95**	–	224751,33	3462429.93
Graph-90-1-1-1	90	304	8257415	2730.90	**2105.35**	39822,1	206860.77
Graph-100-1-1-1	100	345	46865151	10839.7	**8891.98**	171796,33	1186575.70

The bold that indicates that the values is the best execution times

6.2 Results

We use two sets of data. The first set contains synthetic graphs generated using the R3MAT graph generator [36]. It generates graphs that resemble the properties observed in real-world graphs following a power law distribution. The description of these graphs and the execution times of our implementation compared with the execution times of Johnson's algorithm and the parallel multi-core implementation of V-FEC are presented in Table 1.

Column $|C|$ represents the number of elementary circuits found by the compared algorithms. For each graph, we ran the algorithms 20 times and calculate the average time taken. The execution times are presented in milliseconds.

For the graph "graph-80-1-1-0," algorithm T-FEC could not enumerate all circuits of the graph. This is due to the large number of intermediate results (paths) that could not fit in the GPU memory. As an alternative, we succeeded to list all the circuits of size equal or less than 16.

For some graphs, where the number of circuits is very important, it is not possible to enumerate all the circuits, due to memory limitations. The search for elementary circuits with Johnson's algorithm for these graphs result in an out of memory error (OOM). On the other hand, our algorithms proceed by levels. At each iteration, circuits of length l are found. This gives the possibility to enumerate circuits of a given length without having to enumerate all the circuits of the graph. For such graphs it is possible to proceed by enumerating the circuits until a maximum depth is reached. Table 3 shows these results.

In addition to the synthetic graphs generated by the R3MAT model [36], we use real-world datasets from four different sources: Konect [37] (moreno_taro, moreno_sheep), SNAP Library [38] (p2p-Gnutella04, p2p-Gnutella08 and

Table 2 Experimental comparison of V-FEC, T-FEC and the algorithm by Johnson for real-world graphs

Graph	$|V|$	$|E|$	$|C|$	Count circuits		Enumerate circuits		
				V-FEC	T-FEC	V-FEC	T-FEC	Johnson
Moreno_taro	22	78	21671	50.66	35.83	51.52	**36.24**	1618.73
Moreno_sheep	28	250	19727891	27740.69	27450.56	**26723.50**	26982.67	703917.13
EPA	4772	8965	142	6.91	7.48	**7.14**	**7.38**	48.20
p2p-Gnu-tella04	10879	39994	l = 6: 1768	452.17	l = 5	**483.95**	l = 5	OOM
p2p-Gnu-tella08	6301	20777	l = 7: 15732	379.44	l = 6	**371.91**	l = 6	OOM
Email-Eu-core	1005	24929	l = 3: 124765	154.11	213.42	**166.26**	214.28	OOM
Baywet	128	2106	l = 8: 7322229	l = 6	2673.55	l = 6	**l = 7**	OOM

The bold that indicates that the values is the best execution times

email-Eu-core), networkrepository [39] (EPA) and Pajek [40] (Baywet). Characteristics of these graphs are given in Table 2.

We adapt our algorithms to count circuits without enumerating them. This significantly reduces the memory storage needed since we do not keep track of all the visited vertices. The *Paths* matrix is replaced by vectors that stores the *source-vertex* and the last vertex for each path (in addition to the *in-vertex* in the triplets approach). This permits to explore larger graphs and to explore more levels as shown in Table 3.

6.3 Analysis

The results presented in Tables 1, 2, and 3 show that our algorithms can significantly reduce the running time to find all elementary circuits compared to Johnson's algorithm and the multi-core implementation of V-FEC. The speed-ups over the algorithm of Johnson vary from 8.13 (graph-40-1-1-0), to 190.23 (graph-80-1-1-0), and increase with the number of circuits. The recorded speed-ups for the parallel multi-core implementation are between 1.15 (graph-40-1-1-0) and 19.32 (graph-100-1-1-1). Figure 5 shows the speed-ups achieved by V-FEC and T-FEC compared to the algorithm by Johnson.

