

Computing the Largest Bond and the Maximum Connected Cut of a Graph *

Gabriel L. Duarte¹, Hiroshi Eto², Tesshu Hanaka³, Yasuaki Kobayashi⁴, Yusuke Kobayashi⁴, Daniel Lokshantov⁵, Lehilton L. C. Pedrosa⁶, Rafael C. S. Schouery⁶, and Uéverton S. Souza¹

¹Fluminense Federal University, Niterói, RJ, Brazil [†]

²Kyushu University, Fukuoka, Japan [‡]

³Chuo University, Tokyo, Japan [§]

⁴Kyoto University, Kyoto, Japan [¶]

⁵University of California Santa Barbara, CA, USA ^{||}

⁶University of Campinas, SP, Brazil ^{**}

Abstract

The cut-set $\partial(S)$ of a graph $G = (V, E)$ is the set of edges that have one endpoint in $S \subset V$ and the other endpoint in $V \setminus S$, and whenever $G[S]$ is connected, the cut $[S, V \setminus S]$ of G is called a connected cut. A bond of a graph G is an inclusion-wise minimal disconnecting set of G , i.e., bonds are cut-sets that determine cuts $[S, V \setminus S]$ of G such that $G[S]$ and $G[V \setminus S]$ are both connected. Contrasting with a large number of studies related to maximum cuts, there exist very few results regarding the largest bond of general graphs. In this paper, we aim to reduce this gap on the complexity of computing the largest bond, and the maximum connected cut of a graph. Although cuts and bonds are similar, we remark that computing the largest bond and the maximum connected cut of a graph tends to be harder than computing its maximum cut. We show that it does not exist a constant-factor approximation algorithm to compute the largest bond, unless $P = NP$. Also, we show that LARGEST BOND and MAXIMUM CONNECTED CUT are NP-hard even for planar bipartite graphs, whereas MAXIMUM CUT is trivial on bipartite graphs and polynomial-time solvable on planar graphs. In addition, we show that LARGEST BOND and MAXIMUM CONNECTED CUT are NP-hard on split graphs, and restricted to graphs of clique-width w they can not be solved in time $f(w) \times n^{o(w)}$ unless the Exponential Time Hypothesis fails, but they can be solved in time $f(w) \times n^{O(w)}$. Finally, we show that both problems are fixed-parameter tractable when parameterized by the size of the solution, the treewidth, and the twin-cover number.

Keywords— bond; cut; maximum cut; connected cut; FPT; treewidth; clique-width.

*This work is partially supported by JST CREST JPMJCR1401, and JSPS KAKENHI grant numbers JP17H01788, JP16K16010, JP17K19960, and JP19K21537, and by São Paulo Research Foundation (FAPESP) grant number 2015/11937-9, and Rio de Janeiro Research Foundation (FAPERJ) grant number E-26/203.272/2017, and by National Council for Scientific and Technological Development (CNPq-Brazil) grant numbers 308689/2017-8, 425340/2016-3, 313026/2017-3, 422829/2018-8, 303726/2017-2. The Japanese authors thank Akitoshi Kawamura and Yukiko Yamauchi for giving an opportunity to discuss in the Open Problem Seminar in Kyushu University, Japan. Preliminary versions of this paper appeared in [27] and [22].

[†]{gabrieled, ueverton}@ic.uff.br

[‡]h-eto@econ.kyushu-u.ac.jp

[§]hanaka.91t@g.chuo-u.ac.jp

[¶]kobayashi@iip.ist.i.kyoto-u.ac.jp, yusuke@kurims.kyoto-u.ac.jp

^{||}daniello@ucsb.edu

^{**}{lehilton, rafael}@ic.unicamp.br

1 Introduction

Let $G = (V, E)$ be a simple, connected, undirected graph. A *disconnecting set* of G is a set of edges $F \subseteq E(G)$ whose removal disconnects G . The edge-connectivity of G is $\kappa'(G) = \min\{|F| : F \text{ is a disconnecting set of } G\}$. A cut $[S, T]$ of G is a partition of V into two subsets S and $T = V \setminus S$. The cut-set $\partial(S)$ of a cut $[S, T]$ is the set of edges that have one endpoint in S and the other endpoint in T ; these edges are said to cross the cut. In a connected graph, each cut-set determines a unique cut. Note that every cut-set is a disconnecting set, but the converse is not true. An inclusion-wise minimal disconnecting set of a graph is called a *bond* (or a *minimal cut*). It is easy to see that every bond is a cut-set, but there are cut-sets that are not bonds. More precisely, a nonempty set of edges F of G is a bond if and only if F determines a cut $[S, T]$ of G such that $G[S]$ and $G[T]$ are both connected. Let $s, t \in V(G)$. An *st*-bond of G is a bond whose removal disconnects s and t .

A minimum (maximum) cut of a graph G is a cut with cut-set of minimum (maximum) size. Every minimum cut is a bond, thus a minimum bond is also a minimum cut of G , and it can be found in polynomial time using the classical Edmonds–Karp algorithm [25]. Besides that, minimum *st*-bonds are well-known structures, since they are precisely the *st*-cuts involved in the Gomory–Hu trees [36].

MAXIMUM CUT is one of the most fundamental problems in theoretical computer science. Given a graph G and an integer k , the problem asks for a subset S of vertices such that $|\partial(S)| \geq k$. Whenever the subgraph of G induced by S is connected, the cut $[S, V \setminus S]$ is called *connected*. Recall that every bond is a connected cut, but the converse is not true.

In this paper, we are interested in the complexity aspects of the following problem.

MAXIMUM CONNECTED CUT

Instance: A graph $G = (V, E)$; a positive integer k .

Question: Is there a proper subset $S \subset V(G)$ such that $G[S]$ is connected and $|\partial(S)| \geq k$?

LARGEST BOND

Instance: A graph $G = (V, E)$; a positive integer k .

Question: Is there a proper subset $S \subset V(G)$ such that $G[S]$ and $G[V \setminus S]$ are both connected and $|\partial(S)| \geq k$?

We also consider the versions of both problems where the removal of the cut-set must disconnect a given pair of vertices.

MAXIMUM CONNECTED *st*-CUT

Instance: A graph $G = (V, E)$; vertices $s, t \in V(G)$; a positive integer k .

Question: Is there a proper subset $S \subset V(G)$ with $s \in S$ and $t \notin S$, such that $G[S]$ is connected and $|\partial(S)| \geq k$?

LARGEST *st*-BOND

Instance: A graph $G = (V, E)$; vertices $s, t \in V(G)$; a positive integer k .

Question: Is there a proper subset $S \subset V(G)$ with $s \in S$ and $t \notin S$, such that $G[S]$ and $G[V \setminus S]$ are both connected and $|\partial(S)| \geq k$?

Such problems can be seen as variants of MAXIMUM CUT, which was shown to be NP-hard in Karp’s seminal work [43]. To overcome this intractability, a lot of researches have been done from various view points, such as approximation algorithms [35], fixed-parameter tractability [52], and special graph classes [6, 24, 38, 39, 50].

The MAXIMUM CONNECTED CUT problem was defined in [40] and it is known to be NP-complete even on planar graphs [41]. Regarding bonds on planar graphs, a folklore theorem states that if G is a connected planar graph, then a set of edges is a cycle in G if and only if it corresponds to a bond in the dual graph of G [34]. Note that each cycle separates the faces of G into the faces

in the interior of the cycle and the faces of the exterior of the cycle, and the duals of the cycle edges are exactly the edges that cross from the interior to the exterior [51]. Consequently, the girth of a planar graph equals the edge connectivity of its dual [14]. Although cuts and bonds are similar, computing the largest bond of a graph seems to be harder than computing its maximum cut. MAXIMUM CUT is NP-hard in general [32], but becomes polynomial for planar graphs [39]. On the other hand, finding a longest cycle in a planar graph is NP-hard, implying that finding a largest bond of a planar multigraph (or of a simple edge-weighted planar graph) is NP-hard. In addition, it is well-known that if a simple planar graph is 3-vertex-connected, then its dual is a simple planar graph. In 1976, Garey, Johnson, and Tarjan [33] proved that the problem of establishing whether a 3-vertex-connected planar graph is Hamiltonian is NP-complete, thus, as also noted by Haglin and Venkatesan [40], finding the largest bond of a simple planar graph is also NP-hard, contrasting with the polynomial-time solvability of MAXIMUM CUT on planar graphs. Recently, Chaourar proved that LARGEST BOND can be solved in polynomial time on series parallel graphs and graphs without $K_5 \setminus e$ as a minor in [12, 13].

Computing the largest bond of a graph G corresponds to computing the maximum minimal cut of G . Graph problems about finding a maximum minimal (or minimum maximal) solutions such as MAXIMUM MINIMAL VERTEX COVER [9, 60], MAXIMUM MINIMAL DOMINATING SET [2], MAXIMUM MINIMAL EDGE COVER [44], MAXIMUM MINIMAL SEPARATOR [42], MINIMUM MAXIMAL MATCHING [32, 58], and MINIMUM MAXIMAL INDEPENDENT SET [20], have been long studied.

From the point of view of parameterized complexity, it is well known that MAXIMUM CUT can be solved in FPT time when parametrized by the size of the solution [46], and since every graph has a cut with at least half the edges [26], it follows that it has a linear kernel. Concerning approximation algorithms, a $1/2$ -approximation algorithm can be obtained by randomly partitioning the set vertices into two parts, which induces a cut-set whose expected size is at least half of the number of edges [47]. The best-known result is the seminal work of Goemans and Williamson [35], who gave a 0.878 -approximation based on semidefinite programming. This has the best approximation factor unless the Unique Games Conjecture fails [45]. To the best of our knowledge, there is no algorithmic study regarding the parameterized complexity of computing the largest bond of a graph as well as the approximability of the problem. Observe that a bond induces a feasible solution of MAXIMUM CONNECTED CUT, but not the other way around, since $G[T]$ may be disconnected. Indeed, the size of a largest bond can be arbitrarily smaller than the size of the maximum connected cut; take, e.g., a star with n leaves. For MAXIMUM CONNECTED CUT on general graphs, there exists a $\Omega(1/\log n)$ -approximation [30], where n is the number of vertices. Also, there is a constant-factor approximation with factor $1/2$ for graphs of bounded treewidth [56], and a polynomial-time approximation scheme for graphs of bounded genus [41].

Recently, Saurabh and Zehavi [55] considered a generalization of MAXIMUM CONNECTED CUT, named MULTI-NODE HUB. In this problem, given numbers l and k , the objective is to find a cut $[S, T]$ of G such that $G[S]$ is connected, $|S| = l$ and $|\partial(S)| \geq k$. They observed that the problem is $W[1]$ -hard when parameterized on l , and gave the first parameterized algorithm for the problem with respect to the parameter k . We remark that the $W[1]$ -hardness also holds for LARGEST BOND parameterized by $|S|$.

Since every nonempty bond determines a cut $[S, T]$ such that $G[S]$ and $G[T]$ are both connected, every bond of G has size at most $|E(G)| - |V(G)| + 2$. A graph G has a bond of size $|E(G)| - |V(G)| + 2$ if and only if $V(G)$ can be partitioned into two parts such that each part induces a tree. Such graphs are known as *Yutsis graphs*. The set of planar Yutsis graphs is exactly the dual class of Hamiltonian planar graphs. According to Aldred, Van Dyck, Brinkmann, Fack, and McKay [1], cubic Yutsis graphs appear in the quantum theory of angular momenta as a graphical representation of general recoupling coefficients. They can be manipulated following certain rules in order to generate the so-called summation formulae for the general recoupling coefficient (see [4, 23, 59]).

There are very few results about the largest bond size in general graphs. In 2008, Aldred, Van

Dyck, Brinkmann, Fack, and McKay [1] showed that if a Yutsis graph is regular with degree 3, the partition of the vertex set from the largest bond will result in two sets of equal size. In 2015, Ding, Dziobiak and Wu [21] proved that any simple 3-connected graph G will have a largest bond with size at least $\frac{2}{17}\sqrt{\log n}$, where $n = |V(G)|$. In 2017, Flynn [28] verified the conjecture that any simple 3-connected graph G has a largest bond with size at least $\Omega(n^{\log_3 2})$ for a variety of graph classes including planar graphs.

Even though there are many important applications of MAXIMUM CONNECTED CUT and LARGEST BOND such as image segmentation [57], forest planning [11], and computing a market splitting for electricity markets [37], the known results are much fewer than those for MAXIMUM CUT due to the difficult nature of simultaneously maximizing its size and handling the connectivity of a cut.

1.1 Our contributions

In this paper, we complement the state of the art on the problems of computing the largest bond and the maximum connected cut of a graph. Preliminarily, we present general reductions that allows us to observe that LARGEST BOND and MAXIMUM CONNECTED CUT are NP-hard for several graph classes for which MAXIMUM CUT is NP-hard. Using this framework, we are able to show that LARGEST BOND and MAXIMUM CONNECTED CUT on graphs of clique-width w cannot be solved in time $f(w) \times n^{o(w)}$ unless the ETH fails. We also prove that both MAXIMUM CONNECTED CUT and LARGEST BOND are NP-complete even on planar bipartite graphs. Interestingly, although MAXIMUM CUT can be solved in polynomial time on planar graphs [39, 50] and it is trivial on bipartite graphs, our both problems are intractable even on the intersection of these classes. Also, we show that these problems are NP-complete on split graphs. Moreover, we show that LARGEST BOND does not admit a constant-factor approximation algorithm, unless $P = NP$, and thus is asymptotically harder to approximate than MAXIMUM CUT.

To tackle this difficulty, we study both problems from the perspective of the parameterized complexity. Using win/win approaches, we consider the strategy of preprocessing the input in order to bound the treewidth of the resulting instance. After that, we give $O^*(2^{O(tw \log tw)})$ -time algorithms for both problems¹, where tw is the tree-width of the input graph. Moreover, we can improve this running time using the rank-based approach [5] to $O^*(c^{tw})$ for some constant c and using the Cut & Count technique [19] to $O^*(3^{tw})$ for MAXIMUM CONNECTED CUT and $O^*(4^{tw})$ for LARGEST BOND, using randomization. Let us note that our result generalizes the polynomial time algorithms for LARGEST BOND on series parallel graphs and graphs without $K_5 \setminus e$ as a minor due to Chaourar [12, 13] since such graphs are tree-width bounded [54]. Based on these algorithms, we give $O^*(2^{O(k)})$ -time algorithms for both problems. Also, we remark that the problems do not admit polynomial kernels, unless $NP \subseteq coNP/poly$.

Finally, we consider different structural graph parameters. We design tight (assuming ETH) XP-time algorithms for both problems when parameterized by clique-width cw . Also, we give $O^*(2^{2^{tc}+tc})$ -time and $O^*(2^{tc}3^{2^{tc}})$ -time FPT algorithms for MAXIMUM CONNECTED CUT and LARGEST BOND, respectively, where tc is the minimum size of a twin-cover of the input graph.

2 Intractability results

In this section, we discuss aspects of the hardness of computing the largest bond and the maximum connected cut of a graph. Notice that LARGEST BOND and MAXIMUM CONNECTED CUT are Turing reducible to LARGEST st -BOND and MAXIMUM CONNECTED st -CUT, respectively. Therefore, the hardness of LARGEST BOND and MAXIMUM CONNECTED CUT presented in this section also holds for LARGEST st -BOND and MAXIMUM CONNECTED st -CUT as well, unless $P = NP$.

¹The $O^*(\cdot)$ notation suppresses polynomial factors in the input size.

Next, we present a general framework for reducibility from MAXIMUM CUT to LARGEST BOND, by defining a special graph operator ψ such that MAXIMUM CUT on a graph class \mathcal{F} is reducible to LARGEST BOND on the image of \mathcal{F} via ψ . An interesting particular case occurs when \mathcal{F} is closed under ψ (for instance, chordal graphs are closed under ψ).

Definition 1. Let G be a graph and let $n = V(G)$. The graph $\psi(G)$ is constructed as follows:

- (i) create n disjoint copies G_1, G_2, \dots, G_n of G ;
- (ii) add vertices v_a and v_b ;
- (iii) add an edge between v_a and v_b ;
- (iv) add all possible edges between $V(G_1 \cup G_2 \cup \dots \cup G_n)$ and $\{v_a, v_b\}$.

Definition 2. A set of graphs \mathcal{G} is closed under operator ψ if whenever $G \in \mathcal{G}$, then $\psi(G) \in \mathcal{G}$.

Theorem 1. LARGEST BOND is NP-complete for any graph class \mathcal{G} such that:

- \mathcal{G} is closed under operator ψ ; and
- MAXIMUM CUT is NP-complete for graphs in \mathcal{G} .

