



Contents lists available at ScienceDirect

Journal of Computer and System Sciences

journal homepage: www.elsevier.com/locate/jcssOn maximal k -edge-connected subgraphs of undirected graphsLoukas Georgiadis^{a,1}, Giuseppe F. Italiano^{b,2}, Evangelos Kosinas^{a,*,3},
Debasish Pattanayak^b^a University of Ioannina, Greece^b LUISS Guido Carli University, Roma, Italy

ARTICLE INFO

Article history:

Received 19 September 2023

Accepted 31 March 2026

Available online 9 April 2026

Keywords:

Graph algorithms

Graph connectivity

Data structures

Dynamic algorithms

ABSTRACT

We provide the following new results on maximal k -edge-connected subgraphs of undirected graphs. (1) A general framework for maintaining the maximal k -edge-connected subgraphs upon insertions of edges or vertices, by successively partitioning the graph into its k -edge-connected components. This defines a decomposition tree, which can be maintained by using algorithms for the incremental maintenance of the k -edge-connected components as black boxes at every level of the tree. As a concrete application of this framework, we provide two algorithms for the incremental maintenance of the maximal 3-edge-connected subgraphs. These algorithms allow for vertex and edge insertions, interspersed with queries asking whether two vertices belong to the same maximal 3-edge-connected subgraph, and there is a trade-off between their time- and space-complexity. Specifically, the first algorithm has $O(m\alpha(m, n) + n^2 \log^2 n)$ total running time and uses $O(n)$ space, where m is the number of edge insertions and queries, and n is the total number of vertices inserted starting from an empty graph. The second algorithm performs the same operations in faster $O(m\alpha(m, n) + n^2 \alpha(n, n))$ time in total, using $O(n^2)$ space. (2) We provide efficient constructions of (almost) sparse spanning subgraphs that have the same maximal k -edge-connected subgraphs as the original graph. We refer to such subgraphs as k -certificates. We use those certificates to speed up the computation of the maximal k -edge-connected subgraphs in the static and the fully-dynamic setting. (3) Finally, we give a simple reduction for computing the maximal k -edge-connected subgraphs to a fully dynamic mincut algorithm.

© 2026 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

* Corresponding author.

E-mail addresses: loukas@cs.uoi.gr (L. Georgiadis), gitaliano@luiss.it (G.F. Italiano), ekosinas@cs.uoi.gr (E. Kosinas), debasish@iiti.ac.in (D. Pattanayak).¹ Research supported by the Hellenic Foundation for Research and Innovation (H.F.R.I.), under the "First Call for H.F.R.I. Research Projects to support Faculty members and Researchers and the procurement of high-cost research equipment grant", Project FANTA (eFFicient Algorithms for NeTwork Analysis), number HFRI-FM17-431.² Partially supported by MUR, the Italian Ministry of University and Research, under PRIN Project n. 2022TS4Y3N - EXPAND (scalable algorithms for EXploratory Analyses of heterogeneous and dynamic Networked Data).³ The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the 3rd Call for HFRI PhD Fellowships (Fellowship Number: 6547).




Fig. 1. A 2-edge-connected graph G , its maximal 3-edge-connected subgraphs, and its 3-edge-connected components. Note that while any two vertices x_i and y_j are 3-edge-connected, they do not belong to the same maximal 3-edge-connected subgraph.

1. Introduction

A dynamic graph algorithm aims to maintain the solution of a given problem after each update faster than recomputing it from scratch. An algorithm is *fully dynamic* if it supports both insertions and deletions of edges, while it is *incremental* (resp. *decremental*) if it only supports insertions (resp. deletions) of edges.

Let $G = (V, E)$ be a connected undirected multigraph with m edges and n vertices. We let $V(G)$ and $E(G)$ denote the vertex set and the edge set of G , respectively, i.e., $V(G) = V$ and $E(G) = E$. Let $S \subseteq V$ be any subset of vertices of G . The *subgraph induced by S* , denoted by $G[S]$, is defined as the graph having vertex set S and edge set $\{e \in E \mid \text{both endpoints of } e \text{ lie in } S\}$. Let X, Y be two subsets of V . $[X, Y]$ denotes the set of edges that have one endpoint in X and the other in Y . In other words, $[X, Y] = \{(x, y) \in E \mid x \in X \text{ and } y \in Y\}$. Let $S \subset V$ be a non-empty (proper) subset of vertices of G , and let $\bar{S} = V \setminus S \neq \emptyset$. If both $G[S]$ and $G[\bar{S}]$ are connected, then $[S, \bar{S}]$ is called an *edge cut* of G , i.e., a minimal set of edges of G such that its removal disconnects the graph. An edge cut of cardinality k is called a *k -edge cut* of G . As a special case, 1-edge cuts are also called *bridges*. G is said to be *k -edge-connected* if it contains no k' -edge cuts, for $k' < k$.

Two vertices of G are said to be *k -edge-connected* if there are k edge-disjoint paths between them. The *edge-connectivity* of a pair of vertices $x, y \in V$, denoted as $\lambda_G(x, y)$ (or simply $\lambda(x, y)$, when G is clear from context), is the maximum k such that x and y are k -edge-connected. The relation of “ k -edge-connectivity” between pairs of vertices of G is an equivalence relation on $V(G)$ (see, e.g., [46]). The equivalence classes induced by this relation are called the *k -edge-connected components* of G . Thus, a *k -edge-connected component* of G is a maximal subset of vertices $C \subseteq V$ such that any pair of vertices in C is k -edge-connected. Notice that every k -edge-connected component lies entirely within a k' -edge-connected component, for every $k' < k$.

Now let $S \subseteq V$ be a subset of vertices in G . We say that the induced subgraph $G[S]$ is a *maximal k -edge-connected subgraph* of G if (1) $G[S]$ is k -edge-connected and (2) no proper superset of S has this property. Unlike 2-edge connectivity, for $k \geq 3$ the k -edge-connected components of G do not necessarily correspond to maximal k -edge-connected subgraphs. Indeed, for $k \geq 3$, two vertices in the subgraph induced by a k -edge-connected component may not be k -edge-connected in this subgraph, as some of the k edge-disjoint paths may go outside of the component; see Fig. 1. Also, Fig. 2 shows a graph where almost all its vertices are 3-edge-connected but has only trivial maximal 3-edge-connected subgraphs. Notice that, if S is a subset of V , then $\lambda_{G[S]}(x, y) \leq \lambda_G(x, y)$ for any pair of vertices $x, y \in S$, since every path in $G[S]$ is also a path in G . Thus, every maximal k -edge-connected subgraph lies within the subgraph induced by a k -edge-connected component.

Determining or testing various notions of edge connectivity of undirected graphs, as well as computing edge-connected components or subgraphs, is a fundamental graph problem, and indeed k -edge connectivity has received much attention in the literature due to its importance in applications. As a matter of fact, finding the maximal k -edge-connected subgraphs of a graph is of significant interest in several areas, such as in the field of databases, social networks, graph visualization etc. [3,7,9,8,42,52,58,60] (some of those papers refer to the maximal k -edge-connected subgraphs as *k -edge-connected components*).




Fig. 2. A 2-edge-connected graph G with only trivial maximal 3-edge-connected subgraphs, despite that almost all vertices are 3-edge-connected.

The problem of computing maximal k -edge-connected subgraphs appears to be harder than computing k -edge-connected components. In more detail, it is known how to compute the k -edge-connected components in linear time for $k \leq 4$ [25, 27, 30, 36, 44, 45, 53, 57]. On the other hand, linear-time bounds for computing the maximal k -edge-connected subgraphs are known only in the trivial case $k \leq 2$, simply because in these cases the k -edge-connected components coincide with the maximal k -edge-connected subgraphs. As mentioned in [10], the maximal 3-edge-connected subgraphs can be computed with a simple-minded algorithm in $O(mn)$ time in the worst case. Henzinger et al. [31] provided a deterministic algorithm for computing the maximal k -edge-connected subgraphs, for constant k , that runs in $O(n^2 \log n)$ time. Chechik et al. [10] presented an algorithm with running time $O(k^{O(k)}(m + n \log n)\sqrt{n})$ ($O(m\sqrt{n})$ for $k = 3$), which improves the bound of [31] for sparse graphs. Forster et al. [21] gave a Las Vegas randomized algorithm for computing the maximal k -edge-connected subgraphs in $O(km \log^2 n + k^3 n \sqrt{n} \log n)$ expected running time. Recently, Nalam and Saranurak [48] provided a randomized algorithm for computing the maximal k -edge-connected subgraphs in *weighted* undirected graphs in $\tilde{O}(m \cdot \min\{m^{3/4}, n^{4/5}\})$ time. Also, very recently, Saranurak and Yuan [51] gave a deterministic $O(m + n^{1+o(1)})$ -time algorithm for computing the maximal k -edge-connected subgraphs of an undirected graph, for any $k = \log^{o(1)} n$. To achieve this result, Saranurak and Yuan [51] provide a black-box reduction to any decremental algorithm for answering k -edge-connectivity queries, and then they apply the fully dynamic algorithm of Jin and Sun [37]. In the special case of planar graphs, Holm et al. [35] showed how to compute the maximal 3-edge-connected subgraphs in $O(n)$ time.

1.1. Overview of our results

In this paper we present several new results on maximal k -edge-connected subgraphs of undirected graphs. In particular, we provide the following results.

- (1) A general framework for maintaining the maximal k -edge-connected subgraphs upon insertions of edges or vertices, by successively partitioning the graph into its k -edge-connected components. This defines a decomposition tree, which can be maintained by using algorithms for the incremental maintenance of the k -edge-connected components as black boxes at every level of the tree. As an application of our new framework, we provide two algorithms for the incremental maintenance of the maximal 3-edge-connected subgraphs, which constitute the main results of this article. These algorithms allow for vertex and edge insertions, interspersed with queries asking whether two vertices belong to the same maximal 3-edge-connected subgraph, and provide a trade-off between time- and space-complexity. The first algorithm has $O(m\alpha(m, n) + n^2 \log^2 n)$ total running time and uses asymptotically optimal $O(n)$ space, where m is the number of edge insertions and queries, and n is the total number of vertices inserted starting from an empty graph. The second algorithm improves the total running time to $O(m\alpha(m, n) + n^2 \alpha(n, n))$ (i.e., almost optimal for dense graphs) at the expense of using $O(n^2)$ space. We note that those are the first incremental algorithms for this problem, and thus provide significant improvements over recomputing the maximal 3-edge-connected subgraphs after every insertion.
- (2) Building up on results from Benczúr and Karger [5], we provide efficient constructions of (almost) sparse spanning subgraphs that have the same maximal k -edge-connected subgraphs as the original graph. We refer to such subgraphs as *k-certificates*. These are useful in speeding up computations involving the maximal k -edge-connected subgraphs in dense undirected graphs. In particular, we use those certificates to speed up the computation of the maximal k -edge-connected subgraphs. As another application of our *k-certificates*, we use them to provide a fully dynamic algorithm for maintaining information about the maximal k -edge-connected subgraphs for fixed k . Using the sparsification technique of Eppstein et al. [20], we show the existence of a fully dynamic algorithm with update times that are independent of the number of edges and with constant time for maximal k -edge-connected subgraph queries.
- (3) Finally, we give a simple reduction for the computation of the maximal k -edge-connected subgraphs to fully dynamic mincut. By using Thorup's fully dynamic mincut algorithm [56], we obtain a deterministic algorithm that computes the maximal k -edge-connected subgraphs in $O(m + k^{O(1)} n \sqrt{n} \log^{O(1)} n)$ time, for $k = \log^{O(1)} n$.

We believe that our main technical contribution is given by (1). To achieve this result, we provide a structural characterization of the maximal 3-edge-connected subgraphs of an undirected graph by introducing a decomposition tree \mathcal{T} into maximal 3-edge-connected subgraphs, and show how to maintain it efficiently under edge and vertex insertions. \mathcal{T} is a

rooted tree whose root corresponds to the whole graph G and is defined recursively by computing the (subgraphs induced by the) k -edge-connected components, for $k \in \{1, 2, 3\}$, of the graphs of the previous level. We proceed recursively in this decomposition until we reach a graph that is 3-edge-connected, which is a maximal 3-edge-connected subgraph of G (and a leaf of \mathcal{T}). Therefore, the nodes of \mathcal{T} correspond to subgraphs of G , and the parent relation is given by (vertex) set inclusion. (See Section 3.3, and also Fig. 4.)

Although maintaining the entire tree \mathcal{T} may seem more challenging than maintaining only the maximal 3-edge-connected subgraphs, we show that it is in fact easier to maintain the decomposition tree. This is because \mathcal{T} contains enough information to facilitate its efficient update after new insertions to G , as its nodes correspond to all successive partitions of G into k -edge-connected components, for $k \in \{1, 2, 3\}$. In fact, if a new edge e is inserted to G , then we only have to locate the deepest subgraph N on \mathcal{T} that contains the endpoints of e , and all changes to \mathcal{T} due to this insertion apply to the subtree of N . However, we do not explicitly maintain the correspondence between nodes of \mathcal{T} and subgraphs of G , as this would require as much as $\Omega(mn)$ space (where m is the number of edges and n is the number of vertices of G), and a much worse time to maintain \mathcal{T} during a sequence of insertions. Instead, we associate with the nodes of \mathcal{T} some data structures that represent the interconnections between the 2- and 3-edge-connected components of the subgraphs that (abstractly) correspond to the nodes of \mathcal{T} . To be more precise, we associate to every node of \mathcal{T} that corresponds to a connected component C of its parent (a representation of) the tree of the 2-edge-connected components of C ; and to every node of \mathcal{T} that corresponds to a 2-edge-connected component of its parent, we associate (a representation of) the cactus of its 3-edge-connected components. All this information reveals to be useful for maintaining our tree decomposition with the help of previous incremental approaches [28,40,41,59]. We describe in detail, in Sections 5 and 6, how we can augment the previous incremental approaches with enough information to suit our purposes.

Notice that in the bounds provided in (1), there is a trade-off between space- and time-complexity. The second algorithm is more time-efficient (at least asymptotically), due to the following two reasons. First, we find a way to use the more sophisticated data structures of La Poutré [40], in order to efficiently maintain the 2- and 3-edge-connected components of the graphs in the various levels of \mathcal{T} . Secondly, we use an alternative and more intriguing method for answering nearest common ancestor and level ancestor queries, that are needed in the algorithm for maintaining \mathcal{T} . This method relies on the structure of those queries (i.e., they are not completely arbitrary but they depend on modifications already made on \mathcal{T}). However, in order to facilitate the efficient answering of those queries, we pay an extra $O(n^2)$ in space. In any case, the n^2 factor in the time-bound of both algorithms is an inherent bottleneck of the basic procedure that we use, for any sequence of insertions, and a lower bound on the worst-case time for a single insertion (see Fig. 3 and Algorithm 1).

Both our algorithms can efficiently answer queries concerning the maximal 3-edge-connected subgraphs in asymptotically optimal time, plus the time to perform one or two calls to a *find* operation in an underlying disjoint set union (DSU) data structure [54] that maintains the vertex sets of the maximal 3-edge-connected subgraphs. (For the details on this DSU data structure, see Section 4.) For instance, given two query vertices x and y , we can report whether x and y belong to the same maximal 3-edge-connected subgraph using two calls to a *find* operation, or given a query vertex x we can report the maximal 3-edge-connected subgraph that contains x in time proportional to its size plus a call to a *find* operation. (See Section 10.)

We note that these algorithms can be seen as applications of a more general framework for maintaining the maximal k -edge-connected subgraphs, by relying on algorithms that maintain the k -edge-connected components. We think that it is possible that one can develop a similar framework for maintaining the maximal k -vertex-connected subgraphs, by relying on efficient algorithms for maintaining the k -vertex-connected components. (In particular, for the case $k = 3$, one may rely on the algorithms of [13,50] for maintaining the 3-vertex-connected components.)

For (2), we show that it is sufficient to compute (a superset of) all the edges whose endpoints lie in different maximal k -edge-connected subgraphs. Benczúr and Karger [5] provide efficient algorithms that achieve this. From those algorithms we get two different constructions for spanning subgraphs that have the same maximal k -edge-connected subgraphs as the original graph, and there is a trade-off between the time complexity and the size of the output subgraph. Following the terminology of [1], we call such a subgraph a k -certificate. Then, we have a linear-time algorithm for computing a k -certificate with $O(kn \log n)$ edges, and an $O(m \log^2 n)$ -time algorithm for computing a k -certificate with $O(kn)$ edges (where m and n denote the number of edges and vertices of the graph, respectively). A key component in those algorithms is the certificates for k -edge connectivity of Nagamochi and Ibaraki [47]. For the details, see Section 7. We believe that it is an interesting question whether a k -certificate of $O(kn)$ size can be computed in linear time. This would be trivial if we had a linear-time algorithm for computing the maximal k -edge-connected subgraphs of a graph, but it is still an open problem whether this can be done for $k \geq 3$. Thus, we have to perform the construction of the certificates without explicitly computing the maximal k -edge-connected subgraphs, and this seems to be a challenging task. Then, we give two applications of our k -certificates. First, we use them to speed up previous algorithms for computing the maximal k -edge-connected subgraphs in dense undirected graphs. Furthermore, we apply our k -certificates within the sparsification technique of Eppstein et al. [20] to give a fully dynamic algorithm for maintaining information about the maximal k -edge-connected subgraphs for fixed k . This way, we achieve update times that are independent of the number of edges and constant time for maximal k -edge-connected subgraph queries. For the details, see Section 9.

Finally, for (3), we repeatedly find and remove all k' -edge cuts, for $k' < k$. For this purpose, we can rely on the fully dynamic min-cut algorithm of Thorup (Theorem 26 in [56]). Thus, we use this algorithm in order to successively find and remove all k' -edge cuts, for $k' < k$, until we are left with the maximal k -edge-connected subgraphs. This is how we get

a deterministic $O(m + k^{O(1)}n\sqrt{n}\text{polylog}(n))$ -time algorithm for computing the maximal k -edge-connected subgraphs of a graph with m edges and n vertices, for $k = \log^{O(1)}n$. This algorithm is described in Section 8.

1.2. More related work

Dynamic graph algorithms have been extensively studied for several decades, and many important results have been achieved for fundamental problems such as dynamic connectivity [22,29,33,43], minimum spanning trees [20,23,49], minimum cut [39,40,56], shortest paths [2,6,12,55], and transitive closure [32]. Regarding dynamic algorithms for maintaining k -connectivity information (for small values of k) has also received a lot of attention [16,17,19,20,26,34,37,38,40,41,59]. This previous work, however, involves maintaining a graph under edge insertions and/or deletions, so that we can efficiently answer queries of whether two vertices are k -edge- or k -vertex-connected (i.e., belong to the same k -edge- or k -vertex-connected component). Very recently, Aamand et al. [1] provided a Monte Carlo randomized algorithm for the decremental maintenance of the maximal c -edge-connected subgraphs, that has $O(m + n^{3/2+o(1)})$ total update time and $O(1)$ query time, for $c = O(n^{o(1)})$. To the best of our knowledge, no previous non-trivial bounds were known for the dynamic maintenance of maximal 3-edge-connected subgraphs with a deterministic algorithm.

1.3. Organization

The rest of the paper is organized as follows. First, we provide some preliminaries in Section 2. We describe our decomposition tree of the maximal 3-edge-connected subgraphs in Section 3.3. This can be seen as an application of a general framework for maintaining the maximal k -edge-connected subgraphs by using algorithms for maintaining the k -edge-connected components, which is described in Sections 3.1 and 3.2. The details on maintaining the decomposition tree of the maximal 3-edge-connected subgraphs under insertions of new edges and vertices are given in Section 4. In Sections 5 and 6 we present efficient implementations for the data structures that are associated with the nodes of the decomposition tree, in order to efficiently update it. In Section 7 we present constructions of sparse certificates for maximal k -edge-connected subgraphs. From this we derive a deterministic $O(m + k^{O(k)}n\sqrt{n}\log n)$ -time algorithm for computing the maximal k -edge-connected subgraphs. In Section 8 we provide an $O(m + k^{O(1)}n\sqrt{n}\text{polylog}(n))$ deterministic algorithm for computing the maximal k -edge-connected subgraphs, using Thorup's fully dynamic mincut algorithm. Section 9 presents our fully dynamic algorithm for maximal k -edge-connected subgraphs. We conclude in Section 10 with suggestions for further applications of our decomposition tree.

2. Preliminaries

In the sequel, we assume that the reader is familiar with standard graph terminologies, for example, as presented in [11,14,46]. For the sake of completeness, we provide some definitions and state some results that will be used throughout concerning the structure of the 2- and 3-edge-connected components in undirected graphs. Consider the quotient graph Q^k of G that is formed by shrinking every k -edge-connected component into a single vertex, maintaining all inter-connection edges and discarding self-loops. Then the quotient map $\nu : V(G) \rightarrow V(Q^k)$ induces a natural correspondence between the edges of Q^k and some edges of G : that is, for every edge (u, v) of Q^k , there is an edge $(x, y) \in E(G)$ such that $\nu(x) = u$ and $\nu(y) = v$. Now, for $k = 2$, we have that Q^2 is a tree T . The nodes of T correspond to the 2-edge-connected components of G , and the edges of T correspond to the bridges of G . Now let C be a 2-edge-connected component of G . Then we have that $G[C]$ is 2-edge-connected, and every k -edge-connected component of G that lies in C , for $k \geq 2$, is also a k -edge-connected component of $G[C]$. For $k = 3$, the quotient graph Q^3 of G is a *cactus* S (that is, a connected graph in which every edge belongs to a unique cycle) [18,28,41]. The nodes of S correspond to the 3-edge-connected components of G and the 2-edge cuts of Q^3 correspond to the 2-edge cuts of $G[C]$. (We note that these properties generalize to the *cactus of the minimum cuts* of an undirected graph [15].)