We analyze the effect of the *SCC* decomposition phase on the overall processing time by comparing the execution time of both approaches with and without decomposing the graph into *SCCs*. Detailed results for V-FEC and T-FEC algorithms are listed in Tables 4 and 5. V-FEC + SCC and T-FEC + SCC present execution times of V-FEC and T-FEC using *SCC* decomposition, while V-FEC - SCC and T-FEC - SCC present the execution times of the algorithms without the *SCC* decomposition phase. The results show that the decomposition of the graph into *SCCs* can reduce processing time compared to when the graph is not decomposed into *SCCs*. In fact, finding the *SCCs* of the graph can reduce the number of

Table 3 The results per level showing the maximum depth for each algorithm, the maximum number of circuits and the number of circuits for each level

Graph	|V|	|E|	Max circuits	Count circuits		Enumerate circuits		Execution time / Circuits
				V-FEC	T-FEC	V-FEC	T-FEC	
Graph-90-1-1-0	90	203	86533017	**l = 16**	l = 14	l = 13	l = 12	T = 40191.56 C = {26 73 316 1203 4457 15307 49849 150785 424416 1106250 2667819 5945511 12234120 23230600 40702285 }
Graph-100-1-1-0	100	235	328908432	**l = 19**	l = 13	l = 14	l = 12	T = 97980.37 C = {25 104 415 1564 5837 20236 64949 193905 536577 1365891 3195804 6859242 13478597 24201006 39647343 59191874 80492814 99652249}
Graph-200-1-0-0	200	526	2746022	**l = 17**	l = 16	l = 15	l = 15	T = 15829.14 C = {2 7 26 57 98 261 644 1482 3511 8600 20790 50882 120272 286170 672475 1580745 }
Graph-300-1-0-0	300	1256	501259	**l = 13**	**l = 13**	l = 12	l = 12	T = 4937.91 C = {5 13 36 90 204 607 1671 4773 13969 40080 114002 325809 }

The bold that indicates that the values is the best execution times

Fig. 5 Experimental comparison of Johnson's algorithm and mCore-V-FEC with V-FEC and T-FEC algorithms

Table 4 Experimental results for the impact of the SCC graph decomposition in V-FEC

Graph	$|V|$	$|SCC|$	Time (ms) V-FEC + SCC			Time (ms) V-FEC - SCC	
			SCC	Exploration	Total	Exploration	Total
p2p-Gnutella04	10879	4317	3.409	419.01	**448.18**	1609.35	1834.1
p2p-Gnutella08	6301	2068	1.819	351.33	**369.19**	804.23	859.62
Email-Eu-core	1005	803	1.082	130.93	**163.44**	133.67	**150.84**
Baywet	128	103	0.100	494.69	**508.25**	1381.08	1487.08
EPA	4772	{2,11,3,10,4,2,5,2,2,6,3,2, 2,2,2,3,6,2,4,2,2,2,2,2,2,2}	0.990	5.08	**6.99**	48.54	52.35
Moreno_taro	22	22	0.015	51.91	**53.31**	51.18	**52.48**
Moreno_sheep	28	{2, 22}	0.024	26778.5	26927.6	–	–
Graph-40-1-1-1	40	12	0.019	23.61	**24.72**	23.98	**24.90**
Graph-40-1-1-0	40	17	0.019	9.14	**9.83**	10.26	11.06
Graph-50-1-1-1	50	14	0.022	40.13	**41.79**	41.26	**42.71**
Graph-50-1-1-0	50	23	0.020	10.51	**11.29**	11.37	**12.09**
Graph-60-1-1-1	60	18	0.021	1474.35	**1502.46**	1522.29	**1551.51**
Graph-60-1-1-0	60	30	0.022	74.70	**77.08**	92.35	94.63
Graph-70-1-1-1	70	16	0.05	204.70	**207.88**	208.76	**211.74**
Graph-70-1-1-0	70	30	0.025	722.09	**739.82**	1116.81	1135.71
Graph-80-1-1-1	80	13	0.027	125.01	**127.63**	122.39	**124.85**
Graph-80-1-1-0	80	36	0.058	16656.3	**17507.3**	29613	30959.6
G_02_03_02_1	64	51	0.025	3350.25	**3365.48**	4698.41	4723.84
G_02_03_02_3	64	52	0.024	3216.52	**3232.97**	4753.63	4779.21
G_06_01_015_1	64	32	0.025	1430.28	**1463.19**	2307.59	2348.33
G_06_01_015_2	64	33	0.026	984.37	**1003.3**	1507.08	1527.84
G_01_02_03_5	64	38	0.027	1593.95	**1601.35**	6859.72	6883.79

The bold that indicates that the values is the best execution times

Table 5 Experimental results for the impact of the SCC graph decomposition in T-FEC