Proof. Let $G \in \mathcal{G}$, $n = |V(G)|$, and $H = \psi(G)$. By (i), $H \in \mathcal{G}$. Suppose that G has a cut $[S, V(G) \setminus S]$ of size k , and let S_1, S_2, \dots, S_n be the copies of S in G_1, G_2, \dots, G_n , respectively. If $S' = \{v_a\} \cup S_1 \cup S_2 \cup \dots \cup S_n$, then $[S', V(H) \setminus S']$ defines a bond $\partial(S')$ of H of size at least $nk + n^2 + 1$. Conversely, suppose H has a bond $\partial(S')$ of size at least $nk + n^2 + 1$. We consider the following cases: (a) If $\{v_a, v_b\} \subseteq S'$, then for all copies G_i but one we have $V(G_i) \subseteq S'$, as otherwise the graph induced by $V(H) \setminus S'$ would not be connected, and $\partial(S')$ would not be a bond. Thus, $V(H) \setminus S' \subseteq V(G_j)$ for some j , then the size of $\partial(S')$ is smaller than $nk + n^2 + 1$, a contradiction. (b) If $v_a \in S'$ and $v_b \notin S'$, then $\{v_a, v_b\}$ is incident with exactly $n^2 + 1$ edges crossing $[S', V(H) \setminus S']$, which implies that at least one copy G_i has k or more edges crossing $[S', V(H) \setminus S']$. Therefore, G has a cut of size at least k . \square

Corollary 1. LARGEST BOND is NP-complete for the following classes:

1. chordal graphs;
2. co-comparability graphs;
3. P_5 -free graphs;
4. AT-free graphs.

Proof. Bodlaender and Jansen [7] proved that MAXIMUM CUT is NP-complete when restricted to split and co-bipartite graphs. Since split graphs are chordal and co-bipartite graphs are P_5 -free, AT-free and co-comparability graphs, the NP-completeness also holds for these classes. Now we have to show that the classes are closed under ψ .

(1.) A graph is chordal if every cycle of length at least 4 has a chord. Let G be a chordal graph. Notice that the disjoint union of G_1, G_2, \dots, G_n is also chordal. In addition, no chordless cycle of length at least 4 may contain either v_a or v_b because both vertices are universal. Therefore, $\psi(G)$ is chordal.

(2.) A graph is a co-comparability graph if it is the intersection graph of curves from a line to a parallel line. Let G be a co-comparability graph. Notice that the class of co-comparability graphs is closed under disjoint union. Thus, in order to conclude that $\psi(G)$ is co-comparability, it is enough to observe that from a representation of curves (from a line to a parallel line) of the disjoint union of G_1, G_2, \dots, G_n , one can construct a representation of $\psi(G)$ by adding two concurrent lines (representing v_a and v_b) crossing all curves.

(3.) The disjoint union of P_5 -free graphs is also P_5 -free. In addition, no induced P_5 contains either v_a or v_b because both vertices are universal. Then, the class of P_5 -free graphs is closed under ψ .

(4.) Three vertices of a graph form an asteroidal triple if every two of them are connected by a path avoiding the neighbourhood of the third. A graph is AT-free if it does not contain any asteroidal triple. Since an asteroidal triple does not contain universal vertices and it is a connected subgraph, the class of AT-free graphs is closed under ψ . \square

Now we consider a similar result for MAXIMUM CONNECTED CUT.

Definition 3. Let G be a graph and let $n = V(G)$. The graph $\phi(G)$ is constructed as follows:

- (i) create n disjoint copies G_1, G_2, \dots, G_n of G ;
- (ii) add a new vertex v_a ;
- (iv) add exactly one edge from v_a to a vertex of G_i , $1 \leq i \leq n$.

At this point, it is easy to see that a graph G has a cut $[S, V(G) \setminus S]$ of size k if and only if $\phi(G)$ has a bond $\partial(S')$ of size at least nk . Thus, the following theorem also holds.

Theorem 2. MAXIMUM CONNECTED CUT is NP-complete for any graph class \mathcal{G} such that:

- \mathcal{G} is closed under operator ϕ ; and
- MAXIMUM CUT is NP-complete for graphs in \mathcal{G} .

2.1 Algorithmic lower bound for clique-width parameterization

The *clique-width* of a graph G , denoted by $cw(G)$, is defined as the minimum number of labels needed to construct G , using the following four operations:

1. Create a single vertex v with an integer label ℓ (denoted by $\ell(v)$);
2. Take the disjoint union (i.e., co-join) of two graphs (denoted by \oplus);
3. Join by an edge every vertex labeled i to every vertex labeled j for $i \neq j$ (denoted by $\eta(i, j)$);
4. Relabel all vertices with label i by label j (denoted by $\rho(i, j)$).

An algebraic term that represents such a construction of G and uses at most w labels is said to be a w -*expression* of G , and the clique-width of G is the minimum w for which G has a w -expression.

Graph classes with bounded clique-width include cographs, distance-hereditary graphs, graphs of bounded treewidth, graphs of bounded branchwidth, and graphs of bounded rank-width.

In the '90s, Courcelle, Makowsky, and Rotics [16] proved that all problems expressible in MS_1 -logic are fixed-parameter tractable when parameterized by the clique-width of a graph and the logical expression size. The applicability of this meta-theorem has made clique-width become one of the most studied parameters in parameterized complexity. However, although several problems are MS_1 -expressible, this is not the case with MAXIMUM CUT.

In 2014, Fomin, Golovach, Lokshtanov and Saurabh [29] showed that MAXIMUM CUT on a graph of clique-width w cannot be solved in time $f(w) \times n^{o(w)}$ for any function f of w unless Exponential Time Hypothesis (ETH) fails. Using operators ψ and ϕ , we are able to extend this result to LARGEST BOND and MAXIMUM CONNECTED CUT.

Lemma 1. LARGEST BOND and MAXIMUM CONNECTED CUT on graphs of clique-width w cannot be solved in time $f(w) \times n^{o(w)}$ unless the ETH fails.

Proof. MAXIMUM CUT cannot be solved in time $f(w) \times n^{o(w)}$ on graphs of clique-width w , unless Exponential Time Hypothesis (ETH) fails [29]. Therefore, by the polynomial-time reduction presented in Theorem 1 and Theorem 2, it is enough to show that the clique-width of $\psi(G)$ and $\phi(G)$ is upper bounded by a linear function of the clique-width of G .

If G has clique-width $w \geq 2$, then the disjoint union $H_1 = G_1 \oplus G_2 \oplus \dots \oplus G_n$ has clique-width w .

For $\psi(G)$, suppose that all vertices in H_1 have label 1. Now, let H_2 be the graph isomorphic to a K_2 such that $V(H) = \{v_a, v_b\}$, and v_a, v_b are labeled with 2. In order to construct $\psi(G)$ from $H_1 \oplus H_2$ it is enough to apply the join $\eta(1, 2)$. Thus, $\psi(G)$ has clique-width equals w .

For $\phi(G)$, note that from a w -expression of H_1 we can obtain a $w + 1$ -expression of H_1 resulting into a labelled graph such that the neighborhood of v_a are exactly the vertices with label $w + 1$ and all the other vertices of H_1 have label 1. Thus we can add v_a with label 2 and apply the join $\eta(1, 2)$, which implies that $\phi(G)$ has clique-width at most $w + 1$. \square

2.2 On planar bipartite graphs

Although MAXIMUM CUT is trivial for bipartite graphs, we first observe that the same does not apply to compute the largest bond.

Theorem 3. LARGEST BOND is NP-complete even on planar bipartite subcubic graphs.

Proof. In [40], Haglin and Venkatesan proved that LARGEST BOND remains NP-complete on planar cubic graphs. Since subdivision of edges does not increase the size of the largest bond, by subdividing each edge of a planar cubic graph G we obtain planar bipartite subcubic graph G' such that G has a bond of size k if and only if G' has a bond of size k . \square

Theorem 4. Let G be a simple bipartite graph and $\ell \in \mathbb{N}$. To determine the largest bond $\partial(S)$ of G with $|S| = \ell$ is $W[1]$ -hard with respect to ℓ .

Proof. From an instance H of k -INDEPENDENT SET on regular graphs we first construct a multi-graph G' by adding an edge between any pair of vertices. Finally, we obtain a simple graph G by subdividing every edge of G' . Notice that H has an independent set of size k if and only if G has a bond $\partial(S)$ of size $dk + k(n - k)$ with $|S| = k + \binom{k}{2}$, where d is the vertex degree of H . \square

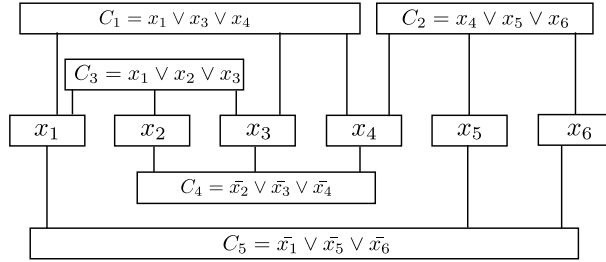
Now, we consider MAXIMUM CONNECTED CUT on planar bipartite graphs.

Theorem 5. MAXIMUM CONNECTED CUT is NP-complete on planar bipartite graphs.


Proof. The reduction is based on the proof of Theorem 4 in [41], which proves that MAXIMUM CONNECTED CUT is NP-hard on planar graphs. We transform the planar reduced graph in [41] into our planar bipartite graph by using additional vertices, called *bridge vertices*.

In this proof, we reduce an instance of PLANAR MONOTONE RECTILINEAR 3-SAT, which is known to be NP-complete [3], to a planar bipartite instance of MAXIMUM CONNECTED CUT. An instance of PLANAR MONOTONE RECTILINEAR 3-SAT consists of a 3-CNF formula ϕ satisfies the following properties: (a) each clause contains either only positive literals or only negative literals, (b) the bipartite incidence graph G_ϕ is planar, and (c) G_ϕ has a *monotone rectilinear representation*. In a monotone rectilinear representation of G_ϕ , the variable vertices are drawn on a straight line in the order of their indices and each positive (resp., negative) clause vertex is drawn in the “positive side” (resp., “negative side”) of the plane defined by the straight line (See Figure 1).

The Reduction. Given a formula ϕ of PLANAR MONOTONE RECTILINEAR 3-SAT with a monotone rectilinear representation as in Figure 1 (a), let $X = \{x_1, x_2, \dots, x_n\}$ be a set of variables and $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ be a set of clauses. Let $K > 4m^2$ and $m > 2$. Then we create the graph $H_\phi = (V, E)$ as follows (see Fig.1). For each variable $x_i \in X$, we create two *literal vertices* $v(x_i)$ and $v(\bar{x}_i)$ corresponding to the literals x_i and \bar{x}_i , respectively. Moreover, we add K *helper vertices* h_1^i, \dots, h_K^i and connect h_k^i to $v(x_i)$ and $v(\bar{x}_i)$ for each $k = 1, \dots, K$. For every clause $C_j \in \mathcal{C}$, we



a) A monotone rectilinear representation of an instance of PLANAR MONOTONE RECTILINEAR 3-SAT.



b) The reduced graph H_ϕ .

Figure 1: An example illustrating (a) the rectilinear representation of a formula $\phi = (x_1 \vee x_3 \vee x_4) \wedge (x_4 \vee x_5 \vee x_6) \wedge (x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_2 \vee \bar{x}_3 \vee \bar{x}_4) \wedge (\bar{x}_1 \vee \bar{x}_5 \vee \bar{x}_6)$ of PLANAR MONOTONE RECTILINEAR 3-SAT and (b) the reduced graph H_ϕ .

create a *clause vertex* $v(C_j)$ and connect $v(x_i)$ (resp., $v(\bar{x}_i)$) to $v(C_j)$ if C_j contains x_i (resp., \bar{x}_i). Moreover, we attach \sqrt{K} pendant vertices to each $v(C_j)$.

Then we attach K pendant vertices to each helper vertex h_k^i . Finally, we add a *bridge vertex* $b_{i,i+1}$ with K pendant vertices that we make adjacent to each $v(x_i)$, $v(\bar{x}_i)$, $v(x_{i+1})$, and $v(\bar{x}_{i+1})$ for $1 \leq i \leq n-1$. We denote by H_ϕ the graph we obtained. Notice that we can draw H_ϕ in the plane according to a monotone rectilinear representation. Moreover, H_ϕ is bipartite since we only add helper and bridge vertices, which have a neighbor only in literal vertices, and pendant vertices to the planar drawing of G_ϕ .

Clearly, this reduction can be done in polynomial time. To complete the proof, we prove the following claim.

Claim 1. *A formula ϕ is satisfiable if and only if there is a maximum connected cut of size at least $m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$ in H_ϕ .*

Proof. Let V_X , V_C , V_H , V_B , and V_P be the set of literal vertices, clause vertices, helper vertices, bridge vertices, and pendant vertices, respectively.

(\Rightarrow) We are given a satisfiable assignment α for ϕ . For α , we denote a true literal by l_i . We also call $v(l_i)$ a *true literal vertex*. Let $S = \bigcup_{i=1}^n \{v(l_i)\} \cup V_C \cup V_H \cup V_B$. That is, S consists of the set of true literal vertices, all the clause vertices, all helper vertices and bridge vertices. Observe that the induced subgraph by S is connected. This follows from the facts that each clause has at least one true literal and literal vertices are connected by bridge vertices.

Finally, we show that $|\partial(S)| \geq m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$. Since each clause vertex has \sqrt{K} pendant vertices and each helper vertex K pendant vertices, there are $m\sqrt{K} + nK^2$ cut edges. Moreover, each bridge vertex has K cut edges incident to its pendant vertices and two cut edges incident to literal vertices not in S . Finally, since either $v(x_i)$ or $v(\bar{x}_i)$ is not in S , there are nK cut edges between literal vertices and helper vertices. Therefore, we have $|\partial(S)| \geq m\sqrt{K} + nK^2 + (n-1)(K+2) + nK = m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$.

(\Leftarrow) We are given a connected cut S in H_ϕ such that $|\partial(S)| \geq m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$. Here, we assume without loss of generality that S is an optimal connected cut of H_ϕ . Suppose, for contradiction, that neither of $v(x_i)$ and $v(\bar{x}_i)$ is contained in S for some variable x_i . Then, all helper vertices h_k^i cannot be contained in S due to the connectivity of S . There are $m\sqrt{K} + 3m + 2nK + (K+4)(n-1)$ edges except for those between helper vertices and its pendant vertices. Thus, it follows that $|\partial(S)| \leq m\sqrt{K} + 3m + 2nK + (K+4)(n-1) + (n-1)K^2$. Since $K > 4m^2$ and $m > 2$, this contradicts the assumption that $|\partial(S)| \geq m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$. Thus, at least one literal vertex must be contained in S for each x_i .

Next, we show that every helper vertex must be contained in S . Suppose a helper vertex h_k^i is not contained in S . Then, all K pendant vertices attached to h_k^i is not contained in S due to the connectivity of S . Since at least one literal vertex of x_i is contained in S , we can increase the size of the cut by moving h_k^i to S , contradicting the optimality of S . Therefore, we assume that every helper vertex is contained in S . Similar to helper vertices, we can prove that every bridge vertex is contained in S .

Then, we observe that exactly one literal vertex must be contained in S for each x_i . Suppose that both $v(x_i)$ and $v(\bar{x}_i)$ are contained in S for some x_i . Since all helper vertices and bridge vertices are contained in S , we may increase the size of the cut by moving either of $v(x_i)$ or $v(\bar{x}_i)$ to $V \setminus S$. However, there are some issues we have to consider carefully. Suppose that $v(x_i)$ is moved to $V \setminus S$. Then, some clause vertices $v(C_j)$ in S can be disconnected in $G[S]$. If so, we also move $v(C_j)$ together with its pendant vertices to $V \setminus S$. Since there are at least $K+1$ cut edges newly introduced but at most $m(\sqrt{K}+3)$ edges removed from the cutset, the size of the cutset is increased, also contradicting the optimality of S .

Finally, we show that every clause vertex is, in fact, contained in S . Suppose that $|S \cap V_C| = m' < m$. If $v(C_j)$ is not in S , its pendant vertices are also not in S . Due to the optimality of S , the pendant vertices of every helper vertex and every bridge vertex is in $V \setminus S$. Thus, we have




Figure 2: An instance of MAXIMUM CONNECTED CUT on split graphs reduced from an instance of EXACT 3-COVER where $X = \{x_1, x_2, x_3, x_4, x_5, x_6, x_7, x_8, x_9\}$ and $\mathcal{F} = \{\{x_1, x_2, x_3\}, \{x_1, x_3, x_4\}, \{x_2, x_4, x_5\}, \{x_5, x_8, x_9\}, \{x_3, x_6, x_7\}, \{x_6, x_7, x_8\}, \{x_7, x_8, x_9\}, \{x_6, x_8, x_9\}, \{x_4, x_8, x_9\}, \{x_2, x_7, x_9\}\}$.

