

# Minimum Stable Cut and Treewidth

Michael Lampis  

Université Paris-Dauphine, PSL University, CNRS, LAMSADE, 75016, Paris, France

---

## Abstract

A stable or locally-optimal cut of a graph is a cut whose weight cannot be increased by changing the side of a single vertex. Equivalently, a cut is stable if all vertices have the (weighted) majority of their neighbors on the other side. Finding a stable cut is a prototypical PLS-complete problem that has been studied in the context of local search and of algorithmic game theory.

In this paper we study MIN STABLE CUT, the problem of finding a stable cut of minimum weight, which is closely related to the Price of Anarchy of the MAX CUT game. Since this problem is NP-hard, we study its complexity on graphs of low treewidth, low degree, or both. We begin by showing that the problem remains weakly NP-hard on severely restricted trees, so bounding treewidth alone cannot make it tractable. We match this hardness with a pseudo-polynomial DP algorithm solving the problem in time  $(\Delta \cdot W)^{O(\text{tw})} n^{O(1)}$ , where  $\text{tw}$  is the treewidth,  $\Delta$  the maximum degree, and  $W$  the maximum weight. On the other hand, bounding  $\Delta$  is also not enough, as the problem is NP-hard for unweighted graphs of bounded degree. We therefore parameterize MIN STABLE CUT by both  $\text{tw}$  and  $\Delta$  and obtain an FPT algorithm running in time  $2^{O(\Delta \text{tw})} (n + \log W)^{O(1)}$ . Our main result for the weighted problem is to provide a reduction showing that both aforementioned algorithms are essentially optimal, even if we replace treewidth by pathwidth: if there exists an algorithm running in  $(nW)^{o(\text{pw})}$  or  $2^{o(\Delta \text{pw})} (n + \log W)^{O(1)}$ , then the ETH is false. Complementing this, we show that we can, however, obtain an FPT *approximation scheme* parameterized by treewidth, if we consider almost-stable solutions, that is, solutions where no single vertex can unilaterally increase the weight of its incident cut edges by more than a factor of  $(1 + \varepsilon)$ .

Motivated by these mostly negative results, we consider UNWEIGHTED MIN STABLE CUT. Here our results already imply a much faster exact algorithm running in time  $\Delta^{O(\text{tw})} n^{O(1)}$ . We show that this is also probably essentially optimal: an algorithm running in  $n^{o(\text{pw})}$  would contradict the ETH.

**2012 ACM Subject Classification** Mathematics of computing → Graph algorithms; Theory of computation → Parameterized complexity and exact algorithms

**Keywords and phrases** Treewidth, Local Max-Cut, Nash Stability

**Digital Object Identifier** 10.4230/LIPIcs.ICALP.2021.92

**Category** Track A: Algorithms, Complexity and Games

**Related Version** *Full Version*: <https://arxiv.org/abs/2104.13097>

**Funding** The author is supported by ANR JCJC project “ASSK” (ANR-18-CE40-0025-01).

**Acknowledgements** Part of this work was conducted while the author was on sabbatical at IRIF, UMR 8243, Université de Paris.

## 1 Introduction

In this paper we study problems related to *stable cuts* in graphs. A *stable* cut of an edge-weighted graph  $G = (V, E)$  is a partition of  $V$  into two sets  $V_0, V_1$  that satisfies the following property: for each  $i \in \{0, 1\}$  and  $v \in V_i$ , the total weight of edges incident on  $v$  whose other endpoint is in  $V_{1-i}$  is at least half the total weight of all edges incident on  $v$ . In other words, a cut is stable if all vertices have the (weighted) majority of their incident edges cut.

The notion of stable cuts has been very widely studied from two different points of view. First, in the context of local search, a stable cut is a locally optimal cut: switching the side of any single vertex cannot increase the total weight of the cut. Hence, stable cuts have



© Michael Lampis;

licensed under Creative Commons License CC-BY 4.0

48th International Colloquium on Automata, Languages, and Programming (ICALP 2021).

Editors: Nikhil Bansal, Emanuela Merelli, and James Worrell; Article No. 92; pp. 92:1–92:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



been studied with the aim to further our understanding of the basic local search heuristic for MAX CUT. Second, in the context of algorithmic game theory a MAX CUT game has often been considered, where each vertex is an agent whose utility is the total weight of edges connecting it to the other side. In this game, a stable cut corresponds exactly to the notion of a Nash equilibrium, that is, a state where no agent has an incentive to change her choice. The complexity of producing a Nash stable or locally optimal cut of a given edge-weighted graph has been heavily studied under the name LOCAL MAX CUT. The problem is known to be PLS-complete, under various restrictions (we give detailed references below).

In this paper we focus on a different but closely related optimization problem: given an edge-weighted graph we would like to produce a stable cut *of minimum total weight*. We call this problem MIN STABLE CUT. In addition to being a fairly natural problem on its own, we believe that MIN STABLE CUT is interesting from the perspective of both local search and algorithmic game theory. In the context of local search, MIN STABLE CUT is the problem of bounding the performance of the local search heuristic on a particular instance. It is folklore (and easy to see) that in general there exist graphs where the smallest stable cut has size half the maximum cut (e.g. consider a  $C_4$ ) and this is tight since any stable cut must cut at least half the total edge weight. However, for most graphs this bound is far from tight. MIN STABLE CUT therefore essentially asks to estimate the ratio between the largest and smallest stable cut for a given specific instance. Similarly, in the context of algorithmic game theory, solving MIN STABLE CUT is essentially equivalent to calculating the Price of Anarchy of the MAX CUT game on the given instance, that is, the ratio between the smallest stable cut and the maximum cut. Since we will mostly focus on cases where MAX CUT is tractable, MIN STABLE CUT can, therefore, be seen as the problem of computing either the approximation ratio of local search or the Price of Anarchy of the MAX CUT game on a given graph.

**Our results.** It appears that little is currently known about the complexity of MIN STABLE CUT. However, since finding a (not necessarily minimum) stable cut is PLS-complete, finding the minimum such cut would be expected to be hard. Our focus is therefore to study the parameterized complexity of MIN STABLE CUT using structural parameters such as treewidth and the maximum degree of the input graph<sup>1</sup>. Our results are the following.