Graph	Time (ms) T-FEC + SCC					Time (ms) T-FEC - SCC			
	Triplets	SCC	T-time	Exploration	Total	Triplets	T-time	Exploration	Total
p2p-Gnutella04	29576	3.511	0.807	220.87	**230.42**	69411	0.95	466.27	476.12
p2p-Gnutella08	17656	1.963	0.932	336.11	346.54	41752	–	–	–
Email-Eu-core	585882	1.133	37.954	156.36	212.24	609644	39.535	160.07	**215.75**
Baywet	1928	0.108	0.346	424.99	**439.27**	2547	0.397	1128.53	1142.53
EPA	126	1.014	0.027	4.27	**7.39**	6467	2.243	21.22	24.24
Moreno_taro	54	0.019	0.028	35.43	**36.58**	54	0.029	36.34	37.21
Moreno_sheep	323	0.026	0.114	26629.4	26758.8	497	–	–	–
Graph-40-1-1-1	131	0.020	0.051	18.64	**19.87**	131	0.049	18.59	19.62
Graph-40-1-1-0	85	0.019	0.045	9.69	**10.59**	128	0.053	13.77	14.51
Graph-50-1-1-1	125	0.023	0.063	26.90	**28.52**	125	0.061	27.68	29.04
Graph-50-1-1-0	152	0.022	0.079	16.83	**17.79**	223	0.089	25.39	26.26
Graph-60-1-1-1	179	0.026	0.088	1137.43	**1163.57**	179	0.085	1155.23	1180.97
Graph-60-1-1-0	220	0.025	0.070	374.19	**377.08**	272	0.073	465.84	468.48
Graph-70-1-1-1	156	0.028	0.052	138.01	**141.21**	156	0.051	132.68	134.97
Graph-70-1-1-0	297	0.027	0.091	3665.96	**3684.68**	388	0.132	6509.8	6528.48
Graph-80-1-1-1	183	0.032	0.055	91.48	**93.93**	183	0.053	90.86	93.07
Graph-80-1-1-0	427	0.032	0.195	16791.4	16939.6	598	–	–	–
G_02_03_02_1	107	0.028	0.034	3200.82	**3212.58**	127	0.029	4679.11	4691.74
G_02_03_02_3	110	0.027	0.025	2551.72	**2563.72**	122	0.027	4413.85	4427.28
G_06_01_015_1	280	0.032	0.079	4770.51	**4797.8**	351	0.086	7998.3	8029.54
G_06_01_015_2	252	0.032	0.075	3290.73	**3307.99**	305	0.102	5596.88	5624.02
G_01_02_03_6	135	0.027	0.031	24933.2	**25784.7**	139	0.038	33428.3	34310.3
G_01_02_03_5	102	0.033	0.031	723.22	**726.70**	135	0.038	3059.05	3063.8

The bold that indicates that the values is the best execution times

vertices to be explored and avoid fruitless paths, thus, minimizing the time of the exploration phase, which justify the use of the preprocessing phase in our algorithms. Experiments demonstrate that time achieved by the decomposition phase is negligible compared to the overall running time, as presented in Fig. 6.

We note that while V-FEC algorithm performs better in some graphs, T-FEC gives better results in others. In an attempt to identify what variables affect the results, we run a series of experiments on synthetic graphs, using the R-Mat graph generator model [41]. The generator takes as an input $|V|$, $|E|$ and the parameters a b c and d which represents the probabilities of an edge falling into partitions. We fix $|V| = 64$ and vary $|E|$ and the values of a, b, c and d and observe the behavior of both V-FEC and T-FEC approaches. The results are presented in Table 6. We observe that the degree distribution of the graph is a relevant parameter in the algorithms' performance. We divide Table 6 into three sets based on the degree distribution. We represent the execution times of both approaches on these graphs in Fig. 7. The second set in the table represent graphs where the degrees of vertices are almost uniform. For these graphs, we remark that V-FEC and T-FEC