$|\partial(S) \cap \partial(V_H \cup V_B)| = nK(K+1) + (K+2)(n-1) = nK^2 + (2n-1)K + 2(n-1)$. Therefore, $|\partial(S)| = |\partial(S) \cap \partial(V_C)| + |\partial(S) \cap \partial(V_H \cup V_B)| \leq m'\sqrt{K} + 3(m-m') + nK^2 + (2n-1)K + 2(n-1) < m\sqrt{K} + nK^2 + (2n-1)K + 2(n-1)$ as $K > 4m^2$. This is also contradicting to the assumption of the size of the cut.

To summarize, exactly one literal vertex is in S for each variable and every clause vertex is in S . Since $G[S]$ is connected, every clause vertex is adjacent to a literal vertex included in S . Given this, we can obtain a satisfying assignment for ϕ . \square

This completes the proof of Theorem 5. \square

2.3 On split graphs

Theorem 6. MAXIMUM CONNECTED CUT is NP-complete on split graphs.

Proof. We reduce the following problem called EXACT 3-COVER, which is known to be NP-complete: Given a set $X = \{x_1, x_2, \dots, x_{3n}\}$ and a family $\mathcal{F} = \{F_1, F_2, \dots, F_m\}$, where each $F_i = \{x_{i_1}, x_{i_2}, x_{i_3}\}$ has three elements of X , the objective is to find a subfamily $\mathcal{F}' \subseteq \mathcal{F}$ such that every element in X is contained in exactly one of the subsets \mathcal{F}' . By making some copies of 3-element sets if necessary, we may assume that $|\{F \in \mathcal{F} \mid x \in F\}| \geq 3(n+2)$ for each $x \in X$, which implies that m is sufficiently large compared to n .

Given an instance of EXACT 3-COVER with $|\{F \in \mathcal{F} \mid x \in F\}| \geq 3(n+2)$ for each $x \in X$, we construct an instance of MAXIMUM CONNECTED CUT in a split graph as follows. We introduce m vertices u_1, u_2, \dots, u_m , where each u_i corresponds to F_i , and introduce $m-2n$ vertices $u_{m+1}, u_{m+2}, \dots, u_{2(m-n)}$. Let $U := \{u_1, u_2, \dots, u_{2(m-n)}\}$. For $i = m+1, m+2, \dots, 2(m-n)$, introduce a vertex set Y_i of size M , where M is a sufficiently large integer compared to n (e.g. $M = 3n+1$). Now, we construct a graph $G = (U \cup X \cup Y, E)$, where $Y := \bigcup_{m+1 \leq i \leq 2(m-n)} Y_i$, $E_U := \{\{u, u'\} \mid u, u' \in U, u \neq u'\}$, $E_X := \{\{u_i, x_j\} \mid 1 \leq i \leq m, 1 \leq j \leq 3n, x_j \in F_i\}$, $E_Y := \{\{u_i, y\} \mid m+1 \leq i \leq 2(m-n), y \in Y_i\}$, and $E := E_U \cup E_X \cup E_Y$. Then, G is a split graph in which U induces a clique and $X \cup Y$ is an independent set. We now show the following claim.

Claim 2. The original instance of EXACT 3-COVER has a solution if and only if the obtained graph G has a connected cut of size at least $(m-n)^2 + 3m - 3n + (m-2n)M$.

Proof. Suppose that the original instance of EXACT 3-COVER has a solution \mathcal{F}' . Then $S := \{u_i \mid F_i \in \mathcal{F}'\} \cup \{u_i \mid m+1 \leq i \leq 2(m-n)\} \cup X$ is a desired connected cut, because $|\partial(S) \cap E_U| = (m-n)^2$, $|\partial(S) \cap E_X| = \sum_{i=1}^m |F_i| - |X| = 3m - 3n$, and $|\partial(S) \cap E_Y| = (m-2n)M$.

Conversely, suppose that the obtained instance of MAXIMUM CONNECTED CUT has a connected cut S such that $|\partial(S)| \geq (m-n)^2 + 3m - 3n + (m-2n)M$. Since $|\partial(S) \cap E_U| \leq (m-n)^2$, $|\partial(S) \cap E_X| \leq 3m$, and $|\partial(S) \cap E_Y| \leq |S \cap \{u_{m+1}, \dots, u_{2(m-n)}\}| \cdot M$, we obtain $|S \cap \{u_{m+1}, \dots, u_{2(m-n)}\}| = m-2n$, that is, $\{u_{m+1}, \dots, u_{2(m-n)}\} \subseteq S$. Let $t = |S \cap \{u_1, \dots, u_m\}|$, $X_0 = \{x \in X \mid N(x) \cap S = \emptyset\}$ the vertices in X that has no neighbor in S , $X_{\text{all}} = \{x \in X \mid N(x) \subseteq S\}$ the vertices in X whose neighbor is entirely included in S , and $X_{\text{part}} = X \setminus (X_0 \cup X_{\text{all}})$ all the other vertices in X . Recall that every element in X is contained in at least $3(n+2)$ subsets of \mathcal{F} . Then, since $|\partial(S) \cap E_U| = (m-t)(m-2n+t) = (m-n)^2 - (t-n)^2$, $|\partial(S) \cap E_X| \leq |E_X| - |X_{\text{part}}| - |\partial(X_0)| \leq 3m - (3n - |X_{\text{all}}| - |X_0|) - 3(n+2)|X_0|$, $|\partial(S) \cap E_Y| \leq (m-2n)M$, and $|\partial(S)| \geq (m-n)^2 + 3m - 3n + (m-2n)M$, we obtain

$$|X_{\text{all}}| - (3n+5)|X_0| - (t-n)^2 \geq 0. \quad (1)$$

By counting the number of edges between $S \cap \{u_1, u_2, \dots, u_m\}$ and X , we obtain $3t \geq |\partial(X_{\text{all}})| \geq 3(n+2)|X_{\text{all}}|$, which shows that $t \geq (n+2)|X_{\text{all}}|$. If $|X_{\text{all}}| \geq 1$, then $t \geq (n+2)|X_{\text{all}}| \geq n+2|X_{\text{all}}|$, and hence $|X_{\text{all}}| - 3(n+5)|X_0| - (t-n)^2 \leq |X_{\text{all}}| - (2|X_{\text{all}}|)^2 < 0$, which contradicts (1). Thus, we obtain $|X_{\text{all}}| = 0$, and hence we have $t = n$ and $X_0 = \emptyset$ by (1). Therefore, $\mathcal{F}' := \{F_i \mid 1 \leq i \in m, u_i \in S\}$ satisfies that $|\mathcal{F}'| = n$ and $\bigcup_{F \in \mathcal{F}'} F = X$. This shows that \mathcal{F}' is a solution of the original instance of EXACT 3-COVER. \square

This shows that EXACT 3-COVER is reduced to MAXIMUM CONNECTED CUT in split graphs, which completes the proof. \square

Theorem 7. LARGEST BOND is NP-complete on split graphs.

Proof. We give a reduction from MAXIMUM CUT. Given a graph $G = (V, E)$ with n vertices, we create a split graph $G' = (V \cup V_E, E')$ where V is a clique, $V_E = \{e^\ell \mid e \in E, 1 \leq \ell \leq n^3\}$ is an independent set, and $E' = \{\{u, e^\ell\}, \{v, e^\ell\} \mid e = \{u, v\} \in E, 1 \leq \ell \leq n^3\}$. We show that G has a cut of size at least k if and only if G' has a bond of size at least kn^3 . Without loss of generality, we assume that $n > 1$ and $k > 2$.

Let $[S_1, S_2]$ be a cut of G of size k . We define a cut $[S'_1, S'_2]$ of G' with $S_i \subseteq S'_i$ for $i \in \{1, 2\}$. For each $e \in E$ and $1 \leq \ell \leq n^3$, we set $e^\ell \in S'_2$ if both endpoints of e are in S_2 in G , and otherwise $e^\ell \in S'_1$. It is straightforward to verify that $G'[S'_1]$ and $G'[S'_2]$ are connected. If $e = \{u, v\}$ contributes to the cut $[S_1, S_2]$, there are n^3 edges $(\{u, e_\ell\}$ or $\{v, e_\ell\})$ in G' that contribute to $[S'_1, S'_2]$. Therefore, the size of $[S'_1, S'_2]$ is at least kn^3 .

Conversely, let $[S'_1, S'_2]$ be a bond of size kn^3 in G' . Let $S_i = S'_i \cap V$ for $i \in \{1, 2\}$. For each $e = \{u, w\}$ and $i \in \{1, 2\}$, we can observe that $e^\ell \in S_i$ if $u, w \in S_i$ due to the connectivity of S_i and $k > 2$. Since V forms a clique in G' , there are at most n^2 edges between vertices of V in the cut $[S'_1, S'_2]$. Thus, at least $kn^3 - n^2 > kn^3 - n^3 = (k-1)n^3$ edges between V and V_E belong to the cutset. This implies that there are at least k pairs $\{u, w\}$ with $u \in S'_1 \cap V$ and $w \in S'_2 \cap V$, and hence G has a cut of size at least k . \square

2.4 Inapproximability for Largest Bond

While the maximum cut of a graph has at least a constant fraction of the edges, the size of the largest bond can be arbitrarily smaller than the number of edges; take, e.g., a cycle on n edges, for which a largest bond has size 2. This discrepancy is also reflected on the approximability of the problems. Indeed, we show that LARGEST BOND is strictly harder to approximate than MAXIMUM CUT. To simplify the presentation, we consider a weighted version of the problem in which edges are allowed to have weights 0 or 1; the hardness results will follow for the unweighted case as well. In the BINARY WEIGHTED LARGEST BOND, the input is given by a connected weighted graph H with weights $w : E(H) \rightarrow \{0, 1\}$. The objective is to find a bond whose total weight is maximum.

Let G be a graph on n vertices and whose maximum cut has size k . Next, we define the G -edge embedding operator ξ_G . Given a connected weighted graph H , the weighted graph $\xi_G(H)$ is

constructed by replacing each edge $\{u, v\} \in E(H)$ with weight 1 by a copy of G , denoted by G_{uv} , whose edges have weight 1, and, for each vertex t of G_{uv} , new edges $\{u, t\}$ and $\{v, t\}$, both with weight 0.

We can also apply the G -edge embedding operation on the graph $\xi_G(H)$, then on $\xi_G(\xi_G(H))$, and so on. In what follows, for an integer $h \geq 0$, denote by $\xi_G^h(H)$ the graph resulting from the operation that receives a graph H and applies ξ_G successively h times. Notice that $\xi_G^h(H)$ can be constructed in $\mathcal{O}(|V(G)|^{h+1})$ time. For some j , $0 \leq j \leq h-1$, observe that an edge $\{u, v\} \in E(\xi_G^j(H))$ will be replaced by a series of vertices added in iterations $j+1, j+2, \dots, h$. These vertices will be called the *descendants* of $\{u, v\}$, and will be denoted by V_{uv} .

Let K_2 be the graph composed of a single edge $\{u, v\}$, and consider the problem of finding a bond of $\xi_G(K_2)$ with maximum weight. Since edges connecting u or v have weight 0, one can assume that u and v are in different sides of the bond, and the problem reduces to finding a maximum cut of G . In other words, the operator ξ_G embeds an instance G of MAXIMUM CUT into an edge $\{u, v\}$ of K_2 .

This suggests the following strategy to solve an instance of MAXIMUM CUT. For some constant integer $h \geq 1$, calculate $H = \xi_G^h(K_2)$, and obtain a bond F of H with maximum weight. Note that, to solve H , one must solve embedded instances of MAXIMUM CUT in multiple levels simultaneously. For a level j , $1 \leq j \leq h-1$, each edge $\{u, v\} \in E(\xi_G^j(K_2))$ with weight 1 will be replaced by a graph G_{uv} which is isomorphic to G . In Lemma 2 below, we argue that F is such that either $V(G_{uv}) \cup \{u, v\}$ are all in the same side of the cut, or u and v are in distinct sides. In the latter case, the edges of F that separate u and v will induce a cut of G .

In the remaining of this section, we consider a constant integer $h \geq 0$. Then, we define $H^j = \xi_G^j(K_2)$ for every j , $0 \leq j \leq h$, and $H = H^h$. Also, we write $[S, T]$ to denote the cut induced by a bond F of H .

Definition 4. Let F be a bond of H with cut $[S, T]$. We say that an edge $\{u, v\} \in E(H^j)$ with weight 1 is nice for F if either

- $|\{u, v\} \cap S| = 1$, or
- $(\{u, v\} \cup V_{uv}) \subseteq S$, or
- $(\{u, v\} \cup V_{uv}) \subseteq T$.

Also, we say that F is nice if, for every j , $0 \leq j \leq h-1$, and every edge $\{u, v\} \in E(H^j)$ with weight 1, $\{u, v\}$ is nice for F .

Lemma 2. There is a polynomial-time algorithm that receives a bond F , and finds a nice bond F' such that $w(F') = w(F)$.

Proof. Let $[S, T]$ be the cut induced by F and let j^* be the minimum value such that there exists an edge $\{u, v\} \in H^{j^*}$ with weight 1 which is not nice for F . Then $|\{u, v\} \cap S| \neq 1$. Assume, without loss of generality that $u, v \in S$. In this case, $U := V_{uv} \cap T$ is not empty. Since removing vertices $\{u, v\}$ disconnects U , and T must be connected, it follows that $U = T$. This implies that $N(T) \subseteq (V_{uv} \setminus T) \cup \{u, v\}$.

We will construct a bond F' of H with cut $[S', T']$. Let S' be the set of vertices in the connected component of $H[S \setminus \{v\}]$ which contains u , and $T' = V(H) \setminus S'$. Since $H[S]$ is connected, so must be $H[S \setminus S']$. Also, each vertex of U is adjacent to v , thus $H[(S \setminus S') \cup U]$ is connected. Observe that $T' = (S \setminus S') \cup U$, so indeed the cut $[S', T']$ induces a bond $F' = \partial(S')$. Observe that any edge that appears only in F or only in F' is adjacent to v . Since such edges have weight 0, this implies $w(F) = w(F')$.

To complete the proof, we claim that if for some j , $0 \leq j \leq h$, there exists an edge $\{u, v\} \in H^j$ with weight 1 which is not nice for F' , then $j > j^*$. If this claim holds, then we need to repeat the previous procedure at most h times before obtaining a nice bond F' .

To prove the claim, consider an edge $\{s, t\} \in H^j$ which is not nice for F' . Suppose, for a contradiction, that $V_{st} \cap V_{uv} = \emptyset$. There are two possibilities. If $s, t \in S'$, then $V_{st} \subseteq S'$; if $s, t \in T'$, then $V_{st} \subseteq S \setminus S' \subseteq T'$. In either possibility, $\{s, t\}$ is nice for F' . This is a contradiction, and thus $V_{uv} \cap V_{st} \neq \emptyset$.

The statement $V_{uv} \cap V_{st} \neq \emptyset$ can only happen if $V_{uv} \subseteq V_{st}$ or $V_{st} \subseteq V_{uv}$. If $V_{uv} \subseteq V_{st}$, then $U \subseteq V_{st}$ and $s, t \in S$. This implies that $\{s, t\}$ is not nice for F . But in this case $j < j^*$, contradicting the choice of j^* . Therefore, $V_{st} \subseteq V_{uv}$, and $j > j^*$, proving our claim. \square

In the following, assume that F is a nice bond with cut $[S, T]$. Consider a level j , $0 \leq j \leq h$, and an edge $\{u, v\} \in E(H^j)$ with weight 1 such that $|\{u, v\} \cap S| = 1$. If $j < h$, then we define F_{uv} to be the subset of edges in F which are incident with some vertex of V_{uv} ; if $j = h$, then we define $F_{uv} = \{\{u, v\}\}$. Note that, because F is nice, if $|\{u, v\} \cap S| \neq 1$, then no edge of F is incident with V_{uv} .

Suppose now that $|\{u, v\} \cap S| = 1$ for some edge $\{u, v\} \in E(H^j)$ with weight 1 and $0 \leq j \leq h - 1$.