- First, we show that bounding only one of the two mentioned parameters is not sufficient to render the problem tractable. This is not surprising for the maximum degree  $\Delta$ , where a reduction from MAX CUT allows us to show the problem is NP-hard for  $\Delta \leq 6$  even in the unweighted case (Theorem 4). It is, however, somewhat more disappointing that bounded treewidth also does not help, as the problem remains weakly NP-hard on trees of diameter 4 (Theorem 1) and bipartite graphs of vertex cover 2 (Theorem 3).
- These hardness results point to two directions for obtaining algorithms for MIN STABLE CUT: first, since the problem is “only” weakly NP-hard for bounded treewidth one could hope to obtain a pseudo-polynomial time algorithm in this case. We show that this is indeed possible and the problem is solvable in time  $(\Delta \cdot W)^{O(\text{tw})} n^{O(1)}$ , where  $W$  is the maximum edge weight (Theorem 5). Second, one may hope to obtain an FPT algorithm when both  $\text{tw}$  and  $\Delta$  are parameters. We show that this is also possible and obtain an algorithm with complexity  $2^{O(\Delta \text{tw})} (n + \log W)^{O(1)}$  (Theorem 6).
- These two algorithms lead to two further questions. First, can the  $(\Delta \cdot W)^{O(\text{tw})} n^{O(1)}$  algorithm be improved to an FPT dependence on  $\text{tw}$ , that is, to running time  $f(\text{tw})(nW)^{O(1)}$ ? And second, can the  $2^{\Delta \text{tw}}$  parameter dependence of the FPT algorithm be improved,

---

<sup>1</sup> We assume familiarity with the basics of parameterized complexity as given in standard textbooks [21].

for example to  $2^{O(\Delta+tw)}$  or even  $\Delta^{O(tw)}$ ? We show that the answer to both questions is negative, even if we replace treewidth with pathwidth: under the ETH there is no algorithm running in  $(nW)^{o(pw)}$  or  $2^{o(\Delta tw)}(n + \log W)^{O(1)}$  (Theorem 8).

- Complementing the above, we show that the problem does become FPT by treewidth alone if we allow the notion of approximation to be used in the concept of stability: there exists an algorithm which, for any  $\varepsilon > 0$ , runs in time  $(tw/\varepsilon)^{O(tw)}(n + \log W)^{O(1)}$  and produces a cut with the following properties: all vertices are  $(1 + \varepsilon)$ -stable, that is, no vertex can unilaterally increase its incident cut weight by more than a factor of  $(1 + \varepsilon)$ ; the cut has weight at most equal to that of the minimum stable cut.
- Finally, motivated by the above mostly negative results, we also consider UNWEIGHTED MIN STABLE CUT, the restriction of the problem where all edge weights are uniform. Our previous results give a much faster algorithm with parameter dependence  $\Delta^{O(tw)}$ , rather than  $2^{\Delta tw}$  (Corollary 12). However, this poses the natural question if in this case the problem finally becomes FPT by treewidth alone. Our main result in this part is to answer this question in the negative and show that, under the ETH, UNWEIGHTED MIN STABLE CUT cannot be solved in time  $n^{o(pw)}$  (Theorem 13).

Taken together, our results paint a detailed picture of the complexity of MIN STABLE CUT parameterized by  $tw$  and  $\Delta$ . All our exact algorithms (Theorems 5, 6) are obtained using standard dynamic programming on tree decompositions, the only minor complication being that for Theorem 6 we edit the decomposition to make sure that for each vertex some bag contains all of its neighborhood (this helps us verify that a cut is stable). The main technical challenge is in proving our complexity lower bounds. It is therefore perhaps somewhat surprising that the lower bounds turn out to be essentially tight, as this indicates that for MIN STABLE CUT and UNWEIGHTED MIN STABLE CUT, the straightforward DP algorithms are essentially optimal, if one wants to solve the problem exactly.

For the approximation algorithm, we rely on two rounding techniques: one is a rounding step similar to the one that gives an FPTAS for KNAPSACK by truncating weights so that the maximum weight is polynomially bounded. However, MIN STABLE CUT is more complicated than KNAPSACK, as an edge which is light for one of its endpoints may be heavy for the other. We therefore define a more general version of the problem, allowing us to decouple the contribution each edge makes to the stability of each endpoint. This helps us bound the largest stability-weight by a polynomial, but is still not sufficient to obtain an FPT algorithm, as the lower bound of Theorem 8 applies to polynomially bounded weights. We then go on to apply a technique introduced in [46] (see also [2, 10, 43, 44]) which allows us to obtain FPT approximation algorithms for problems which are W-hard by treewidth by applying a different notion of rounding to the dynamic program. This allows us to produce a solution that is simultaneously of optimal weight (compared to the best stable solution) and almost-stable, using essentially the same algorithm as in Theorem 5. However, it is worth noting that in general there is no obvious way to transform almost-stable solutions to stable solutions [11, 17], so our algorithm is not immediately sufficient to obtain an FPT approximation for MIN STABLE CUT if we insist on obtaining a cut which is exactly stable.

**Related work.** From the point of view of local search algorithms, there is an extensive literature on the LOCAL MAX CUT problem, which asks us to find a stable cut (of any size). The problem has long been known to be PLS-complete [42, 52]. It remains PLS-complete for graphs of maximum degree 5 [27], but becomes polynomial-time solvable for graphs of maximum degree 3 [48, 51]. The problem remains PLS-complete if weights are assigned to vertices, instead of edges, and the weight of an edge is defined simply as the product of the

weights of its endpoints [31]. Even though the problem is PLS-complete, it has long been observed that local search quickly finds a stable solution in most practical instances. One theoretical explanation for this phenomenon was given in a recent line of work which showed that LOCAL MAX CUT has quasi-polynomial time smoothed complexity [3, 12, 18, 29]. LOCAL MAX CUT is of course polynomial time solvable if all weights are polynomially bounded in  $n$ , as local improvements always increase the size of the cut.

In algorithmic game theory much work has been done on the complexity of computing Nash equilibria for the cut game and the closely related *party affiliation game*, in which players, represented by vertices, have to pick one of two parties and edge weights indicate how much two players gain if they are in the same party [6, 7, 19, 30, 35]. Note that for general graphical games finding an equilibrium is PPAD-hard on trees of constant pathwidth [25]. Because computing a stable solution is generally intractable, approximate equilibria have also been considered [11, 17]. Note that the notion of approximate equilibrium corresponds exactly to the approximation guarantee given by Theorem 11, but unlike the cited works, Theorem 11 produces a solution that is both approximately stable and as good as the optimal.