Now let us consider how the insertion of a new edge (x, y) to G affects its k -edge-connected components, for $k \leq 3$. (In general, observe that the edge-connectivity for any pair of vertices increases at most by one.) We distinguish four different cases, depending on whether $\lambda(x, y) \in \{0, 1, 2\}$ or $\lambda(x, y) \geq 3$, prior to the insertion of (x, y) .

- (1) If $\lambda(x, y) = 0$, then x and y belong to two different connected components C_1 and C_2 , respectively. Thus the only change that occurs is that $\lambda(u, v) = 1$, for any two vertices $u \in C_1$ and $v \in C_2$.
- (2) If $\lambda(x, y) = 1$, then x and y belong to the same connected component C , and lie in two different 2-edge-connected components X and Y , respectively. Let T be the tree of the 2-edge-connected components of $G[C]$, and let $P = X_1, \dots, X_k$ be the simple path on T with endpoints X and Y (i.e., we have $X_1 = X$ and $X_k = Y$). Then, after inserting (x, y) , we have $\lambda(u, v) \geq 2$, for any pair of vertices $u, v \in X_1 \cup \dots \cup X_k$, and $\lambda(u, v)$ stays the same for any other pair of vertices. In particular, this implies that $X_1 \cup \dots \cup X_k$ is a new 2-edge-connected component. Now, for every edge (X_i, X_{i+1}) , $i \in \{1, \dots, k-1\}$, of T , let (x_i, y_i) be the edge of G that corresponds to (X_i, X_{i+1}) , and let $(x, y) = (y_0, x_k)$ (this is for notational convenience). Then, for every $i \in \{1, \dots, k\}$, we have $y_{i-1}, x_i \in X_i$, and the remaining changes in the k -edge-connected components of G , after inserting (x, y) , are given by supposing that we introduce a virtual edge (y_{i-1}, x_i) to G ; thus they are described sufficiently in the following two cases.

- (3) If $\lambda(x, y) = 2$, then x and y belong to the same 2-edge-connected component C , and lie in two different 3-edge-connected components of $G[C]$. Let S be the cactus of the 3-edge-connected components of $G[C]$. Then, after inserting (x, y) , we have that at least X , and Y are now united into a larger 3-edge-connected component (together with some other nodes of S). To describe precisely the nodes of S that get united into a new 3-edge-connected component, we introduce the concept of the *cycle-path* on S connecting X and Y . Let Z_1, \dots, Z_k be a simple path on S with $Z_1 = X$ and $Z_k = Y$. Then the cycle-path Q connecting X and Y on S is the set of nodes consisting of X, Y , and all Z_i , $i \in \{2, \dots, k-1\}$, such that the edges $(Z_{i-1}, Z_i), (Z_i, Z_{i+1})$ belong to different cycles of S . (Notice that this definition is independent of the choice of the simple path Z_1, \dots, Z_k .) Then, after inserting (x, y) , all the nodes of Q are merged into a new maximal 3-edge-connected component. The edge-connectivity between any pair of vertices u and v , not both of which lie in nodes of Q , stays the same. Observe that the new cactus of the 3-edge-connected components of $G[C]$ is given by S where we have “squeezed” every cycle c that contains two consecutive nodes Z and W of Q , by merging Z and W into a new node U . (If Z and W are not connected with an edge on S , then the squeezing of c at Z and W produces two new cycles that meet at U .)
- (4) If $\lambda(x, y) \geq 3$, then all the k -edge-connected components of G , for $k \leq 3$, stay the same [16].

3. The decomposition tree of the maximal k -edge-connected subgraphs

In this section we provide a general framework for maintaining the maximal k -edge-connected subgraphs of an undirected graph upon insertions of edges or vertices. This framework relies on a structural characterization of the maximal k -edge-connected subgraphs that is derived from the repeated decomposition of the graph into its k -edge-connected components. Thus we get a decomposition tree, which can be maintained by using any algorithm for the incremental maintenance of the k -edge-connected components (that satisfies some properties) as a black box. We provide a concrete application of this framework: two algorithms for maintaining the maximal 3-edge-connected subgraphs. These algorithms provide a trade-off between time- and space-bounds, because they rely on different algorithms for maintaining the 3-edge-connected components, and they utilize the structure of the decomposition tree in different ways.

3.1. A general framework for maintaining the k -edge-connected subgraphs

In what follows we let “ k -ecc” mean “ k -edge-connected component”. We will not distinguish between the vertex set of a k -ecc and the subgraph induced by it. Let G be an undirected multigraph with m edges and n vertices. We consider the decomposition tree \mathcal{T} of the maximal k -edge-connected subgraphs of G , which is a rooted tree whose nodes correspond to subgraphs of G . It is defined recursively as follows. (1) The root of \mathcal{T} corresponds to G . (2) Every node of \mathcal{T} that corresponds to a k -edge-connected subgraph of G is a leaf. (3) The children of every node N of \mathcal{T} that is not a leaf correspond to the k -eccs of the subgraph of G corresponding to N . This completes the recursive definition of \mathcal{T} . Observe that the size of \mathcal{T} is $O(n)$, because the leaves of \mathcal{T} correspond precisely to the maximal k -edge-connected subgraphs of G (which are at most n), and every node of \mathcal{T} is either a leaf or it has at least two children.

We observe that \mathcal{T} has the following useful property.

Property 3.1. *Let G be an undirected graph, and consider the quotient graph Q that is formed by shrinking a maximal k -edge-connected subgraph of G into a single vertex. Then Q has the same decomposition tree into maximal k -edge-connected subgraphs as G .*

We use this decomposition tree because we want to maintain the maximal k -edge-connected subgraphs of G while new edges or vertices are inserted to it. The definition using successive partitions into k -edge-connected components is convenient, because every maximal k -edge-connected subgraph lies entirely within a k -edge-connected component, and we can utilize algorithms for the incremental maintenance of the k -edge-connected components (that satisfy some properties) in order to maintain the decomposition tree under new insertions.

Now let us describe the changes that take place in \mathcal{T} after a new edge is inserted to G . Suppose a new edge $e = (x, y)$ is inserted to G . Let X and Y be the maximal k -edge-connected subgraphs of G that contain x and y , respectively. If $X = Y$, there is nothing to do. Otherwise, we find the nearest common ancestor N of X and Y on \mathcal{T} , and perform the insertion of e to N . Observe that N is the deepest node of \mathcal{T} that contains both x and y . It is sufficient to perform the insertion of e to N , because N is separated from its siblings by a k' -edge cut, for some $k' < k$, and this does not change even if we insert e to N . Thus, all changes in \mathcal{T} after the insertion of e to G take place in the subtree of N . If x and y do not become k -edge-connected in N , then no change takes place in \mathcal{T} . Otherwise, some k -eccs of N get merged into a larger k -ecc. If all vertices of N become k -edge-connected, then N becomes a leaf node. Otherwise, we have to find its children that correspond to the k -eccs of N that get merged, and substitute them with a single child that corresponds to the bigger k -ecc that is formed. Then we have to push down to the new child all the edges that have their endpoints in different k -eccs that got merged. In effect, we re-insert those edges to the child of N that represents the new k -ecc that is formed. This completes the recursive description of the changes that take place in \mathcal{T} after inserting a new edge to G .

In order to efficiently maintain the decomposition tree, we will have to attach more information to it. We discuss this in the following section, where we use as a black box any algorithm that maintains the k -edge-connected components (and

satisfies some conditions). In any case, it is interesting to note - and it is essential in proving our time bounds - that the number of edge insertions that can affect the decomposition tree, in any sequence of edges insertions, is $O(kn)$. This is formally stated and proved in Theorem 3.3.

Lemma 3.2. [1] Let G be a graph with n vertices and m edges, such that every maximal k -edge-connected subgraph of G is trivial. Then $m \leq (k-1)(n-1)$.

Theorem 3.3. Let $k \geq 3$, and let $G_0 = (V, E_0)$ be an empty graph with n vertices and no edges ($E_0 = \emptyset$). Let e_1, \dots, e_m be a sequence of edges joining vertices of V , and define $E_i = \{e_1, \dots, e_i\}$ and $G_i = (V, E_i)$, for $i = 1, \dots, m$. Then, $|\{i \in \{1, \dots, m\} \mid e_i \text{ joins two different maximal } k\text{-edge-connected subgraphs of } G_{i-1}\}| \leq k(n-1)$.

Proof. We proceed by induction on the number n of vertices. For $n = 1$, the theorem trivially holds, as there are no possible edge insertions. Assume that the theorem holds for every graph with at most n vertices, for some $n \geq 1$. Given an empty graph with $n+1$ vertices, we will show that at most kn edges can affect the decomposition tree during any sequence e_1, \dots, e_m of edge insertions.

Let ℓ , $0 \leq \ell \leq m-1$, be the lowest index such that $G_{\ell+1} = (V, E_{\ell+1})$ contains a non-trivial maximal k -edge-connected subgraph (i.e., $G_i = (V, E_i)$, $0 \leq i \leq \ell$ contain only trivial maximal k -edge-connected subgraphs). If no such index ℓ exists, then the final graph G_m has only trivial maximal k -edge-connected subgraphs and the theorem holds by Lemma 3.2. Otherwise, let S be the non-trivial maximal k -edge-connected subgraph that appears in $G_{\ell+1}$ after the edge $e_{\ell+1}$ has been inserted to G_ℓ , and let S contain d vertices ($d \geq 2$) and t edges. Note that $e_{\ell+1}$ must necessarily be in S . Let $S' = S \setminus e_{\ell+1}$: then S' has $d \geq 2$ vertices and $t-1$ edges, and contains only trivial maximal k -edge-connected subgraphs. By Lemma 3.2 applied to S' , we have $t-1 \leq (k-1)(d-1)$, and therefore $t \leq (k-1)(d-1) + 1$.

Now let us consider the quotient graph of $G_{\ell+1}/S$ that is formed by contracting S into one single vertex. This quotient graph has $\ell+1-t$ edges and $(n+1)-d+1 = n-d+2 \leq n$ vertices (since $d \geq 2$). Since we are contracting a maximal k -edge-connected subgraph, the quotient graph has exactly the same decomposition tree as $G_{\ell+1}$. By induction, there are at most $k((n-d+2)-1) - (\ell+1-t) = kn - kd + k - \ell - 1 + t$ new edges that can be added to the quotient graph and affect its decomposition tree. This number is at most $kn - \ell - d + 1$ (since t is at most $kd - k - d + 2$). Since we have already added $\ell+1$ edges to the initial graph, we have that, in total, at most $kn - d + 2$ edges can affect the decomposition tree. The fact that $d \geq 2$ yields the desired result. \square

3.2. Maintaining the decomposition tree

In order to be able to update \mathcal{T} after a new insertion to the graph, we need a k -edge-connected components algorithm that satisfies the following properties. The algorithm works on an incremental graph G and maintains a unique label for every k -ecc of G . Furthermore, it supports the following two operations.

1. When a new edge is inserted to G , return the labels of the k -eccs of G that get merged due to this insertion, the label of the new k -ecc that is formed, and the set of the interconnection edges between those k -eccs that get merged.
2. Given a collection of vertex-disjoint graphs G_1, \dots, G_t , join the underlying data structures that the algorithm uses for those graphs into a new data structure for the graph $G = G_1 \cup \dots \cup G_t$.

We call an algorithm (plus the underlying data structure) that supports these operations an “incremental k -edge-connected components algorithm” (or “a data structure for maintaining the k -edge-connected components”).

Now suppose that we are equipped with an incremental k -edge-connected components algorithm. In order to efficiently maintain the decomposition tree \mathcal{T} , we associate with every non-leaf node N a data structure for maintaining the k -edge-connected components. Furthermore, for every k -ecc of N , we have a two-way correspondence between its label maintained by the data structure and the child of N that corresponds to it. Now suppose that an edge e is inserted to N . Then we use operation (1) in order to find the labels C_1, \dots, C_t of the k -eccs of N that get merged due to the insertion of e . If there is only one such k -ecc returned, there is nothing to do. If all the k -eccs of N are returned, then we make N a leaf node. Otherwise, we get from those labels the corresponding children of N , we merge those children into a new child, and we establish a two-way correspondence between this child and the new k -ecc of N that is formed. Furthermore, we retrieve the associated data structures of the children of N that get merged, we join them into a new data structure using operation (2), and we associate the new data structure with the new child of N . Finally, we insert the interconnection edges between C_1, \dots, C_t that we got from operation (1) into the new child of N . This completes the recursive description of the update of \mathcal{T} using an incremental k -edge-connected components algorithm as a black box.

We will describe two different methods with which we can perform the merging of nodes of \mathcal{T} , with a trade-off in time- and space-complexity. The first method is the simplest one. Let N be a node of \mathcal{T} , and let C_1, \dots, C_t be the children of N that we have to merge. Then we will use the node with the greatest number of children among C_1, \dots, C_t as the new node in which C_1, \dots, C_t get merged. So let C be a node among C_1, \dots, C_t with maximum number of children. Then we discard all nodes in $\{C_1, \dots, C_t\} \setminus \{C\}$, and we redirect the parent pointer of their children to C . Since this involves at most n nodes




Fig. 3. An example of a sequence of insertions that can force $\Omega(n)$ 3-interconnection edges to increase their level by one $\Omega(n)$ times. The sequence of insertions, starting from the empty graph with vertices in $\{1, \dots, n\}$ is $(1, 2), (1, 2), (3, 1), (3, 2), (4, 2), (4, 3), \dots, (i, i-2), (i, i-1), \dots, (n, n-2), (n, n-1)$. After the insertion of $(i, i-1)$, we have that $(i, i-1)$ and $(i, i-2)$ are the only 1-level edges, and the level of the other edges $(1, 2), \dots, (i-1, i-2)$ increases by one. Note that $\{1, \dots, i-1\}$ is a 3-edge-connected component at this point, whereas all the maximal 3-edge-connected subgraphs are trivial.

at every level, and a parent pointer can be redirected at most $\log n$ times, we get at most $O(n \log n)$ redirections of parent pointers at every level, and $O(n^2 \log n)$ redirections in total. In the second method we maintain all nodes of \mathcal{T} (possibly in a “deactivated” mode) throughout the sequence of all insertions, and thus we may need as much as $\Omega(n^2)$ space (since there are at most $O(n)$ insertions that can affect the decomposition tree). The idea is to use a disjoint-set union data structure DSU_i [54] on the nodes of level i of the tree, for all $i \geq 1$. In order to merge nodes at level i we unite all of them using DSU_i and we choose one of them as the representative. We consider the representative to be an *active* node, in the sense that we can use its parent pointer to move to level $i-1$, whereas the parent pointer of the other nodes that got united will never be used again. Thus, in order to access the parent of a node N at level i , we use $find_{i-1}(parent(N))$, where $parent$ denotes the parent pointer in \mathcal{T} , and $find_{i-1}$ is the *find* operation of DSU_{i-1} . Notice that this approach has the advantage that it does not explicitly redirect the parent pointers. Using an optimal implementation for DSU_i [54], we can perform any sequence of m operations $find_i$ or $unite_i$ in $O(m\alpha(m, n))$ time in total. Since at most n *unite* operations can take place at each level, we have an $O(n\alpha(n, n))$ time-bound at every level for the mergings, and an $O(n^2\alpha(n, n))$ time-bound in total.

We note that the n^2 expression in the time-bound is an inherent bottleneck of this approach for maintaining the maximal k -edge-connected subgraphs, for any $k \geq 3$. In order to demonstrate this, let us introduce some concepts. We call the edges that connect two different maximal k -edge-connected subgraphs of a graph k -interconnection edges. Then we partition the k -interconnection edges into levels as follows. If an edge participates in a k' -edge cut of the graph, for $k' < k$, we define the level of this edge to be 1. Now suppose that we have defined all d -level edges, for some $d \geq 1$. Then an edge that participates in a k' -edge cut of (a connected component of) the graph that remains after we remove all d' -level edges, for $d' \leq d$, is assigned level $d+1$. Now we observe that our method for updating the decomposition tree explicitly maintains the levels of the k -interconnection edges (because every d -level k -interconnection edge is essentially maintained in the associated data structure of a node of depth d of the decomposition tree). There are sequences of insertions of edges in which $\Omega(n)$ insertions may force $\Omega(n)$ edges to change their level, and thus we need at least $\Omega(n^2)$ time in total to maintain the decomposition tree. An example for $k=3$ is shown in Fig. 3. (This generalizes easily to any $k > 3$.)

3.3. The decomposition tree of the maximal 3-edge-connected subgraphs

Here we apply the framework outlined in the previous section for the special case $k=3$. We focus on the case $k=3$ for the following reasons. (1) There exist very efficient/asymptotically optimal algorithms for the incremental maintenance of the 3-edge-connected components that we can use for maintaining the decomposition tree for the case $k=3$. (2) Although there exist very efficient algorithms for maintaining also the 4- and 5-edge-connected components [19,16], and they seem applicable in our framework, they involve a lot of technicalities and it would be very tedious to show how to extend them in order to apply them for our purposes. For $k > 5$ we are not even aware of any efficient incremental algorithms that we could apply.⁴ (3) The case $k=3$ is the simplest one to consider. And yet, even here, we have to extend appropriately the existing data structures and algorithms, and this involves various technicalities.

For the case $k=3$ we incorporate the algorithms of [28,41,40,59], for maintaining the 3-eccs, in the decomposition tree. Thus, since these algorithms basically maintain the 2-eccs of the 1-eccs, and the 3-eccs of the 2-eccs, we augment the decomposition tree with more levels in order to capture the decomposition into all k' -eccs, for $k' \leq 3$. To be more precise, we consider the decomposition tree \mathcal{T} of the maximal 3-edge-connected subgraphs of G that is formed by repeatedly removing all 1- and 2-edge cuts. The nodes of \mathcal{T} correspond to subgraphs of G . First, the root of \mathcal{T} corresponds to the entire graph. If G is 3-edge-connected, then this is the only node of \mathcal{T} , and the decomposition is over. Otherwise, the children of the root correspond to the connected components of G , which we call 1-ecc nodes. Then the children of every 1-ecc node correspond to its 2-edge-connected components, and we call them 2-ecc nodes. Finally, the children of every 2-ecc node correspond to its 3-edge-connected components and we call them 3-ecc nodes. Now, if a 3-ecc node corresponds to a 3-edge-connected subgraph of G , then it is a leaf of \mathcal{T} , and a maximal 3-edge-connected subgraph of G . Otherwise, it becomes the root of a subtree of \mathcal{T} that is produced recursively with the same procedure. In this way, we compute the decomposition tree \mathcal{T} of the maximal 3-edge-connected subgraphs of G . (See Fig. 4 for an example.)

⁴ The fully dynamic algorithm of Jin and Sun [37] seems very difficult to be usable in our framework, because it does not explicitly maintain the k -edge-connected components.

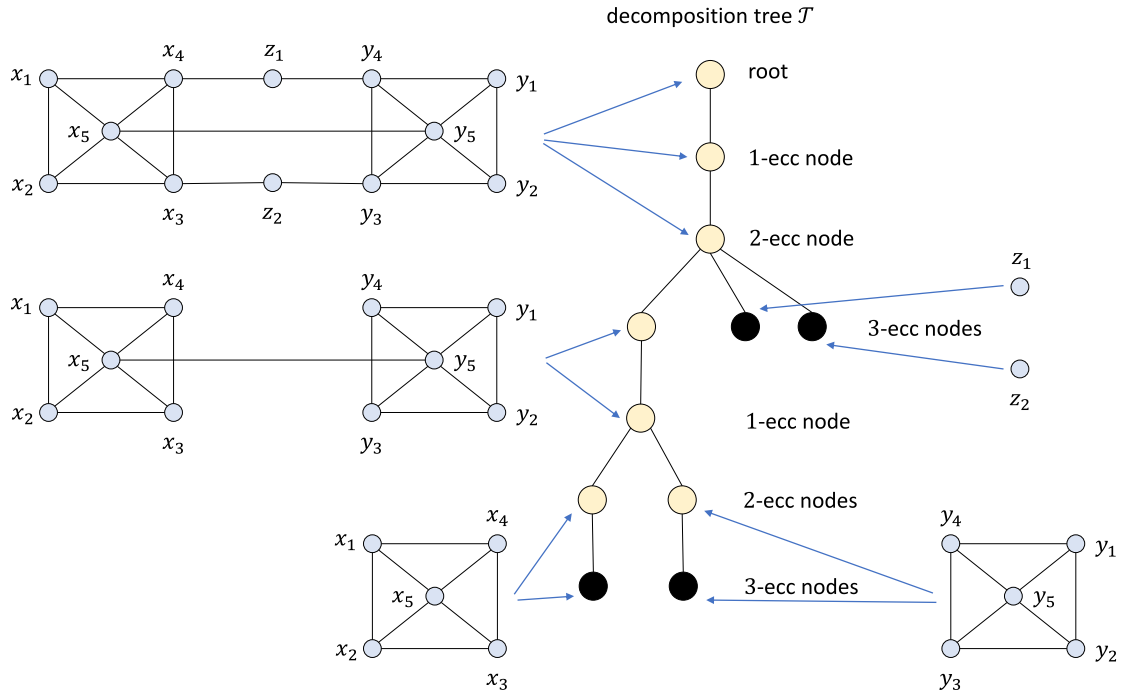


Fig. 4. The decomposition tree \mathcal{T} of the graph G of Fig. 1. Notice the correspondence between the nodes of \mathcal{T} and some subgraphs of G . The leaves of \mathcal{T} (coloured in black) correspond to the maximal 3-edge-connected subgraphs of G .