In this case, F induces a cut-set of G_{uv} . Namely, define $\hat{S}_{uv} := S \cap V(G_{uv})$ and $\hat{T}_{uv} := T \cap V(G_{uv})$ and let \hat{F}_{uv} be the cut-set of G_{uv} corresponding to cut $[\hat{S}_{uv}, \hat{T}_{uv}]$.

Observe that for distinct edges $\{u, v\}$ and $\{s, t\}$, it is possible that $|\hat{F}_{uv}| \neq |\hat{F}_{st}|$. We will consider bonds F for which all induced cut-sets \hat{F}_{uv} have the same size.

Definition 5. Let ℓ be a positive integer. A bond F of H with cut $[S, T]$ is said to be ℓ -uniform if, (i) F is nice, and (ii) for every j , $0 \leq j \leq h - 1$, and every edge $\{u, v\} \in E(H^j)$ with weight 1 such that $|\{u, v\} \cap S| = 1$, $|\hat{F}_{uv}| = \ell$.

An ℓ -uniform bond induces a cut-set of G of size ℓ .

Lemma 3. Suppose F is an ℓ -uniform bond of H . One can find in polynomial time a cut-set L of G with $|L| = \ell$.

Proof. Let u, v be the vertices of K_2 to which ξ_G was applied. Since F is ℓ -uniform, $|\hat{F}_{uv}| = \ell$. Note that \hat{F}_{uv} induces a cut-set of size ℓ on G . \square

In the opposite direction, a cut of G induces an ℓ -uniform bond of H .

Lemma 4. Suppose L is a cut-set of G with $|L| = \ell$. One can find in polynomial time an ℓ -uniform bond F of H with $w(F) = \ell^h$.

Proof. For each $j \geq 0$, we construct a bond F^j of H^j . For $j = 0$, let F^0 be the set containing the unique edge of $H^0 = K_2$. Suppose now that we already constructed a bond F^{j-1} of H^{j-1} . For each edge $\{u, v\} \in F^{j-1}$, let L_{uv} be the set of edges of G_{uv} corresponding to L . Define $F^j := \cup_{\{u, v\} \in F^{j-1}} L_{uv}$. One can verify that indeed F^j is a bond of H^j , and that $w(F_j) = |L| \times w(F_{j-1}) = \ell^j$. \square

Lemma 5. There is a polynomial-time algorithm that receives a bond F of H , and finds an ℓ -uniform bond F' of H such that $w(F') = \ell^h \geq w(F)$.

Proof. Let $[S, T]$ be the cut corresponding to F . First, find the largest cut-set of a graph G_{uv} over cut-sets \hat{F}_{uv} . More precisely, define \hat{F} to be the cut-set \hat{F}_{uv} with maximum $|\hat{F}_{uv}|$ over all edges $\{u, v\} \in E(H^j)$ with weight 1 such that $|\{u, v\} \cap S| = 1$, and over all j , $0 \leq j \leq h - 1$. Let $\ell := |\hat{F}|$.

We claim that for every j , $0 \leq j \leq h$, and every edge $\{u, v\} \in E(H^j)$ with weight 1 such that $|\{u, v\} \cap S| = 1$, $w(F_{uv}) \leq \ell^{h-j}$. The proof is by (backward) induction on j . For $j = h$, $F_{uv} = \{u, v\}$, so $w(F_{uv}) = 1$. Next, let $j < h$, and assume the claim holds for $j + 1$.

Let F_{uv}^0 be the subset of edges in F_{uv} incident with u or v . The set F_{uv} can be partitioned into F_{uv}^0 and sets F_{st} for $\{s, t\} \in \hat{F}_{uv}$. To see this, observe that each edge $\{x, y\} \in F_{uv} \setminus F_{uv}^0$ must be incident with descendants of $\{u, v\}$, and thus $\{x, y\}$ is incident with vertices of V_{st} , for some edge

$\{s, t\} \in E(G_{uv})$. Since $|\{x, y\} \cap S| = 1$, neither $V_{st} \cup \{s, t\} \subseteq S$, nor $V_{st} \cup \{s, t\} \subseteq T$. Because F is nice, it follows that $|\{s, t\} \cap S| = 1$, then $\{s, t\} \in \hat{F}_{uv}$, and thus $\{x, y\} \in F_{st}$. To complete the claim, observe that, by the induction hypothesis, $w(F_{st}) \leq \ell^{h-j-1}$ for each $\{s, t\} \in \hat{F}_{uv}$, and recall that $|\hat{F}_{uv}| \leq |\hat{F}|$. Therefore

$$w(F) = w(F_{uv}^0) + \sum_{\{s,t\} \in \hat{F}_{uv}} w(F_{st}) \leq |\hat{F}| \times \ell^{h-j-1} = \ell^{h-j}.$$

Using Lemma 4 for \hat{F} , we construct a bond F' for H with $w(F') = \ell^h$. □

Lemma 6. *Let F^* be a bond of H with maximum weight. Then $w(F^*) = k^h$.*

Proof. We assume that F^* is ℓ -uniform such that $w(F^*) = \ell^h$ for some ℓ ; if this is not the case, then use Lemma 5.

Since F^* is ℓ -uniform, using Lemma 4 one obtains a cut-set L of G with size ℓ , then $\ell \leq k$, and thus $w(F^*) \leq k^h$.

Conversely, let L be a cut-set of G with size k . Using Lemma 4 for L , we obtain a bond F of H with weight k^h , and thus $w(F^*) \geq k^h$. □

Lemma 7. *If there exists a constant-factor approximation algorithm for WEIGHTED LARGEST BOND, then $P = NP$.*

Proof. Consider a graph G whose maximum cut has size k . Construct graph H and obtain a bond F of H using an α -approximation, for some constant $0 < \alpha < 1$. Using the algorithm of Lemma 5, obtain an ℓ -uniform bond F' of H such that $w(F') = \ell^h \geq w(F)$. Using Lemma 6 and the fact that F' is an α -approximation, $\ell^h \geq \alpha \times k^h$. Using Lemma 3, one can obtain a cut-set L of G with size $\ell \geq \alpha^{\frac{1}{h}} k$.

For any constant ε , $0 < \varepsilon < 1$, we can set $h = \lceil \log_{1-\varepsilon} \alpha \rceil$, such that the cut-set L has size at least $\ell \geq (1 - \varepsilon)k$. Since MAXIMUM CUT is APX-hard, this implies $P = NP$. □

Theorem 8. *If there exists a constant-factor approximation algorithm for LARGEST BOND, then $P = NP$.*

Proof. We show that if there exists an α -approximation algorithm for LARGEST BOND, for constant $0 < \alpha < 1$, then there is an $\alpha/2$ -approximation algorithm for the BINARY WEIGHTED LARGEST BOND, so the theorem will follow from Lemma 7.

Let H be a weighted graph whose edge weights are all 0 or 1. Let m be the number of edges with weight 0, and let l be the weight of a bond of H with maximum weight. Assume $l \geq 2/\alpha$, as otherwise, one can find an optimal solution in polynomial time by enumerating sets of up to $2/\alpha$ edges.

Construct an unweighted graph G as follows. Start with a copy of H and, for each edge $\{u, v\} \in E(H)$ with weight 1, replace $\{u, v\} \in E(G)$ by m parallel edges. Finally, to obtain a simple graph, subdivide each edge of G . If F is a bond of G , then one can construct a bond F' of H by undoing the subdivision and removing the parallel edges. Each edge of F' has weight 1, with exception of at most m edges. Thus, $w(F') \geq (|F| - m)/m$.

Observe that an optimal bond of H induces a bond of G with size at least ml . Thus, if F is an α -approximation for G , then $|F| \geq \alpha ml$ and therefore

$$w(F') \geq \frac{\alpha ml - m}{m} = \alpha l - 1 \geq \alpha l - \alpha l/2 = \alpha l/2.$$

We conclude that F' is an $\alpha/2$ -approximation for H . □

3 Parameterized algorithms

In this section, we present parameterized algorithms for LARGEST BOND, LARGEST *st*-BOND, MAXIMUM CONNECTED CUT, and MAXIMUM CONNECTED *st*-CUT.

3.1 Algorithmic upper bounds for clique-width parameterization

Lemma 1 shows that LARGEST BOND on graphs of clique-width w cannot be solved in time $f(w) \times n^{o(w)}$ unless the ETH fails. Now, we show that given an expression tree of width at most w , LARGEST BOND can be solved in $f(w) \times n^{O(w)}$ time.

An expression tree \mathcal{T} is *irredundant* if for any join node $\eta(i, j)$, the vertices labeled by i and j are not adjacent in the graph associated with its child. It was shown by Courcelle and Olariu [17] that every expression tree \mathcal{T} of G can be transformed into an irredundant expression tree \mathcal{T} of the same width in time linear in the size of \mathcal{T} . Therefore, without loss of generality, we can assume that \mathcal{T} is irredundant.

Our algorithm is based on dynamic programming over the expression tree of the input graph. We first describe what we store in the tables corresponding to the nodes in the expression tree.

Given a w -labeled graph G , two connected components of G has the same *type* if they have the same set of labels. Thus, a w -labeled graph G has at most $2^w - 1$ types of connected components.

Now, for every node X_ℓ of \mathcal{T} , denote by G_{X_ℓ} the w -labeled graph associated with this node, and let $L_1(X_\ell), \dots, L_w(X_\ell)$ be the sets of vertices of G_{X_ℓ} labeled with $1, \dots, w$, respectively. We define a table where each entry is of the form

$$c[\ell, s_1, \dots, s_w, r, e_1, \dots, e_{2^w-1}, d_1, \dots, d_{2^w-1}],$$

such that: $0 \leq s_i \leq |L_i(X_\ell)|$ for $1 \leq i \leq w$; $0 \leq r \leq |E(G_{X_\ell})|$; $0 \leq e_i \leq \min\{2, |L_i(X_\ell)|\}$ for $1 \leq i \leq 2^w - 1$; and $0 \leq d_i \leq \min\{2, |L_i(X_\ell)|\}$ for $1 \leq i \leq 2^w - 1$.

Each entry of the table represents whether there is a partition V_1, V_2 of $V(G_{X_\ell})$ such that: $|V_1 \cap L_i(G_{X_\ell})| = s_i$; the cut-set of $[V_1, V_2]$ has size at least r ; $G_{X_\ell}[V_1]$ has e_i connected components of type i ; $G_{X_\ell}[V_2]$ has d_i connected components of type i , where $e_i = 2$ means that $G_{X_\ell}[V_1]$ has *at least* two connected components of type i . The same holds for d_i .

Notice that this table contains $f(w) \times n^{O(w)}$ entries. If X_ℓ is the root node of \mathcal{T} (that is, $G = G_{X_\ell}$), then the size of the largest bond of G is equal to the maximum value of r for which the table for X_ℓ contains a valid entry (true value), such that there are j and k such that $e_i = 0$, $e_j = 1$ for $1 \leq i, j \leq 2^w - 1$, $i \neq j$; and $d_i = 0$, $d_k = 1$ for $1 \leq i, k \leq 2^w - 1$, $i \neq k$. Similarly, the maximum connected cut of G can be found in a valid entry for r such that some $e_j = 1$, and $e_i = 0$ for every $i \neq j$.

It is easy to see that we store enough information to compute a largest bond. Note that a w -labeled graph is connected if and only if it has exactly one type of connected components and exactly one component of such a type.

Now we provide the details of how to construct and update such tables. The construction for introduce nodes of \mathcal{T} is straightforward.

Relabel node: Suppose that X_ℓ is a relabel node $\rho(i, j)$, and let $X_{\ell'}$ be the child of X_ℓ . The table for X_ℓ contains a valid entry $c[\ell, s_1, \dots, s_w, r, e_1, \dots, e_{2^w-1}, d_1, \dots, d_{2^w-1}]$ if and only if the table for $X_{\ell'}$ contains an entry

$$c[\ell', s'_1, \dots, s'_w, r, e'_1, \dots, e'_{2^w-1}, d'_1, \dots, d'_{2^w-1}] = \text{true},$$

where:

$$\begin{aligned} s_i &= 0; \\ s_j &= s'_i + s'_j; \\ s_p &= s'_p \text{ for } 1 \leq p \leq w, p \neq i, j; \\ e_p &= e'_p \text{ for any type that contain neither } i \text{ nor } j; \end{aligned}$$

$e_p = 0$ for any type that contains i ;
 and for any type e_p that contains j , it holds that $e_p = \min\{2, e'_p + e'_q + e'_r\}$ where
 e'_q represent the set of labels $(C_p \setminus \{j\}) \cup \{i\}$,
 e'_r represent the set of labels $C_p \cup \{i\}$, and
 C_p is the set of labels associated to p . The same holds for d_1, \dots, d_{2^w-1} .

Union node: Suppose that X_ℓ is a union node with children $X_{\ell'}$ and $X_{\ell''}$. It holds that $c[\ell, s_1, \dots, s_w, r, e_1, \dots, e_{2^w-1}, d_1, \dots, d_{2^w-1}]$ equals true if and only if there are valid entries

$$c[\ell', s'_1, \dots, s'_w, r', e'_1, \dots, e'_{2^w-1}, d'_1, \dots, d'_{2^w-1}]$$

and

$$c[\ell'', s''_1, \dots, s''_w, r'', e''_1, \dots, e''_{2^w-1}, d''_1, \dots, d''_{2^w-1}],$$

having:

$$\begin{aligned}
 s_i &= s'_i + s''_i \text{ for } 1 \leq i \leq w; \\
 r' + r'' &\geq r; \\
 e_k &= \min\{2, e'_k + e''_k\}, \text{ and} \\
 d_k &= \min\{2, d'_k + d''_k\} \text{ for } 1 \leq k \leq 2^w - 1.
 \end{aligned}$$

Join node: Finally, let X_ℓ be a join node $\eta(i, j)$ with the child $X_{\ell'}$. Remind that since the expression tree is irredundant then the vertices labeled by i and j are not adjacent in the graph $G_{X_{\ell'}}$. Therefore, the entry $c[\ell, s_1, \dots, s_w, r, e_1, \dots, e_{2^w-1}, d_1, \dots, d_{2^w-1}]$ equals true if and only if there is a valid entry

$$c[\ell', s_1, \dots, s_w, r', e'_1, \dots, e'_{2^w-1}, d'_1, \dots, d'_{2^w-1}]$$

where

$$r' + s_i \times (|L_j(X_{\ell'})| - s_j) + s_j \times (|L_i(X_{\ell'})| - s_i) \geq r,$$

and $e_p = e'_p$, case p is associated to a type that contains neither i nor j ; $e_p = 1$, case p is associated to $C_{i,j}^{\ell'} \setminus \{i\}$, where $C_{i,j}^{\ell'}$ is the set of labels obtained by the union of the types of $G_{X_{\ell'}}$ with some connected component having either label i or label j ; $e_p = 0$, otherwise. The same holds for d_1, \dots, d_{2^w-1} .

The correctness of the algorithm follows from the description of the procedure. Since for each ℓ , there are $\mathcal{O}((n+1)^w \times m \times (3^{2^w-1})^2)$ entries, the running time of the algorithm is $f(w) \times n^{\mathcal{O}(w)}$. This algorithm together with Lemma 1 concludes the proof of the Theorem 9.

Theorem 9. *Given an expression tree of width at most w , LARGEST BOND and MAXIMUM CONNECTED CUT can be solved in time $f(w) \times n^{\mathcal{O}(w)}$, but they cannot be solved in time $f(w) \times n^{\mathcal{O}(w)}$ unless ETH fails.*

In order to extend this result to LARGEST *st*-BOND and MAXIMUM CONNECTED *st*-CUT, it is enough to observe that given a tree expression \mathcal{T} of G with width w , it is easy to construct a tree expression \mathcal{T}' with width equals $w + 2$, where no vertex of $V(G)$ has the same label than either s or t . Let $w + 1$ be the label of s , and let $w + 2$ be the label of t . By fixing, for each ℓ , $s_{w+1} = |L_{w+1}(X_\ell)|$ and $s_{w+2} = 0$, one can solve LARGEST *st*-BOND in time $f(w) \times n^{\mathcal{O}(w)}$.

In addition, we can improve the clique-width dependence for a clique-width parameterization of LARGEST BOND and MAXIMUM CONNECTED CUT, using another parameter, called "module width", proposed by [53]. However, since the algorithms are complicated, for simplicity and readability, we present them only in Appendix.

Theorem 10. *Given an expression tree of width at most w , LARGEST BOND and MAXIMUM CONNECTED CUT can be solved in time $n^{\mathcal{O}(w)}$.*

3.2 Bounding the treewidth of G to a linear function of k

Now, we deal with the problems parameterized by the size of the solution (k). We consider the strategy of preprocessing the input in order to bound the treewidth of the resulting instance. We start our analysis with LARGEST BOND.