Table 6 Results of the impact of the graph degree distribution on execution time

Graph	$|V|$	$|E|$	$|C|$	V-FEC	Count circuits time (ms)	
					Triplets	T-FEC
G_02_02_045_2	64	147	816347	**1827.96**	65	2102.32
G_01_02_03_4	64	179	29404	**342.46**	88	**331.87**
G_01_02_03_5	64	179	255064	1708.35	102	**748.35**
G_01_02_03_6	64	179	22280272	67482.4	135	**26221.7**
G_02_02_045_3	64	179	3558059	13296.96	102	**8240.42**
G_02_02_045_4	64	230	6903928 (l = 16)	31656.84	175	**14577.16**
G_02_03_02_1	64	147	770263	**3503.28**	107	3298.44
G_02_03_02_3	64	147	783545	3415.25	110	**2675.1**
G_035_015_035_2	64	147	938758	**4067.80**	91	**4123.50**
G_025_025_025_2	64	147	2873222	**8648.52**	99	**8105**
G_02_03_02_4	64	179	32980789 (l = 25)	78766.44	157	**79135.08**
G_02_03_02_5	64	179	11348731	**28826.70**	127	**25192.58**
G_02_03_02_6	64	179	37618516 (l = 28)	**68731.56**	159	87724.42
G_025_025_025_3	64	179	17920050 (l = 25)	81767.37	118	**43598.87**
G_035_015_035_3	64	179	22603921 (l = 26)	88475.50	134	**72638.87**
G_025_025_025_4	64	230	6669043 (l = 15)	**15742.1**	277	**15594.58**
G_035_015_035_4	64	230	4563399 (l = 16)	**12570.66**	215	14345.82
G_05_02_015_2	64	147	988967	**503.95**	241	2354.7
G_06_01_015_1	64	147	3180582	**1536.35**	280	5055.02
G_06_01_015_2	64	147	1766087	**1060.39**	252	3434.8
G_06_01_015_3	64	147	419801	806.74	172	**748.73**
G_05_02_015_3	64	179	42767911 (l = 21)	**17457.02**	289	84379.7
G_06_01_015_4	64	179	48576117 (l = 19)	**24837.16**	451	72551.22
G_06_01_015_5	64	179	35867587	**22945.82**	345	56180.6
G_06_01_015_6	64	179	16792958	**10489.28**	324	20738
G_05_02_015_4	64	230	13016638 (l = 13)	**4715.12**	506	17352.32

The bold that indicates that the values is the best execution times


(a) V_FEC

(b) T_FEC

Fig. 6 Experimental results for the impact of the SCC graph decomposition

(a) High degree vertices having higher Ids

(b) High degree vertices having lower Ids

(c) Uniform degree distribution

Fig. 7 Results of the impact of the graph degree distribution

algorithms are equivalent (Fig. 7c). The first and the third set are graphs with few vertices having higher degree than other vertices with a difference in the ordering of these vertices. In the first set the high degree vertices are processed in the end, while in the last set they are processed in the beginning. We observe that V-FEC algorithm outperforms the T-FEC algorithm for the graphs in the last set of the table (Fig. 8b). In this case, the number of triplets that are generated is important for those vertices with high degree. Indeed, a lot of combinations are generated, which means more paths to be explored and some of them are fruitless. On the contrary, the triplets approach outperforms the vertex-based one for the first set of graphs (Fig. 7a).

To further analyze the impact of the distribution order of the vertices' degree, we use a graph labeling function based on the in and out degrees of the vertices of the graph. We know that the more *in-neighbors* and *out-neighbors* a vertex has the more probable it is to belong to circuits. Based on that, we order the vertices according to the product of their in and out degrees. We use two graph labeling functions: the first one l_{Asc} organizes the vertices in an increasing order of *in-degree * out-degree*. The second function l_{Desc}, lists the vertices in a decreasing order of the vertices' *in-degree * out-degree*. We apply this labeling to some of

Fig. 8 Results of the impact of the degree distribution order

the graphs in Table 6 and run algorithm V-FEC. We compare the results with and without the labeling functions. The results are represented in Fig. 8. We observe that algorithm V-FEC gives better results when the l_{desc} label is applied. Indeed, by visiting vertices with more neighbors in early iterations we have higher probability to find circuits earlier and without exploring much of the fruitless paths. Experimentally observing, we can tell that the degree distribution order is a crucial parameter in the algorithm execution.

7 Conclusion

In this paper, we proposed parallel algorithms for enumerating all elementary circuits of a directed graph. Algorithms V-FEC and T-FEC first detect whether a circuit exists, by searching for strongly connected components, then, explore possible paths in parallel to find elementary circuits. In addition, they enumerate circuits of a given length and circuits going through a given vertex. We have provided theoretical guarantees on the correctness of V-FEC and T-FEC. The presented algorithms were implemented and tested in a GPU environment. To the best of our knowledge, these are the first parallel GPU-based algorithms for finding all the elementary circuits of a graph. Conducted experiments showed a significant improvement in execution times of V-FEC and T-FEC compared with the algorithm by Johnson, due to the

massive parallelism offered by the GPU. The speed-up over the sequential algorithm was from ≈ 8 to 190 times.