The problem we consider in this paper is more closely related to the problem of computing the *worst* (or best) Nash equilibrium, which in turn is closely linked to the notion of Price of Anarchy. For most problems in algorithmic game theory this type of question is usually NP-hard [13, 20, 26, 32, 34, 37, 53] and hard to approximate [5, 16, 22, 39, 49]. Even though these results show that finding a Nash equilibrium that maximizes an objective function is NP-hard under various restrictions (e.g. graphical games of bounded degree), to the best of our knowledge the complexity of finding the worst equilibrium of the MAX CUT game (which corresponds to the MIN STABLE CUT problem of this paper) has not been considered.

Finally, another topic that has recently attracted attention in the literature is that of MinMax and MaxMin versions of standard optimization problems, where we search the worst solution which cannot be improved using a simple local search heuristic. The motivation behind this line of research is to provide bounds and a refined analysis of such basic heuristics. Problems that have been considered under this lens are MAX MIN DOMINATING SET [8, 24], MAX MIN VERTEX COVER [15, 54], MAX MIN SEPARATOR [38], MAX MIN CUT [28], MIN MAX KNAPSACK [4, 33, 36], MAX MIN EDGE COVER [45], MAX MIN FEEDBACK VERTEX SET [23]. Some problems in this area also arise naturally in other forms and have been extensively studied, such as MIN MAX MATCHING (also known as EDGE DOMINATING SET [41]) and GRUNDY COLORING, which can be seen as a MAX MIN version of COLORING [1, 9].

## 2 Definitions – Preliminaries

We generally use standard graph-theoretic notation and consider edge-weighted graphs, that is, graphs  $G = (V, E)$  supplied with a weight function  $w : E \rightarrow \mathbb{N}$ . For a vertex  $v \in V$ , The weighted degree of a vertex  $v \in V$  is  $d_w(v) = \sum_{uv \in E} w(uv)$ . A cut of a graph is a partition of  $V$  into  $V_0, V_1$ . A cut is *stable* for vertex  $v \in V_i$  if  $\sum_{vu \in E \wedge u \in V_{1-i}} w(vu) \geq \frac{d_w(v)}{2}$ , that is, if the total weight of edges incident on  $v$  crossing the cut is at least half the weighted degree of  $v$ . In the MIN STABLE CUT problem we are given an edge-weighted graph and are looking for a cut that is stable for all vertices that minimizes the sum of weights of cut edges (that is, edges with endpoints on both sides of the cut). In UNWEIGHTED MIN STABLE CUT we restrict the problem so that the  $w$  function returns 1 for all edges. When describing stable cuts we will sometimes say that we “assign” value 0 (or 1) to a vertex; by this we mean that we place this vertex in  $V_0$  (or  $V_1$  respectively).

For the definitions of treewidth, pathwidth, and the related (nice) decompositions we refer to [21]. We will use as a complexity assumption the Exponential Time Hypothesis (ETH) [40] which states that there exists a constant  $c > 1$  such that 3-SAT with  $n$  variables and  $m$  clauses cannot be solved in time  $c^{n+m}$ . In fact, we will use the slightly weaker and simpler form of the ETH which states that 3-SAT cannot be solved in time  $2^{o(n+m)}$ .

### 3 Weighted Min Stable Cut

In this section we present our results on exact algorithms for (weighted) MIN STABLE CUT. We begin with some basic NP-hardness reductions in Section 3.1, which establish that the problem remains (weakly) NP-hard when either the treewidth or the maximum degree are bounded. These set the stage for two algorithms, given in Section 3.2, solving the problem in pseudo-polynomial time for constant treewidth; and in FPT time parameterized by  $\text{tw} + \Delta$ . In Section 3.3 we present a more fine-grained hardness argument, based on the ETH, which shows that the dependence on  $\text{tw}$  and  $\Delta$  of our two algorithms is essentially optimal.

#### 3.1 Basic Hardness Proofs

► **Theorem 1.** *MIN STABLE CUT is weakly NP-hard on trees of diameter 4.*

**Proof.** We describe a reduction from PARTITION. Recall that in this problem we are given  $n$  positive integers  $x_1, \dots, x_n$  such that  $\sum_{i=1}^n x_i = 2B$  and are asked if there exists  $S \subseteq [n]$  such that  $\sum_{i \in S} x_i = B$ . We construct a star with  $n$  leaves and subdivide every edge once. For each  $i \in [n]$  we select a distinct leaf of the tree and set the weight of both edges in the path from the center to this leaf to  $x_i$ . We claim that the graph has a stable cut of weight  $3B$  if and only if there is a partition of  $x_1, \dots, x_n$  into two sets with the same sum.

For the first direction, suppose  $S \subseteq [n]$  is such that  $\sum_{i \in S} x_i = B$ . For each  $i \in S$  we select a degree two vertex of the tree whose incident edges have weight  $x_i$  and assign it value 1. We assign all other degree two vertices value 0 and assign to all leaves the opposite of the value of their neighbor. We give the center value 0. This partition is stable as the center has edge weight exactly  $B$  towards each side, and all degree two vertices have a leaf attached that is placed on the other side and contributes half their total incident weight. The total weight cut is  $2B$  from edges incident on leaves, plus  $B$  from half the weight incident on the center.

For the converse direction, observe that in any stable solution all edges incident on leaves are cut, contributing a weight of  $2B$ . As a result, in a stable cut of size  $3B$ , the weight of cut edges incident on the center is at most  $B$ . However, this weight is also at least  $B$ , since the edge weight incident on the center is  $2B$ . We conclude that the neighborhood of the center must be perfectly balanced. From this we can infer a solution to the PARTITION instance. ◀

► **Remark 2.** Theorem 1 is tight, because MIN STABLE CUT is trivial on trees of diameter at most 3.

► **Theorem 3.** *MIN STABLE CUT is weakly NP-hard on bipartite graphs with vertex cover 2.*

► **Theorem 4.** *UNWEIGHTED MIN STABLE CUT is strongly NP-hard and APX-hard on bipartite graphs of maximum degree 6.*

#### 3.2 Algorithms

► **Theorem 5.** *There is an algorithm which, given an instance of MIN STABLE CUT with  $n$  vertices, maximum weight  $W$ , and a tree decomposition of width  $\text{tw}$ , finds an optimal solution in time  $(\Delta \cdot W)^{O(\text{tw})} n^{O(1)}$ .*

► **Theorem 6.** *There is an algorithm which, given an instance of MIN STABLE CUT with  $n$  vertices, maximum weight  $W$ , maximum degree  $\Delta$  and a tree decomposition of width  $\text{tw}$ , finds an optimal solution in time  $2^{O(\Delta \text{tw})}(n + \log W)^{O(1)}$ .*


**Proof.** We describe an algorithm which works in a way similar to the standard algorithm for MAX CUT parameterized by treewidth, except that we work in a tree decomposition that is essentially a decomposition of the square of  $G$ . More precisely, before we begin, we do the following: for each  $v \in V$  we add to every bag of the decomposition that contains  $v$  all the vertices of  $N(v)$ . It is not hard to see that we now have a decomposition of width at most  $(\Delta + 1)(\text{tw} + 1)$  and also that the new decomposition is still a valid tree decomposition. Crucially, we now also have the following property: for each  $v \in V$  there exists at least one bag of the decomposition that contains all of  $N[v]$ .