Definition 6. A graph H is called a minor of a graph G if H can be formed from G by deleting edges, deleting vertices, and by contracting edges. For each vertex v of H , the set of vertices of G that are contracted into v is called a branch set of H .

Lemma 8. Let G be a simple connected undirected graph, and k be a positive integer. If G contains $K_{2,k}$ as a minor then G has a bond of size at least k .

Proof. Let H be a minor of G isomorphic to $K_{2,k}$. Since G is connected and each branch set of H induces a connected subgraph of G , from H it is easy to construct a bond of G of size at least k . \square

Combined with Lemma 8, the following results show that, without loss of generality, our study on k -bonds can be reduced to graphs of treewidth $\mathcal{O}(k)$.

Lemma 9. [8] Every graph $G = (V, E)$ contains $K_{2,k}$ as a minor or has treewidth at most $2k - 2$.

Lemma 10. [8] There is a polynomial-time algorithm that either concludes that the input graph G contains $K_{2,k}$ as a minor, or outputs a tree-decomposition of G of width at most $2k - 2$.

From Lemma 8 and Lemma 10, the following holds.

Corollary 2. There is a polynomial-time algorithm that either concludes that the input graph G contains a bond of size k , or outputs a tree-decomposition of G of width at most $2k - 2$.

Since k -bonds are also connected cuts, it holds that there is a polynomial-time algorithm that either concludes that the input graph G contains a bond of size k , or outputs a tree-decomposition of G of width at most $2k - 2$. Such a bound can be improved to $k - 1$ by replace $K_{2,k}$ with $K_{1,k}$ (see [8]) as a minor for MAXIMUM CONNECTED CUT.

3.2.1 The st -bond case

Let $S \subseteq V(G)$ and let $\partial(S)$ be a bond of a connected graph G . Recall that a block is a 2-vertex-connected subgraph of G which is inclusion-wise maximal, and a block-cut tree of G is a tree whose vertices represent the blocks and the cut vertices of G , and there is an edge in the block-cut tree for each pair of a block and a cut vertex that belongs to that block. Then, $\partial(S)$ intersects at most one block of G . More precisely, for any two distinct blocks B_1 and B_2 of G , if $S \cap V(B_1) \neq \emptyset$ and $S \cap V(B_2) \neq \emptyset$, then either $V(B_2) \subseteq S$, or $V(B_2) \subseteq V \setminus S$. Indeed, if this is not the case, then either $G[S]$ or $G[V \setminus S]$ would be disconnected. Thus, to solve LARGEST st -BOND, it is enough to consider, individually, each block on the path between s and t in the block-cut tree of G . Also, if a block is composed of a single edge, then it is a bridge in G , which is not a solution for the problem unless $k = 1$. Thus, we may assume without loss of generality that G is 2-vertex-connected.

Lemma 11. Let G be a 2-vertex-connected graph. For all $v \in V(G) \setminus \{s, t\}$, there is an sv -path and a tv -path which are internally disjoint.

Proof. Since G is 2-vertex-connected, there are two disjoint sv -paths P_s and P'_s and there is a tv -path P'_t which does not include s . Let x be the first vertex of P'_t which belongs to $V(P_s \cup P'_s)$ and assume, w.l.o.g., that $x \in P'_s$. Let P''_t be the sub-path of P'_t from t to x and P''_s the sub-path of P'_s from x to v . Now define P_t as $tP''_t x P''_s v$ and notice that P_t is a tv -path disjoint from P_s . \square

Lemma 12. Let G be a 2-vertex-connected graph. If G contains $K_{2,2k}$ as a minor, then there exists $S \subseteq V(G)$ such that $\partial(S)$ is a bond of size at least k .

Proof. Let G be a graph containing a $K_{2,2k}$ as a minor. If $k = 1$, the statement holds trivially, thus assume $k \geq 2$. Also, since G is connected, one can assume that this minor was obtained by contracting or removing edges only, and thus its branch sets contain all vertices of G . Let A and B be the branch sets corresponding to first side of $K_{2,2k}$, and let X_1, X_2, \dots, X_{2k} be the remaining branch sets.

First, suppose that s and t are in distinct branch sets. If this is the case, then there exist distinct indices $a, b \in \{1, \dots, 2k\}$ such that $s \in A \cup X_a$ and $t \in B \cup X_b$. Now observe that $G[A \cup X_a]$ and $G[B \cup X_b]$ are connected, which implies an st -bond with at least $2k - 1 \geq k$ edges. Now, suppose that s and t are in the same branch set. In this case, one can assume without loss of generality that $s, t \in A \cup X_{2k}$.

Define $U = A \cup X_{2k}$ and $Q = V(G) \setminus U$. Observe that $G[U]$ and $G[Q]$ are connected. Consider an arbitrary vertex v in the set Q . Since G is 2-vertex-connected, Lemma 11 implies that there exist an sv -path P_s and a tv -path P_t which are internally disjoint. Let P'_s and P'_t be maximal prefixes of P_s and P_t , respectively, whose vertices are contained in U .

We partition the set U into parts U_s and U_t such that $G[U_s]$ and $G[U_t]$ are connected. Since $G[U]$ is connected, there exists a tree T spanning U . Direct all edges of T towards s and partition U as follows. Every vertex in P'_s belongs to U_s and every vertex in P'_t belongs to U_t . For a vertex $u \notin V(P'_s \cup P'_t)$, let w be the first ancestor of u (accordingly to T) which is in $P'_s \cup P'_t$. Notice that w is well-defined since $u \in V(T)$ and the root of T is $s \in V(P'_s \cup P'_t)$. Then u belongs to U_s if $w \in V(P'_s)$, and u belongs to U_t if $w \in V(P'_t)$.

Observe that there are at least $2k - 1$ edges between U and Q , and thus there are at least k edges between U_s and Q , or between U_t and Q . Assume the former holds, as the other case is analogous. It follows that $G[U_s]$ and $G[U_t \cup Q]$ are connected and induce a bond of G with at least k edges. \square

Lemma 10 and Lemma 12 imply that there is an algorithm that either concludes that the input graph G has a bond of size at least k , or outputs a tree-decomposition of an equivalent instance G' of width $\mathcal{O}(k)$.

Corollary 3. *Given a graph G , vertices $s, t \in V(G)$, and an integer k , there exists a polynomial-time algorithm that either concludes that G has an st -bond of size at least k or outputs a subgraph G' of G together with a tree decomposition of G' of width equals $\mathcal{O}(k)$, such that G' has an st -bond of size at least k if and only if G has an st -bond of size at least k .*

Proof. Find a block-cut tree of G in linear time [15], and let B_s and B_t be the blocks of G that contain s and t , respectively. Remove each block that is not in the path from B_s to B_t in the block-cut tree of G . Let G' be the remaining graph. For each block B of G' , consider the vertices s' and t' of B which are nearest to s and t , respectively. Using Lemmas 10 and 12 one can in polynomial time either conclude that B has an $s't'$ -bond, in which case G is a yes-instance, or compute a tree decomposition of B with width at most $\mathcal{O}(k)$.

Now, construct a tree decomposition of G' as follows. Start with the union of the tree decompositions of all blocks of G' . Next, create a bag $\{u\}$ for each cut vertex u of G' . Finally, for each cut vertex u and any bag corresponding to a block B connected through u , add an edge between $\{u\}$ and one bag of the tree decomposition of B containing u . Note that this defines a tree decomposition of G' and that each bag has at most $\mathcal{O}(k)$ vertices. \square

Since st -bonds are solutions for MAXIMUM CONNECTED st -CUT, the results presented in Corollary 3 naturally apply to such a problem as well.

3.3 Taking the treewidth as parameter

In the following, given a tree decomposition \mathcal{T} , we denote by ℓ one node of \mathcal{T} and by X_ℓ the vertices contained in the *bag* of ℓ . We assume w.l.o.g that \mathcal{T} is an extended version of a *nice* tree decomposition (see [18]), that is, we assume that there is a special root node r such that $X_r = \emptyset$

and all edges of the tree are directed towards r and each node ℓ has one of the following five types: *Leaf*; *Introduce vertex*; *Introduce edge*; *Forget vertex*; and *Join*. Moreover, define G_ℓ to be the subgraph of G which contains only vertices and edges that have been introduced in ℓ or in a descendant of ℓ .

The number of partitions of a set of k elements is the k -th *Bell number*, which we denote by $B(k)$ ($B(k) \leq k!$ [49]).

Theorem 11. *Given a nice tree decomposition of G with width tw , one can find a bond of maximum size in time $2^{\mathcal{O}(tw \log tw)} \times n$ where n is the number of vertices of G .*

Proof. Let $\partial_G(U)$ be a bond of G , and $[U, V \setminus U]$ be the cut defined by such a bond. Set $S_U^\ell = U \cap X_\ell$. The removal of $\partial_G(U)$ partitions $G_\ell[U]$ into a set C_U^ℓ of connected components, and $G_\ell[V \setminus U]$ into a set $C_{V \setminus U}^\ell$ of connected components. Note that C_U^ℓ and $C_{V \setminus U}^\ell$ define partitions of S_U^ℓ and $X_\ell \setminus S_U^\ell$, denoted by ρ_1^ℓ and ρ_2^ℓ respectively, where the intersection of each connected component of C_U^ℓ with S_U^ℓ corresponds to one part of ρ_1^ℓ . The same holds for $C_{V \setminus U}^\ell$ with respect to $X_\ell \setminus S_U^\ell$ and ρ_2^ℓ .

We define a table for which an entry $c[\ell, S, \rho_1, \rho_2]$ is the size of a largest cut-set (partial solution) of the subgraph G_ℓ , where S is the subset of X_ℓ to the left part of the bond, $X_\ell \setminus S$ is the subset to the right part, and ρ_1, ρ_2 are the partitions of S and $X_\ell \setminus S$ representing, after the removal of the partial solution, the intersection with the connected components to the left and to the right, respectively. If there is no such a partial solution then $c[\ell, S, \rho_1, \rho_2] = -\infty$.

For the case that S is empty, two special cases may occur: either $U \cap V(G_\ell) = \emptyset$, in which case there are no connected components in C_U^ℓ , and thus $\rho_1 = \emptyset$; or C_U^ℓ has only one connected component which does not intersect X_ℓ , i.e., $\rho_1 = \{\emptyset\}$, this case means that the connected component in C_U^ℓ was completely forgotten. Analogously, we may have $\rho_2 = \emptyset$ and $\rho_2 = \{\emptyset\}$. Note that we do not need to consider the case $\{\emptyset\} \subsetneq \rho_i$ since it would imply in a disconnected solution. The largest bond of a connected graph G corresponds to the root entry $c[r, \emptyset, \{\emptyset\}, \{\emptyset\}]$.

To describe a dynamic programming algorithm, we only need to present the recurrence relation for each node type.

Leaf: In this case, $X_\ell = \emptyset$. There are a few combinations for ρ_1 and ρ_2 : either $\rho_1 = \emptyset$, or $\rho_1 = \{\emptyset\}$, and either $\rho_2 = \emptyset$, or $\rho_2 = \{\emptyset\}$. Since for this case G_ℓ is empty, there can be no connected components, so having $\rho_1 = \emptyset$ and $\rho_2 = \emptyset$ is the only feasible choice.

$$c[\ell, S, \rho_1, \rho_2] = \begin{cases} 0 & \text{if } S = \emptyset, \rho_1 = \emptyset \text{ and } \rho_2 = \emptyset, \\ -\infty & \text{otherwise.} \end{cases}$$

Introduce vertex: We have only two possibilities in this case, either v is an isolated vertex to the left ($v \in S$) or it is an isolated vertex to the right ($v \notin S$). Thus, a partial solution on ℓ induces a partial solution on ℓ' , excluding v from its part.

$$c[\ell, S, \rho_1, \rho_2] = \begin{cases} c[\ell', S \setminus \{v\}, \rho_1 \setminus \{\{v\}\}, \rho_2] & \text{if } \{v\} \in \rho_1, \\ c[\ell', S, \rho_1, \rho_2 \setminus \{\{v\}\}] & \text{if } \{v\} \in \rho_2, \\ -\infty & \text{if } \{v\} \notin \rho_1 \cup \rho_2. \end{cases}$$

Introduce edge: In this case, either the edge $\{u, v\}$ that is being inserted is incident with one vertex of each side, or the two endpoints are at the same side. In the former case, a solution on ℓ corresponds to a solution on ℓ' with the same partitions, but with value increased. In the latter case, edge $\{u, v\}$ may connect two connected components of a partial solution on ℓ' .

$$c[\ell, S, \rho_1, \rho_2] = \begin{cases} c[\ell', S, \rho_1, \rho_2] + 1 & \text{if } u \in S \text{ and } v \notin S \text{ or } u \notin S \text{ and } v \in S, \\ \max_{\rho'_1} \{c[\ell', S, \rho'_1, \rho_2]\} & \text{if } u \in S \text{ and } v \in S, \\ \max_{\rho'_2} \{c[\ell', S, \rho_1, \rho'_2]\} & \text{if } u \notin S \text{ and } v \notin S. \end{cases}$$

Here, ρ'_1 spans over all refinements of ρ_1 such that the union of the parts containing u and v results in the partition ρ_1 . The same holds for ρ'_2 .

Forget vertex: In this case, either the forgotten vertex v is in the left side of the partial solution induced on ℓ , or is in the right side. Thus, v must be in the connected component which contains some part of ρ_1 , or some part of ρ_2 . We select the possibility that maximizes the value

$$c[\ell, S, \rho_1, \rho_2] = \max_{\rho'_1, \rho'_2} \{c[\ell', S \cup \{v\}, \rho'_1, \rho_2], c[\ell', S, \rho_1, \rho'_2]\}.$$

Here, ρ'_1 spans over all partitions obtained from ρ_1 by adding v in some part of ρ_1 (if $\rho_1 = \{\emptyset\}$ then $\rho'_1 = \{v\}$). The same holds for ρ'_2 .

Join: This node represents the join of $G_{\ell'}$ and $G_{\ell''}$ where $X_\ell = X_{\ell'} = X_{\ell''}$.

By counting the bond edges contained in $G_{\ell'}$ and in $G_{\ell''}$, each edge is counted at least once, but edges in X_ℓ are counted twice. Thus

$$c[\ell, S, \rho_1, \rho_2] = \max \{c[\ell', S, \rho'_1, \rho'_2] + c[\ell', S, \rho''_1, \rho''_2]\} - |\{\{u, v\} \in E, u \in S, v \in X_\ell \setminus S\}|.$$

In this case, we must find the best combination between the two children. Namely, for $i \in \{1, 2\}$, we consider combinations of ρ'_i with ρ''_i which merge into ρ_i . If $\rho_i = \{\emptyset\}$ then either $\rho'_i = \{\emptyset\}$ and $\rho''_i = \emptyset$; or $\rho'_i = \emptyset$ and $\rho''_i = \{\emptyset\}$. Also, if $\rho_i = \emptyset$ then $\rho'_i = \emptyset$ and $\rho''_i = \emptyset$.

The running time of the dynamic programming algorithm can be estimated as follows. The number of nodes in the decomposition is $\mathcal{O}(tw \times n)$ [18]. For each node ℓ , the parameters ρ_1 and ρ_2 induce a partition of X_ℓ ; the number of partitions of X_ℓ is given by the corresponding Bell number, $B(|X_\ell|) \leq B(tw + 1)$. Each such a partition ρ corresponds to a number of choice of parameter S that corresponds to a subset of the parts of ρ ; thus the number of choices for S is not larger than $2^{|\rho|} \leq 2^{|X_\ell|} \leq 2^{tw+1}$. Therefore, we conclude that the table size is at most $\mathcal{O}(B(tw + 1) \times 2^{tw} \times tw \times n)$. Since each entry can be computed in $2^{\mathcal{O}(tw \log tw)}$ time, the total complexity is $2^{\mathcal{O}(tw \log tw)} \times n$. The correctness of the recursive formulas is straightforward. \square

The reason for the $2^{\mathcal{O}(tw \log tw)}$ dependence on treewidth is because we enumerate all partitions of a bag to check connectivity. However, one can obtain single exponential-time dependence by modifying the presented algorithm using rank-based approach as described in Section 3.3.1.

Theorem 12. LARGEST *st*-BOND, MAXIMUM CONNECTED CUT, and MAXIMUM CONNECTED *st*-CUT can be solved in time $2^{\mathcal{O}(tw \log tw)} \times n$.