Observe that the size of \mathcal{T} is $O(n)$,⁵ since the children of every node N of \mathcal{T} correspond to subgraphs of the graph corresponding to N that partition its vertex set (and although some nodes have only one child, by construction this child must have at least two children or two grandchildren, unless it is a leaf). For convenience, and since no ambiguity arises, we identify every node of \mathcal{T} with the subgraph of G that corresponds to it. (We do not maintain this correspondence explicitly in any data structure, as this would be impractically expensive in terms of space (and therefore time).) It is only a conceptual convention that we make, and we will use it extensively in the sequel.) Furthermore, since \mathcal{T} is rooted, we let the nodes of \mathcal{T} be partitioned into levels according to their depth, where the level of the root is 0. Thus we may speak of a subgraph of G at some level of \mathcal{T} . Observe that \mathcal{T} contains $3k$ levels, for some $k \geq 0$, and that every maximal 3-edge-connected subgraph of G is contained at some level $3k'$, for $k' \leq k$, as a leaf. It is useful to note that \mathcal{T} satisfies the following.

Property 3.4. *If we shrink a maximal 3-edge-connected subgraph of G into a single vertex, then the decomposition tree of the resulting graph is given again by \mathcal{T} .*

Now we will give a high-level description of the changes that \mathcal{T} undergoes when a new edge (x, y) is inserted to G . Theorem 3.3 implies that at most $3n - 3$ insertions of edges may affect the decomposition tree, in any sequence of edge insertions, where n is the total number of vertices inserted. By Property 3.4, we have that a newly inserted edge can affect the decomposition tree only if its endpoints lie in different maximal 3-edge-connected subgraphs of G . Thus we consider the nearest common ancestor N of the maximal 3-edge-connected subgraphs of G that contain x and y . This corresponds to the deepest node in the decomposition tree that contains both x and y , and it should be clear that the insertion of (x, y) affects only the decomposition of N . We distinguish three different cases, depending on whether (1) N is the root or a 3-ecc node, or (2) N is a 1-ecc node, or (3) N is a 2-ecc node. The whole procedure is summarized in Algorithm 1.

3.3.1. *N is the root or a 3-ecc node*

In this case (x, y) joins two different connected components C_1 and C_2 of N , and so it becomes a bridge of N that connects them. Thus we only have to merge C_1 and C_2 into a new 1-ecc node (see Fig. 5).

3.3.2. *N is a 1-ecc node*

In this case (x, y) joins two different 2-edge-connected components C_1 and C_2 of N . Without loss of generality, assume that $x \in C_1$ and $y \in C_2$. Let T be tree of the 2-edge-connected components of N , and let $P = X_1, \dots, X_k$ be the path on

⁵ However, in order to efficiently maintain \mathcal{T} after the insertion of new edges to G , we may have to attach more information to \mathcal{T} , so that its total size will be $O(n^2)$. (See Section 4.2.)




Fig. 5. Here we show the changes to the decomposition tree on the insertion of the edge (x, y) . The nearest common ancestor N of (the maximal 3-edge-connected subgraphs that contain) x and y is a 3-*ecc* node. Therefore, we need to know the two children of N that contain x and y ; then we retrieve the associated trees, join them properly with (an equivalent of) the edge (x, y) , and then we merge the two nodes. Finally, we associate with it the new tree that has been formed. In this figure, the middle graphs show the decomposition of the graph into 3-edge-connected components at every level. Also, in the decomposition trees, we can see the trees associated with the nodes we have to merge and the tree associated with the merged node. (The new edge (x, y) is stored inside this tree.) The 3-edge-connected components are highlighted with pink colour. Lower intensity signifies greater depth in the decomposition tree. (For interpretation of the colours in the figure(s), the reader is referred to the web version of this article.)

T with endpoints C_1 and C_2 . (Thus we have $X_1 = C_1$ and $X_k = C_2$.) Then the vertices that are contained in X_1, \dots, X_k become 2-edge-connected, and for every pair of vertices not both of which belong to $X_1 \cup \dots \cup X_k$ the edge-connectivity remains the same. Furthermore, for every $i \in \{1, \dots, k - 1\}$, let (x_i, y_i) be the bridge of N that corresponds to the edge (X_i, X_{i+1}) of P , and let also $(x, y) = (y_0, x_k)$ (this is for notational convenience). Then, every pair of edges in $\{(x_i, y_i) \mid i \in \{1, \dots, k - 1\}\} \cup \{(x, y)\}$ is a 2-edge cut of N . Thus we have to merge all 2-*ecc* nodes X_1, \dots, X_k into a new 2-*ecc* node C , and no change takes place on \mathcal{T} outside the subtree of C .

Now we have to consider the changes that possibly take place in the subtree of C . For every $i \in \{1, \dots, k\}$, let S_i be the cactus of the 3-edge-connected components of X_i , and let Q_i be the cycle-path on S_i with endpoints the 3-*ecc* of X_i that contains y_{i-1} and the 3-*ecc* of X_i that contains x_i . Let also $D_{(i,1)}, \dots, D_{(i,t(i))}$ be the subgraphs of X_i that correspond to the nodes of Q_i . Then, the vertices that are contained in $D_{(i,1)}, \dots, D_{(i,t(i))}$ become 3-edge-connected; furthermore, for every pair of vertices of X_i not both of which lie in $D_{(i,1)} \cup \dots \cup D_{(i,t(i))}$ the edge-connectivity remains the same. Thus we have to merge all $D_{(i,1)}, \dots, D_{(i,t(i))}$ into a new 3-*ecc* node D_i ; furthermore, no change in the subtree of X_i takes place outside the subtree of D_i . Now we have to consider the graphs $D_{(i,1)}, \dots, D_{(i,t(i))}$ as the connected components of the new graph D_i . However, in order to maintain the decomposition subtree of D_i on \mathcal{T} , we have to take care of two things. First, some graphs $D_{(i,j)}$, for $j \in \{1, \dots, t(i)\}$, might be leaves (prior to the insertion of (x, y)). This means that they are 3-edge-connected subgraphs, and so we have to include them in the subtree of D_i by expanding their corresponding nodes on \mathcal{T} with the addition of three intermediary nodes. And secondly, there might exist edges in X_i that connect some of $D_{(i,1)}, \dots, D_{(i,t(i))}$. (Note that these edges correspond naturally to those of S_i that connect the nodes of Q_i .) These edges are now included in D_i , and so we have to consider their effect on the decomposition subtree of D_i . Thus, in order to capture fully the effect on \mathcal{T} of the insertion of (x, y) , we have to re-insert those edges to G . Observe that these edges constitute parts of 1- or 2-edge cuts of D_i , and so they will be re-inserted only once in order to perform the insertion of (x, y) . However, their re-insertion may force other edges of G , that lie in graphs on deeper levels, to be re-inserted to G . Nevertheless, the nearest common ancestor involved in each such re-insertion will either be a 3-*ecc* or a 1-*ecc* node. As a consequence, we note that no new maximal 3-edge-connected subgraphs will be formed after the insertion of (x, y) to G . Fig. 6 is an example of this case.

We refer to the three step process of (a) merging all $D_{(i,1)}, \dots, D_{(i,t(i))}$ into a single node D_i , (b) expanding every node $D_{(i,j)}$ that is a leaf, for $j \in \{1, \dots, t(i)\}$, with the addition of three intermediary nodes, and (c) re-inserting into G the inter-edges between the subgraphs $D_{(i,1)}, \dots, D_{(i,t(i))}$ of X_i , as *merging* the 3-*ecc* nodes $D_{(i,1)}, \dots, D_{(i,t(i))}$ into D_i . This procedure is shown in Algorithm 2. (The variable “parent” in Lines 6, 7 and 8 denotes the parent relation of \mathcal{T} .)




Fig. 6. Here we insert the edge (x, y) to the graph. This joins two different 2-edge-connected components (at the first decomposition level), which lie in the tree T associated with the nearest common ancestor of x and y . The path of those 2-eccs on T consists of three nodes, which we have to merge. So first we have to retrieve their associated cactuses (which are shown in this figure), to merge some of their nodes using `merge3ecc`, and then to join them on a new cactus using (essentially) the edges of the cycle. Here we can see the effect of `merge3ecc` on the three cactuses. With grey are shown the nodes that we have to merge. The resulting three cactuses get joined on their grey (compressed) nodes along the edges of the triangle, one of which is (essentially) the new one, and the other two lay on a higher level (inside the tree T of the 2-eccs). Thus, those two edges got pushed down one level. Also notice that the height of the decomposition tree increased.

3.3.3. N is a 2-ecc node

In this case (x, y) joins two different 3-edge-connected components C_1 and C_2 of N . We note that this is the only case in which the formation of new maximal 3-edge-connected subgraphs of G may take place (by merging together smaller maximal 3-edge-connected subgraphs of G). Let S be cactus of the 3-edge-connected components of N , and let $Q = X_1, \dots, X_k$ be the cycle-path on S with endpoints C_1 and C_2 . Then the vertices that are contained in X_1, \dots, X_k become 3-edge-connected, and for every pair of vertices not both of which belong to $X_1 \cup \dots \cup X_k$ the edge-connectivity remains the same. Thus, if Q contains all the nodes of S , then N becomes 3-edge-connected. Therefore, if N prior to the insertion of (x, y) was the only 2-edge-connected component of its grandparent R on \mathcal{T} , then this means that R becomes 3-edge-connected, and so its subtree on T is condensed into R . Otherwise, all the proper descendants of N are condensed into a new single 3-ecc node of \mathcal{T} , which corresponds to the new maximal 3-edge-connected subgraph that has been formed.

Now suppose that Q does not contain all the nodes of S . Then we have to perform a merging of the 3-ecc nodes X_1, \dots, X_k into a new node D , and then repeat the insertion of the edge (x, y) . Let X and Y be the maximal 3-edge-connected subgraphs of G that contain x and y , respectively. Then the nearest common ancestor of X and Y on \mathcal{T} is a descendant of D , and it can either be a 3-ecc node or a 1-ecc node (in which case we are work as previously), or a 2-ecc node again. In the last case, either the formation of a new maximal 3-edge-connected subgraph of G will take place (and we are done), or we will have to merge again some 3-ecc nodes of \mathcal{T} into a new node and then repeat the insertion of (x, y) . Observe that, eventually, this process must terminate, since every time that we have to merge some 3-ecc nodes into a new node, we leave some of their siblings on the same level, and then the computation involves only the subtree of the new node.

Fig. 7 is a simple example of this case, where the whole graph becomes 3-edge-connected.

4. Maintaining the decomposition tree after insertions

As seen in Algorithms 1 and 2, in order to efficiently update \mathcal{T} after inserting an edge (x, y) to G , we have to provide efficient implementations for the following procedures:

- (1) Find the leaves X and Y of \mathcal{T} that contain x and y , respectively.
- (2) Find the nearest common ancestor N of X and Y on \mathcal{T} .




Fig. 7. This is a simple example where the insertion of a new edge (x, y) makes the whole graph 3-edge-connected. This happens immediately, because the nearest common ancestor N of x and y is a 2-ecc node, where the associated cactus is the cycle-path connecting the nodes that contain x and y . Thus we merge all leaves of N , we discard the whole subtree N , and assign as a child of N the new maximal 3-edge-connected subgraph that has been formed. Notice that the parent of N has only one child, and the parent of the parent of N has again only one child. Thus, we replace the grandparent of N with the child of N , and we discard N , its parent, and its grandparent.

- (3) Find the grandchild of N that contains a particular vertex $v \in N$. (This is needed in Line 20 of Algorithm 1, in order to find the 3-eccs of X_i that contain y_{i-1} and x_i .)
- (4) Merge sets of nodes of \mathcal{T} (lying on the same level) into a new single node.
- (5) Find the nodes of \mathcal{T} that have to get merged, and the inter-connection edges between the subgraphs that correspond to them.

We will provide two different solutions for this set of procedures. In Section 4.1 we rely on the ideas and the data structures of [28,41,59], in order to provide an algorithm that uses $O(n)$ space and runs in $O(n^2 \log^2 n + m\alpha(m, n))$ time, for any sequence of m edge and n vertex insertions. In Section 4.2 we rely on [41] and [40], in order to provide an algorithm that uses $O(n^2)$ space and runs in $O(n^2\alpha(n, n) + m\alpha(m, n))$ time, for any sequence of m edge and n vertex insertions. We note that the first algorithm is more readily amenable to implementations. The second algorithm, although it is asymptotically more time-efficient, it uses the sophisticated data structures of [40] (summarized and expanded in Section 6), and it may be less efficient than the first algorithm in practice.

In both algorithms we use the same solution for (1): an optimal disjoint set union data structure DSU_{3ecs} [54], that operates on $V(G)$, and uses as representatives the leaves of T (which correspond bijectively to the maximal 3-edge-connected subgraphs of G). DSU_{3ecs} supports the operations $find_{3ecs}(x)$ and $unite_{3ecs}(x, y)$. $find_{3ecs}(x)$ returns a pointer to the maximal 3-edge-connected subgraph of G that contains x , and $unite_{3ecs}(x, y)$ unites the sets of vertices of the maximal 3-edge-connected subgraphs of G that contain x and y . DSU_{3ecs} uses $O(n)$ space, and it can perform any sequence of m operations $find_{3ecs}$ and $unite_{3ecs}$ in $O(m\alpha(m, n))$ time in total [54]. (Thus we get this expression in the time-bounds of both algorithms.)

Furthermore, for (5) we rely on data structures associated with the nodes of \mathcal{T} , that represent trees or cactuses. In particular, recall that (the graph corresponding to) every 1-ecc node X is a connected component of (the graph corresponding to) its parent, and the children of X correspond to its 2-edge-connected components. Thus, we associate with every 1-ecc node a data structure that represents the tree of its 2-edge-connected components. This data structure is used in order to find the children of X that we have to merge (in Line 13 of Algorithm 1), and the inter-connection edges between (the graphs that correspond to) those children (in Line 15 of Algorithm 1). Similarly, every 2-ecc node X has an associated data structure that represents the cactus of its 3-edge-connected components. This data structure is used in order to find the children of X that we have to merge (in Lines 20 and 27 of Algorithm 1), and the inter-connection edges between (the graphs that correspond to) those children (in Line 12 of Algorithm 2). More details about the operations supported by these data structures will be given in Section 4.1.

Finally, we note that the $O(n^2)$ expression in the time-bounds of both algorithms is a bottleneck in their total running time. In fact, Lines 34 and 35 of Algorithm 1 and Line 15 of Algorithm 2 may create a sequence of recursive calls, so that, even for a single edge insertion, we may have to perform $O(n^2)$ re-insertions of already inserted edges. However, since by Theorem 3.3 there are only $O(n)$ edge insertions that can affect the decomposition tree, and every re-insertion (performed internally by the algorithm) can only affect deeper levels than before, and there can be only $O(n)$ levels in the decomposition tree, we have that any sequence of edge insertions to G can initiate at most $O(n^2)$ calls to procedure `insert` of Algorithm 1.

Algorithm 1: Update the decomposition tree \mathcal{T} after inserting the edge (x, y) to G .

```

1 procedure insert( $x, y$ )
2 begin
3    $X \leftarrow$  the leaf of  $\mathcal{T}$  that contains  $x$ 
4    $Y \leftarrow$  the leaf of  $\mathcal{T}$  that contains  $y$ 
5    $N \leftarrow$  the nearest common ancestor of  $X$  and  $Y$  on  $\mathcal{T}$ 
6    $C_1 \leftarrow$  the child of  $N$  that contains  $x$ 
7    $C_2 \leftarrow$  the child of  $N$  that contains  $y$ 
8   if  $N$  is the root of  $\mathcal{T}$  or a 3-ecc node then
9     merge  $C_1$  and  $C_2$  into a new 1-ecc node
10  end
11  else if  $N$  is a 1-ecc node then
12     $T \leftarrow$  the tree of the 2-edge-connected components of  $N$ 
13     $X_1, \dots, X_k \leftarrow$  the path on  $T$  with endpoints  $C_1$  and  $C_2$ 
14    foreach  $i \in \{1, \dots, k-1\}$  do
15       $(x_i, y_i) \leftarrow$  the edge of  $N$  that corresponds to  $(X_i, X_{i+1})$  of  $T$ 
16    end
17     $(y_0, x_k) \leftarrow (x, y)$ 
18    foreach  $i \in \{1, \dots, k\}$  do
19       $S_i \leftarrow$  the cactus of the 3-edge-connected components of  $X_i$ 
20       $D_{(i,1)}, \dots, D_{(i,t(i))} \leftarrow$  the cycle-path on  $S_i$  with endpoints the 3-ecc of  $X_i$  that contains  $y_{i-1}$  and the 3-ecc of  $X_i$  that contains  $x_i$ 
21      merge3ecc( $D_{(i,1)}, \dots, D_{(i,t(i))}$ )
22    end
23    merge all  $X_1, \dots, X_k$  into a new 2-ecc node of  $\mathcal{T}$ 
24  end
25  else if  $N$  is a 2-ecc node then
26     $S \leftarrow$  the cactus of the 3-edge-connected components of  $N$ 
27     $D_1, \dots, D_k \leftarrow$  the cycle-path on  $S$  with endpoints  $C_1$  and  $C_2$ 
28    if  $V(S) = \{D_1, \dots, D_k\}$  then
29       $R \leftarrow$  the grandparent of  $N$  on  $\mathcal{T}$ 
30      if  $N$  is the only 2-ecc of  $R$  then condense the subtree of  $R$  into  $R$ 
31      else condense all the proper descendants of  $N$  into a new 3-ecc node
32    end
33    else
34      merge3ecc( $D_1, \dots, D_k$ )
35      insert( $x, y$ )
36    end
37  end
38 end

```

Algorithm 2: Merge the children D_1, \dots, D_k of a 2-ecc node X .

```

1 procedure merge3ecc( $D_1, \dots, D_k$ )
2 begin
3   for  $i \in \{1, \dots, k\}$  do
4     if  $D_i$  is a leaf of  $\mathcal{T}$  then
5       let  $D'_i, D''_i$  and  $D'''_i$  be a new 1-ecc, 2-ecc and 3-ecc node, respectively, where  $D'_i, D''_i$  and  $D'''_i$  correspond to the same graph as  $D_i$ 
6       parent( $D'''_i$ )  $\leftarrow D'_i$ 
7       parent( $D''_i$ )  $\leftarrow D'_i$ 
8       parent( $D'_i$ )  $\leftarrow D_i$ 
9     end
10  end
11   $S \leftarrow$  the cactus of the 3-edge-connected components of  $X$ 
12   $\mathcal{E} \leftarrow$  the edge set of  $S[D_1 \cup \dots \cup D_k]$ 
13  merge all  $D_1, \dots, D_k$  into a new 3-ecc node
14  foreach edge  $e$  in  $\mathcal{E}$  do
15    insert( $e$ )
16  end
17 end

```

4.1. An $O(n^2 \log^2 n + m\alpha(m, n))$ -time algorithm for the incremental maintenance of \mathcal{T}

An efficient solution for (2) is to use the top-trees data structure of Alstrup et al. [4]. With this data structure we can perform nearest common ancestor queries in dynamic trees with n nodes in amortized $O(\log n)$ time per query. Furthermore, with top-trees we can also answer level ancestor queries within the same time-bounds. Thus we can solve (3) efficiently with a query for the $d+2$ level ancestor of V , where V is the leaf of \mathcal{T} that contains v , and d is the level of N . (We may assume that every node of \mathcal{T} has an attribute for its level on \mathcal{T} .)

The merging of nodes in (4) can be performed by redirecting the parents of the children of the nodes with the least number of children to the node with the most number of children (breaking ties arbitrarily), and then discarding these nodes. (We may assume that every node of \mathcal{T} has an attribute for its number of children.) In other words, suppose that we have to merge the nodes X_1, \dots, X_k , where X_1 has the greatest number of children among X_1, \dots, X_k . Then we redirect the parents of the children of X_2, \dots, X_k to X_1 , and then we discard X_2, \dots, X_k (and all the information associated with them). Since the nodes of every level of \mathcal{T} are $O(n)$, this procedure ensures that at most $O(n \log n)$ redirections can take place in total in every level. However, for every such redirection, a deletion and an insertion of an edge of \mathcal{T} must take place, and this takes $O(\log n)$ amortized time using the top-trees. Thus we get an $O(n \log^2 n)$ time-bound for every level of \mathcal{T} , and since there are at most $O(n)$ levels in \mathcal{T} , we get the $O(n^2 \log^2 n)$ time-bound in total.