The algorithm now performs dynamic programming by storing for each bag the value of the best solution for each partition of  $B_t$ . As a result, the size of the DP table is  $2^{O(\Delta \text{tw})}$ . The only difference with the standard MAX CUT algorithm (beyond the fact that we are looking for a cut of minimum weight) is that when we consider a bag that contains all of  $N[v]$ , for some  $v \in V$ , we discard all partitions which are unstable for  $v$ . Since the bag contains all of  $N[v]$ , this can be checked in time polynomial in  $n$  and  $\log W$  (assuming weights are given in binary). ◀

### 3.3 Tight ETH-based Hardness

We first give a reduction from 3-SET SPLITTING to MIN STABLE CUT whose main properties are laid out in Lemma 7. This reduction gives the lower bound of Theorem 8.

► **Lemma 7.** *There is a polynomial-time algorithm which, given a 3-SET SPLITTING instance  $H = (V, E)$  with  $n$  elements, produces a MIN STABLE CUT instance  $G$  with the following properties: (i)  $G$  is a Yes instance if and only if  $H$  is a Yes instance; (ii) if  $\Delta$  is the maximum degree of  $G$  and  $\text{pw}$  its pathwidth, then  $\Delta = O(\log n)$  and  $\text{pw} = O(n / \log n)$ ; (iii) the maximum weight of  $G$  is  $W = O(2^\Delta)$ .*



■ **Figure 1** Sketch of the construction of Lemma 7. On the left, the general architecture:  $m$  columns, each with  $n$  vertices, partitioned into groups of size  $\log n$ . On each column we add a checker vertex (on top). Between the same groups of consecutive columns we add propagator vertices. On the right, more details about the exponentially increasing weights of edges incident on propagators.



**Proof.** Let  $H = (V, E)$  be the given 3-SET SPLITTING instance,  $V = \{v_0, \dots, v_{n-1}\}$  and suppose that  $E$  contains  $e_2$  sets of size 2 and  $e_3$  sets of size 3, where  $|E| = e_2 + e_3$  will be denoted by  $m$ . Assume without loss of generality that  $n$  is a power of 2 (otherwise add some dummy elements to  $V$ ). Let  $\delta = \log n$ . We construct a graph by first making  $m$  copies of  $V$ , call them  $V_j, j \in [m]$  and label their vertices as  $V_j = \{v_{(i,j)} \mid i \in \{0, \dots, n-1\}\}$ . Intuitively, the vertices  $\{v_{(i,j)} \mid j \in [m]\}$  are all meant to represent the element  $v_i$  of  $H$ . We now add to the graph the following:

1. Checkers: Suppose that the  $j$ -th set of  $E$  contains elements  $v_{i_1}, v_{i_2}, v_{i_3}$ . Then we construct a vertex  $c_j$  and connect it to  $v_{(i_1,j)}, v_{(i_2,j)}, v_{(i_3,j)}$  with edges of weight 1. If the  $j$ -th set has size two, we do the same (ignoring  $v_{i_3}$ ).
2. Propagators: For each  $j \in [m-1]$  we construct  $\rho = \lceil n/\delta \rceil$  vertices labeled  $p_{(i,j)}, i \in \{0, \dots, \rho-1\}$ . Each  $p_{(i,j)}$  is connected to (at most)  $\delta$  vertices of  $V_j$  and  $\delta$  vertices of  $V_{j+1}$  with edges of exponentially increasing weight. Specifically, for  $i \in \{0, \dots, \rho-1\}, \ell \in \{0, \dots, \delta-1\}$ , we connect  $p_{(i,j)}$  to  $v_{(i\delta+\ell,j)}$  and to  $v_{(i\delta+\ell,j+1)}$  (if they exist) with an edge of weight  $2^\ell$ .
3. Stabilizers: For each  $j \in [m], i \in \{0, \dots, n-1\}$  we attach to  $v_{(i,j)}$  a leaf. The edge connecting this leaf to  $v_{(i,j)}$  has weight  $3 \cdot 2^{(i \bmod \delta)}$ .

This completes the construction of the graph. Let  $L_w$  be the total weight of edges incident on leaves and  $P$  be the total weight of edges incident on Propagator vertices  $p_{(i,j)}$ . We set  $B = L_w + \frac{P}{2} + e_2 + 2e_3$  and claim that the new instance has a stable cut of weight  $B$  if and only if  $H$  can be split.

For the forward direction, suppose that  $H$  can be split by the partition of  $V$  into  $L, R = V \setminus L$ . We assign the following values for our new instance: for each  $j \in [m]$  odd, we set  $v_{(i,j)}$  to value 0 if and only if  $v_i \in L$ ; for each  $j \in [m]$  even, we set  $v_{(i,j)}$  to value 0 if and only if  $v_i \in R$ . In other words, we use the same partition for all copies of  $V$ , but flip the roles of 0, 1 between consecutive copies. We place leaves on the opposite side from their neighbors and greedily assign values to all other vertices of the graph to obtain a stable partition. Observe that all vertices  $v_{(i,j)}$  are stable with the values we assigned, since the edge connecting each such vertex to a leaf has weight at least half its total incident weight.

In the partition we have, we observe that (i) all edges incident on leaves are cut (total weight  $L_w$ ) (ii) all Propagator vertices have balanced neighborhoods, so exactly half of their incident weight is cut (total weight  $P/2$ ) (iii) since  $L, R$  splits all sets of  $E$ , each checker vertex will have exactly one neighbor on the same side (total weight  $e_2 + 2e_3$ ). So, the total weight of the cut is  $B$ .

For the converse direction, suppose we have a stable cut of size  $B$  in the constructed instance. Because of the stability condition, this solution must cut all edges incident on leaves (total weight  $L_w$ ); at least half of the total weight of edges incident on Propagators (total weight  $P/2$ ); and for each checker vertex all its incident edges except at most one (total weight at least  $e_2 + 2e_3$ ). We conclude that, in order to achieve weight  $B$ , the cut must properly balance the neighborhood of all Propagators and make sure that each Checker vertex has one neighbor on its own side.