Proof. The solution of LARGEST *st*-BOND can be found by a dynamic programming as presented in Theorem 11 where we add s and t in all the nodes and we fix $s \in S$ and $t \notin S$. To obtain similar algorithms for MAXIMUM CONNECTED CUT and MAXIMUM CONNECTED *st*-CUT we just do not have to take care with the connectivity information for S_2 and simply drop it in the above computation. \square

The dynamic programming algorithms in Theorems 11 and 12 can be seen as ones for *connectivity problems* such as finding a Hamiltonian cycle, a feedback vertex set, and a Steiner tree. For such problems, we can improve the running time $2^{\mathcal{O}(tw \log tw)}$ to $2^{\mathcal{O}(tw)}$ using two techniques called the *rank-based approach* due to Bodlaender et al. [5] and the *cut \mathcal{E} count technique* due to Cygan et al. [19]. In the next two subsections, we improve the running time of the algorithms described in this section using these techniques.

3.3.1 Rank-based approach

In this subsection, we provide faster $2^{\mathcal{O}(tw)}$ -time deterministic algorithms parameterized by treewidth. To show this, we use the rank-based approach proposed by Bodlaender et al. [5]. The key idea of the rank-based approach is to keep track of *small* representative sets of size $2^{\mathcal{O}(tw)}$ that

capture partial solutions of an optimal solution instead of $2^{O(tw \log tw)}$ partitions. Indeed, we can compute small representative sets within the claimed running time using reduce algorithm [5].

We begin with some definition used in the Rank-based approach.

Definition 7 (Set of weighted partitions [5]). *Let $\Pi(U)$ be the set of all partitions of some set U . A set of weighted partitions is $\mathcal{A} \subseteq \Pi(U) \times \mathbb{N}$, i.e., a family of pairs, each consisting of a partition of U and a non-negative integer weight.*

The weight of a partition corresponds to the size of a partial solution. For $p, q \in \Pi(U)$, let $J(p, q)$ denote the join of the partition. We say that a set of weighted partitions $\mathcal{A}' \subseteq \Pi(U) \times \mathbb{N}$ represents another set $\mathcal{A} \subseteq \Pi(U) \times \mathbb{N}$, if for all $q \in \Pi(U)$ it holds that $\max\{w \mid (p, w) \in \mathcal{A}' \wedge J(p, q) = \{U\}\} = \max\{w \mid (p, w) \in \mathcal{A} \wedge J(p, q) = \{U\}\}$. Then Bodlaender et al. [5] provided reduce algorithm that computes a *small* representative set of weighted partitions.

Theorem 13 (reduce [5]). *There exists an algorithm reduce that given a set of weighted partitions $\mathcal{A} \subseteq \Pi(U) \times \mathbb{N}$, outputs in time $|\mathcal{A}|2^{(\omega-1)|U|}|U|^{O(1)}$ a set of weighted partitions $\mathcal{A}' \subseteq \mathcal{A}$ such that \mathcal{A}' represents \mathcal{A} and $|\mathcal{A}'| \leq 2^{|U|-1}$, where $\omega < 2.3727$ denotes the matrix multiplication exponent.*

The reduce algorithm allows us to compute an optimal solution without keeping all weighted partitions. We apply reduce algorithm to the set of partitions at each node in the $O^*(2^{O(tw \log tw)})$ -time algorithm of the previous section.

Theorem 14. *Given a tree decomposition of width tw , there are $O^*((1+2^{\omega+1})^{tw})$ -time deterministic algorithms for MAXIMUM CONNECTED CUT and MAXIMUM CONNECTED st -CUT.*

Proof. For a bag X_ℓ , we compute the value $c[\ell, S_i, \rho_1]$ for each $S_i \subseteq X_\ell$ and a partition ρ_1 of S_i . For each S_i and T_i , we apply the reduce algorithm to a set of weighted partitions $(\rho_1, c[\ell, S_i, \rho_1])$ that are obtained by recursive formulas as described in the previous section. At each node i , the reduce algorithm outputs only $2^{|S_i|-1}$ weighted partitions for each S_i . Thus, at each node except join nodes, the running time of evaluating the recursive formula is $\sum_{S_i \subseteq X_i} O^*(2^{|S_i|}) = O^*(3^{tw})$ and of the reduce algorithm is $\sum_{S_i \subseteq X_i} O^*(2^{|S_i|}2^{(\omega-1)|S_i|}) = \sum_{S_i \subseteq X_i} O^*(2^{\omega|S_i|}) = O^*((1+2^\omega)^{tw})$. At each join node, since the output of evaluating the recursive formula may contain $O^*(2^{2|S_i|})$ weighted partitions for each S_i . Thus, the total running time at join node i is

$$\sum_{S_i \subseteq X_i} O^*(2^{2|S_i|}2^{(\omega-1)|S_i|}) = O^*((1+2^{\omega+1})^{tw}).$$

Hence, the theorem follows. \square

Note that if a tree decomposition has no join nodes, namely a path decomposition, the overall running time is $O^*((1+2^\omega)^{pw})$.

Theorem 15. *Given a tree decomposition of width tw , there are $O^*(2^{(\omega+2)tw})$ -time deterministic algorithms for LARGEST BOND and LARGEST st -BOND.*

Proof. For a bag X_ℓ , we compute the value $c[\ell, S_i, \rho_1, \rho_2]$ for each $S_i \subseteq X_\ell$ and ρ_1, ρ_2 being partitions of S_i and $X_\ell \setminus S_i$, respectively. Similar to Theorem 14, for each S_i , we apply the reduce algorithm to a set of weighted partitions $(\rho_1, c[\ell, S_i, \rho_1, \rho_2])$ and then apply it again to weighted partitions $(\rho_2, c[\ell, S_i, T_i, \rho_1', \rho_2])$ for each S_i, T_i , and for each remaining ρ_1 of the first application. Since there are at most $2^{|S_i|-1}2^{|T_i|-1} = 2^{|X_i|-2}$ weighted partitions in the representative set for each $S_i \subseteq X_i$, the total running time is $\sum_{S_i \subseteq X_i} O^*(2^{2(|X_i|-2)}2^{(\omega-1)|X_i|}) = O^*(2^{(\omega+2)tw})$. \square

Corollary 4. LARGEST BOND, LARGEST st -BOND, MAXIMUM CONNECTED CUT and MAXIMUM CONNECTED st -CUT can be solved in $2^{O(k)} \times n^{O(1)}$ time.

Proof. It follows directly from Corollary 2, Corollary 3, Theorem 14 and Theorem 15. \square

3.3.2 Cut & Count

In this subsection, we design much faster randomized algorithms by using Cut & Count, which is the framework for solving the connectivity problems faster [19]. In Cut & Count, we count the number of *relaxed* solutions modulo 2 on a tree decomposition and determine whether there exists a connected solution by cancellation tricks.

Definition 8 ([19]). *A cut $[V_1, V_2]$ of $V' \subseteq V$ such that $V_1 \cup V_2 = V'$ and $V_1 \cap V_2 = \emptyset$ is consistent if $v_1 \in V_1$ and $v_2 \in V_2$ implies $(v_1, v_2) \notin E$.*

In other words, a cut $[V_1, V_2]$ of V' is consistent if there are no edge between V_1 and V_2 .

Fix an arbitrary vertex v in V_1 . If $G[V]$ has k components, then there exist 2^{k-1} consistent cuts of V . Thus, when $G[V]$ is connected, there only exists one consistent cut $[V_1, V_2] = [V, \emptyset]$. From this observation, $G[V]$ is connected if and only if the number of consistent cuts is odd. Therefore, in order to compute “connected solutions”, it seems to suffice to count the number of consistent cuts modulo two at first glance. However, this computation may fail to count the number of “connected solutions” since there can be even number of such solutions. To overcome this obstacle, Cygan et al. [19] used the Isolation Lemma [48], which ensures with high probability that the problem has a unique minimum solution. For the detail of the Isolation Lemma, see [18, 48].

We follow the Cut & Count framework in [18, 19]: We apply it to determining whether there exists a minimal st -cut, a cut that separates s and t , of size k , namely st -bond. Recall that $[S, V \setminus S]$ is a st -bond of a connected graph $G = (V, E)$ if both $G[S]$ and $G[V \setminus S]$ are connected, $s \in S$, and $t \in V \setminus S$.

Let i be the index of a node of a nice tree decomposition of T of G .

Definition 9. *Let $r \in \{1, 2, \dots, |E|\}$. Let $S_i^l, S_i^r, T_i^l, T_i^r$ be pairwise disjoint (possibly) subsets of X_i such that $S_i^l \cup S_i^r \cup T_i^l \cup T_i^r = X_i$. A partial solution for $(S_i^l, S_i^r, T_i^l, T_i^r, r)$ is a cut $[S, V_i \setminus S]$ of G_i such that:*

- $S \cap X_i = S_i^l \cup S_i^r$ and $(V_i \setminus S) \cap X_i = T_i^l \cup T_i^r$,
- $[S_i^l, S_i^r]$ and $[T_i^l, T_i^r]$ are consistent cuts of $S_i^l \cup S_i^r$ and $T_i^l \cup T_i^r$, respectively,
- there are exactly r cut edges between S and $V_i \setminus S$ in G_i , and
- $s \in V_i \implies s \in S_i^l$ and $t \in V_i \implies t \in T_i^l$.

Before proceeding to our dynamic programming, we assign a weight w_v to each vertex $v \in V$ by choosing an integer from $\{1, \dots, 2n\}$ independently and uniformly at random. We also use the following preprocessing: add s and t to each node of T and remove the bags introduce s or t from T . In our dynamic programming algorithm, for each node i and for $0 \leq w \leq 2n^2$ and $0 \leq r \leq |E|$, we count the number of partial solutions $[S, V_i \setminus S]$ for $(S_i^l, S_i^r, T_i^l, T_i^r, r)$ such that the total weight of S is exactly w , which we denote by $c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w]$. By the Isolation Lemma, with high probability, there is a minimal st -cut of G of size exactly k if and only if $c(i, \{s\}, \emptyset, \{t\}, \emptyset, k, w)$ is odd for some $0 \leq w \leq 2n^2$ in the root node $r(T) = i$. In the following, we describe the recursive formula for our dynamic programming.

Leaf node: In a leaf node i , since $X_i = \{s, t\}$, we have $c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w] = 1$ if $S_i^l = \{s\}$, $T_i^l = \{t\}$, $S_i^r = T_i^r = \emptyset$, $r = 0$, and $w = 0$. Otherwise, $c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w] = 0$.

Introduce vertex v node: In an introduce vertex node i , we consider the following four cases:

$$c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w] = \begin{cases} c[j, S_i^l \setminus \{v\}, S_i^r, T_i^l, T_i^r, r, w - w(v)] & \text{if } v \in S_i^l, \\ c[j, S_i^l, S_i^r \setminus \{v\}, T_i^l, T_i^r, r, w - w(v)] & \text{if } v \in S_i^r, \\ c[j, S_i^l, S_i^r, T_i^l \setminus \{v\}, T_i^r, r, w] & \text{if } v \in T_i^l, \\ c[j, S_i^l, S_i^r, T_i^l, T_i^r \setminus \{v\}, r, w] & \text{if } v \in T_i^r. \end{cases}$$

As $v \in X_i$, exactly one of the above cases is applied.

Introduce edge (u, v) node: Let i be an introduce node of T . Let $S_i^l, S_i^r, T_i^l, T_i^r$ be disjoint subsets of X_i whose union covers X_i . If exactly one of u and v belongs to $S_i^l \cup S_i^r$ (i.e. the other one belongs to $T_i^l \cup T_i^r$), the edge is included in the cutset. Suppose otherwise, that is, either $u, v \in S_i^l \cup S_i^r$ or $u, v \in T_i^l \cup T_i^r$. If u and v belong to different sets, say $u \in S_i^l$ and $v \in S_i^r$, then $[S_i^l, S_i^r]$ is not consistent. Therefore, there is no partial solutions in this case. To summarize these facts, we have the following:

$$c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w] = \begin{cases} c[j, S_i^l, S_i^r, T_i^l, T_i^r, r-1, w] & \text{if } |(S_i^l \cup S_i^r) \cap \{u, v\}| = 1, \\ c[j, S_i^l, S_i^r, T_i^l, T_i^r, r, w] & \text{if } u, v \text{ are in the same set,} \\ 0 & \text{otherwise.} \end{cases}$$

Forget v node: In a forget node i , we just sum up the number of partial solutions:

$$c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w] = c[j, S_i^l \cup \{v\}, S_i^r, T_i^l, T_i^r, r, w] + c[j, S_i^l, S_i^r \cup \{v\}, T_i^l, T_i^r, r, w] \\ + c[j, S_i^l, S_i^r, T_i^l \cup \{v\}, T_i^r, r, w] + c[j, S_i^l, S_i^r, T_i^l, T_i^r \cup \{v\}, r, w].$$

Join node: Let i be a join node and j_1 and j_2 its children. As $X_i = X_{j_1} = X_{j_2}$, it should hold that $S_i^l = S_{j_1}^l = S_{j_2}^l$, $S_i^r = S_{j_1}^r = S_{j_2}^r$, $T_i^l = T_{j_1}^l = T_{j_2}^l$, and $T_i^r = T_{j_1}^r = T_{j_2}^r$.

The size of a partial solution S_i at i is the sum of the size of partial solutions S_{j_1} and S_{j_2} at its children, minus the number of edges from $S_i \cap X_i$ to $X_i \setminus S_i$, since such edges are in both S_{j_1} and S_{j_2} . Also, the total weight of S_i is the the sum of the weight of S_{j_1} and S_{j_2} minus the the total weight of $S_i \cap X_i$.

Thus, we have

$$c[i, S_i^l, S_i^r, T_i^l, T_i^r, r, w]$$

equals to:

$$\sum_{r_{j_1} + r_{j_2} - \alpha = r} \sum_{w_{j_1} + w_{j_2} - \beta = w} c[j_1, S_i^l, S_i^r, T_i^l, T_i^r, r_{j_1}, w_{j_1}] c[j_2, S_i^l, S_i^r, T_i^l, T_i^r, r_{j_2}, w_{j_2}].$$

where α is the set of edges from $S_i^l \cup S_i^r$ to $T_i^l \cup T_i^r$, and β is the sum of the weights of vertices in $S_i^l \cup S_i^r$.

The running time of evaluating the recursive formulas is $O^*(4^{|X_i|})$ for each node i . Therefore, the total running is $O^*(4^{tw})$. We can also solve LARGEST BOND in time $O^*(4^{tw})$ by applying the algorithm for LARGEST *st*-BOND for all combinations of s and t .

Theorem 16. *Given a tree decomposition of width tw , there is a Monte-Carlo algorithm that solves LARGEST BOND and LARGEST *st*-BOND in time $O^*(4^{tw})$. It cannot give false positives and may give false negatives with probability at most $1/2$.*

We can also solve MAXIMUM CONNECTED CUT and MAXIMUM CONNECTED *st*-CUT. Since it suffices to keep track of consistent cuts of S , the running time is $O^*(3^{tw})$.

Theorem 17. *Given a tree decomposition of width tw , there is a Monte-Carlo algorithm that solves MAXIMUM CONNECTED CUT and MAXIMUM CONNECTED *st*-CUT in time $O^*(3^{tw})$. It cannot give false positives and may give false negatives with probability at most $1/2$.*

By combining Corollary 2 and Theorems 16 and 17, we have the following theorem.

Theorem 18. *There are Monte-Carlo algorithms that solve LARGEST BOND and MAXIMUM CONNECTED CUT in time $O^*(16^k)$ and $O^*(3^k)$, respectively. It cannot give false positives and may give false negatives with probability at most $1/2$.*

3.4 Twin-cover

Two vertices u, v are called *twins* if both u and v have the same closed/open neighbourhood. Moreover, if twins u, v have edge $\{u, v\}$, they are called *true twins* and the edge is called a *twin edge*. Then a *twin-cover* of G is defined as follows.

Definition 10 ([31]). *A set of vertices X is a twin-cover of G if every edge $\{u, v\} \in E$ satisfies either*

- $u \in X$ or $v \in X$, or
- u, v are true twins.

The twin-cover number of G , denoted by $tc(G)$, is defined as the size of minimum twin-cover in G .

An important observation is that the complement of a twin-cover X induces disjoint cliques. Moreover, for each clique Z of $G[V \setminus X]$, $N(u) \cap X = N(v) \cap X$ for every $u, v \in Z$ [31].

MAXIMUM CUT is FPT when parameterized by twin-cover number [31]. In this section, we show that MAXIMUM CONNECTED CUT and LARGEST BOND are also FPT when parameterized by the twin-cover number.