Now, for (5), we assume, as above, that every 1-ecc node has an associated data structure that represents the tree of its 2-edge-connected components, and every 2-ecc node has an associated data structure that represents the cactus of its 3-edge-connected components. We use the associated data structures for the trees of the 2-edge-connected components in order to be able to perform efficiently the following. In Line 13 of Algorithm 1, we access the associated data structure T of N , in order to find its children X_1, \dots, X_k that we have to merge in Line 23. Thus we assume that the data structure for trees supports the operation `compressPath(T, x, y)`, which, given a pointer T to a tree, and pointers x and y to nodes of this tree, it finds the simple path P that connects these nodes on the tree, returns pointers to the nodes of P and the edges of P , and then merges all nodes of P into a new node and also returns a pointer to this node. We assume that the nodes and the edges of T have some information associated with them. Specifically, since every edge e of T corresponds essentially to a bridge of N , and this bridge is an edge (x, y) of G , we let e point to the edge (x, y) of G . This information is needed for Lines 15 and 20. Furthermore, since every node v of T corresponds essentially to a child V of N , we let v point to V . This is used precisely in order to find the children X_1, \dots, X_k of N that we have to merge. Conversely, every child of N must point to its corresponding node of T , in order to be able to find the endpoints of the path on T that we have to find with the call to `compressPath`. Finally, since we may have to merge some children of the root or of a 3-ecc node in Line 9 of Algorithm 1, we also have to link the associated (representations of the) trees of those children, with the addition of an edge that corresponds to the one that has been (re-)inserted to G . Specifically, suppose that an edge (x, y) is (re-)inserted to G , and that the nearest common ancestor N of X and Y is the root or a 3-ecc node, where X and Y are the leaves of \mathcal{T} that contain x and y , respectively. Let C_1 and C_2 the children of N that contain x and y , respectively. Then, in Line 9, we have to merge C_1 and C_2 into a new 1-ecc node C . Now, the data structure associated with C must represent a tree that is formed by linking the tree represented by the associated data structure to C_1 with the tree represented by the associated data structure to C_2 , through the addition of an edge $e = (u, v)$ with endpoints corresponding to the children of C_1 and C_2 that contain x and y , respectively. (These children of C_1 and C_2 can be found with a query for the $d + 2$ level ancestors of X and Y , respectively, where d is the level of N .) Furthermore, since e corresponds essentially to (x, y) , the associated information to e must be a pointer to (x, y) . Thus, the data structure for trees must support the operation `joinTrees($T_1, T_2, (u, v)$)`, which, given a pointer T_1 to a (representation of a) tree \tilde{T}_1 , a pointer T_2 to a (representation of a) tree \tilde{T}_2 , a pointer u to a node \tilde{u} of \tilde{T}_1 , a pointer v to a node \tilde{v} of \tilde{T}_2 , and some extra information to be associated with the edge (\tilde{u}, \tilde{v}) , links the trees \tilde{T}_1 and \tilde{T}_2 through the addition of the edge (\tilde{u}, \tilde{v}) .

The associated data structures for the cactuses of the 3-edge-connected components of 2-ecc nodes are used to perform the analogous operations that are performed by the data structures for trees. Specifically, in Lines 20 and 27 of Algorithm 1, we access the associated data structure for the 3-edge-connected components of a node X of \mathcal{T} , in order to find its children that we have to merge later on (with a call to procedure `merge3ecc`). Thus we assume that the data structure for cactuses supports the operation `compressCyclePath(S, x, y)`, which, given a pointer S to a cactus, and pointers x and y to nodes of this cactus, it finds the cycle-path Q on S with endpoints these nodes, it returns pointers to the nodes of Q and the edges of S between the nodes of Q , and then merges all nodes of Q into a new node and also returns a pointer to this node. We assume that the nodes and the edges of S have some information associated with them. Specifically, since every edge e of S corresponds essentially to an edge of X , and this edge is an edge (x, y) of G , we let e point to (x, y) . This information is needed for Line 12 of Algorithm 2. Furthermore, since every node v of S corresponds essentially to a child V of X , we let v point to V . This is used precisely in order to find the children of X that we have to merge. Conversely, every child of X must point to its corresponding node of S , in order to be able to find the endpoints of the cycle-path on S that we have to find with the call to `compressCyclePath`. Finally, since we may have to merge the children X_1, \dots, X_k of the 1-ecc node N in Line 23 of Algorithm 1, we also have to link the associated (representations of the) cactuses of those children, with the addition of a cycle that corresponds to the one that has been formed by the (re-)insertion of (x, y) to G . Specifically, when we merge X_1, \dots, X_k into a new 2-ecc node X , the associated cactus of X must be that which is formed by joining the cactuses associated with X_1, \dots, X_k with a cycle corresponding to D_1, \dots, D_k , where D_i , for $i \in \{1, \dots, k\}$, is the new 3-ecc node that has been formed by merging $D_{(i,1)}, \dots, D_{(i,t(i))}$ through the call `merge3ecc` in Line 21. Furthermore, the edges of the new cactus that correspond to the edges $(D_1, D_2), \dots, (D_{k-1}, D_k), (D_k, D_1)$ of X , must correspond to the edges $(x_1, y_1), \dots, (x_{k-1}, y_{k-1}), (y, x)$. This information is needed in Line 12 of Algorithm 2. Thus, the data structure for cactuses must support the operation `joinCactuses($S_1, \dots, S_k, (d_1, d_2), \dots, (d_k, d_1)$)`, which, given pointers S_1, \dots, S_k to (representations of) cactuses $\tilde{S}_1, \dots, \tilde{S}_k$, pointers d_1, \dots, d_k to nodes $\tilde{d}_1, \dots, \tilde{d}_k$ of $\tilde{S}_1, \dots, \tilde{S}_k$, respectively, and some extra information to be associated with the edges $(\tilde{d}_1, \tilde{d}_2), \dots, (\tilde{d}_k, \tilde{d}_1)$, links the cactuses $\tilde{S}_1, \dots, \tilde{S}_k$ through the addition of the cycle $(\tilde{d}_1, \tilde{d}_2), \dots, (\tilde{d}_k, \tilde{d}_1)$.

Algorithm 3: Update the decomposition tree \mathcal{T} after inserting a new vertex v to G .

```

1 procedure insert( $v$ )
2 begin
3   let  $C, C'$  and  $C''$  be a new 1-ecc, 2-ecc and 3-ecc node, respectively
4    $parent(C) \leftarrow root$  of  $\mathcal{T}$ 
5    $parent(C') \leftarrow C$ 
6    $parent(C'') \leftarrow C'$ 
7   let the representative of  $v$  in  $DSU_{3ecc}$  be  $C''$ 
8   initialize a new trivial tree  $T$ , and associate it with  $C$ 
9   initialize a new trivial cactus  $S$ , and associate it with  $C'$ 
10 end

```

In Section 5 we provide efficient implementations for the associated data structures for trees and cactuses. These implementations use size $O(n)$ for a collection of trees or cactuses with n nodes, and can perform any sequence of operations in $O(n \log n)$ time in total. We use a data structure for trees in every $3k + 1$ level of the tree, and a data structure for cactuses in every $3k + 2$ level of the tree. Since the number of nodes of the trees (of the 2-eccs) or of the cactuses (of the 3-eccs) that correspond to the nodes of every level of \mathcal{T} can be at most n , and there are at most $O(n)$ levels on \mathcal{T} , we thus have that the operations in the associated data structures can take time $O(n^2 \log n)$ in total for any sequence of edge insertions to G .

Finally, Algorithm 3 shows how we can handle the insertion of a new vertex v to G . We simply introduce three new nodes C, C' , and C'' (an 1-ecc node, a 2-ecc node, and a 3-ecc node, respectively), that correspond to v , and we set $parent(C) \leftarrow root$, $parent(C') \leftarrow C$, and $parent(C'') \leftarrow C'$. Since the leaf of \mathcal{T} that contains v is C'' , we also set the representative $find_{3ecc}(v) \leftarrow C''$. We also associate a (representation of a) trivial tree to C , and a (representation of a) trivial cactus to C' . The DSU data structure and the data structures in Section 5 allow for new node insertions, and the same time-bounds hold (where n is interpreted as the total number of vertices that will have been inserted to G at the moment we estimate the time that it took to perform all the operations of the algorithm so far).

4.2. An $O(n^2\alpha(n, n) + m\alpha(m, n))$ -time algorithm for the incremental maintenance of \mathcal{T}

Now we will describe a more time-efficient algorithm to handle insertions to the graph. (This algorithm, however, uses $O(n^2)$ space.) First, we can handle the operations in (5) exactly as we did in Section 4.1, although here we use the more sophisticated data structures for trees and cactuses that are described in Section 6. The implementations given in Section 6 use size $O(n)$ for a collection of trees or cactuses with n nodes, and can perform any sequence of operations in $O(n\alpha(n, n))$ time in total. Since we use a data structure for trees in every $3k + 1$ level, and a data structure for cactuses in every $3k + 2$ level, and there are at most $O(n)$ levels in \mathcal{T} , we thus we get the $O(n^2\alpha(n, n))$ expression in the total time-bound.

To perform the merging of nodes in (4) we use an optimal disjoint set union data structure DSU_i [54], on every level i of \mathcal{T} , that operates on the node set of that level and uses as representatives nodes on that level that we consider to be “active”. The function of an active node is to represent all the nodes that have been merged with it; initially, all nodes of \mathcal{T} are active, and, throughout, the parent pointer on \mathcal{T} exists only for the active nodes. DSU_i supports the operations $find_i(x)$ and $unite_i(x, y)$. $find_i(x)$ returns a pointer to the active node that has been merged with x , and $unite_i(x, y)$ merges the sets of nodes that are represented by the active nodes x and y , and sets as representative one of x or y (while deactivating the other). Thus, in order to find the parent of an active node X of \mathcal{T} on level i , we use $find_{i-1}(parent(X))$. DSU_i uses $O(n)$ space, and it can perform any sequence of m operations $find_i$ and $unite_i$ in $O(m\alpha(m, n))$ time in total [54].

To perform efficiently the operations (2) and (3), we basically maintain, for every inter-connection edge $e = (x, y)$ of G (that is, for every edge e whose endpoints, at the time of its insertion, lie in different maximal 3-edge-connected subgraphs of G), two paths of \mathcal{T} whose nodes contain an endpoint of e , that start from the leaves of \mathcal{T} that contain x and y , and go up until at least their nearest common ancestor. In order to achieve this, we augment the information associated with \mathcal{T} as follows. For every node Z of \mathcal{T} we associate a linked list L_Z , where every element of L_Z corresponds to an endpoint of an inter-connection edge that lies in an ancestor of Z . Thus, the existence of those lists means that we may need as much as $O(n^2)$ space (since the number of inter-connection edges can be at most $O(n)$, and the number of levels of \mathcal{T} is $O(n)$). Every element of L_Z is a pointer that points to an element of L_C , for some child C of Z . Also, every element of L_Z has a pointer to Z (so that we can find in constant time the node in whose associated list this pointer lies). Thus we may say that a pointer of L_Z lies in Z . Finally, for every inter-connection edge $e = (x, y)$, we maintain two pointers e_x and e_y , that point to the elements of a list L_Z , for some node Z that contains both x and y .

Now we work as follows. Let $e = (x, y)$ be an edge that is inserted to G for the first time. (I.e., this is the first time that we call $insert(e)$ of Algorithm 1. Notice, of course, that an edge $e' = (x, y)$ may have previously been inserted to G , since we allow for multiple edges. But now we consider e as a new insertion. Furthermore, we can easily determine whether $insert(e)$ is the first time we call $insert$ on e , by using a flag every time we call $insert$ from Line 35 of Algorithm 1 or Line 15 of Algorithm 13, in order to signify that this is a re-insertion.) First we find the leaves X and Y of \mathcal{T} that contain x and y , respectively, by calling $find_{3ecc}(x)$ and $find_{3ecc}(y)$, respectively. Suppose that $X \neq Y$ (for otherwise there is nothing to do). Then we crawl up the tree, starting from X and Y , by following the parents alternately, marking all the nodes that

we meet, and keeping them in a stack in order to unmark them later. When we meet a node N that is already marked, we have that this is the nearest common ancestor of X and Y on \mathcal{T} , and we unmark all the nodes that we traversed. Now we start again from X , and we follow the parents until we reach N . For every node Z that we traverse, including N , we append a new pointer $Z_{(e,x)}$ to L_Z . If $Z = X$, then we let $Z_{(e,x)}$ point to *null*. Otherwise, let C be the child of Z that we traversed before reaching Z ; then we let $Z_{(e,x)}$ point to $C_{(e,x)}$. We repeat the same process starting from Y . Finally, we let e_x point to $N_{(e,x)}$ and we let e_y point to $N_{(e,y)}$. Observe that this process may need as many as $O(n)$ calls to the parent function, and $O(n)$ pointer manipulation operations. Thus, since the total number of first-time edge insertion in any sequence of edge insertions is $O(n)$, and every call to the parent function involves a call to $find_i$, for some i , the total cost of this process in the running time of the algorithm is $O(n^2\alpha(n, n))$.

Now suppose that we re-insert an edge $e = (x, y)$. Then we can find the nearest common ancestor of the leaves X and Y of \mathcal{T} that contain x and y , respectively, by following the path of pointers starting from e_x and e_y , and moving on to deeper levels in parallel, until we reach two pointers that lie on two nodes C and D that have not got merged together, in which case the common parent of C and D is the nearest common ancestor of X and Y on \mathcal{T} . Then, since C and D will have to get merged, we let e_x point to $C_{(e,x)}$ and e_y point to $D_{(e,y)}$, in order to avoid traversing the same path of pointers again. Since the total number of edges that can affect the decomposition tree, in any sequence of edge insertions, is $O(n)$, and every time an edge is re-inserted it is moved into deeper levels, observe that the total cost of this process in the running time of the algorithm is $O(n^2\alpha(n, n))$. Now suppose that we want to find the grandchild of N that contains x . (We can follow the analogous procedure in order to find the grandchild of N that contains y .) Then we can use the pointer $N_{(e,x)}$ stored in L_N . This points to a pointer $C_{(e,x)}$ in L_C , for a node C that is merged with a child C' of N . And then $C_{(e,x)}$ points to $D_{(e,x)}$ in L_D , for a node D that is merged with a child D' of C' . Thus we can find D in constant time, and then we can find D' with a call $find_i(D)$, where i is the level of D . Again, since there can be at most $O(n^2)$ re-insertions of edges, the total cost of this process in the running time of the algorithm is $O(n^2\alpha(n, n))$.

All the data structures with which \mathcal{T} is equipped (the DSU data structures, and the associated data structures for trees and cactuses), allow for new node insertions, and the same time-bounds hold. Thus, the insertion of new vertices to G is performed as shown in Algorithm 3.

5. Data structures for trees and cactuses

To maintain the decomposition tree of the maximal 3-edge-connected subgraphs, we will need the following two data structures, which we may use as black boxes in the algorithm. They operate on collections of trees and cactuses, respectively. We assume that there is some information associated with every edge of every tree or cactus.⁶

The first is a data structure that works on a collection of trees and supports the following operations.

- $compressPath(T, x, y)$. Given a tree T and nodes x and y on T , this operation finds the simple path P that connects x and y on T and compresses it into a new node z (converting T into a smaller tree). Furthermore, it returns pointers to all nodes and (information associated with the) edges on P , as well as to the new node z . Finally, if (u, v) was an edge of T , prior to the call of this operation, with $u \notin P$ and $v \in P$, then (u, z) (in the new tree) must contain the same information that was previously associated with (u, v) .
- $joinTrees(T_1, T_2, (x, y))$. Given two trees T_1, T_2 and nodes $x \in T_1$ and $y \in T_2$, this operation links the trees T_1 and T_2 by introducing a new edge (x, y) .

The second data structure works on a collection of cactuses and supports the following operations.

- $compressCyclePath(S, x, y)$. Given a cactus S and nodes x and y on S , this operation finds the “cycle-path” P that connects x and y on S . It compresses all nodes on P into a new node z by properly squeezing the involved cycles. Furthermore, it returns pointers to all nodes on P , to all (information associated with the) edges that connect pairs of those nodes (if there are any), as well as to the new node z . Finally, if (u, v) was an edge of S , prior to the call of this operation, with $u \notin P$ and $v \in P$, then (u, z) (in the new cactus) must contain the same information that was previously associated with (u, v) .
- $joinCactuses(S_1, \dots, S_k, (x_1, x_2), \dots, (x_k, x_1))$. Given a collection of cactuses S_1, \dots, S_k and nodes $x_i \in S_i$, $i \in \{1, \dots, k\}$, this operation links the cactuses S_1, \dots, S_k into a larger cactus by introducing the edges $(x_1, x_2), \dots, (x_k, x_1)$ (which induce a new cycle).

In Sections 5.1 and 5.2, we provide implementations for data structures maintaining trees and cactuses where we can perform any sequence of operations in $O(n \log n)$ time, assuming that the total number of nodes is at most n . To achieve this time-bound, we will draw ideas from [28], [41], and [59]. Using these datastructures, we spend an additional $O(n^2 \log n)$ time for the incremental maintenance of the decomposition tree. To see this, recall that we use representations for the trees of the 2-edge-connected components at levels $3k + 1$ of the decomposition tree (where $0 \leq k < n$), and representations for

⁶ For our algorithm, this information is a pointer to an edge of the original graph (whose maximal 3-edge-connected subgraphs we want to maintain).

the cactuses of the 3-edge-connected components at levels $3k+2$ (where $0 \leq k < n$). Then, every operation `compressPath`, `joinTrees`, `compressCyclePath` and `joinCactuses` is performed on structures on the same level. Since there are $O(n)$ levels on which these operations may be performed, and the total number of nodes at any level is $O(n)$, we get the $O(n^2 \log n)$ time bound.

In Section 6, we use the more sophisticated implementations for those data structures provided by [40], in order to get a $O(n\alpha(n, n))$ -time bound for performing any sequence of operations on trees or cactuses (on the same level), which gives $O(n^2\alpha(n, n))$ in the total time bound of the algorithm for maintaining the decomposition tree.

5.1. An implementation for trees

(a) Representation

An idea is to represent the trees as rooted trees. For every node x we let $p(x)$ denote its parent (if x is the root, we leave $p(x)$ undefined). Also, every node x has a pointer $x.edge$ to the edge $(x, p(x))$ that connects x with its parent. We use a DSU data structure to perform mergings of nodes. The DSU data structure supports the operations `find(x)` and `unite(x, y)`. Thus, in order to access the parent of node x we actually use `find(p(x))` (although we will not write this explicitly in what follows).

(b) Operations

We execute `compressPath(T, x, y)` as follows. We start from x and y , and we alternately climb up the tree following the parents, marking all nodes that we traverse, until we meet an already marked node z . (This ensures that we traverse at most $2d - 1$ nodes, where d is the number of nodes of the simple path connecting x and y .) Then there are two possibilities: either (1) $z = x$ or $z = y$, or (2) z is neither x nor y . In the first case, assuming w.l.o.g. that $z = y$, the path that we have to compress consists of the nodes $x, p(x), p(p(x)), \dots, y$. In the second case, the path consists of two parts: $x, p(x), \dots, z$ and $y, p(y), \dots, z$. In case (1) we work as follows. We start from x ; we retrieve the edge $(x, p(x))$ (using the pointer stored in x), we call `unite(x, p(x))`, setting as representative of the new set $p(x)$, and we repeat, until we reach $z = y$. In case (2) we work as in case (1) twice, starting from both x and y , until we reach z each time.

The `joinTrees($T_1, T_2, (x, y)$)` operation works as follows. First, we determine which of T_1, T_2 is smaller (w.r.t. its number of nodes). This can be done easily if we maintain for every tree an attribute `size` indicating the number of its nodes. Assume w.l.o.g. that the smallest tree is T_1 (or that its size is equal to that of T_2). Then we reroot T_1 at x , and we set $p(x) \leftarrow y$, $x.edge \leftarrow (x, y)$, and $T_2.size \leftarrow T_2.size + T_1.size$. The rerooting of T_1 at x is performed as follows. First we find the path P connecting x with the root r of T_1 , by following the parents. Then we process the nodes of P in reverse order, starting from r until we reach x (which we do not process), and we simply hand over the information of every child to its parent. To be precise, let v be a node that we process, and let u be the node on P that has $p(u) = v$. Then we set $p(v) \leftarrow u$ and $v.edge \leftarrow u.edge$.

(c) Analysis

For convenience in the time analysis, we may use a simple implementation for the DSU data structure which performs all unions in $O(n \log n)$ time and every `find` in $O(1)$ worst-case time. Now, the total time it takes to perform any sequence of `compressPath` operations is $O(n \log n)$. This is because, every time `compressPath(T, x, y)` is performed, it takes $O(|P|)$ time to find the simple path P that connects x and y , using $O(|P|)$ calls to the `find` operation. Then we merge the nodes of P (thereby deleting them essentially), so they will not be traversed again. Thus, the total time to find all paths, in all `compressPath` calls, is $O(n)$. Therefore, the time bound for all `compressPath` operations is dominated by that for performing the mergings (using the `union` operation), and thus it is $O(n \log n)$.

When we perform a `joinTrees` operation on two trees, we may have to reroot the smallest tree T , which means that we may have to traverse all of its nodes. This takes $O(|T|)$ time and uses $O(|T|)$ calls to `find`. However, a node can only be accessed during a rerooting (that is, it may lie on the smallest tree) at most $O(\log n)$ times (because the trees are subsequently joined to produce a bigger tree whose size is at least twice that of the smallest). Thus, we need at most $O(n \log n)$ steps to perform all rerootings, using at most $O(n \log n)$ calls to the `find` operation. Thus, the total time, for any sequence of `compressPath` and `joinTrees` operations, is at most $O(n \log n)$.

5.2. An implementation for cactuses

(a) Representation

In order to efficiently determine the cycle-path connecting two nodes of a cactus, we represent the cactus as a rooted tree with “real nodes” and “cycle nodes”. Every real node corresponds to a cactus node, and every cycle node corresponds to a cycle. Thus, for every cactus S , there is a *cactus tree* T on the set of real and cycle nodes corresponding to S , which is formed by adding edges corresponding to the incidence relation of the real nodes to the cycle nodes. Specifically, there is an edge (x, C) in T , where x is a real node and C is a cycle node, if and only if the cactus node corresponding to x lies on the cactus cycle corresponding to C . We root this tree on an arbitrary real node. Thus, the parent of a real node (different from the root) is a cycle node, and the parent of a cycle node is a real node. The parent of a (real or cycle) node v is denoted as

$p(v)$. We use a DSU data structure to perform mergings of real nodes. This data structure supports $find(x)$ and $unite(x, y)$. Thus, in order to access the parent of a cycle node C , we use $find(p(C))$.