We now argue that because the neighborhood of each Propagator is balanced we have for all  $i \in \{0, \dots, n-1\}, j \in [m-1]$  that  $v_{(i,j)}, v_{(i,j+1)}$  are on different sides of the partition. To see this, suppose for contradiction that for two such vertices this is not the case and to ease notation consider the vertices  $v_{(i\delta+\ell,j)}, v_{(i\delta+\ell,j+1)}$ , where  $0 \leq \ell \leq \delta-1$ . Among all such pairs select one that maximizes  $\ell$ . Both vertices are connected to the Propagator  $p_{(i,j)}$  with edges of weight  $2^\ell$ . But now  $p_{(i,j)}$  has strictly larger edge weight connecting it to the side of the

partition that contains  $v_{(i\delta+\ell,j)}$  and  $v_{(i\delta+\ell,j+1)}$  than to the other side because (i) for neighbors of  $p_{(i,j)}$  connected to it with edges of higher weight, the neighborhood of  $p_{(i,j)}$  is balanced by the maximality of  $\ell$  (ii) the total weight of all other edges is  $2 \cdot (2^{\ell-1} + 2^{\ell-2} + \dots + 1) < 2 \cdot 2^\ell$ .

We thus have that for all  $i, j$ ,  $v_{(i,j)}, v_{(i,j+1)}$  must be on different sides, and therefore all  $V_j$  are partitioned in the same way (except some have the roles of 0 and 1 reversed). From this, we obtain a partition of  $V$ . To conclude this direction, we argue that this partition of  $V$  must split all sets. Indeed, if not, there will be a checker vertex such that all its neighbors are on the same side, which, as we argued, means that the cut must have weight strictly more than  $B$ .

Finally, let us show that the constructed instance has the claimed properties. The maximum degree is  $\Delta = 2\delta = O(\log n)$  in the Propagators vertices (all other vertices have degree at most 4); the maximum weight is  $O(2^\delta) = O(2^\Delta)$ . Let us also consider the pathwidth of the constructed graph. Let  $G_j$  be the subgraph induced by  $V_j$  and its attached leaves, the Checker  $c_j$ , and all Propagators adjacent to  $V_j$ . We claim that we can build a path decomposition of  $G_j$  that contains all Propagators adjacent to  $V_j$  in all bags and has width  $O(n/\log n)$ . Indeed, if we place all the (at most  $\lceil 2n/\delta \rceil$ ) Propagators and  $c_j$  in all bags, we can delete them from  $G_j$ , and all that is left is a union of isolated edges, which has pathwidth 1. Now, since the union of all  $G_j$  covers all vertices and edges, we can construct a path decomposition of the whole graph of width  $O(n/\log n)$  by gluing together the decompositions of each  $G_j$ , that is, by connecting the last bag of the decomposition of  $G_j$  to the first bag of the decomposition of  $G_{j+1}$ . ◀

► **Theorem 8.** *If the ETH is true then (i) there is no algorithm solving MIN STABLE CUT in time  $(nW)^{o(\text{pw})}$  (ii) there is no algorithm solving MIN STABLE CUT in time  $2^{o(\Delta \text{pw})}(n + \log W)^{O(1)}$ . These statements apply even if we restrict the input to instances where weights are written in unary and the maximum degree is  $O(\log n)$ .*

## 4 Approximately Stable Cuts

In this section we present an algorithm which runs in FPT time parameterized by treewidth and produces a solution that is  $(1 + \varepsilon)$ -stable and has weight upper bounded by the weight of the optimal stable cut. Before we proceed, we will need to define a more general version of our problem. In EXTENDED MIN STABLE CUT we are given as input: a graph  $G = (V, E)$ ; a cut-weight function  $w : E \rightarrow \mathbb{N}$ ; and a stability-weight function  $s : E \times V \rightarrow \mathbb{N}$ . For  $v \in V$  we denote  $d_s(v) = \sum_{vu \in E} s(vu, v)$ , which we call the stability degree of  $v$ . If we are also given an error parameter  $\rho > 1$ , we will then be looking for a partition of  $V$  into  $V_0, V_1$  which satisfies the following: (i) each vertex is  $\rho$ -stable, that is, for each  $i \in \{0, 1\}$  and  $v \in V_i$  we have  $\sum_{vu \in E \wedge u \in V_{1-i}} s(vu, v) \geq \frac{d_s(v)}{2\rho}$  (ii) the total cut weight  $\sum_{u \in V_0, v \in V_1, uv \in E} w(uv)$  is minimum. Observe that this extended version of the problem contains MIN STABLE CUT as a special case if  $\rho = 1$  and for all  $uv \in E$  we have  $s(uv, v) = s(uv, u) = w(uv)$ .

The generalization of MIN STABLE CUT is motivated by three considerations. First, the algorithm of Theorem 5 is inefficient because it has to store exact weight values to satisfy the stability constraints; however, it can efficiently store the total weight of the cut. We therefore decouple the contribution of an edge to the size of the cut (given by  $w$ ) from a contribution of an edge to the stability of its endpoints (given by  $s$ ). Second, our strategy will be to truncate the values of  $s$  so that the DP of the algorithm of Theorem 5 can be run more efficiently. To do this we will first simply divide all stability-weights by an appropriate value. However, a problem we run into if we do this is that the edge  $uv$  could simultaneously be one of the heavier edges incident on  $u$  and one of the lighter edges incident on  $v$ , so it



is not clear how we can adjust its weight in a way that minimizes the distortion for both endpoints. As a result it is simpler if we allow edges to contribute different amounts to the stability of their endpoints. In this sense,  $s(uv, u)$  is the amount that the edge  $uv$  contributes to the stability of  $u$  if the edge is cut. Observe that with the new definition, if we set a new stability-weight function for a specific vertex  $u$  as  $s'(uv, v) = c \cdot s(uv, v)$  for all  $v \in N(u)$ , that is, if we multiply the stability-weight of all edges incident on  $u$  by a constant  $c$  and leave all other values unchanged, we obtain an equivalent instance, and this does not affect the stability of other vertices. Finally, the parameter  $\rho$  allows us to consider solutions where a vertex is stable if its cut incident edges are at least a  $(\frac{1}{2\rho})$ -fraction of its stability degree.