We should note that the size of the GPU memory can be a limiting factor that prevents from handling large graphs with millions of circuits. For future work, we intend to scale to larger graphs by working on multi-GPU approaches and parallel strategies on cluster of GPUs.

Data availability All data generated or analyzed during this study have been deposited in the Open Science Framework repository (OSF) (https://osf.io/74dve/?view_only=a5291e6e898843c48686b06d49c03ed5).

References

1. Fronzetti Colladon A, Remondi E (2017) Using social network analysis to prevent money laundering. Expert Syst Appl 67:49–58. https://doi.org/10.1016/j.eswa.2016.09.029
2. Dunne JA, Williams RJ, Martinez ND (2002) Network structure and biodiversity loss in food webs: robustness increases with connectance. Ecol Lett 5(4):558–567. https://doi.org/10.1046/j.1461-0248.2002.00354.x
3. Bodaghi A, Teimourpour B (2018) Automobile insurance fraud detection using social network analysis. Applications of data management and analysis case studies in social networks and beyond. Springer, Berlin, pp 11–16. https://doi.org/10.1007/978-3-319-95810-1_2
4. Safar M, Mahdi K, Kassem A (2009) Universal cycles distribution function of social networks. In: 2009 First International Conference on Networked Digital Technologies, pp 354–359. https://doi.org/10.1109/NDT.2009.5272805
5. Giscard P-L, Rochet P, Wilson RC (2017) Evaluating balance on social networks from their simple cycles. J Complex Netw 5(5):750–775. https://doi.org/10.1093/comnet/cnx005
6. Kwon Y-K, Cho K-H (2007) Analysis of feedback loops and robustness in network evolution based on boolean models. BMC Bioinform 8(1):1–9. https://doi.org/10.1186/1471-2105-8-430
7. Klamt S, von Kamp A (2009) Computing paths and cycles in biological interaction graphs. BMC Bioinform 10(1):1–11. https://doi.org/10.1186/1471-2105-10-181
8. Chitturi B, Bein D, Grishin NV (2010) Complete enumeration of compact structural motifs in proteins. In: Proceedings of the International Symposium on Biocomputing, pp 1–8. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/1722024.1722047
9. Parasar M, Farrokhbakht H, Enright Jerger N, Gratz PV, Krishna T, San Miguel J (2020) Drain: deadlock removal for arbitrary irregular networks. In: 2020 IEEE International Symposium on High Performance Computer Architecture (HPCA), pp 447–460. https://doi.org/10.1109/HPCA47549.2020.00044
10. Tiernan JC (1970) An efficient search algorithm to find the elementary circuits of a graph. Commun ACM 13(12):722–726. https://doi.org/10.1145/362814.362819
11. Tarjan R (1973) Enumeration of the elementary circuits of a directed graph. SIAM J Comput 2(3):211–216. https://doi.org/10.1137/0202017
12. Johnson DB (1975) Finding all the elementary circuits of a directed graph. SIAM J Comput 4(1):77–84. https://doi.org/10.1137/0204007
13. Lu W, Zhao Q, Zhou C (2018) A parallel algorithm for finding all elementary circuits of a directed graph. In: 2018 37th Chinese Control Conference (CCC), pp 3156–3161. https://doi.org/10.23919/ChiCC.2018.8482857. IEEE

14. Giscard P-L, Kriege N, Wilson RC (2019) A general purpose algorithm for counting simple cycles and simple paths of any length. Algorithmica 81(7):2716–2737. https://doi.org/10.1007/s00453-019-00552-1

15. Gupta A, Suzumura T (2021) Finding all bounded-length simple cycles in a directed graph. CoRR. https://doi.org/10.48550/ARXIV.2105.10094

16. Mahdi F, Safar M, Mahdi K (2011) Detecting cycles in graphs using parallel capabilities of gpu. In: International Conference on Digital Information and Communication Technology and Its Applications, pp 193–205. Springer

17. Rungta S, Srivastava S, Yadav US, Rastogi R (2014) A comparative analysis of new approach with an existing algorithm to detect cycles in a directed graph. In: ICT and Critical Infrastructure: Proceedings of the 48th Annual Convention of Computer Society of India-Vol II, Springer, pp 37–47