In order to retrieve the information associated with the cactus edges and to efficiently perform the “cycle-squeezing” operation, we use doubly-linked circular lists to represent the cycles of the cactuses. Specifically, for every cycle C of a cactus, there is a doubly-linked circular list L consisting of as many nodes as there are in C . The nodes of L correspond to the nodes of C and are ordered accordingly. Every node x of L has pointers $x.left$ and $x.right$ to its neighbouring nodes, and pointers $x.leftEdge$ and $x.rightEdge$ to the respective cactus edges.

Finally, we have the following correspondence between the nodes of the cactus trees and the nodes of the circular lists. Let x be a real node of a cactus tree with parent C . Then C corresponds to a cactus cycle and therefore to a circular list L . Then we have a pointer for x to its corresponding node on L , and vice versa. Furthermore, $p(C)$ also corresponds to a node on L , but we do not keep a pointer from $p(C)$ to this node (because it cannot be uniquely determined from $p(C)$, as this can be the parent of many cycle nodes). Instead, we have a pointer from C to the node of L corresponding to $p(C)$. For notational convenience, we may use the same name to denote a real node and its corresponding circular list node. Also, for a cycle node C , we may denote the circular list node pointed to by C as $p(C)$.

(b) Operations

compressCyclePath

To perform $compressCyclePath(S, x, y)$ we work as follows. First, we determine the cycle-path P connecting x and y . To do this, we work on the cactus tree associated with S ; we start from x and y , and we alternately climb up the tree following the parents, marking all the nodes that we traverse, until we meet an already marked node z . Then there are two possibilities: either (1) $z = x$ or $z = y$, or (2) z is neither x nor y . In the first case, assuming w.l.o.g. that $z = x$, P is given by the real nodes on the ascended path starting from x . In the second case, P is given by two parts: by the real nodes on the path starting from x and ending in z , and by the real nodes on the path starting from y and ending in z . (Note that z can be either a cycle node or a real node.) In case (1), assuming w.l.o.g. that $z = x$, let $u_1, C_1, \dots, C_k, u_{k+1}$, where $u_1 = x$ and $u_{k+1} = y$, be the cycle-path connecting x and y , including the cycles nodes that are connected with every two consecutive real nodes. Then we only have to perform the operations $squeezeCycle(u_i, u_{i+1}, C_i)$, for every $i \in \{1, \dots, k\}$, successively. In case (2), we work as in case (1) for every one of the two parts that form P . Furthermore, in the case that z is a cycle node, we also have to perform $squeezeCycle(u, v, z)$, where u and v are the second-last nodes of the two parts that form P . In particular, let u_1, C_1, \dots, u_k, z and v_1, C'_1, \dots, v_l, z , where $u_1 = x$ and $v_1 = y$, be the two paths that we ascended, starting from x and y , respectively. Then, in addition to all operations $squeezeCycle(u_i, u_{i+1}, C_i)$, for every $i \in \{1, \dots, k-1\}$, and $squeezeCycle(v_i, v_{i+1}, C'_i)$, for every $i \in \{1, \dots, l-1\}$, we also have to perform $squeezeCycle(u_k, v_l, z)$.

Now let us describe in detail the operation $squeezeCycle(u, v, C)$, where u and v are real nodes connected with a cycle node C . (Intuitively, we want this operation to “squeeze” the cycle C by merging its two nodes u and v .) We distinguish two cases, depending on whether (i) u and v are related as ancestor and descendant, or (ii) u and v are siblings on the rooted cactus tree (the later occurs in case (2) above, when z is a cycle node).

(i) u and v are related as ancestor and descendant

We may assume w.l.o.g. that v is an ancestor of u . Thus we have $p(u) = C$ and $p(C) = v$. We distinguish three cases, depending on whether u has no siblings, or u is the leftmost or rightmost child of C (to be defined shortly), or neither of the previous two cases holds. First, if u has no siblings, then we simply return the two edges that connect u and v , we discard the cycle C , and we merge u and v , setting as representative of the union v . To return the two edges we use the pointer from u to the node in the circular list corresponding to C (let us call it u again). Then the two edges are precisely $u.leftEdge$ and $u.rightEdge$. Now let us consider the case that u is the leftmost or rightmost child of C . This is determined by checking whether $u.left$ or $u.right$, respectively, is the node pointed to by C (which essentially corresponds to v). So let us assume that $u.left = v$ (the case that $u.right = v$ is treated analogously). The first thing to do is to return the edge that connects u and v . This is simply $u.leftEdge$. Then we discard the node u from the circular list and we update it accordingly. That is, we set $u.right.left \leftarrow v$, $v.right \leftarrow u.right$, and $v.rightEdge \leftarrow u.rightEdge$. Finally, we merge u and v (the real nodes on the cactus tree), setting as representative of the union v .

Now we consider the third case. That is, we have that $u.left \neq v$ and $u.right \neq v$. (Notice that, in this case, there is no cactus edge connecting u and v .) Now the first thing to do is to determine the smallest segment of the circular list L corresponding to C that contains u and v . To achieve this efficiently, we start from u and v , and we alternately follow the $left$ pointers, marking all the nodes that we traverse, until we meet an already marked node z . Then we have that either $z = v$ or $z = u$. Let us assume that $z = v$ (the other case is treated analogously). Now we update the cactus tree as follows. Let P be the internal path on L connecting u and v following the $left$ pointers starting from u (that is, $P = u.left, u.left.left, \dots, v.right$). We introduce a new cycle node C' , and we let C' be the parent of all the real nodes that correspond to nodes of P . The parent of C' is set to be v , and we introduce a new circular list node \tilde{v} , corresponding to v on the cycle C' . Thus, there is a pointer from C' to \tilde{v} . Now we detach P from L and we attach it appropriately to the circular list corresponding to C' . This is done by setting $v.right.left \leftarrow \tilde{v}$, $\tilde{v}.right \leftarrow v.right$, $u.left.right \leftarrow \tilde{v}$, and $\tilde{v}.left \leftarrow u.left$. To maintain the cactus edges, we also set $\tilde{v}.rightEdge \leftarrow v.rightEdge$ and $\tilde{v}.leftEdge \leftarrow u.leftEdge$. In the list L we simply discard u and we update L accordingly. That is, we set $u.right.left \leftarrow v$ and $v.right \leftarrow u.right$. Again, to maintain the cactus edge, we also set $v.rightEdge \leftarrow u.rightEdge$. Finally, we merge u and v (the real nodes on the cactus tree), setting as representative

of the union v .

(ii) u and v are related as siblings

Let us assume first that (the circular list nodes corresponding to) u and v are neighbours. This means that either $u.left = v$ or $u.right = v$. Suppose that $u.left = v$ (the other case is treated similarly). First we return $u.leftEdge$: the edge that connects u and v . Then we discard u from the circular list and we update it accordingly. That is, we set $v.right \leftarrow u.right$, $u.right.left \leftarrow v$, and $v.rightEdge \leftarrow u.rightEdge$. Finally, we merge u and v (on the cactus tree), setting as representative of the union v .

Now suppose that u and v are not neighbours. This means that $u.left \neq v$ and $u.right \neq v$. (Notice that, in this case, there is no cactus edge connecting u and v .) Let C be the common parent of u and v , and let $w = p(C)$. Let also L be the circular list corresponding to C , and denote as w the node of L pointed to by C . Now the first thing to do is to determine the smallest part of L that contains u and v . To achieve this efficiently, we start from u and v , and we alternately follow the *left* pointers, marking all the nodes that we traverse, until we meet an already marked node z . Then we have that either $z = v$ or $z = u$. Let us assume that $z = v$ (the other case is treated analogously). Let P be the internal path on L connecting u and v following the *left* pointers starting from u (that is, $P = u.left, u.left.left, \dots, v.right$). Here we have that either $w \notin P$ or $w \in P$. If $w \notin P$, then we work precisely as in the second paragraph of case (ii) above. So let us assume that $w \in P$. Then we introduce a new cycle node C' , we let C' be the parent of all the real nodes that correspond to nodes of $P \setminus w$, and we let C' point to w . Then we introduce a new circular list node \tilde{v} (corresponding to the node that will be formed by merging u and v), we let C point to \tilde{v} , and we attach the segment $L \setminus (P \cup \{u, v\})$ of L to \tilde{v} . Again, this is done precisely as in the second paragraph of case (ii) above. Furthermore, we discard the circular list node u , and we properly update the list corresponding to C' . That is, we set $v.right \leftarrow u.right$, $u.right.left \leftarrow v$, and $v.rightEdge \leftarrow u.rightEdge$. Finally, we let the parent of C be v , and we merge u and v (the real nodes on the cactus tree), setting as representative of the union v .

`joinCactuses`

To perform `joinCactuses`($S_1, \dots, S_k, (x_1, x_2), \dots, (x_k, x_1)$), we first have to find the largest (w.r.t. its number of nodes) cactus among S_1, \dots, S_k . This can be done easily if we keep for every cactus an attribute *size*, signifying its number of nodes. Now let us assume, w.l.o.g., that one of the cactuses with the greatest size is S_k . Then we have to do two things: to reroot the cactuses S_1, \dots, S_{k-1} on their nodes that are to be connected on a new cycle, and then to form the cycle $(x_1, x_2), \dots, (x_k, x_1)$.

Let us describe how to perform a rerooting of a cactus S at a node $x \in S$. We assume that x is not the root of the cactus tree of S (for otherwise there is nothing to do). First we find the path on the cactus tree that connects x with the root by following the parents. This has the form $u_1, C_1, \dots, C_{t-1}, u_t$, where u_1, \dots, u_t are real nodes with $u_1 = x$, and C_1, \dots, C_{t-1} are cycle nodes. Then we process this path in reverse order, and for every triple (u, C, v) , where u, v are real nodes and C is a cycle node, we do the following. (Due to reverse processing, notice that we have $p(v) = C$ and $p(C) = u$.) Let L be the circular list corresponding to C . Then we simply change the pointer of C , so that it points to the node of L corresponding to v , and we reverse the parent relation between the nodes of (u, C, v) . (Thus, we set $p(u) \leftarrow C$ and $p(C) \leftarrow v$.) Finally, we update the status of x so that it is recognized as the root.

Thus we may assume that the cactus tree of every S_i , $i \in \{1, \dots, k-1\}$, is rooted at x_i . Now we introduce a new cycle node C , and we set the parent of every x_i , $i \in \{1, \dots, k-1\}$, to be C , and the parent of C to be x_k . Then we introduce new circular list nodes $\tilde{x}_1, \dots, \tilde{x}_k$, pointers between x_i and \tilde{x}_i (in both directions), for every $i \in \{1, \dots, k-1\}$, and we let C point to \tilde{x}_k . All nodes \tilde{x}_i are linked in a circular structure according to their ordering on the new cycle. That is, we set $\tilde{x}_i.left \leftarrow \tilde{x}_{i-1}$, for $i \in \{2, \dots, k\}$, $\tilde{x}_1.left \leftarrow \tilde{x}_k$, $\tilde{x}_i.right \leftarrow \tilde{x}_{i+1}$, for $i \in \{1, \dots, k-1\}$, and $\tilde{x}_k.right \leftarrow \tilde{x}_1$. Furthermore, we fix the pointers to the newly introduced edges. That is, we set $\tilde{x}_i.leftEdge \leftarrow (x_{i-1}, x_i)$, for $i \in \{2, \dots, k\}$, $\tilde{x}_1.leftEdge \leftarrow (x_k, x_1)$, $\tilde{x}_i.rightEdge \leftarrow (x_i, x_{i+1})$, for $i \in \{1, \dots, k-1\}$, and $\tilde{x}_k.rightEdge \leftarrow (x_k, x_1)$. Finally, we set $S_k.size \leftarrow S_1.size + \dots + S_{k-1}.size$, and the description of `joinCactuses`($S_1, \dots, S_k, (x_1, x_2), \dots, (x_k, x_1)$) is complete.

(c) Analysis

First, if we use a simple implementation for the DSU data structure which performs all unions in $O(n \log n)$ time and every *find* in $O(1)$ worst-case time, we can argue as in the analysis for the data structure for trees, that the total time it takes to find and merge all cycle-path during any sequence of `compressCyclePath` operations is $O(n \log n)$. Thus we only have to consider the total time it takes to perform all `squeezeCycle` operations. Recall that `squeezeCycle`(u, v, C) has to find the smallest segment P of the circular list corresponding to C that contains u and v (all other operations of `squeezeCycle`(u, v, C) take $O(1)$ time in total). This is performed in $O(|P|)$ time, and then P forms a new circular list on its own (plus at most one more node). We say that the two new cycles into which C was squeezed are *formed by* C . (Thus, we may observe that after any repeated application of `squeezeCycle` operations on C and on cycles formed by C , we have that all cycles formed by C form a cactus.) We overload our terminology by saying that a node of a cycle formed by C is also a node of C . Now we may argue as follows. Fix a cycle C of size k and a node $x \in C$. Then we observe that in any sequence of `squeezeCycle` operations on C or on cycles formed by C , x can be part of the smallest segment (on a circular list) explored by `squeezeCycle` at most $O(\log k)$ times. Thus, the total time to perform any sequence of `squeezeCycle` operations on C or on cycles formed by C is $O(k \log k)$. Now observe that the cycles are introduced into the collection of cactuses by the operation `joinCactuses`; and that once a cycle has been introduced, it can only get

squeezed to form smaller cactuses. Thus, let k_1, \dots, k_t be the sizes of all cycles that have been introduced in the data structure by the operation *joinCactuses*. Then we have that the total time to perform all *squeezeCycle* operations is $O(k_1 \log k_1 + \dots + k_t \log k_t) = O(k_1 \log n + \dots + k_t \log n) = O((k_1 + \dots + k_t) \log n) = O(n \log n)$, since $k_i = O(n)$, for every $i \in \{1, \dots, t\}$, and $k_1 + \dots + k_t = O(n)$.

When we perform a *joinCactuses* operation on some cactuses, we may have to reroot the smallest ones, which means that we may have to access all of their nodes. However, we can argue as in the analysis of the data structure for trees, in order to show that the total time-bound for all rerootings is $O(n \log n)$. All the other operations performed in *joinCactuses* take time analogous to the number of the cactuses involved. Since the total size of all cycles that we can introduce (while maintaining a collection of cactuses on n nodes) is $O(n)$, this shows that the total time for any sequence of *joinCactuses* operations is $O(n \log n)$.

6. Improved data structures for trees and cactuses

Here we provide data structures and algorithms for the operations *compressPath*, *joinTrees*, *compressCyclePath* and *joinCactuses*, on collections of trees and cactuses with at most n nodes, with time bounds better than $O(n \log n)$. To be precise, any sequence of m operations on collections of trees or cactuses with at most n nodes can be performed in $O((n+m)\alpha(m, n))$ time in total. The analysis for this time bound is essentially contained in [41] and [40]; our own additions in the algorithms of [41] and [40] involve only a worst-case $O(1)$ calls to some DSU operations and a worst-case $O(1)$ pointer manipulations, for every operation performed. Also, the arguments that establish correctness are sufficiently contained in the description of the algorithms.

Now, to achieve this time bound, we will basically use the data structures provided by La Poutré et al. [41], [40], with some minor additions to suit our purposes. These data structures rely on the so-called “fractionally rooted trees” (FRT), introduced in [40]. In the next section we give a brief overview of the operations supported by FRTs. Then we describe the implementations for collections of trees and cactuses. We only give as many details of the data structures and algorithms of [41] and [40] as are needed in order to show where our own additions fit in and establish our results.

6.1. Fractionally rooted trees

The FRT data structure operates on a forest F equipped with a partition of its edges such that the classes corresponding to this partition induce subtrees of F . Before we describe the operations supported by FRT, we introduce some terminology. We say that a node x of F belongs to an edge class if it is incident to at least one edge of that class. Also, we say that an edge class intersects a path P if P contains at least one edge of that class. Now let x, y be two nodes of a tree of F , and let P be the simple path that connects them. We call a node z on P a *boundary node* of P if it is either one of x, y or it is incident to two edges of P which belong to different classes. A *boundary edge set* for a boundary node z on P is a set of (0, 1 or 2) edges incident to z , one from each edge class to which z belongs and which intersects P . (We do not demand that the edges in a boundary edge set for z lie on P ; however, one of their endpoints must be z .) A *boundary list* for P is a list consisting of the boundary nodes of P , where each boundary node has a sublist that contains a boundary edge set for it on P . Now let L be a list of nodes where each node has a sublist of edges. We say that an edge class *occurs* in L , or that L *contains* an edge class, if there is an edge of that class in some sublist of L . The edge classes of L are precisely those that occur in it. A *joining list* J is a list of nodes with sublists of edges such that the union of the classes occurring in J induces some subtree in F . (We note that this is always the case for a boundary list of a path.) In addition, the nodes in J must be the nodes belonging to at least two edge classes occurring in J , and the sublist for each node contains an edge for each class in J to which this node belongs. (Thus we have that a boundary list of a path is a joining list if we remove the endpoints of the path from the list.)

Now the operations supported by FRT are the following.

- *link*(x, y). Let x, y be two nodes lying on different trees of F . Then join the two trees by introducing a new edge (x, y) in the FRT data structure.
- *boundary*(x, y). Let x, y be two nodes lying on the same tree of F . Then return a boundary list for the simple path with endpoints x and y .
- *joinClasses*(J). Let J be a joining list. Then join all the edge classes of which an edge occurs in the list.

We say that a call *boundary*(x, y) is *essential* if there are at least two different edge classes that intersect the simple path that connects x and y . (Equivalently, a call *boundary*(x, y) is essential if the boundary list that it provides contains more than two nodes.)

By [40], we have the following result.

Theorem 6.1. (Theorem 9.2 in [40]) *There is an implementation of fractionally rooted trees with the following guarantee. Suppose that we start with an empty forest, and we perform n insertions of nodes and m operations *link*, *boundary*, and *joinClasses*, where every essential call of *boundary* is immediately followed by *joinClasses* on a joining list that contains all the edge classes that occur in the list provided by *boundary*. Then all these operations can be performed in total $O((n+m)\alpha(m, n))$ time.*

6.2. An implementation for trees

(a) Representation

For the operations `compressPath` and `joinTrees` on collections of trees, we essentially use the solution of La Poutré for the incremental maintenance of the 2-edge-connected components in general graphs with asymptotically optimal time complexity [40].

Thus we represent the collection of trees \mathcal{C} as a forest F , which is implemented as an FRT data structure. Every tree $T \in \mathcal{C}$ corresponds to a tree T_{FRT} of F . The edges of F are partitioned into edge classes, where every edge class induces a subtree of F . Some edges of F are called *bridges* and belong to singleton classes which are called *quasi* classes. All other edge classes are called *real*. There is an one-to-one correspondence between the edges of a tree $T \in \mathcal{C}$ and the bridges of T_{FRT} . Thus every edge of T_{FRT} that is a bridge contains a pointer to its corresponding edge on T . Every node x of T_{FRT} belongs to at most one real class. For every node x that belongs to a real class, we maintain an edge x_{assoc} of that class. By shrinking the subtrees that are induced by the real classes into single nodes, we get a natural isomorphism between T and T_{FRT} . (The idea here is that the subtrees induced by the real classes correspond to maximal sets of nodes that got merged due to the operation `compressPath`.) Thus, every node $x \in T$ corresponds to a subset S of T_{FRT} , and we associate x with a node x_{FRT} of S ; conversely, every node in S is associated with x . To implement the later, we use a DSU data structure DSU_F on the nodes of F , where the representatives of the sets are nodes of the trees of \mathcal{C} . DSU_F supports the operations $find_F$ and $unite_F$. Thus, in order to find which node of \mathcal{C} corresponds to a node u of F , we use $find_F(u)$. DSU_F (and every other DSU data structure in the sequel) is implemented using rooted trees, with path compression and union by size, thus achieving the asymptotically optimal time complexity [54].

We can summarize the above properties as follows.

Property 6.2. *Let x, y be two nodes of a tree $T \in \mathcal{C}$, and let P be the simple path on T that connects x and y . Let also \tilde{P} be the simple path on T_{FRT} that connects x_{FRT} and y_{FRT} . Then there is an one-to-one correspondence between the edges of P and the edges of \tilde{P} that are bridges. This correspondence is compatible with that between the nodes of T and the nodes of T_{FRT} . In other words, for every bridge (u, v) on \tilde{P} , there is an edge $(find_F(u), find_F(v))$ on P .*

(b) Operations

`compressPath`

Let x and y be two distinct nodes of a tree $T \in \mathcal{C}$, and let P be the simple path on T that connects x and y . To perform $compressPath(T, x, y)$ we work as follows. First we call $boundary(x_{FRT}, y_{FRT})$ to obtain a boundary list L for the simple path \tilde{P} that connects x_{FRT} and y_{FRT} on T_{FRT} . According to Property 6.2, in order to retrieve the edges of P , it is sufficient to identify the edges of \tilde{P} that are bridges. Since bridges belong to singleton (quasi) classes, we have that every bridge of \tilde{P} must lie in the sublist of some node in L . Thus we can retrieve the bridges of \tilde{P} by scanning the sublists of L . (Incidentally, we note here that the efficiency of the operation $boundary$ lies in the fact that it only computes boundary nodes and corresponding edge classes, without always traversing the entire path.) Furthermore, using again Property 6.2, we can retrieve the nodes of P using the operation $find_F$ of DSU_F on the endpoints of every bridge of \tilde{P} .