Armed with this intuition we can now explain our approach to obtaining our FPT approximation algorithm. Given an instance of the extended problem, we first adjust the  $s$  function so that its maximum value is bounded by a polynomial in  $n$ . We achieve this by dividing  $s(uv, u)$  by a value that depends only on  $d_s(u)$  and  $n$ . This allows us to guarantee that near-stable solutions are preserved. Then, given an instance where the maximum value of  $s$  is polynomially bounded, we apply the technique of [46], using the algorithm of Theorem 5 as a base, to obtain our approximation. We give these separate steps in the Lemmas below.

► **Lemma 9.** *There is an algorithm which, given a graph  $G = (V, E)$  on  $n$  vertices and a stability-weight function  $s : E \times V \rightarrow \mathbb{N}$  with maximum value  $S$ , runs in time polynomial in  $n + \log S$  and produces a stability-weight function  $s' : E \times V \rightarrow \mathbb{N}$  with the following properties: (i) the maximum value of  $s'$  is  $O(n^2)$  (ii) for all partitions  $V$  into  $V_0, V_1$ ,  $i \in \{0, 1\}$ ,  $v \in V_i$  we have*

$$\left( \frac{\sum_{vu \in E, u \in V_{1-i}} s(vu, v)}{d_s(v)} \right) / \left( \frac{\sum_{vu \in E, u \in V_{1-i}} s'(vu, v)}{d_{s'}(v)} \right) \in [1 - 1/n, 1 + 1/n]$$

Using Lemma 9 we can assume that all stability-weights are bounded by  $O(n^2)$ . The most important part is that Lemma 9 guarantees us that almost-optimal solutions are preserved in both directions, as for any cut and for each vertex the ratio of stability weight going to the other side over the total stability-degree of the vertex does not change by more than a factor  $(1 + \frac{1}{n})$ . Let us now see the second ingredient of our algorithm.

► **Lemma 10.** *There is an algorithm which takes as input a graph  $G = (V, E)$ , a cut-weight function  $w : E \rightarrow \mathbb{N}$  with maximum  $W$ , a stability-weight function  $s : E \times V \rightarrow \mathbb{N}$  with maximum  $S$ , a tree decomposition of  $G$  of width  $\text{tw}$ , and an error parameter  $\varepsilon > 0$  and returns a  $(1+2\varepsilon)$ -stable solution that has cut-weight at most equal to that of the minimum  $(1+\varepsilon)$ -stable solution. If  $S = O(n^2)$ , then the algorithm runs in time  $(\text{tw}/\varepsilon)^{O(\text{tw})}(n + \log W)^{O(1)}$ .*

**Proof.** We use the methodology of [46]. Before we proceed, let us explain that we are actually aiming for an algorithm with running time roughly  $(\log n/\varepsilon)^{O(\text{tw})}$ . This type of running time implies the time stated in the lemma using a standard Win/Win argument: if  $\text{tw} \leq \sqrt{\log n}$  then  $(\log n)^{O(\text{tw})}$  is  $n^{o(1)}$ , so the  $(\log n)^{O(\text{tw})}$  factor is absorbed in the  $n^{O(1)}$  factor; while if  $\log n \leq \text{tw}^2$ , then an algorithm running in  $(\log n)^{\text{tw}}$  actually runs in time  $(\text{tw})^{O(\text{tw})}$ .

To be more precise, if the given tree decomposition has height  $H$ , then we will formulate an algorithm with running time  $(H \log S/\varepsilon)^{O(\text{tw})}(n + \log W)^{O(1)}$ . This running time achieves parameter dependence  $(\log n/\varepsilon)^{O(\text{tw})}$  if we use the fact that  $S = O(n^2)$  and a theorem due to [14] which proves that any tree decomposition can be edited (in polynomial time) so that its height becomes  $O(\log n)$ , without increasing its width by more than a constant factor.

The basis of our algorithm will be the algorithm of Theorem 5, appropriately adjusted to the extended version of the problem. Let us first sketch the modifications to the algorithm of Theorem 5 that we would need to do to solve this more general problem, since the details

are straightforward. First, we observe that in solution signatures we would now take into account stability-weights, and signatures would have values going up to  $S$ . Second, in Forget nodes, since we are happy with a  $(1 + \varepsilon)$ -stable solution, we would only discard solutions which violate this constraint. With these modifications, we can run this exact algorithm to return the minimum  $(1 + \varepsilon)$ -stable solution in time  $(2S)^{O(\text{tw})}(n + \log W + \log(1/\varepsilon))^{O(1)}$ .


The idea is to modify this algorithm so that the DP tables go from size  $(2S)^{\text{tw}}$  to roughly  $(H \log S)^{\text{tw}}$ . To do this, we define a parameter  $\delta = \frac{\varepsilon}{5H}$ . We intend to replace every value  $x$  that would be stored in the signature of a solution in the DP table, with the next larger integer power of  $(1 + \delta)$ , that is, to construct a DP table where  $x$  is replaced by  $(1 + \delta)^{\lceil \log_{(1+\delta)} x \rceil}$ .

More precisely, the invariant we maintain is the following. Consider a node  $t$  of the decomposition at height  $h$ , where  $h = 0$  corresponds to leaves. We maintain a collection of solution signatures such that: (i) each signature contains a partition of  $B_t$  and for each  $v \in B_t$  an integer that is upper-bounded by  $\lceil \log_{(1+\delta)} d_s(v) \rceil$ ; (ii) Soundness: for each stored signature there exists a partition of  $B_t^\downarrow$  which approximately corresponds to it. Specifically, the partition and the signature agree exactly on the assignment of  $B_t$  and the total cut-weight; the partition is  $(1 + 2\varepsilon)$ -stable for all vertices of  $B_t^\downarrow \setminus B_t$ ; and for each  $v \in B_t$ , if the signature stores the value  $x(v)$  for  $v$ , that is, it states that  $v$  has approximate stability-weight  $(1 + \delta)^{x(v)}$  towards its own side in  $B_t^\downarrow \setminus B_t$ , then in the actual partition the stability-weight of  $v$  to its own side of  $B_t^\downarrow \setminus B_t$  is at most  $(1 + \delta)^h (1 + \delta)^{x(v)}$ . (iii) Completeness: conversely, for each partition of  $B_t^\downarrow$  that is  $(1 + \varepsilon)$ -stable for all vertices of  $B_t^\downarrow \setminus B_t$  there exists a signature that approximately corresponds to it. Specifically, the partition and signature agree on the assignment of  $B_t$  and the total cut-weight; and for each  $v \in B_t$ , if the stability-weight of  $v$  towards its side of the partition of  $B_t^\downarrow \setminus B_t$  is  $y(v)$ , and the signature stores the value  $x(v)$ , then  $(1 + \delta)^{x(v)} \leq (1 + \delta)^h y(v)$ .