18. Rocha RC, Thatte BD (2015) Distributed cycle detection in large-scale sparse graphs. In: Proceedings of Simpósio Brasileiro de Pesquisa Operacional (SBPO'15), pp 1–11

19. Cui H, Niu J, Zhou C, Shu M (2017) A multi-threading algorithm to detect and remove cycles in vertex-and arc-weighted digraph. Algorithms 10(4):115

20. Xu-guang L, Da-ming Z (2010) An approximation algorithm for the shortest cycle in an undirected unweighted graph. In: 2010 International Conference on Computer, Mechatronics, Control and Electronic Engineering, vol. 1, pp 297–300. IEEE

21. Yuster R (2011) A shortest cycle for each vertex of a graph. Inf Process Lett 111(21–22):1057–1061

22. Karimi M, Banihashemi AH (2012) Message-passing algorithms for counting short cycles in a graph. IEEE Trans Commun 61(2):485–495

23. Paulusma D, Yoshimito K (2007) Cycles through specified vertices in triangle-free graphs. Discuss Math Graph Theory 27(1):179–191

24. Li B, Zhang S (2011) Heavy subgraph conditions for longest cycles to be heavy in graphs. arXiv preprint arXiv:1109.4675

25. Li B, Xiong L, Yin J (2016) Large degree vertices in longest cycles of graphs, i. Discuss Math Graph Theory 36(2)

26. Li B, Xiong L, Yin J (2019) Large degree vertices in longest cycles of graphs, ii. Electron J Graph Theory Appl (EJGTA) 7(2):277–299

27. Weinblatt H (1972) A new search algorithm for finding the simple cycles of a finite directed graph. J ACM 19(1):43–56. https://doi.org/10.1145/321679.321684

28. Liu H, Wang J (2006) A new way to enumerate cycles in graph. In: Advanced Int'l Conference on Telecommunications and Int'l Conference on Internet and Web Applications and Services (AICT-ICIW'06), pp 57–57. https://doi.org/10.1109/AICT-ICIW.2006.22. IEEE

29. Sankar K, Sarad A (2007) A time and memory efficient way to enumerate cycles in a graph. In: 2007 International Conference on Intelligent and Advanced Systems, pp 498–500. https://doi.org/10.1109/ICIAS.2007.4658438. IEEE

30. Reif JH (1985) Depth-first search is inherently sequential. Inf Process Lett 20(5):229–234. https://doi.org/10.1016/0020-0190(85)90024-9

31. Qing Z., Yuan L, Chen Z, Lin J, Ma G (2020) Efficient parallel cycle search in large graphs. In: International Conference on Database Systems for Advanced Applications. Springer, pp 349–367

32. Williamson, EABSG. Lists, Decisions and Graphs. S. Gill Williamson. https://books.google.dz/books?id=vaXv_yhefG8C

33. Tarjan R (1972) Depth-first search and linear graph algorithms. SIAM J Comput 1(2):146–160. https://doi.org/10.1137/0201010

34. Harris M, Sengupta S, Owens JD (2007) Parallel prefix sum (scan) with cuda. GPU Gems 3(39):851–876

35. Gui C-Y, Zheng L, He B, Liu C, Chen X-Y, Liao X-F, Jin H (2019) A survey on graph processing accelerators: challenges and opportunities. J Comput Sci Technol 34(2):339–371. https://doi.org/10.1007/s11390-019-1914-z

36. Angles R, Paredes R, García R (2020) R3MAT: a rapid and robust graph generator. IEEE Access 8:130048–130065. https://doi.org/10.1109/ACCESS.2020.3009577

37. Kunegis J (2013) KONECT – The Koblenz Network Collection. http://konect.cc/networks/

38. Leskovec J, Krevl A (2014) SNAP Datasets: stanford large network dataset collection. http://snap.stanford.edu/data

39. Rossi RA, Ahmed NK (2015) The network data repository with interactive graph analytics and visualization. http://networkrepository.com
40. Batagelj V, Mrvar A (2006) Pajek datasets. http://vlado.fmf.uni-lj.si/pub/networks/data/
41. Chakrabarti D, Zhan Y, Faloutsos C (2004) R-mat: a recursive model for graph mining. In: Proceedings of the 2004 SIAM International Conference on Data Mining (SDM), pp 442–446. https://doi.org/10.1137/1.9781611972740.43. SIAM

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.