Now, for every node u in L that is incident to a real class, we append u_{assoc} to the sublist of u . This is to ensure that the real class to which u belongs will get joined to all classes that occur in L . (Because there is a possibility that the edge classes to which u belongs and that intersect \tilde{P} are quasi classes, and so L does not contain any edge from the real edge class to which u belongs.) If either x_{FRT} or y_{FRT} still has only one edge in its sublist, then it is removed from L . (This is to ensure that L is a joining list.) If L is non-empty, then we call $joinClasses(L)$. Otherwise, we have that x_{FRT} and y_{FRT} are connected with a bridge (x_{FRT}, y_{FRT}) on T_{FRT} . Then we just change the status of (x_{FRT}, y_{FRT}) so that it is no longer marked as a bridge, and we convert the quasi class that contains (x_{FRT}, y_{FRT}) to a real class. We let the associated edge of every node on \tilde{P} be any edge on \tilde{P} . (For this purpose, we may keep in a variable e_{temp} the edge that was in the boundary edge set of x_{FRT} after the call $boundary(x_{FRT}, y_{FRT})$; so now we set $u_{assoc} \leftarrow e_{temp}$ for every node u on \tilde{P} .) We introduce a new node z on T that replaces the entire path P , and we let z_{FRT} be any node on \tilde{P} . Finally, we unite all nodes on \tilde{P} using DSU_F , and we let z be the representative of the resulting set. This completes the description of $compressPath$, which is summarized in Algorithm 4.

`joinTrees`

Let T_1 and T_2 be two distinct trees of \mathcal{C} , and let $x \in T_1$ and $y \in T_2$. Then $joinTrees(T_1, T_2, (x, y))$ is performed by simply calling $link(x_{FRT}, y_{FRT})$, where (x_{FRT}, y_{FRT}) is marked as a bridge and the corresponding singleton edge class is marked as a quasi class. Finally we let (x_{FRT}, y_{FRT}) point to (x, y) . (See Algorithm 5.)

6.3. An implementation for cactuses

(a) Representation

Following [40], we represent the cactuses with a data structure that generalizes the concept of the tree of cycles. Recall that the tree of cycles of a cactus is the graph that represents the incidence relation of the nodes of the cactus to its cycles [28], [41]. By the definition of the cactus, we have that this graph is a tree. Now we generalize the tree of cycles as follows.

Algorithm 4: $compressPath(T, x, y)$.

```

1  $P \leftarrow \emptyset$  // the set of nodes of the simple path on  $T$  that connects  $x$  and  $y$ 
2  $\mathcal{E} \leftarrow \emptyset$  // the set of edges of  $T$  to be returned
3  $L \leftarrow \text{boundary}(x_{FRT}, y_{FRT})$ 
4 let  $e_{temp}$  be the edge in the sublist of  $x_{FRT}$  in  $L$ 
5 foreach node  $u$  in  $L$  do
6   foreach edge  $e$  in the sublist of  $u$  do
7     if  $e$  is a bridge then
8        $\tilde{e} \leftarrow$  edge of  $T$  pointed to by  $e$ 
9        $\mathcal{E} \leftarrow \mathcal{E} \cup \{\tilde{e}\}$ 
10      // let  $e = (z_1, z_2)$ 
11       $x_1 \leftarrow \text{find}_F(z_1)$ 
12       $x_2 \leftarrow \text{find}_F(z_2)$ 
13       $P \leftarrow P \cup \{x_1, x_2\}$ 
14     end
15   end
16   if  $u_{assoc} \neq \emptyset$  then append  $u_{assoc}$  to the sublist of  $u$ 
17   if the sublist of  $u$  contains only one edge then remove  $u$  from  $L$ 
18 end
19 if  $L \neq \emptyset$  then  $\text{joinclasses}(L)$ 
20 else
21   unmark  $(x_{FRT}, y_{FRT})$  as a bridge
22   mark  $\{(x_{FRT}, y_{FRT})\}$  as a real class
23 end
24 foreach node  $u$  in  $L \cup \{x_{FRT}, y_{FRT}\}$  do set  $u_{assoc} \leftarrow e_{temp}$ 
25 let  $z$  be a new node on  $T$  that replaces the path  $P$ 
26 unite all nodes in  $L \cup \{x_{FRT}, y_{FRT}\}$  using  $DSU_F$ ; let  $z$  be the representative
27 set  $z_{FRT} \leftarrow x_{FRT}$ 
28 return  $\{P, \mathcal{E}, z\}$ 

```

Algorithm 5: $joinTrees(T_1, T_2, (x, y))$.

```

1  $\text{link}(x_{FRT}, y_{FRT})$ 
2 mark  $(x_{FRT}, y_{FRT})$  as a bridge
3 mark  $\{(x_{FRT}, y_{FRT})\}$  as a quasi class
4 make  $(x_{FRT}, y_{FRT})$  point to  $(x, y)$ 

```

- We partition the cactus S into subcactuses. Based on this partition, we have the graph T that represents the incidence relation of the nodes of S to the subcactuses. (Again, we have that T is a tree.)
- We allow T to be extended with extra nodes, that serve as copies of the nodes of S . However, we demand that the resulting graph is also a tree. All nodes of T that serve as copies of the same node of S are said to belong to the same *cluster*. (Thus, the set of clusters is a partition of the set of nodes of T that correspond to nodes of S .)

We call this type of representation of a cactus S a *tree of cactuses*. See Fig. 8 for an example of this kind of representation. (We note that, contrary to the tree of cycles, a tree-of-cactuses representation of a cactus is not unique.) We distinguish two types of nodes in a tree of cactuses: those that correspond to the nodes of S , and they are called *real nodes*, and those that correspond to the subcactuses of S , and they are called *cactus nodes*. Thus we have a correspondence between the nodes of S and the real nodes of a tree-of-cactuses representation T of S . Furthermore, we may speak of the *nodes* of a cactus node C , and by that we mean the nodes of the subcactus of S that corresponds to C .

In our implementation, every node of S has a pointer to a unique corresponding node on T ; whenever no confusion arises, we will denote these two nodes with the same symbol. Conversely, every node of T has a pointer to its corresponding node on S . Specifically, to maintain the later correspondence we use a DSU data structure DSU_{cls} on the real nodes of T whose sets coincide with the clusters. This data structure supports the operations find_{cls} and unite_{cls} . Every cluster has a representative real node z in it, which can be found with a call $\text{find}_{cls}(x)$ on any node x in this cluster. Then we only need to maintain the pointer of z to its corresponding node on S . Finally, every edge (x, C) , where x is a real node and C is a cactus node, has pointers to its endpoints x and C .

Now let T be a tree of cactuses of a cactus S . We equip the set of edges of T with the following equivalence relation. If (x, C_1) and (x, C_2) are two edges of T , where x is a real node and C_1, C_2 are two cactus nodes, then (x, C_1) and (x, C_2) are equivalent. Furthermore, if (x, C) and (x', C') are two edges of T , where x and x' belong to the same cluster and C, C' are cactus nodes, then (x, C) and (x', C') are equivalent. It should be clear that this is indeed an equivalence relation, and that its equivalence classes induce subtrees of T . Thus, on a set of trees of cactuses (representing a set of cactuses), each one equipped with its respective equivalence relation, we can perform the operations of an FRT data structure. Furthermore, it should also be clear that every node x (in a non-trivial tree of cactuses) belongs to a unique edge class, and we associate to x an edge x_{assoc} of that class.




Fig. 8. An example of a tree-of-cactuses representation of a cactus. The square nodes are the cactus nodes. In figures A and C we can see the corresponding subcactuses of some cactus nodes. The edges that belong to the same edge class are painted with the same colour. The edges in black colour constitute singleton edge classes. We can see that some real nodes correspond to the same nodes of the cactus (and thus they belong to the same cluster). For example, real nodes 6a and 6b, in figure A, correspond to node 6 of the cactus. In figure B we can see an intermediate step of the call `compressCyclePath(4, 10)`: this is where we have to determine the path on the tree that connects 4 and 10. In figure C we can see the result of this call. The nodes 4, 3, 6 and 10 were merged into a new node, which is called 16. Observe that, in order to determine that 6 was also part of the cycle path connecting 4 and 10, we were guided by the associated data structure of an intermediary cactus node. Notice also that the operation `compressCyclePath` does not affect the structure of the tree of cactuses, but only its edge classes, the names of its real nodes, and the associated data structures of the involved cactus nodes.

Before we describe how we can use the FRT operations to perform `compressCyclePath` and `joinCactuses`, let us give an overview of what is involved in using a tree-of-cactuses representation T of a cactus S in order to determine cycle paths on S . So let x, y be two nodes of S . Let also Q be the cycle path on S connecting x and y , and let $P = z_1, C_1, \dots, C_{k-1}, z_k$ be the simple path on T with endpoints x and y , where $z_1 = x, z_k = y$, and C_1, \dots, C_{k-1} are the intermediary cactus nodes. Then we have that Q consists of (the nodes of S that correspond to) all the real nodes on P , plus (the nodes of S that correspond to) some nodes w_1, \dots, w_t that are incident to the cactuses C_1, \dots, C_{k-1} . The nodes w_1, \dots, w_t are precisely those that appear on the cycle paths of the subcactuses of S corresponding to C_1, \dots, C_{k-1} that form part of Q . Thus, in order to determine w_1, \dots, w_t , we are guided by some data structures associated with C_1, \dots, C_{k-1} that represent cactuses. (See also Fig. 8.) For this purpose we use the data structures of [41], that represent the cactuses

as rooted trees of cycles equipped with *circular split-find* structures. (These representations are used to solve the problem of the incremental maintenance of the 3-edge-connected components of 2-edge-connected graphs in asymptotically optimal time. We will not give a full exposition of this representation here, but later on we will describe how to augment it in order to facilitate the retrieval of the cactus edges.) Furthermore, in order to retrieve the edges of S that connect any pair of consecutive nodes on Q , we use precisely the data structures associated with C_1, \dots, C_{k-1} , as the information concerning the cactus edges is stored in the intermediary cactuses.

Now we recall that a circular split-find data structure operates on a collection of circular lists, where every list is related with a representative element, and supports the operations *find* and *split*. *find*(x) on a list node x returns the representative of the list containing x . *split*(x, y) on two distinct nodes x and y of the same list creates two new circular lists, consisting of the nodes from x (resp. y) and up to - but excluding - y (resp. x), following the same direction in both cases (e.g., the left one), and relates each new list with a representative element. [39] provides an implementation for the circular split-find problem, where we can perform m operations *split* and *find* on a collection of circular lists on n element in $O(n + m\alpha(m, n))$ time in total.

In the tree-of-cycles representation of cactuses used by [41], the circular lists correspond to the cycles of the cactus. Specifically, let C be a cycle of a cactus. Then, for every node x of C , there is an element $\text{repr}(x, C)$ that corresponds to the edge (x, C) of the tree of cycles. Now, the circular list corresponding to C consists precisely of the elements $\text{repr}(x, C)$, for every node x of C , ordered in the same way in which the nodes of C occur on C .

(We note that, in order to retrieve x or C from an element $\text{repr}(x, C)$, we use a *find* operation on a DSU data structure that operates on the collection of all *repr* elements. In fact, we use two different DSU data structures: one for retrieving the node x from an element $\text{repr}(x, C)$, and one for retrieving C from $\text{repr}(x, C)$. However, we will not mention explicitly those DSU data structures, because with an optimal implementation [54] they do not affect the asymptotical time bounds.)

Now we have to augment the information on the circular lists, in order to be able to return the cactus edges that connect consecutive nodes on cycle paths on a cactus. To achieve this, we simply store four pointers *left*, *right*, *leftEdge* and *rightEdge* on every *repr* element. Specifically, let C be a cycle of the cactus, and suppose that we have a specific orientation of its nodes (that is, a way to determine, given two consecutive nodes x and y on C , which one is on the left and which one is on the right). Now let x and y be two consecutive nodes of C , where x is on the left of y . Then we have $\text{repr}(y, C).\text{left} = \text{repr}(x, C)$, $\text{repr}(x, C).\text{right} = \text{repr}(y, C)$, $\text{repr}(y, C).\text{leftEdge} = (x, y)$ and $\text{repr}(x, C).\text{rightEdge} = (x, y)$.

Finally, for every edge (x, C) on a tree of cactuses, where x is a real node and C is a cactus node, we have a pointer from (x, C) to the node in the associated structure of C that corresponds to x , and reversely. (To be more precise, in order to maintain the reverse correspondence we use a DSU data structure that operates internally on the nodes of the tree of cycles representing C . This is because the operations on this data structure merge nodes into larger sets from time to time; and thus we maintain representatives of those sets, and pointers from those representatives to their corresponding nodes on the tree of cactuses. However, we will not mention explicitly the calls to this DSU data structure in what follows, since they do not affect the asymptotical time complexity overall.)

(b) Operations

`compressCyclePath`

To perform `compressCyclePath`(S, x, y) (with $x \neq y$), we first have to determine the cycle path connecting x and y on S . Let us use the same symbols, x and y , to denote the real nodes on the tree-of-cactuses representation T of S that correspond to x and y , respectively. As we noted in subsection (a), we have to determine two things: first, the simple path $P = z_1, C_1, \dots, C_{k-1}, z_k$ on T that connects x and y , where $z_1 = x$, $z_k = y$, and C_1, \dots, C_{k-1} are the intermediary cactus nodes, and second, the real nodes w_1, \dots, w_t that are incident to some of the cactuses C_1, \dots, C_{k-1} and that have to get merged too with the other nodes. The property that characterizes the nodes w_1, \dots, w_t can be explained as follows. For every edge (z, C) on P , where z is a real node and C is a cactus node, we let \tilde{z} denote the node in the associated tree of cycles of C that is pointed to by (z, C) (and corresponds essentially to z , and to any other node in the cluster of z). Now suppose that, for some $i \in \{1, \dots, k-1\}$, (z_i, C_i) and (C_i, z_{i+1}) belong to different edge classes. Then \tilde{z}_i and \tilde{z}_{i+1} are different nodes of (the associated structure to) C , and so we have to find and merge the cycle path on C with endpoints \tilde{z}_i and \tilde{z}_{i+1} (and also return the cactus edges that connect consecutive nodes on this path). The nodes on this cycle path are a subset of w_1, \dots, w_t ; and with this procedure, applied to all $i \in \{1, \dots, k-1\}$, we get precisely all w_1, \dots, w_t .

Now we notice that it is sufficient to determine only the subset of cactuses C_1, \dots, C_{k-1} that consists of all C_i , $i \in \{1, \dots, k-1\}$, such that (z_i, C_i) and (C_i, z_{i+1}) belong to different edge classes. (This is because, if (z_i, C_i) and (C_i, z_{i+1}) belong to the same edge class, then z_i and z_{i+1} correspond to the same node of S , and therefore no real nodes incident to C_i will be involved, and no operations on the associated data structure of C_i will have to be performed.) We can determine those cactuses with a call *boundary*(x, y) on T , by discarding the nodes x and y from the resulting boundary list L . (Observe that the intermediary real nodes of P will not appear in L , since all the edges that they are incident to belong to the same edge class.) Let D_1, \dots, D_l be the nodes that appear in L . Then, for every $i \in \{1, \dots, l\}$, there is an $a(i)$ such that $D_i = C_{a(i)}$. Now, for every $i \in \{1, \dots, l\}$, D_i contains a sublist in L consisting of two edges (z'_{i_1}, D_i) and (z'_{i_2}, D_i) , such that (z'_{i_1}, D_i) and $(z_{a(i)}, C_{a(i)})$ (resp. (z'_{i_2}, D_i) and $(C_{a(i)}, z_{a(i)+1})$) belong to the same edge class. This implies that z'_{i_1} corresponds to the same node as $z_{a(i)}$, and z'_{i_2} corresponds to the same node as $z_{a(i)+1}$. Thus, the node (in the associated data structure) of D_i pointed to by (z'_{i_1}, D_i) (resp. (z'_{i_2}, D_i)) is precisely the node that is pointed to by $(z_{a(i)}, C_{a(i)})$ (resp. $(C_{a(i)}, z_{a(i)+1})$), and so we can get $\tilde{z}_{a(i)}$ (resp. $\tilde{z}_{a(i)+1}$) from this pointer.

Algorithm 6: updateCactus(D, z_1, z_2).

```

// compress the cycle path on  $D$  that connects  $z_1$  and  $z_2$  using the associated data structure; return the
// set of cactus edges that connect consecutive nodes on this path, and a list of edges on the tree of
// cactuses that contains  $D$  that correspond to the nodes of this path
1  $\mathcal{E} \leftarrow \emptyset$  // the set of cactus edges to be returned
2  $edgelist \leftarrow \emptyset$  // the list of corresponding edges to be returned
3 find the cycle path  $P = \{u_1, \dots, u_k\}$  of  $D$  that connects  $z_1$  and  $z_2$  // for this step we refer to [41]
4 foreach  $i \in \{1, \dots, k-1\}$  do
5   let  $C$  be the cycle that contains  $u_i$  and  $u_{i+1}$ 
6   if  $repr(u_i, C).left = repr(u_{i+1}, C)$  then
7      $\mathcal{E} \leftarrow \mathcal{E} \cup repr(u_i, C).leftEdge$ 
8   end
9   if  $repr(u_i, C).right = repr(u_{i+1}, C)$  then
10     $\mathcal{E} \leftarrow \mathcal{E} \cup repr(u_i, C).rightEdge$ 
11  end
12 end
13 foreach  $u \in \{u_1, \dots, u_k\}$  do
14   get the edge  $(\tilde{u}, D)$  that is pointed to by  $u$ 
15    $edgelist \leftarrow edgelist \cup \{(\tilde{u}, D)\}$ 
16 end
17 merge the nodes on  $P$  and properly update the data structure // again, for this step we refer to [41]
18 fix the pointers  $left, right, leftEdge$  and  $rightEdge$  of the  $repr$  elements, as described in the text
19 return  $\{\mathcal{E}, edgelist\}$ 

```

Now, for every $i \in \{1, \dots, l\}$, we have to do the following things: (1) determine the cycle path on D_i whose nodes we have to merge, (2) find the cactus edges that connect consecutive nodes on this path, (3) find the corresponding nodes on the tree of cactuses T , and (4) merge the nodes on this path and properly update the associated data structure of D_i . Let z_1 and z_2 be nodes on D_i that are the endpoints of the cycle path we have to determine. (We recall that z_1 and z_2 are given by the pointers stored in (z'_1, D_i) and (z'_2, D_i) .) We will not provide the details on how to determine this path (i.e., how to perform (1)), as these can be found in [41]. (In particular, we use the procedure $TreePath_3$ in [41].) To perform (2), we use the pointers that we introduced in the $repr$ elements of the circular split-find data structures. To be more precise, once we have determined the cycle path that connects z_1 and z_2 , we then have to perform a *split* operation on every cycle involved in this path. (For the full details of what is involved in this step, see the procedure $AdjustCycles$ in [41].) So let C be a cycle that we have to split on nodes u_1 and u_2 . (We have that u_1 and u_2 are nodes on the cycle path connecting z_1 and z_2 .) Then we check whether $repr(u_1, C).left = repr(u_2, C)$ or $repr(u_1, C).right = repr(u_2, C)$, or both. If either case holds, we have to return $repr(u_1, C).leftEdge$ or $repr(u_1, C).rightEdge$, or both, respectively. For (3) we simply use the pointer of every node on the cycle path to the corresponding edge of T . (In particular, notice that for z_1 and z_2 we will get edges that belong to the same edge classes as (z'_1, D_i) and (z'_2, D_i) , respectively. Every other edge (w, D_i) that we get, belongs to a different edge class, and provides one of the extra nodes w_1, \dots, w_t that we have to merge (and which cannot be derived simply from the call *boundary*.) Finally, we will only provide the details to (4) that have to do with the additional information that we have attached to the associated data structure of C ; for the rest, we refer again to [41]. We first have to ensure that the pointers of the $repr$ elements of the circular lists are updated correctly after the splittings. So let C be a cycle that we split on nodes u_1 and u_2 . Then $repr(u_1, C)$ and $repr(u_2, C)$ are assigned to different circular lists; let us call them L_1 and L_2 , respectively. Suppose that L_1 contains more than one element (for otherwise there is nothing to do for $repr(u_1, C)$). Then, one of $repr(u_1, C).left, repr(u_1, C).right$ has been assigned to L_1 , and the other one has been assigned to L_2 . Assume w.l.o.g. that $repr(u_1, C).left$ has been assigned to L_1 (i.e., the same list that $repr(u_1, C)$ has been assigned to). Then we must set $repr(u_1, C).right \leftarrow repr(u_2, C).right$ and $repr(u_1, C).rightEdge \leftarrow repr(u_2, C).rightEdge$. The other case for $repr(u_1, C)$ and the analogous cases for $repr(u_2, C)$ are treated similarly. To conclude (4), we note that the correspondence between the nodes of (the associated data structure to) D_i and the edges of the tree of cactuses T is maintained, because, although the nodes of the cycle path on D_i that connects z_1 and z_2 got merged, this was done using a DSU data structure internal to the associated data structure to D_i . Thus we only have to make sure that, in order to access the node in an associated tree of cycles that corresponds to an edge of the tree of cactuses, we first perform a *find* on the node pointed to by this edge, using the internal DSU data structure. Algorithm 6 shows the operations that are performed in the associated data structure of D_i .

Thus, from the internal operations on the associated structure of every cactus node D_i , $i \in \{1, \dots, l\}$, we get a collection of edges $(w_{i,1}, D_i), \dots, (w_{i,k_i}, D_i)$ that we include in the edge sublist of D_i in L . Furthermore, from the edges $\{(w_{i,j}, D_i) \mid i \in \{1, \dots, l\}, j \in \{1, \dots, k_i\}\}$ we get the nodes of the cycle path on S that connects x and y , by using the pointer from every $(w_{i,j}, D_i)$ to $w_{i,j}$, and then the pointer from $find_{cls}(w_{i,j})$ to the corresponding node of S .