In more simple terms, the signatures in our DP table store values  $x(v)$  so that we estimate that in the corresponding solution  $v$  has approximately  $(1 + \delta)^{x(v)}$  weight towards its own side in  $B_t^\downarrow$ , that is, we estimate that the DP of the exact algorithm would store approximately the value  $(1 + \delta)^{x(v)}$  for this solution. Of course, it is hard to maintain this relation exactly, so we are happy if for a node at height  $h$  the “true” value which we are approximating is at most a factor of  $(1 + \delta)^h$  off from our approximation.

Now, the crucial observation is that the approximate DP tables can be maintained because our invariant allows the error to increase with the height. For example, suppose that  $t$  is a Forget node at height  $h$  and let  $u \in B_t$  be a neighbor of the vertex  $v$  we forget. The exact algorithm would construct the signature of a solution in  $t$  by looking at the signature of a solution in its child node, and then adding to the value stored for  $u$  the weight  $s(vu, u)$  (if  $u, v$  are on the same side). Our algorithm will take an approximate signature from the child node, which may have a value at most  $(1 + \delta)^{h-1}$  the correct value, add to it  $s(vu, u)$  and then, perhaps, round-up the value to an integer power of  $(1 + \delta)$ . The new approximation will be at most  $(1 + \delta)^h$  larger than the value that the exact algorithm would have calculated. Similar argumentation holds for Join nodes. Furthermore, in Forget nodes we will only discard a solution if according to our approximation it is not  $(1 + 2\varepsilon)$ -stable. We may be over-estimating the stability-weight a vertex has to its own side of the cut by a factor of at most  $(1 + \delta)^h \leq (1 + \frac{\varepsilon}{5H})^H \leq 1 + \frac{\varepsilon}{2}$  so if for a signature our approximation says that the solution is not  $(1 + 2\varepsilon)$ -stable, the solution cannot be  $(1 + \varepsilon)$ -stable, because  $(1 + \varepsilon)(1 + \frac{\varepsilon}{2}) < 1 + 2\varepsilon$  (for sufficiently small  $\varepsilon$ ).

Finally, to estimate the running time, the maximum value we have to store for each vertex in a bag is  $\log_{(1+\delta)} S = \frac{\log S}{\log(1+\delta)} \leq O(\frac{\log n}{\delta}) = O(\frac{H \log n}{\varepsilon})$ . Using the fact that  $H = O(\log n)$  we get that the size of the DP table is  $(\log n / \varepsilon)^{O(\text{tw})}$ . ◀



■ **Figure 2** Checker gadget for Theorem 13. On the right two Selector gadgets. This Checker verifies that we have not taken an edge which has endpoints (2, 3), hence  $t^1, t^3$  are connected to the first 2 and 3 vertices of the Selectors.

► **Theorem 11.** *There is an algorithm which, given an instance of MIN STABLE CUT  $G = (V, E)$  with  $n$  vertices, maximum weight  $W$ , a tree decomposition of width  $\text{tw}$ , and a desired error  $\varepsilon > 0$ , runs in time  $(\text{tw}/\varepsilon)^{O(\text{tw})}(n + \log W)^{O(1)}$  and returns a cut with the following properties: (i) for all  $v \in V$ , the total weight of edges incident on  $v$  crossing the cut is at least  $(1 - \varepsilon)\frac{d_w(v)}{2}$  (ii) the cut has total weight at most equal to the weight of the minimum stable cut.*

## 5 Unweighted Min Stable Cut

In this section we consider UNWEIGHTED MIN STABLE CUT. We first observe that applying Theorem 5 gives a parameter dependence of  $\Delta^{O(\text{tw})}$ , since  $W = 1$ . We then show that this algorithm is essentially optimal, as the problem cannot be solved in  $n^{o(\text{pw})}$  under the ETH.

► **Corollary 12.** *There is an algorithm which, given an instance of UNWEIGHTED MIN STABLE CUT with  $n$  vertices, maximum degree  $\Delta$ , and a tree decomposition of width  $\text{tw}$ , returns an optimal solution in time  $\Delta^{O(\text{tw})}n^{O(1)}$ .*

We now first state our hardness result, then describe the construction of our reduction, and finally go through a series of lemmas that establish its correctness.

► **Theorem 13.** *If the ETH is true then no algorithm can solve UNWEIGHTED MIN STABLE CUT on graphs with  $n$  vertices in time  $n^{o(\text{pw})}$ . Furthermore, UNWEIGHTED MIN STABLE CUT is  $W[1]$ -hard parameterized by pathwidth.*

To prove Theorem 13 we will describe a reduction from  $k$ -MULTI-COLORED INDEPENDENT SET, a well-known  $W[1]$ -hard problem that cannot be solved in  $n^{o(k)}$  time under the ETH [21]. In this problem we are given a graph  $G = (V, E)$  with  $V$  partitioned into  $k$  color classes  $V_1, \dots, V_k$ , each of size  $n$ , and we are asked to find an independent set of size  $k$  which selects one vertex from each  $V_i$ . In the remainder we use  $m$  to denote the number of edges of  $E$  and assume that vertices of  $V$  are labeled  $v_{(i,j)}, i \in [k], j \in [n]$ , where  $V_i = \{v_{(i,j)} \mid j \in [n]\}$ .

Before we proceed, let us give some intuition. Our reduction will rely on a  $k \times m$  grid-like construction, where each row represents the selection of a vertex in the corresponding color class of  $G$  and each column represents an edge of  $G$ . The main ingredients will be a Selector gadget, which will represent a choice of an index in  $[n]$ ; a Propagator gadget which will make sure that the choice we make in each row stays consistent throughout; and a Checker gadget which will verify that we did not select the two endpoints of any edge. Each Selector gadget

will contain a path on (roughly)  $n$  vertices such that any reasonable stable cut will have to cut exactly one edge of the path. The choice of where to cut this path will represent an index in  $[n]$  encoding a vertex of  $G$ .

In our construction we will also make use of a simple but important gadget which we will call a “heavy” edge. Let  $A = n^5$ . When we say that we connect  $u, v$  with a heavy edge we will mean that we construct  $A$  new vertices and connect them to both  $u$  and  $v$ . The intuitive idea behind this gadget is that the large number of degree two vertices will force  $u$  and  $v$  to be on different sides of the partition (otherwise too many edges will be cut). We will also sometimes attach leaves on some vertices with the intention of making it easier for this vertex to achieve stability (as its attached leaves will always be on the other side of the partition).

Let us now describe our construction step-by-step.