Theorem 19. MAXIMUM CONNECTED CUT can be solved in time $O^*(2^{2^{tc}+tc})$.

Proof. First compute a minimum twin-cover X of G in time $O^*(1.2738^{tc})$ [31]. Now, we have a twin-cover X of size tc . Recall that $G[V \setminus X]$ consists of vertex disjoint cliques and for each $u, v \in Z$ in a clique Z of $G[V \setminus X]$, $N(u) \cap X = N(v) \cap X$.

We iterate over all possible subsets X' of X and compute the size of a maximum cut $[S, V \setminus S]$ of G with $S \cap X = X'$.

If $X' = \emptyset$, exactly one of the cliques of $G[V \setminus X]$ intersects S as $G[S]$ is connected. Thus, we can compute a maximum cut by finding the best partition for each clique of $G[V \setminus X]$, which can be done in polynomial time.

Suppose otherwise that $X' \neq \emptyset$. We define a *type* of each clique Z of $G[V \setminus X]$. The type of Z , denoted by $T(Z)$, is $N(Z) \cap X$. Note that there are at most $2^{tc} - 1$ types of cliques in $G[V \setminus X]$.

For each type of cliques, we guess that S has an intersection with this type of cliques. There are at most $2^{2^{tc}-1}$ possible combinations of types of cliques. Let \mathcal{T} be the set of types in $G[V \setminus X]$. For each guess $\mathcal{T}' \subseteq \mathcal{T}$, we try to find a maximum cut $[S, V \setminus S]$ such that $G[S]$ is connected, $S \cap X = X'$, for each $T \in \mathcal{T}'$, at least one of the cliques of type T has an intersection with S , and for each $T \notin \mathcal{T}'$, every clique of type T has no intersection with S . We can easily check if $G[S]$ will be connected as S contains a vertex of a clique of type $T \in \mathcal{T}'$. Consider a clique Z of type $T(Z) = X'' \subseteq X$. Since every vertex in Z has the same neighborhood in X , we can determine the number of cut edges incident to Z from the cardinality of $S \cap Z$. More specifically, if $|S \cap Z| = p$, the number of cut edges incident to Z is equal to $p(|Z| - p) + p|X'' \cap (X \setminus X')| + (|Z| - p)|X'' \cap X'|$. Moreover, we can independently maximize the number of cut edges incident to Z for each clique Z of $G[V \setminus X]$.

Overall, for each $X' \subseteq X$ and for each set of types \mathcal{T}' , we can compute a maximum connected cut with respect to X' and \mathcal{T}' in polynomial time. Therefore, the total running time is bounded by $O^*(2^{2^{tc}+tc})$. \square

Theorem 20. LARGEST BOND can be solved in time $O^*(2^{tc}3^{2^{tc}})$.

Proof. We design an $O^*(2^{tc}3^{2^{tc}})$ -time algorithm for LARGEST BOND, where tc is the size of a minimum twin-cover of $G = (V, E)$. This is quite similar to the one for MAXIMUM CONNECTED CUT developed in this section. As with an algorithm for MAXIMUM CONNECTED CUT, we first compute a minimum twin-cover X in time $O^*(1.2738^{tc})$ [31]. Then we guess all 2^{tc} possible subsets $X' \subseteq X$ and compute the size of maximum cut $(S, V \setminus S)$ of G with $S \cap X = X'$.

If $X' = \emptyset$, exactly one of the cliques of $G[V \setminus X]$ intersects S due to the connectivity of $G[S]$. Thus, we can compute a maximum cut in polynomial time. Note that $G[V \setminus S]$ is also connected because $X \subseteq V \setminus S$. We are also done for the case where $X' = X$ by a symmetric argument. Thus, in the following, we assume that our guess X' is non-empty and a proper subset of X .

For each guess $X' \subseteq X$, we further guess each type of cliques in $G[V \setminus X]$ has an intersection with only S , with only $V \setminus S$, or with both S and $V \setminus S$. For each guess, we can easily check S and $V \setminus S$ will be connected and maximize the size of a cut in polynomial time as in Theorem 19. Since there are at most 2^{tc} types of cliques in $G[V \setminus X]$, the total running time is $O^*(2^{tc}3^{2^{tc}})$. \square

4 Infeasibility of polynomial kernels for solution size parameterization

It is not hard to see that MAXIMUM CONNECTED CUT do not admit a polynomial kernel unless $\text{NP} \subseteq \text{coNP}/\text{poly}$, since it is trivially or-compositional; at least one of graphs G_1, G_2, \dots, G_t have a connected cut of size at least k if and only if their disjoint union $G_1 \cup G_2 \cup \dots \cup G_t$ also has a connected cut of size at least k .

Regarding to bonds, as seen previously, any bond $\partial(S)$ of a graph G intersects at most one of its block. Thus, an or-composition for LARGEST BOND parameterized by k can be done from the disjoint union of ℓ inputs, by selecting exactly one vertex of each input graph and contracting them into a single vertex. Now, let $(G_1, k, s_1, t_1), (G_2, k, s_2, t_2), \dots, (G_\ell, k, s_\ell, t_\ell)$ be ℓ instances of LARGEST st -BOND parameterized by k . An or-composition for LARGEST st -BOND parameterized by k can be done from the disjoint union of G_1, G_2, \dots, G_ℓ , by contracting t_i, s_{i+1} into a single vertex, $1 \leq i \leq \ell - 1$, and setting $s = s_1$ and $t = t_\ell$.

Therefore, the following holds.

Theorem 21. LARGEST BOND, LARGEST st -BOND and MAXIMUM CONNECTED CUT do not admit polynomial kernel, unless $\text{NP} \subseteq \text{coNP}/\text{poly}$.

5 Conclusions

In this work, we present a multivariate analysis on the complexity of computing the largest bond and the maximum connected cut of a graph. Some of our contributions is summarized in Table 1.

Also, we present general reductions that allows us to observe that LARGEST BOND and MAXIMUM CONNECTED CUT are NP-hard for several graph classes for which MAXIMUM CUT is NP-hard. Using this frameworks, we are able to show that LARGEST BOND and MAXIMUM CONNECTED CUT on graphs of clique-width w cannot be solved in time $f(w) \times n^{o(w)}$ unless the ETH fails. Moreover, we show that LARGEST BOND does not admit a constant-factor approximation algorithm, unless $\text{P} = \text{NP}$, and thus is asymptotically harder to approximate than MAXIMUM CUT.

Table 1: The summary of the computational complexity of MAXIMUM CUT and its variants.

	Graph Class		Parameter				poly kernel
	Split	Planar Bipartite	cw	tw	tc	k	k
MAXIMUM CUT	NP-c [6]	P [trivial]	$n^{O(cw)}$ [29]	2^{tw} [6]	2^{tc} [31]	1.2418^k [52]	$O(k)$ [40, 46]
CONNECTED CUT	NP-c [Th. 6]	NP-c [Th. 5]	$n^{O(cw)}$ [Th. 10]	3^{tw} [Th. 17]	$2^{2^{tc}+tc}$ [Th. 19]	$2^{O(k)}$ [Cor. 4]	No [Th. 21]
LARGEST BOND	NP-c [Th. 7]	NP-c [Th. 3]	$n^{O(cw)}$ [Th. 10]	4^{tw} [Th. 16]	$2^{tc}3^{2^{tc}}$ [Th. 20]	$2^{O(k)}$ [Cor. 4]	No [Th. 21]

References

- [1] Aldred, R.E.L., V. Dyck, D., Brinkmann, G., Fack, V., McKay, B.D.: Graph structural properties of non-Yutsis graphs allowing fast recognition. *Discrete Applied Mathematics* **157**(2), 377–386 (2009)
- [2] Bazgan, C., Brankovic, L., Casel, K., Fernau, H., Jansen, K., Klein, K.M., Lampis, M., Liedloff, M., Monnot, J., Paschos, V.T.: The many facets of upper domination. *Theoretical Computer Science* **717**, 2–25 (2018)
- [3] de Berg, M., Khosravi, A.: Finding perfect auto-partitions is NP-hard. In: *EuroCG 2009*, pp. 255–258. (2009)
- [4] Biedenharn, L.C., Louck, J.D.: *The Racah-Wigner algebra in quantum theory*. Addison-Wesley (1981)
- [5] Bodlaender, H.L., Cygan, M., Kratsch, S., Nederlof, J.: Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth. *Information and Computation* **243**, 86–111 (2015)
- [6] Bodlaender, H.L., Jansen, K.: On the complexity of the maximum cut problem. *Nordic Journal of Computing* **7**(1), 14–31 (2000)
- [7] Bodlaender, H.L., Jansen, K.: On the complexity of the maximum cut problem. *Nordic Journal of Computing* **7**(1), 14–31 (2000)
- [8] Bodlaender, H.L., Van Leeuwen, J., Tan, R., Thilikos, D.M.: On interval routing schemes and treewidth. *Information and Computation* **139**(1), 92–109 (1997)
- [9] Boria, N., Croce, F.D., Paschos, V.T.: On the max min vertex cover problem. *Discrete Applied Mathematics* **196**, 62–71 (2015)
- [10] Bui-Xuan, B.M., Suchý, O., Telle, J.A., Vatshelle, M.: Feedback vertex set on graphs of low clique-width. *European Journal of Combinatorics* **34**(3), 666–679 (2013)
- [11] Carvajal, R., Constantino, M., Goycoolea, M., Vielma, J.P., Weintraub, A.: Imposing connectivity constraints in forest planning models. *Operations Research* **61**(4), 824–836 (2013)
- [12] Chaourar, B.: A linear time algorithm for a variant of the MAX CUT problem in series parallel graphs. *Advances in Operations Research* pp. 1267108:1–1267108:4 (2017)
- [13] Chaourar, B.: Connected max cut is polynomial for graphs without $K_5 \setminus e$ as a minor. *CoRR abs/1903.12641* (2019)
- [14] Cho, J.J., Chen, Y., Ding, Y.: On the (co)girth of a connected matroid. *Discrete Applied Mathematics* **155**(18), 2456 – 2470 (2007)
- [15] Cormen, T., Leiserson, C., Rivest, R., Stein, C.: *Introduction To Algorithms*. MIT Press (2001). URL https://books.google.co.in/books?id=NLngYyWF1_YC
- [16] Courcelle, B., Makowsky, J.A., Rotics, U.: Linear time solvable optimization problems on graphs of bounded clique-width. *Theory of Computing Systems* **33**(2), 125–150 (2000)
- [17] Courcelle, B., Olariu, S.: Upper bounds to the clique width of graphs. *Discrete Applied Mathematics* **101**(1-3), 77–114 (2000)
- [18] Cygan, M., Fomin, F.V., Kowalik, Ł., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: *Parameterized algorithms*, vol. 4. Springer (2015)

- [19] Cygan, M., Nederlof, J., Pilipczuk, M., Pilipczuk, M., Rooij, J.M.M.v., Wojtaszczyk, J.O.: Solving connectivity problems parameterized by treewidth in single exponential time. In: FOCS 2011, pp. 150–159 (2011)
- [20] Demange, M.: A note on the approximation of a minimum-weight maximal independent set. *Computational Optimization and Applications* **14**(1), 157–169 (1999)
- [21] Ding, G., Dziobiak, S., Wu, H.: Large-or-minors in 3-connected graphs. *Journal of Graph Theory* **82**(2), 207–217 (2016)
- [22] Duarte, G.L., Lokshtanov, D., Pedrosa, L.L.C., Schouery, R.C.S., Souza, U.S.: Computing the Largest Bond of a Graph. In: IPEC 2019, vol. 148, pp. 12:1–12:15 (2019)
- [23] Dyck, D.V., Fack, V.: On the reduction of Yutsis graphs. *Discrete Mathematics* **307**(11), 1506 – 1515 (2007). The Fourth Caracow Conference on Graph Theory
- [24] Díaz, J., Kamiński, M.: Max-cut and max-bisection are NP-hard on unit disk graphs. *Theoretical Computer Science* **377**(1), 271–276 (2007)
- [25] Edmonds, J., Karp, R.M.: Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM* **19**(2), 248–264 (1972)
- [26] Erdős, P.: On some extremal problems in graph theory. *Israel Journal of Mathematics* **3**(2), 113–116 (1965)
- [27] Eto, H., Hanaka, T., Kobayashi, Y., Kobayashi, Y.: Parameterized Algorithms for Maximum Cut with Connectivity Constraints. In: IPEC 2019, vol. 148, pp. 13:1–13:15 (2019)
- [28] Flynn, M.: The largest bond in 3-connected graphs. Ph.D. thesis, The University of Mississippi (2017)
- [29] Fomin, F.V., Golovach, P.A., Lokshtanov, D., Saurabh, S.: Almost optimal lower bounds for problems parameterized by clique-width. *SIAM Journal on Computing* **43**(5), 1541–1563 (2014)
- [30] Gandhi, R., Hajiaghayi, M.T., Kortsarz, G., Purohit, M., Sarpatwar, K.: On maximum leaf trees and connections to connected maximum cut problems. *Information Processing Letters* **129**, 31 – 34 (2018)
- [31] Ganian, R.: Improving Vertex Cover as a Graph Parameter. *Discrete Mathematics and Theoretical Computer Science* **17**(2), 77–100 (2015)
- [32] Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman (1979)
- [33] Garey, M.R., Johnson, D.S., Tarjan, R.E.: The planar hamiltonian circuit problem is NP-complete. *SIAM Journal on Computing* **5**(4), 704–714 (1976)
- [34] Godsil, C.D., Royle, G.F.: *Algebraic Graph Theory*. Graduate texts in mathematics. Springer (2001)
- [35] Goemans, M.X., Williamson, D.P.: Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM* **42**(6), 1115–1145 (1995)
- [36] Gomory, R.E., Hu, T.C.: Multi-terminal network flows. *Journal of the Society for Industrial and Applied Mathematics* **9**(4), 551–570 (1961)

- [37] Grimm, V., Kleinert, T., Liers, F., Schmidt, M., Zöttl, G.: Optimal price zones of electricity markets: a mixed-integer multilevel model and global solution approaches. *Optimization Methods and Software* **34**(2), 406–436 (2019)
- [38] Guruswami, V.: Maximum cut on line and total graphs. *Discrete Applied Mathematics* **92**(2), 217–221 (1999)
- [39] Hadlock, F.: Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing* **4**(3), 221–225 (1975)
- [40] Haglin, D.J., Venkatesan, S.M.: Approximation and intractability results for the maximum cut problem and its variants. *IEEE Transactions on Computers* **40**(1), 110–113 (1991)
- [41] Hajiaghayi, M.T., Kortsarz, G., MacDavid, R., Purohit, M., Sarpatwar, K.: Approximation algorithms for connected maximum cut and related problems. In: *ESA 2015*, pp. 693–704 (2015)
- [42] Hanaka, T., Bodlaender, H.L., van der Zanden, T.C., Ono, H.: On the maximum weight minimal separator. In: *TAMC 2017*, pp. 304–318 (2017)
- [43] Karp, R.M.: *Reducibility among Combinatorial Problems*, pp. 85–103. Springer US, Boston, MA (1972)
- [44] Khoshkhan, K., Ghadikolaie, M.K., Monnot, J., Sikora, F.: Weighted upper edge cover: Complexity and approximability. In: *WALCOM 2019*, pp. 235–247 (2019)
- [45] Khot, S., Kindler, G., Mossel, E., O’Donnell, R.: Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing* **37**(1), 319–357 (2007)
- [46] Mahajan, M., Raman, V.: Parameterizing above guaranteed values: MaxSat and MaxCut. *Journal of Algorithms* **31**(2), 335–354 (1999)
- [47] Mitzenmacher, M., Upfal, E.: *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press (2005)
- [48] Mulmuley, K., Vazirani, U.V., Vazirani, V.: Matching is as easy as matrix inversion. *Combinatorica* **7**(1), 105–113 (1987)
- [49] Odlyzko, A.M.: Asymptotic enumeration methods. *Handbook of combinatorics* **2**(1063), 1229 (1995)
- [50] Orlova, G.I., Dorfman, Y.G.: Finding the maximal cut in a graph. *Engineering Cybernetics* **10**(3), 502–506 (1972)
- [51] Oxley, J.G.: *Matroid theory*, vol. 3. Oxford University Press, USA (2006)
- [52] Raman, V., Saurabh, S.: Improved fixed parameter tractable algorithms for two “edge” problems: MAXCUT and MAXDAG. *Information Processing Letters* **104**(2), 65–72 (2007)
- [53] Rao, M.: Clique-width of graphs defined by one-vertex extensions. *Discrete Mathematics* **308**(24), 6157–6165 (2008)
- [54] Robertson, N., Seymour, P.D.: Graph minors. V. excluding a planar graph. *Journal of Combinatorial Theory, Series B* **41**(1), 92–114 (1986)
- [55] Saurabh, S., Zehavi, M.: Parameterized Complexity of Multi-Node Hubs. In: *IPEC 2018*, pp. 8:1–8:14 (2019)

- [56] Shen, X., Lee, J., Nagarajan, V.: Approximating graph-constrained max-cut. *Mathematical Programming* **172**(1), 35–58 (2018)
- [57] Vicente, S., Kolmogorov, V., Rother, C.: Graph cut based image segmentation with connectivity priors. In: *CVPR 2008*, pp. 1–8 (2008)
- [58] Yannakakis, M., Gavril, F.: Edge dominating sets in graphs. *SIAM Journal on Applied Mathematics* **38**(3), 364–372 (1980)
- [59] Yutsis, A.P., Vanagas, V.V., Levinson, I.B.: *Mathematical apparatus of the theory of angular momentum*. Israel program for scientific translations (1962)
- [60] Zehavi, M.: Maximum minimal vertex cover parameterized by vertex cover. *SIAM Journal on Discrete Mathematics* **31**(4), 2440–2456 (2017)

Appendix

Faster algorithms parameterized by clique-width (Theorem 10)

In this section, we design faster XP algorithms for both MAXIMUM CONNECTED CUT and LARGEST BOND when parameterized by clique-width, which run in time $n^{O(w)}$.