To conclude $compressCyclePath(S, x, y)$ we call $joinclasses(L)$, we return a pointer to a new node z that takes the place of all the nodes on the cycle path on S connecting x and y , we merge all $w_{i,j}$, $i \in \{1, \dots, l\}$, $j \in \{1, \dots, k_i\}$ into a larger cluster (using the DSU_{cls} data structure), and we let z point to the representative of this cluster, and reversely. The operation $compressCyclePath(S, x, y)$ is summarized in Algorithm 7.

Algorithm 7: `compressCyclePath(S, x, y)`.

```

1  $P \leftarrow \emptyset$  // the set of nodes of the cycle path on  $S$  that connects  $x$  and  $y$ 
2  $\mathcal{E} \leftarrow \emptyset$  // the set of edges of  $S$  to be returned
3  $L \leftarrow \text{boundary}(x, y)$ 
4 remove  $x$  and  $y$  from  $L$ 
5 foreach node  $C$  in  $L$  do
6   let  $(z_1, C)$  and  $(z_2, C)$  be the two edges in the sublist of  $C$  in  $L$ 
7   let  $\tilde{z}_1$  be the node pointed to by  $(z_1, C)$  in the associated data structure of  $C$ 
8   let  $\tilde{z}_2$  be the node pointed to by  $(z_2, C)$  in the associated data structure of  $C$ 
9    $\{\mathcal{E}_0, \text{edgelist}\} \leftarrow \text{updateCactus}(C, \tilde{z}_1, \tilde{z}_2)$ 
10   $\mathcal{E} \leftarrow \mathcal{E} \cup \mathcal{E}_0$ 
11  foreach edge  $(w, C)$  in  $\text{edgelist}$  do
12    get  $w$  using the pointer from  $(w, C)$ 
13    get the node  $u$  of  $S$  that corresponds to  $w$ , using the pointer of  $\text{find}_{cls}(w)$ 
14     $P \leftarrow P \cup \{u\}$ 
15  end
16  append  $\text{edgelist}$  to the sublist of  $C$  in  $L$ 
17 end
18  $\text{joinclasses}(L)$ 
19 let  $z$  be a new node on  $S$  substituting the cycle path  $P$ 
20 merge all the real nodes that appear as endpoints of the edges in the sublists of  $L$  using  $\text{DSU}_{cls}$ ; let  $\tilde{z}$  be the representative
21 make  $z$  point to  $\tilde{z}$  and  $\tilde{z}$  to  $z$ 
22 return  $\{P, \mathcal{E}, z\}$ 

```

Algorithm 8: `joinCactuses($S_1, \dots, S_k, (x_1, x_2), \dots, (x_k, x_1)$)`.

```

1 introduce a new cactus node  $C$ 
2 foreach  $x \in \{x_1, \dots, x_k\}$  do  $\text{link}(x, C)$ 
3 foreach  $x \in \{x_1, \dots, x_k\}$  do let  $\{(x, C)\}$  constitute a new edge class
4 foreach  $x \in \{x_1, \dots, x_k\}$  do
5   if  $x_{\text{assoc}} \neq \emptyset$  then
6     let  $L$  be a singleton list consisting of  $x$ 
7     let the sublist of  $x$  be  $\{x_{\text{assoc}}, (x, C)\}$ 
8      $\text{joinclasses}(L)$ 
9   end
10  else  $x_{\text{assoc}} \leftarrow (x, C)$ 
11 end
12 perform  $\text{initialize\_cycle}(C, x_1, \dots, x_k, (x_1, x_2), \dots, (x_k, x_1))$ , and collect the corresponding nodes  $\{\tilde{x}_1, \dots, \tilde{x}_k\}$ 
13 foreach  $i \in \{1, \dots, k\}$  do
14   make  $(x_i, C)$  point to  $\tilde{x}_i$  and  $\tilde{x}_i$  point to  $(x_i, C)$ 
15 end

```

Algorithm 9: `initialize_cycle($C, x_1, \dots, x_k, (x_1, x_2), \dots, (x_k, x_1)$)`.

```

1 create nodes  $\tilde{x}_1, \dots, \tilde{x}_k$  and  $\tilde{C}$ 
2 create a tree of cycles consisting of the edges  $(\tilde{x}_1, \tilde{C}), \dots, (\tilde{x}_k, \tilde{C})$ 
3 root the tree at  $\tilde{x}_1$ 
4 create the elements  $\text{repr}(\tilde{x}_1, \tilde{C}), \dots, \text{repr}(\tilde{x}_k, \tilde{C})$ 
5 initialize a circular split-find data structure on the elements  $\text{repr}(\tilde{x}_1, \tilde{C}), \dots, \text{repr}(\tilde{x}_k, \tilde{C})$  (in this order), and associated it with  $\tilde{C}$  // here we refer to [41]
6 fix the pointers  $\text{left}$ ,  $\text{right}$ ,  $\text{leftEdge}$  and  $\text{rightEdge}$  of the  $\text{repr}$  elements as described in the text
7 return  $\{\tilde{x}_1, \dots, \tilde{x}_k\}$ 

```

`joinCactuses`

To perform `joinCactuses($S_1, \dots, S_k, (x_1, x_2), \dots, (x_k, x_1)$)`, we have to link the nodes x_1, \dots, x_k in a new cycle. Thus we introduce a new cactus node C (that will be made to correspond to the new cycle) and the edges $(x_1, C), \dots, (x_k, C)$, by performing $\text{link}(x_1, C), \dots, \text{link}(x_k, C)$. This links all the trees of cactuses corresponding to S_1, \dots, S_k to C , and produces a larger tree of cactuses. For every $x \in \{x_1, \dots, x_k\}$ that has an associated edge x_{assoc} , we put (x, C) in the edge class of x_{assoc} by a call of `joinclasses`. If an $x \in \{x_1, \dots, x_k\}$ does not have an associated edge, then (x, C) constitutes a new edge class of its own, and the associated edge of x is set to be $x_{\text{assoc}} \leftarrow (x, C)$.

Then we have to construct the associated data structure to the cactus node C . This must be a rooted tree of cycles which contains only one cycle \tilde{C} with nodes corresponding to x_1, \dots, x_k , in this order. Let $\tilde{x}_1, \dots, \tilde{x}_k$ be the nodes of \tilde{C} that correspond to x_1, \dots, x_k , respectively. We root the tree arbitrarily to any one of $\tilde{x}_1, \dots, \tilde{x}_k$. Then we construct the circular list corresponding to \tilde{C} . The nodes of this list are the elements $\text{repr}(\tilde{x}_1, \tilde{C}), \dots, \text{repr}(\tilde{x}_k, \tilde{C})$, in this order. The pointers left and right on the repr elements are easy to fix. Also, we set $\text{repr}(\tilde{x}_i, \tilde{C}).\text{leftEdge} \leftarrow (x_{i-1}, x_i)$, for every $i \in \{2, \dots, k\}$, $\text{repr}(\tilde{x}_1, \tilde{C}).\text{leftEdge} \leftarrow (x_k, x_1)$, $\text{repr}(\tilde{x}_i, \tilde{C}).\text{rightEdge} \leftarrow (x_i, x_{i+1})$, for every $i \in \{1, \dots, k-1\}$, and

Algorithm 10: Compute a certificate for the maximal k -edge-connected subgraphs of G .

- 1 let E' be a set of edges of G that contains all its k -interconnection edges
 - 2 compute a forest decomposition \mathcal{F} of $G \setminus E'$ with k forests
 - 3 return $\mathcal{F} \cup E'$
-

$\text{repr}(\tilde{x}_k, \tilde{C}).\text{rightEdge} \leftarrow (x_k, x_1)$. (This is to be able to retrieve the real cactus edges.) The construction of the associated data structure to C is shown in Algorithm 9.

Finally, in order to establish the correspondence between the tree of cycles associated with C and the tree of cactuses which contains C , for every $i \in \{1, \dots, k\}$ we have a pointer from (x_i, C) to \tilde{x}_i , and reversely. (This is to be able to find the node within the cactus C that corresponds to x_i , and conversely.) The procedure *joinCactuses* is summarized in Algorithm 8.

7. Sparse certificates for the maximal k -edge-connected subgraphs

In this section we discuss constructions of (almost) sparse subgraphs that have the same maximal k -edge-connected subgraphs as the original graph. Following the terminology of [1], we define a k -certificate of a graph to be a spanning subgraph that has the same maximal k -edge-connected subgraphs as the original graph.⁷

Using results from Benczúr and Karger [5], we show that (1) in linear time we can construct a k -certificate of $O(kn \log n)$ size, and (2) in $O(m \log^2 n)$ time we can construct a k -certificate of $O(kn)$ size. Thus, by combining (1) and (2), we can construct a k -certificate of $O(kn)$ size in $O(m + kn \log^3 n)$ time. This is a result analogous to [47], which provides certificates for the k -edge-connected components. In fact, we use the following concept from [47] (defined formally in [46]): a *forest decomposition* with t forests of a graph G is a collection of forests F_1, \dots, F_t , such that F_1 is a spanning forest of G , and F_i is a spanning forest of $G \setminus (F_1 \cup \dots \cup F_{i-1})$, for $i \in \{2, \dots, t\}$. We also note that (1) is used by Saranurak and Yuan [51] in their $O(m + n^{1+o(1)})$ -time algorithm for computing the k -edge-connected subgraphs of an undirected graph, for any $k = \log^{o(1)} n$.

Let G be a graph. An edge of G whose endpoints lie in different maximal k -edge-connected subgraphs of G is called a k -interconnection edge of G . Algorithm 10 describes an algorithm for computing a k -certificate of G .

Lemma 7.1. *Algorithm 10 outputs a certificate for the maximal k -edge-connected subgraphs of G .*

Proof. If we remove all the k -interconnection edges from G , then the connected components of the resulting graph coincide with the maximal k -edge-connected subgraphs of G . Let S be a maximal k -edge-connected subgraph of G and let E' be a set of edges of G that contains all its k -interconnection edges. Then $[S, \tilde{S}] = \emptyset$ in $G \setminus E'$. Now let H be the set of all edges of E' that are contained in S . Then the set \mathcal{F}_S of all the edges of \mathcal{F} that are contained in $S \setminus H$ is a forest decomposition of $S \setminus H$ with k forests⁸ (precisely because S , and therefore $S \setminus H$, is disconnected from the rest of the graph in $G \setminus E'$). Thus, we know from the sparsification paper of Eppstein et al. [20] that \mathcal{F}_S (considered as a graph) is a *strong certificate*⁹ for the k -edge-connectivity of $S \setminus H$. Thus, $\mathcal{F}_S \cup H$ is a strong certificate for the k -edge-connectivity of $(S \setminus H) \cup H = S$. But S is k -edge-connected, and so $\mathcal{F} \cup E' \supseteq \mathcal{F}_S \cup H$ contains enough edges of S to maintain it as a k -edge-connected subgraph of $\mathcal{F} \cup E'$. Since $\mathcal{F} \cup E'$ is a subgraph of G , we thus have that its maximal k -edge-connected subgraphs coincide with those of G . \square

Corollary 7.2. *Let \mathcal{A} be an algorithm that, given a graph G with m edges and n vertices, computes in $T(m, n)$ time a subset E' of $E(G)$ with $S(m, n)$ size that contains all its k -interconnection edges. Then we can construct a k -certificate of G with $O(kn + S(m, n))$ size in $O(m + n + T(m, n))$ time.*

Proof. First we apply algorithm \mathcal{A} in order to compute a subset E' of $E(G)$ with size $S(m, n)$ that contains all its k -interconnection edges. This takes $T(m, n)$ time. Then we apply Algorithm 10: we remove E' from G , we compute a forest decomposition \mathcal{F} of $G \setminus E'$ with k forests, and we return $\mathcal{F} \cup E'$. Using [47], the computation of \mathcal{F} takes time $O(m + n)$. The output has size $O(kn + S(m, n))$. Correctness is guaranteed by Lemma 7.1. \square

In the following discussion we refer to Section 8 of [5]. (Note that in [5] the k -interconnection edges are called k -weak edges, and the maximal k -edge-connected subgraphs are called k -strong components.) In [5], they prove that a forest decomposition with at least $4k \log n$ forests contains all the k -interconnection edges of a graph. Since this decomposition can be computed in linear time using the MA-ordering algorithm of Nagamochi and Ibaraki [47], by Corollary 7.2 we get result (1): we can construct a k -certificate of size $O(kn \log n)$ in linear time.

⁷ This definition can be extended, in order to include objects such as a graph G' from which “we can derive easily” the maximal k -edge-connected subgraphs of the original graph once we have computed those of G' . However, the definition we provided here is enough for our purposes.

⁸ More precisely: it either coincides with $S \setminus H$, or it has k forests.

⁹ For any graph property \mathcal{P} , and graph G , a strong certificate for G is a graph G' on the same vertex set such that, for any graph H , $G \cup H$ has property \mathcal{P} if and only if $G' \cup H$ has the property.

Alternatively, [5] defines the *strength* of an edge e , denoted by k_e , as the largest k' such that e lies entirely within a maximal k' -edge-connected subgraph. Then the k -interconnection edges are precisely those whose strength is less than k . [5] provides an $O(m \log^2 n)$ -time algorithm that assigns a value \tilde{k}_e to every edge e , such that $\tilde{k}_e \leq k_e$ and $\sum_{e \in E} 1/\tilde{k}_e = O(n)$. Now we consider the set E' of all edges e that have $\tilde{k}_e < k$. Observe that E' contains all the k -interconnection edges, since $\tilde{k}_e \leq k_e$ for every edge e . Then we have that $O(n) = \sum_{e \in E} 1/\tilde{k}_e \geq \sum_{e \in E'} 1/\tilde{k}_e > \sum_{e \in E'} 1/k = |E'|/k$. Thus, E' has size $O(kn)$. Therefore, by Corollary 7.2 we get immediately our result (2): we can construct a k -certificate of size $O(kn)$ in $O(m \log^2 n)$ time.

8. Computing the maximal k -edge-connected subgraphs

Let G be an undirected multigraph with m edges and n vertices, and let $k > 2$ be a fixed integer. Our k -certificates developed in Section 7 can be immediately applied to compute the maximal k -edge-connected subgraphs of G . In particular, we first observe that we can apply the local search routine of Chechik et al. [10] (Lemma 5.1) into the framework of Forster et al. [21] (Lemma 7.2), so that we get a deterministic $O(km \log^2 n + (k^{O(k)} + k^{O(1)} \log n)n\sqrt{n})$ -time algorithm for computing the maximal k -edge-connected subgraphs of an undirected graph. By the result of the previous section, we can compute in $O(m + kn \log^3 n)$ time a k -certificate of $O(kn)$ size. More precisely, the certificate has $O(\min\{m, kn\})$ size, since, by construction, it is a subgraph of the original graph. Now we can apply the deterministic variant of the Forster et al. algorithm on this certificate, and thus we get the following.

Theorem 8.1. *There is an $O(m + (k^{O(k)} + k^{O(1)} \log n)n\sqrt{n})$ -time algorithm for computing the maximal k -edge-connected subgraphs of an undirected graph.*

Note that for constant $k \geq 3$, the time-bound provided by Theorem 8.1 is $O(m + n\sqrt{n} \log n)$, which improves the randomized bound of $O(m \log^2 n + n\sqrt{n} \log n)$ given by Forster et al. [21]. Similarly, by applying the $O(km \log^2 n + k^3 n\sqrt{n} \log n)$ -time algorithm of Forster et al. [21] on our k -certificate we obtain our next result.

Theorem 8.2. *There is a randomized Las Vegas algorithm for computing the maximal k -edge-connected subgraphs of an undirected graph that has $O(m + k^3 n^{3/2} \log n)$ expected running time.*

We next describe another algorithm to compute the maximal k -edge-connected subgraphs of an undirected graph G . We can do that by repeatedly removing any k' -edge cut from G , for $k' < k$, until there are no more k' -edge cuts in the graph for $k' < k$. Then the connected components of the resulting graph coincide with the maximal k -edge-connected subgraphs of G . The best known deterministic algorithm for computing a k' -edge cut of an undirected multigraph, for $k' < k$, or concluding that the graph is k -edge-connected, is given by Gabow [24], and it runs in $O(m + k^2 n \log(n/k))$ time. Thus the total running time of this algorithm would be bounded by $O(mn + k^2 n^2 \log(n/k))$.

To improve this bound, we reduce the time that it takes to successively find a min-cut with k' -edges (for $k' < k$). To speed this up, we repeatedly compute a min-cut of a graph and store enough information to facilitate the search for further cuts. To accomplish this task, we rely on Thorup's fully dynamic min-cut algorithm [56]. This algorithm supports edge insertions and deletions in $\tilde{O}(\sqrt{n})$ time per update, and it maintains a min-cut of size up to $k - 1$, for any fixed k (polylogarithmic on the number of vertices). It also demands an additional $O(m + n)$ time to initialize the underlying data structure on a sparse certificate of k -edge-connectivity for G [47]. (Thus, this algorithm maintains dynamically both a sparse certificate for G , using sparsification by Eppstein et al. [20], and a min-cut of size up to $k - 1$ of this certificate.)

Now we can use the fully dynamic min-cut algorithm of [56] on each connected component C of G as follows. As long as there is a k' -edge cut $[S, \bar{S}]$ of C , for some $k' < k$, we remove the edges of $[S, \bar{S}]$ from C , and then we select two arbitrary vertices $x \in S$ and $y \in \bar{S}$, and we reconnect the two resulting components S and \bar{S} by adding k multiple edges between x and y . Thus, the reconnection of the two components S and \bar{S} with k multiple edges ensures that all k' -edge cuts, for $k' < k$, are maintained in each component after the removal of $[S, \bar{S}]$, and that these are all the k' -edge cuts that may appear now in C . Eventually all components of G will become k -edge-connected. In the meantime, we collect all the edge-cuts that we find, in every connected component of G , and in the end we remove all of them from G . It should be clear that the connected components of the resulting graph coincide with the maximal k -edge-connected subgraphs of G . Observe that every search for a k' -edge cut, for $k' < k$, on some connected component of G , is immediately followed by k' deletions and k insertions of edges. Furthermore, the total number of these cuts is $O(n)$. Thus, the total running time of this algorithm is $\tilde{O}(m + k^{O(1)} n\sqrt{n})$. To provide a more refined analysis, we give a lower bound of the dependency of the complexity of this algorithm on k and on the number of log factors involved. To do that, we note that Thorup's algorithm uses a greedy tree packing of the sparse certificate using $k^7 \log^4 n$ trees, and it maintains it using Frederickson's dynamic minimum spanning tree algorithm [23]. Thus, for every k' -edge cut that we find, the cost of the deletions and insertions that follow is at least $\Omega(k^8 \sqrt{n} \log^4 n)$, and the total complexity of the algorithm is at least $\Omega(m + k^8 n\sqrt{n} \log^4 n)$. This yields the following.

Theorem 8.3. *Let $k = \log^{O(1)} n$. Then there is an $O(m + k^{O(1)} n\sqrt{n} \log^{O(1)} n)$ -time algorithm for computing the maximal k -edge-connected subgraphs of an undirected graph.*

Table 1
Previous best time bounds for computing the maximal k -edge-connected subgraphs.

Algorithm	Time	Type	Note
Chechik et al. [10]	$O(k^{O(k)}(m+n\log n)\sqrt{n})$	Det.	$k \in \mathbb{N}$
Henzinger et al. [31]	$O(n^2 \log n)$	Det.	$k = O(1)$
Forster et al. [21]	$O(k^3 n \sqrt{n} \log n + km \log^2 n)$	Las Vegas Rand.	$k \in \mathbb{N}$
Saranurak and Yuan [51]	$O(m + n^{1+o(1)})$	Det.	$k = \log^{o(1)} n$

Table 2
Improved time bounds for computing the maximal k -edge-connected subgraphs obtained in this section.

Algorithm	Time	Type	Note
This paper (Theorem 8.1)	$O(m + (k^{O(k)} + k^{O(1)} \log n)n\sqrt{n})$	Det.	$k \in \mathbb{N}$
This paper (Theorem 8.2)	$O(m + k^3 n \sqrt{n} \log n)$	Las Vegas Rand.	$k \in \mathbb{N}$
This paper (Theorem 8.3)	$O(m + k^{O(1)} n \sqrt{n} \log^{O(1)} n)$	Det.	$k = \log^{O(1)} n$

We remark here that the $O(m + n^{1+o(1)})$ time bound of Saranurak and Yuan [51] requires that $k = \log^{o(1)} n$, while the bounds in Theorems 8.1 and 8.2 hold for any k , and the bound of Theorem 8.3 holds for any $k = \log^{O(1)} n$. As another remark, we note that for constant k the bound provided by Theorem 8.3 is slightly worse than the bound of Theorem 8.1, since there are more log factors involved. However, the time-bound in Theorem 8.3 has polynomial dependency on k (whereas that of Theorem 8.1 has exponential dependency on k). Also, we believe that the algorithm of Theorem 8.3 can still be relevant, as any future improvement in the time bounds for a fully dynamic min-cut algorithm implies improved time bounds for computing the maximal k -edge-connected subgraphs.

We end this section by listing in Table 1 the previously known best time bounds for computing the maximal k -edge-connected subgraphs. The new results obtained in this section are described in Table 2.

9. A fully dynamic algorithm for maximal k -edge-connectivity

In this section we describe our fully dynamic algorithm for maintaining information about the maximal k -edge-connected subgraphs of a dynamic graph. In more detail, we wish to maintain an undirected graph $G = (V, E)$ throughout an intermixed sequence of the following operations:

insert(x, y): Add edge (x, y) to G ;
delete(x, y): Remove edge (x, y) from G (the operation assumes that (x, y) is in G);
max- k -edge(x, y): Return *true* if vertices x and y are in the same maximal k -edge-connected subgraph of G , and *false* otherwise.

Before stating our bounds, let us review some simple minded approaches for the problem. In Section 7 we showed that the maximal k -edge-connected subgraphs can be computed in time $O(m + n\sqrt{n} \log n)$ (for constant k). Note that if we recompute from scratch the maximal k -edge-connected subgraphs after each update, *max- k -edge* queries can be answered in constant time. This yields a simple algorithm that implements *insert* and *delete* operations in time $O(m + n\sqrt{n} \log n)$ and *max- k -edge* queries in constant time. On the other side, one could simply do no extra work during *insert* and *delete* operations, but then answering a *max- k -edge* query would require recomputing the maximal k -edge-connected subgraphs from scratch, yielding constant time per update and $\tilde{O}(m + n\sqrt{n})$ time per query. We next show how to implement *insert* and *delete* operations in better $\tilde{O}(n\sqrt{n})$ time, while still keeping the running time for *max- k -edge* queries constant. To achieve our improved bounds, we exploit the sparsification technique of Eppstein et al. [20]. We start with the following definition.

Definition 9.1. Let $k \geq 3$ be a fixed integer. Given an undirected graph $G = (V, E)$ with m edges and n vertices, a *sparse certificate of maximal k -edge-connectivity* for G is a graph G' defined on the same vertex set as G such that the following holds:

- (i) G' has $O(n)$ edges;
- (ii) For any graph H , any two vertices are in the same maximal k -edge-connected subgraph of $G' \cup H$ if and only if they are in the same maximal k -edge-connected subgraph in $G \cup H$.

The following lemma provides a sufficient condition for inferring that a k -certificate is a sparse certificate of maximal k -edge-connectivity.

Lemma 9.2. *Let G be an undirected graph, and let C be a k -certificate of G with $O(n)$ edges that contains all the k -interconnection edges of G . Then C is a sparse certificate of maximal k -edge-connectivity for G .*

Proof. Condition (i) of Definition 9.1 is satisfied. It remains to establish condition (ii). Since C is a k -certificate of G , it is a subgraph of G that has the same maximal k -edge-connected subgraphs as G . Let Q be the quotient graph that is formed by shrinking every maximal k -edge-connected subgraph of C into a single vertex. Then, since C contains all the k -interconnection edges of G , by Property 3.1 we have that Q has the same decomposition tree into maximal k -edge-connected subgraphs as G . Thus, the changes that this tree undergoes after inserting all the vertices and edges of H into G , are the same as if inserting them to C . Thus, $C \cup H$ has the same maximal k -edge-connected subgraphs as $G \cup H$. This shows that condition (ii) is satisfied too. Therefore C is a sparse certificate of maximal k -edge-connectivity for G . \square

Corollary 9.3. *Let $G = (V, E)$ be an undirected graph with m edges and n vertices. A sparse certificate of maximal k -edge-connectivity for G can be computed in time $O(m \log^2 n)$.*

Proof. In Section 7 we show that we can construct a k -certificate C of G with $O(n)$ edges in time $O(m \log^2 n)$. Furthermore, this k -certificate has the property that it contains all the k -interconnection edges of G . Thus, Lemma 9.2 implies that C is a sparse certificate of maximal k -edge-connectivity for G . \square

We are now ready to apply the sparsification framework of Eppstein et al. [20]:

Theorem 9.4 ([20]). *Let $k \geq 3$ be a fixed integer, let $f(n, m)$ be the time required to compute a sparse certificate of maximal k -edge-connectivity, and let $g(n, m)$ the time required to compute the maximal k -edge-connected subgraphs. Then we can build a fully dynamic data structure that can handle insert and delete operations in time $O(f(n, O(n)) + g(n, O(n)))$ and max- k -edge queries in constant time.*

From Corollary 9.3 we have $f(m, n) = O(m \log^2 n)$. From [10] we have $g(m, n) = O(m\sqrt{n})$ for $k \in \{3, 4\}$, and $g(m, n) = O((m + n \log n)\sqrt{n})$ for fixed $k > 4$. (The improved time-bounds for $k \in \{3, 4\}$ are derived by using either of the algorithms of [30,44] for computing 3-edge cuts in linear time.) By fitting those bounds into Theorem 9.4 we obtain:

Theorem 9.5. *Let $G = (V, E)$ be an undirected graph with n vertices. Then we can build a fully dynamic data structure that can handle insert and delete operations in time $O(n\sqrt{n})$, for $k \in \{3, 4\}$, and in time $O(n\sqrt{n} \log n)$, for fixed $k > 4$, so that it can answer max- k -edge queries in constant time.*

Also, for $k = \log^{o(1)} n$, by [51] we have $g(m, n) = O(m + n^{1+o(1)})$, so now Theorem 9.4 gives:

Theorem 9.6. *Let $G = (V, E)$ be an undirected graph with n vertices, and let $k = \log^{o(1)} n$. Then we can build a fully dynamic data structure that can handle insert and delete operations in time $O(n \log^2 n + n^{1+o(1)})$, so that it can answer max- k -edge queries in constant time.*

10. Conclusions

We presented two algorithms for maintaining a decomposition tree structure of the maximal 3-edge-connected subgraphs of a graph. The first algorithm uses $O(n)$ space and can handle any sequence of m edge insertions and n vertex insertions in total time $O(n^2 \log^2 n + m\alpha(m, n))$. The second algorithm uses $O(n^2)$ space and can handle any sequence of m edge insertions and n vertex insertions in total time $O(n^2\alpha(n, n) + m\alpha(m, n))$.

We note that one can use this data structure to efficiently answer interspersed queries concerning the maximal 3-edge-connected subgraphs, such as:

- Given vertices x and y , report whether x and y belong to the same maximal 3-edge-connected subgraph (using two calls to a DSU-*find* operation)
- Find the maximal 3-edge-connected subgraph that contains x , in time analogous to its size (plus a call to a DSU-*find* operation)
- Report the size (the number of vertices or edges) of the maximal 3-edge-connected subgraph that contains x in constant time (plus a call to a DSU-*find* operation)
- Return all maximal 3-edge-connected subgraphs in time analogous to their size
- Return the number of all maximal 3-edge-connected subgraphs in constant time

The methods used in this algorithm extend previous work in maintaining the 2- and 3-edge-connected components [28,41,59], and may prove useful in solving other similar problems. For instance, it seems possible that we can add more

levels to the decomposition tree, and rely on previous work for maintaining the 4- and 5-edge-connected components, in order to maintain the maximal 4- and 5-edge-connected subgraphs (by properly adjusting the data structures and algorithms in [16,19]). Furthermore, it seems that any improvement in maintaining the k -edge-connected components, for any constant k , would imply (under some conditions) that this solution can be plugged in to our framework for maintaining the decomposition tree, in order to maintain the maximal k -edge-connected subgraphs (with time-bounds analogous to those provided in our paper). It seems possible that this framework could also be useful for maintaining the maximal k -vertex-connected subgraphs, by relying on efficient algorithms for maintaining the k -vertex-connected components (such as those described in [13,50], for the cases $k = 3$).

We also showed that we can construct (almost) sparse certificates for them maximal k -edge-connected subgraphs. Our result implies that the difficulty in designing efficient algorithms for computing the maximal k -edge-connected subgraphs lies essentially in sparse graphs, and we can use it in order to speed up the running time of already known algorithms. In particular, by using the algorithm by Chechik et al. [10], we get an $O(m + k^{O(k)}n\sqrt{n}\log n)$ -time algorithm for computing the maximal k -edge-connected subgraphs in undirected graphs.

We believe that it is an interesting question whether a k -certificate of $O(n)$ size can be computed in linear time. This would be trivial if we had a linear-time algorithm for computing the maximal k -edge-connected subgraphs of a graph, but it is still an open problem whether this can be done for $k \geq 3$. Thus, we have to perform the construction of the certificates without explicitly computing the maximal k -edge-connected subgraphs, and this seems to be a challenging task.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] Anders Aamand, Adam Karczmarz, Jakub Lacki, Nikos Parotsidis, Peter M.R. Rasmussen, Mikkel Thorup, Optimal decremental connectivity in non-sparse graphs, CoRR, arXiv:2111.09376, 2021.
- [2] Ittai Abraham, Shiri Chechik, Sebastian Krininger, Fully dynamic all-pairs shortest paths with worst-case update-time revisited, in: Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, 2017, pp. 440–452.
- [3] Takuya Akiba, Yoichi Iwata, Yuichi Yoshida, Linear-time enumeration of maximal K -edge-connected subgraphs in large networks by random contraction, in: Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM '13, Association for Computing Machinery, New York, NY, USA, October 2013, pp. 909–918.
- [4] Stephen Alstrup, Jacob Holm, Kristian de Lichtenberg, Mikkel Thorup, Maintaining information in fully dynamic trees with top trees, ACM Trans. Algorithms 1 (2) (2005) 243–264.
- [5] András A. Benczúr, David R. Karger, Randomized approximation schemes for cuts and flows in capacitated graphs, SIAM J. Comput. 44 (2) (2015) 290–319.
- [6] Aaron Bernstein, Maximilian Probst, Christian Wulff-Nilsen, Decremental strongly-connected components and single-source reachability in near-linear time, in: Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23–26, 2019, pp. 365–376.
- [7] Endre Boros, Konrad Borys, Vladimir Gurvich, Gábor Rudolf, Generating 3-vertex connected spanning subgraphs, Discrete Math. 308 (24) (2008) 6285–6297.
- [8] Lijun Chang, Zhiyi Wang, A near-optimal approach to edge connectivity-based hierarchical graph decomposition, Proc. VLDB Endow. 15 (6) (Feb 2022) 1146–1158.
- [9] Lijun Chang, Jeffrey Xu Yu, Lu Qin, Xuemin Lin, Chengfei Liu, Weifa Liang, Efficiently computing k -edge connected components via graph decomposition, in: Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, SIGMOD '13, Association for Computing Machinery, New York, NY, USA, June 2013, pp. 205–216.
- [10] Shiri Chechik, Thomas Dueholm Hansen, Giuseppe F. Italiano, Veronika Loitzenbauer, Nikos Parotsidis, Faster algorithms for computing maximal 2-connected subgraphs in sparse directed graphs, in: Philip N. Klein (Ed.), Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, SIAM, 2017, pp. 1900–1918.
- [11] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, Introduction to Algorithms, 2nd edition, The MIT Press, 2001.
- [12] Camil Demetrescu, Giuseppe F. Italiano, A new approach to dynamic all pairs shortest paths, J. ACM 51 (6) (2004) 968–992.
- [13] G. Di Battista, R. Tamassia, On-line maintenance of triconnected components with SPQR-trees, Algorithmica 15 (4) (Apr 1996) 302–318.
- [14] R. Diestel, Graph Theory, 2nd edition, Springer-Verlag, 2000.
- [15] E.A. Dinic, A.V. Karzanov, M.V. Lomonosov, On the structure of the system of minimum edge-cuts in a graph, in: Studies in Discrete Optimization, 1976, pp. 290–306.
- [16] Y. Dinitz, R. Nossenson, Incremental maintenance of the 5-edge-connectivity classes of a graph, in: Proceedings of the 7th Scandinavian Workshop on Algorithm Theory, SWAT '00, Springer-Verlag, Berlin, Heidelberg, 2000, pp. 272–285.
- [17] Y. Dinitz, Z. Nutov, A 2-level cactus model for the system of minimum and minimum+1 edge-cuts in a graph and its incremental maintenance, in: Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing, STOC '95, Association for Computing Machinery, New York, NY, USA, 1995, pp. 509–518.
- [18] Yefim Dinitz, The 3-edge-components and a structural description of all 3-edge-cuts in a graph, in: Ernst W. Mayr (Ed.), Graph-Theoretic Concepts in Computer Science, 18th International Workshop, Proceedings, WG '92, Wiesbaden-Naurod, Germany, June 19–20, 1992, in: Lecture Notes in Computer Science, vol. 657, Springer, 1992, pp. 145–157.
- [19] Yefim Dinitz, Jeffery R. Westbrook, Maintaining the classes of 4-edge-connectivity in a graph on-line, Algorithmica 20 (3) (1998) 242–276.

- [20] David Eppstein, Zvi Galil, Giuseppe F. Italiano, Amnon Nissenzeig, Sparsification - a technique for speeding up dynamic graph algorithms, *J. ACM* 44 (5) (1997) 669–696.
- [21] Sebastian Forster, Danupon Nanongkai, Liu Yang, Thatchaphol Saranurak, Sorrachai Yingchareonthawornchai, Computing and testing small connectivity in near-linear time and queries via fast local cut algorithms, in: *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Proceedings, Society for Industrial and Applied Mathematics, December 2019, pp. 2046–2065.
- [22] Paolo Giulio Franciosa, Giorgio Gambosi, Umberto Nanni, The incremental maintenance of a depth-first-search tree in directed acyclic graphs, *Inf. Process. Lett.* 61 (2) (1997) 113–120.
- [23] Greg N. Frederickson, Data structures for on-line updating of minimum spanning trees, with applications, *SIAM J. Comput.* 14 (4) (1985) 781–798.
- [24] H.N. Gabow, A matroid approach to finding edge connectivity and packing arborescences, *J. Comput. Syst. Sci.* 50 (1995) 259–273.
- [25] H.N. Gabow, Path-based depth-first search for strong and biconnected components, *Inf. Process. Lett.* 74 (2000) 107–114.
- [26] Z. Galil, G.F. Italiano, Fully dynamic algorithms for edge connectivity problems, in: *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing, STOC '91*, Association for Computing Machinery, New York, NY, USA, 1991, pp. 317–327.
- [27] Z. Galil, G.F. Italiano, Reducing edge connectivity to vertex connectivity, *SIGACT News* 22 (1) (March 1991) 57–61.
- [28] Zvi Galil, Giuseppe F. Italiano, Maintaining the 3-edge-connected components of a graph on-line, *SIAM J. Comput.* 22 (1) (February 1993) 11–28.
- [29] Loukas Georgiadis, Thomas Dueholm Hansen, Giuseppe F. Italiano, Sebastian Krinninger, Nikos Parotsidis, Incremental data structures for connectivity and dominators in directed graphs, in: *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017*, July 10–14, 2017, Warsaw, Poland, 2017, 42.
- [30] Loukas Georgiadis, Giuseppe F. Italiano, Evangelos Kosinas, Computing the 4-edge-connected components of a graph in linear time, in: *Petra Mutzel, Rasmus Pagh, Grzegorz Herman (Eds.), 29th Annual European Symposium on Algorithms, ESA 2021*, September 6–8, 2021, Lisbon, Portugal (Virtual Conference), in: *LIPICs*, vol. 204, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 47.
- [31] Monika Henzinger, Sebastian Krinninger, Veronika Loitzenbauer, Finding 2-edge and 2-vertex strongly connected components in quadratic time, in: *Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, Bettina Speckmann (Eds.), Automata, Languages, and Programming - 42nd International Colloquium, Proceedings, Part I, ICALP 2015*, Kyoto, Japan, July 6–10, 2015, in: *Lecture Notes in Computer Science*, vol. 9134, Springer, 2015, pp. 713–724.
- [32] Monika Rauch Henzinger, Valerie King, Fully dynamic biconnectivity and transitive closure, in: *36th Annual Symposium on Foundations of Computer Science, Milwaukee, Wisconsin, USA, 23–25 October 1995*, 1995, pp. 664–672.
- [33] M.R. Henzinger, V. King, Randomized fully dynamic graph algorithms with polylogarithmic time per operation, *J. ACM* 46 (4) (1999) 502–516.
- [34] J. Holm, K. de Lichtenberg, M. Thorup, Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity, *J. ACM* 48 (4) (July 2001) 723–760.
- [35] Jacob Holm, Giuseppe F. Italiano, Adam Karczmarz, Jakub Lacki, Eva Rotenberg, Piotr Sankowski, Contracting a planar graph efficiently, in: *Kirk Pruhs, Christian Sohler (Eds.), 25th Annual European Symposium on Algorithms, ESA 2017*, September 4–6, 2017, Vienna, Austria, in: *LIPICs*, vol. 87, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, 50.
- [36] J.E. Hopcroft, R.E. Tarjan, Dividing a graph into triconnected components, *SIAM J. Comput.* 2 (3) (1973) 135–158.
- [37] Wenyu Jin, Xiaorui Sun, Fully dynamic s-t edge connectivity in subpolynomial time (extended abstract), in: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, 2022, pp. 861–872.
- [38] A. Kanevsky, R. Tamassia, G. Di Battista, J. Chen, On-line maintenance of the four-connected components of a graph, in: *Proceedings 32nd Annual Symposium of Foundations of Computer Science (FOCS 1991)*, 1991, pp. 793–801.
- [39] J.A. La Poutré, *Dynamic Graph Algorithms and Data Structures*, Ph.D. thesis, Utrecht University, The Netherlands, September 1991.
- [40] J.A. La Poutré, Maintenance of 2- and 3-edge-connected components of graphs II, *SIAM J. Comput.* 29 (5) (2000) 1521–1549.
- [41] J.A. La Poutré, J. van Leeuwen, M.H. Overmars, Maintenance of 2- and 3-edge-connected components of graphs I, *Discrete Math.* 114 (1–3) (1993) 329–359.
- [42] Yuan Li, Guoren Wang, Yuhai Zhao, Feida Zhu, Yubao Wu, Towards k-vertex connected component discovery from large networks, *World Wide Web* 23 (2) (March 2020) 799–830.
- [43] William Di Luigi, Loukas Georgiadis, Giuseppe F. Italiano, Luigi Laura, Nikos Parotsidis, 2-Connectivity in directed graphs: an experimental study, in: *Proceedings of the Seventeenth Workshop on Algorithm Engineering and Experiments, ALENEX 2015*, San Diego, CA, USA, January 5, 2015, 2015, pp. 173–187.
- [44] Wojciech Nadara, Mateusz Radecki, Marcin Smulewicz, Marek Sokolowski, Determining 4-edge-connected components in linear time, in: *Petra Mutzel, Rasmus Pagh, Grzegorz Herman (Eds.), 29th Annual European Symposium on Algorithms, ESA 2021*, September 6–8, 2021, Lisbon, Portugal (Virtual Conference), in: *LIPICs*, vol. 204, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 71.
- [45] H. Nagamochi, T. Ibaraki, A linear time algorithm for computing 3-edge-connected components in a multigraph, *Jpn. J. Ind. Appl. Math.* 9 (163) (1992).
- [46] H. Nagamochi, T. Ibaraki, *Algorithmic Aspects of Graph Connectivity*, 1st edition, Cambridge University Press, 2008.
- [47] Hiroshi Nagamochi, Toshihide Ibaraki, A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph, *Algorithmica* 7 (5&6) (1992) 583–596.
- [48] Chaitanya Nalam, Thatchaphol Saranurak, Maximal k-edge-connected subgraphs in weighted graphs via local random contraction, in: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 2023, pp. 183–211.
- [49] Danupon Nanongkai, Thatchaphol Saranurak, Christian Wulff-Nilsen, Dynamic minimum spanning forest with subpolynomial worst-case update time, in: *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017*, Berkeley, CA, USA, October 15–17, 2017, 2017, pp. 950–961.
- [50] Johannes A. La Poutré, Maintenance of triconnected components of graphs (extended abstract), in: *Proceedings of the 19th International Colloquium on Automata, Languages and Programming, ICALP '92*, Springer-Verlag, Berlin, Heidelberg, 1992, pp. 354–365.
- [51] Thatchaphol Saranurak, Wuwei Yuan, Maximal k-edge-connected subgraphs in almost-linear time for small k, in: *Proceedings of the 31st Annual European Symposium on Algorithms (ESA 2023)*, 2023.
- [52] Heli Sun, Jianbin Huang, Yang Bai, Zhongmeng Zhao, Xiaolin Jia, Fang He, Yang Li, Efficient k-edge connected component detection through an early merging and splitting strategy, *Knowl.-Based Syst.* 111 (November 2016) 63–72.
- [53] R.E. Tarjan, Depth-first search and linear graph algorithms, *SIAM J. Comput.* 1 (2) (1972) 146–160.
- [54] Robert Endre Tarjan, Efficiency of a good but not linear set union algorithm, *J. ACM* 22 (2) (1975) 215–225.
- [55] Mikkel Thorup, Fully-dynamic all-pairs shortest paths: faster and allowing negative cycles, in: *Algorithm Theory - SWAT 2004*, 9th Scandinavian Workshop on Algorithm Theory, Proceedings, Humlebaek, Denmark, July 8–10, 2004, 2004, pp. 384–396.
- [56] Mikkel Thorup, Fully-dynamic min-cut, *Combinatorica* 27 (1) (2007) 91–127.
- [57] Yung H. Tsin, Yet another optimal algorithm for 3-edge-connectivity, *J. Discrete Algorithms* 7 (1) (2009) 130–146.
- [58] Dong Wen, Lu Qin, Ying Zhang, Lijun Chang, Ling Chen, Enumerating k-vertex connected components in large graphs, in: *35th IEEE International Conference on Data Engineering, ICDE 2019*, Macao China, April 8–11, 2019, 2019, pp. 52–63.
- [59] Jeffery R. Westbrook, Robert Endre Tarjan, Maintaining bridge-connected and biconnected components on-line, *Algorithmica* 7 (5&6) (1992) 433–464.
- [60] Long Yuan, Lu Qin, Xuemin Lin, Lijun Chang, Wenjie Zhang, I/O efficient ECC graph decomposition via graph reduction, *Proc. VLDB Endow.* 9 (7) (March 2016) 516–527.