Here, we rather use a different graph parameter and its associated decomposition closely related to clique-width. We believe that this decomposition is more suitable to describe our dynamic programming.

Definition 11. Let $X \subseteq V(G)$. We say that $M \subseteq X$ is a twin-set of X if for any $v \in V(G) \setminus X$, either $M \subseteq N(v)$ or $M \cap N(v) = \emptyset$ holds. A twin-set M is called a twin-class of X if it is maximal subject to being a twin-set of X . X can be partitioned into twin-classes of X .

Definition 12. Let w be an integer. We say that $X \subseteq V(G)$ is a w -module of G if X can be partitioned into w twin-classes $\{X_1, X_2, \dots, X_w\}$. A decomposition tree of G is a pair of a rooted binary tree T and a bijection ϕ from the set of leaves of T to $V(G)$. For each node v of T , we denote by L_v the set of leaves, each of which is either v or a descendant of v . The width of a decomposition tree (T, ϕ) of G is the minimum w such that for every node v in T , the set $\bigcup_{l \in L_v} \phi(l)$ is a w_v -module of G with $w_v \leq w$. The module-width of G is the minimum t such that there is a decomposition tree of G of width w .

Rao [53] proved that clique-width and module-width are linearly related to each other. Let cw and mw be the clique-width and the module-width of G , respectively. We note that a similar terminology “modular-width” has been used in the literature, but module-width used in this paper is different from it.

Theorem 22 ([53]). For every graph G , $mw \leq cw \leq 2mw$.

Moreover, given a w -expression tree of G , we can in time $O(n^2)$ compute a decomposition tree (T, ϕ) of G of width at most w and $w_v \leq w$ twin-classes of $\bigcup_{l \in L_v} \phi(l)$ for each node v in T [10].

Fix a decomposition tree (T, f) of G whose width is w . Our dynamic programming algorithm runs over the nodes of the decomposition tree in a bottom-up manner. For each node v in T , we let $\{X_1^v, X_2^v, \dots, X_{w_v}^v\}$ be the twin-classes of $\bigcup_{l \in L_v} \phi(l)$. From now on, we abuse the notation to denote $\bigcup_{l \in L_v} \phi(l)$ simply by L_v . A tuple of $4w_v$ integers $t = (p_1, \bar{p}_1, p_2, \bar{p}_2, \dots, p_{w_v}, \bar{p}_{w_v}, c_1, \bar{c}_1, c_2, \bar{c}_2, \dots, c_{w_v}, \bar{c}_{w_v})$ is valid for v if it holds that $0 \leq p_i, \bar{p}_i \leq |X_i^v|$ with $p_i + \bar{p}_i = |X_i^v|$ and $c_i, \bar{c}_i \in \{0, 1\}$ for each $1 \leq i \leq w_v$. For a valid tuple t for v , we say that a cut $(S, L_v \setminus S)$ of $G[L_v]$ is t -legitimate if for each $1 \leq i \leq w_v$, it satisfies the following conditions:

- $p_i = |S \cap X_i^v|$,
- $\bar{p}_i = |(L_v \setminus S) \cap X_i^v|$,
- $G[S \cap X_i^v]$ is connected if $c_i = 1$, and
- $G[(L_v \setminus S) \cap X_i^v]$ is connected if $\bar{c}_i = 1$.

The size of a t -legitimate cut is defined accordingly. In this section, we allow each side of a cut to be empty and the empty graph is considered to be connected. Our algorithm computes the value $mc(v, t)$ that is the maximum size of a t -legitimate cut for each valid tuple t and each node v in the decomposition tree.

Leaves (Base step): For each valid tuple t for a leaf v , $mc(v, t) = 0$. Note that there is only one twin-class $X_1^v = \{v\}$ for v in this case.

Internal nodes (Induction step): Let v be an internal node of T and let a and b be the children of v in T . Consider twin-classes $\mathcal{X}^v = \{X_1^v, X_2^v, \dots, X_{w_v}^v\}$, $\mathcal{X}^a = \{X_1^a, X_2^a, \dots, X_{w_a}^a\}$, and $\mathcal{X}^b = \{X_1^b, X_2^b, \dots, X_{w_b}^b\}$ of L_v , L_a , and L_b , respectively. Note that $\mathcal{X}^a \cup \mathcal{X}^b$ is a partition of L_v .

Observation .1. \mathcal{X}^v is a partition of L_v coarser than $\mathcal{X}^a \cup \mathcal{X}^b$.

To see this, consider an arbitrary twin-class X_i^a of L_a . By the definition of twin-sets, for every $z \in V(G) \setminus L_a$, either $X_i^a \subseteq N(z)$ or $X_i^a \cap N(z) = \emptyset$ holds. Since $V(G) \setminus L_v \subseteq V(G) \setminus L_a$, X_i^a is also a twin-set of L_v , which implies X_i^a is included in some twin-class X_j^v of L_v . This argument indeed holds for twin-classes of L_b . Therefore, we have the above observation.

The intuition of our recurrence is as follows. By Observation .1, every twin-class of L_v can be obtained by merging some twin-classes of L_a and of L_b . This means that every t_v -legitimate cut of $G[L_v]$ for a valid tuple t_v for v can be obtained from some t_a -legitimate cut and t_b -legitimate cut for valid tuples for a and b , respectively. Moreover, for every pair of twin-classes X_i^a of L_a and X_j^b of L_b , either there are no edges between them or every vertex in X_i^a is adjacent to every vertex in X_j^b as X_i^a is a twin-set of L_v . Therefore, the number of edges in the cutset of a cut $(S, L_v \setminus S)$ between X_i^a and X_j^b depends only on the cardinality of $X_i^a \cap S$ and $X_j^b \cap S$ rather than actual cuts $(S \cap X_i^a, (L_a \setminus S) \cap X_i^a)$ and $(S \cap X_j^b, (L_b \setminus S) \cap X_j^b)$.

Now, we formally describe this idea. Let X^v be a twin-class of L_v . We denote by $I_a(X^v)$ (resp. $I_b(X^v)$) the set of indices i such that X_i^a (resp. X_i^b) is included in X^v and by $\mathcal{X}^a(X^v)$ (resp. $\mathcal{X}^b(X^v)$) the set $\{X_i^a : i \in I_a(X^v)\}$ (resp. $\{X_i^b : i \in I_b(X^v)\}$). For $X^a \in \mathcal{X}^a(X^v)$ and $X^b \in \mathcal{X}^b(X^v)$, we say that X^a is adjacent to X^b if every vertex in X^a is adjacent to every vertex in X^b and otherwise X^a is not adjacent to X^b . This adjacency relation naturally defines a bipartite graph whose vertex set is $\mathcal{X}^a(X^v) \cup \mathcal{X}^b(X^v)$. We say that a subset of twin-classes of $\mathcal{X}^a(X^v) \cup \mathcal{X}^b(X^v)$ is *non-trivially connected* if it induces a connected bipartite graph with at least twin-classes. Let $S \subseteq X^v$. To make $G[S]$ (and $G[X^v \setminus S]$) connected, the following observation is useful.

Observation .2. Suppose $S \subseteq X^v$ has a non-empty intersection with at least two twin-classes of $\mathcal{X}^a(X^v) \cup \mathcal{X}^b(X^v)$. Then, $G[S]$ is connected if and only if the twin-classes having a non-empty intersection with S are non-trivially connected.

This observation immediately follows from the fact that every vertex in a twin-class is adjacent to every vertex in an adjacent twin-class and is not adjacent to every vertex in a non-adjacent twin-class.

Let $t_v = (p_1^v, \bar{p}_1^v, \dots, p_{w_v}^v, \bar{p}_{w_v}^v, c_1^v, \bar{c}_1^v, \dots, c_{w_v}^v, \bar{c}_{w_v}^v)$ be a valid tuple for v . For notational convenience, we use \mathbf{p}^v to denote $(p_1^v, \bar{p}_1^v, \dots, p_{w_v}^v, \bar{p}_{w_v}^v)$ and \mathbf{c}^v to denote $(c_1^v, \bar{c}_1^v, \dots, c_{w_v}^v, \bar{c}_{w_v}^v)$ for each node v in T . For valid tuples $t_a = (\mathbf{p}^a, \mathbf{c}^a)$ for a and $t_b = (\mathbf{p}^b, \mathbf{c}^b)$ for b , we say that t_v is *consistent with the pair* (t_a, t_b) if for each $1 \leq i \leq w_v$,

$$\text{C1 } p_i^v = \sum_{j \in I_a(X_i^v)} p_j^a + \sum_{j \in I_b(X_i^v)} p_j^b;$$

$$\text{C2 } \bar{p}_i^v = \sum_{j \in I_a(X_i^v)} \bar{p}_j^a + \sum_{j \in I_b(X_i^v)} \bar{p}_j^b;$$

\text{C3} if $c_i^v = 1$, either (1) $\{X_j^a : j \in I_a(X^v), p_j^a > 0\} \cup \{X_j^b : j \in I_b(X^v), p_j^b > 0\}$ is non-trivially connected or (2) exactly one of $\{p_j^s : s \in \{a, b\}, 1 \leq j \leq w_s\}$ is positive, say p_j^s , and $c_j^s = 1$;

\text{C4} if $\bar{c}_i^v = 1$, either (1) $\{X_j^a : j \in I_a(X^v), \bar{p}_j^a > 0\} \cup \{X_j^b : j \in I_b(X^v), \bar{p}_j^b > 0\}$ is non-trivially connected or (2) exactly one of $\{\bar{p}_j^s : s \in \{a, b\}, 1 \leq j \leq w_s\}$ is positive, say \bar{p}_j^s , and $\bar{c}_j^s = 1$.

Lemma 13.

$$\text{mc}(v, t_v) = \max_{t_a, t_b} \left(\text{mc}(a, t_a) + \text{mc}(b, t_b) + \sum_{\substack{X_i^a \in \mathcal{X}^a, X_j^b \in \mathcal{X}^b \\ X_i^a, X_j^b: \text{adjacent}}} (p_i^a \bar{p}_j^b + p_j^b \bar{p}_i^a) \right),$$

where the maximum is taken over all consistent pairs (t_a, t_b) .

Proof. We first show that the left-hand side is at most the right-hand side. Suppose $(S, L_v \setminus S)$ be a t_v -legitimate cut of $G[L_v]$ whose size is equal to $\text{mc}(v, t_v)$. Let $S_a = S \cap L_a$ and $S_b = S \cap L_b$. We claim that $(S_a, L_a \setminus S_a)$ is a t_a -legitimate cut of $G[L_a]$ for some valid tuple t_a for a . This is obvious since we set $p_i^a = |S_a \cap X_i^a|$, $\bar{p}_i^a = |(L_a \setminus S_a) \cap X_i^a|$, $c_i^a = 1$ if $G[S_a \cap X_i^a]$ is connected, and $\bar{c}_i^a = 1$ if $G[(L_a \setminus S_a) \cap X_i^a]$ is connected, which yields a valid tuple t_a for a . We also conclude that $(S_b, L_b \setminus S_b)$ is a t_b -legitimate cut of $G[L_b]$ for some valid tuple t_b for b . Moreover, the number of cut edges between twin-class X_i^a of L_a and twin-class X_j^b of L_b is $|S_a \cap X_i^a| \cdot |(L_b \setminus S_b) \cap X_j^b| + |S_b \cap X_j^b| \cdot |(L_a \setminus S_a) \cap X_i^a| = p_i^a \bar{p}_j^b + p_j^b \bar{p}_i^a$ if X_i^a and X_j^b is adjacent, zero otherwise. Therefore, the left-hand side is at most the right-hand side.

To show the converse direction, suppose $(S_a, L_a \setminus S_a)$ is a t_a -legitimate cut of $G[L_a]$ and $(S_b, L_b \setminus S_b)$ is a t_b -legitimate cut of $G[L_b]$, where t_v is consistent with (t_a, t_b) and the sizes of the cuts are $\text{mc}(a, t_a)$ and $\text{mc}(b, t_b)$, respectively. We claim that $(S_a \cup S_b, L_v \setminus (S_a \cup S_b))$ is a t_v -legitimate cut of $G[L_v]$. Since t_v is consistent with (t_a, t_b) , for each $1 \leq i \leq w_v$, we have $p_i^v = \sum_{j \in I_a(X_i^v)} p_j^a + \sum_{j \in I_b(X_i^v)} p_j^b = \sum_{1 \leq j \leq w_a} |S_a \cap X_j^i| + \sum_{1 \leq j \leq w_b} |S_b \cap X_j^i| = |(S_a \cup S_b) \cap X_i^v|$. Symmetrically, we have $\bar{p}_i^v = |(L_v \setminus (S_a \cup S_b)) \cap X_i^v|$. If $\bar{c}_i^v = 1$, by condition C3 of the consistency, either (1) $\{X_j^a : j \in I_a(X_i^v), p_j^a > 0\} \cup \{X_j^b : j \in I_b(X_i^v), p_j^b > 0\}$ is non-trivially connected or (2) exactly one of $\{p_j^s : s \in \{a, b\}, 1 \leq j \leq w_s\}$ is positive, say p_j^s , and $\bar{c}_j^s = 1$. If (1) holds, by Observation .2, $G[(S_a \cap S_b) \cap X_i^v]$ is connected. Otherwise, as $\bar{c}_j^s = 1$, $G[S_s \cap X_i^v] = G[(S_a \cup S_b) \cap X_i^v]$ is also connected. By a symmetric argument, we conclude that $G[(L_v \setminus (S_a \cup S_b)) \cap X_i^v]$ is connected if $\bar{c}_i^v = 1$. Therefore the cut $(S_a \cup S_b, L_v \setminus (S_a \cup S_b))$ is t_v -legitimate. Since the cut edges between two twin-classes of L_a is counted by $\text{mc}(a, t_a)$ and those between two twin-classes of L_b is counted by $\text{mc}(b, t_b)$. Similar to the forward direction, the number of cut edges between a twin-class of L_a and a twin-class of L_b can be counted by the third term in the right-hand side of the equality. Hence, the left-hand side is at least the right-hand side. \square

Proof of Theorem 10. From a w -expression tree of G , we can obtain a decomposition tree (T, ϕ) of width at most w in $O(n^2)$ time using Rao's algorithm [53]. Based on this decomposition, we evaluate the recurrence in Lemma 13 in a bottom-up manner. The number of valid tuples for each node of T is at most $4^w n^w$. For each internal node v and for each valid tuple t_v for v , we can compute $\text{mc}(v, t_v)$ in $(4^w n^w)^2 n^{O(1)}$ time. Overall, the running time of our algorithm is $n^{O(w)}$. Let r be the root of T . For MAXIMUM CONNECTED CUT, by the definition of legitimate cuts, we should take the maximum value among $\text{mc}(r, (i, n - i, 1, j))$ for $1 \leq i < n$ and $j \in \{0, 1\}$. Note that as L_v has only one twin-class, the length of valid tuples is exactly four. For LARGEST BOND, we should take the maximum value among $\text{mc}(r, (i, n - i, 1, 1))$ for $1 \leq i < n$.