1. Construct two “palette” vertices  $p_0, p_1$  and a heavy edge connecting them. Note that all heavy edges we will add will be incident on at least one palette vertex.
2. For each  $i \in [k], j \in [m]$  construct the following Selector gadget:
  - a. Construct a path on  $n + 1$  vertices  $P_{(i,j)}$  and label its vertices  $P_{(i,j)}^1, \dots, P_{(i,j)}^{n+1}$ .
  - b. If  $j$  is odd, then add a heavy edge from  $P_{(i,j)}^1$  to  $p_1$  and a heavy edge from  $P_{(i,j)}^{n+1}$  to  $p_0$ . If  $j$  is even, then add a heavy edge from  $P_{(i,j)}^1$  to  $p_0$  and a heavy edge from  $P_{(i,j)}^{n+1}$  to  $p_1$ .
  - c. Attach 5 leaves to each  $P_{(i,j)}^\ell$  for  $\ell \in \{2, \dots, n\}$ . Attach  $A + 5$  leaves to  $P_{(i,j)}^1$  and  $P_{(i,j)}^{n+1}$ .
3. For each  $i \in [k], j \in [m - 1]$  construct a new vertex connected to all vertices of the paths  $P_{(i,j)}$  and  $P_{(i,j+1)}$ . This vertex is the Propagator gadget.
4. For each  $j \in [m]$  consider the  $j$ -th edge of the original instance and suppose it connects  $v_{(i_1, j_1)}$  to  $v_{(i_2, j_2)}$ . We construct the following Checker gadget (see Figure 2)
  - a. We construct four vertices  $t_j^1, t_j^2, t_j^3, t_j^4$ . These are connected to existing vertices as follows:  $t_j^1$  is connected to  $\{P_{(i_1, j)}^1, \dots, P_{(i_1, j)}^{j_1}\}$  (that is, the first  $j_1$  vertices of the path  $P_{(i_1, j)}$ );  $t_j^2$  is connected to  $\{P_{(i_1, j)}^{j_1+1}, \dots, P_{(i_1, j)}^{n+1}\}$  (that is, the remaining  $n + 1 - j_1$  vertices of  $P_{(i_1, j)}$ ); similarly,  $t_j^3$  is connected to  $\{P_{(i_2, j)}^1, \dots, P_{(i_2, j)}^{j_2}\}$ ; and finally  $t_j^4$  is connected to  $\{P_{(i_2, j)}^{j_2+1}, \dots, P_{(i_2, j)}^{n+1}\}$ .
  - b. We construct four independent sets  $T_j^1, T_j^2, T_j^3, T_j^4$  with respective sizes  $j_1, n + 1 - j_1, j_2, n + 1 - j_2$ . We connect  $t_j^1$  to all vertices of  $T_j^1$ ,  $t_j^2$  to  $T_j^2$ ,  $t_j^3$  to  $T_j^3$ , and  $t_j^4$  to  $T_j^4$ . We attach two leaves to each vertex of  $T_j^1 \cup T_j^2 \cup T_j^3 \cup T_j^4$ .
  - c. We construct three vertices  $a_j, b_j, c_j$ . We connect  $c_j$  to both  $a_j$  and  $b_j$ . We connect  $a_j$  to an arbitrary vertex of  $T_j^1$  and an arbitrary vertex of  $T_j^3$ . We connect  $b_j$  to an arbitrary vertex of  $T_j^2$  and an arbitrary vertex of  $T_j^4$ .

Let  $L_1$  be the number of leaves of the construction we described above and  $L_2$  be the number of degree two vertices which are part of heavy edges. We set  $B = L_1 + L_2 + km + k(m - 1)(n + 1) + m(2n + 6)$ .

► **Lemma 14.** *If  $G$  has a multi-colored independent set of size  $k$ , then the constructed instance has a stable cut of size at most  $B$ .*

► **Lemma 15.** *If the constructed instance has a stable cut of size at most  $B$ , then  $G$  has a multi-colored independent set of size  $k$ .*

► **Lemma 16.** *The constructed graph has pathwidth  $O(k)$ .*

## 6 Conclusions

Our results paint a clear picture of the complexity of MIN STABLE CUT with respect to  $\text{tw}$  and  $\Delta$ . As directions for further work one could consider stronger notions of stability such as demanding that switching sets of  $k$  vertices cannot increase the cut, for constant  $k$ . We conjecture that, since the structure of this problem has the form  $\exists\forall_k$ , its complexity with respect to treewidth will turn out to be double-exponential in  $k$  [47]. Another direction is to consider *hedonic games* where vertices self-partition into an unbounded number of groups. The complexity of finding a stable solution in such games parameterized by  $\text{tw} + \Delta$  has already been considered by Peters [50], whose algorithm runs in time exponential in  $\Delta^5 \text{tw}$ . Can we bridge the gap between this complexity and the  $2^{O(\Delta \text{tw})}$  complexity of MIN STABLE CUT?

---

### References

- 1 Pierre Aboulker, Édouard Bonnet, Eun Jung Kim, and Florian Sikora. Grundy coloring & friends, half-graphs, bicliques. In *STACS*, volume 154 of *LIPICs*, pages 58:1–58:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 2 Eric Angel, Evripidis Bampis, Bruno Escoffier, and Michael Lampis. Parameterized power vertex cover. *Discret. Math. Theor. Comput. Sci.*, 20(2), 2018. URL: <http://dmtcs.episciences.org/4873>.
- 3 Omer Angel, Sébastien Bubeck, Yuval Peres, and Fan Wei. Local max-cut in smoothed polynomial time. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 429–437. ACM, 2017. doi:10.1145/3055399.3055402.
- 4 Esther M. Arkin, Michael A. Bender, Joseph S. B. Mitchell, and Steven Skiena. The lazy bureaucrat scheduling problem. *Inf. Comput.*, 184(1):129–146, 2003.
- 5 Per Austrin, Mark Braverman, and Eden Chlamtac. Inapproximability of NP-complete variants of Nash equilibrium. *Theory Comput.*, 9:117–142, 2013. doi:10.4086/toc.2013.v009a003.
- 6 Baruch Awerbuch, Yossi Azar, Amir Epstein, Vahab S. Mirrokni, and Alexander Skopalik. Fast convergence to nearly optimal solutions in potential games. In Lance Fortnow, John Riedl, and Tuomas Sandholm, editors, *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, pages 264–273. ACM, 2008. doi:10.1145/1386790.1386832.
- 7 Maria-Florina Balcan, Avrim Blum, and Yishay Mansour. Improved equilibria via public service advertising. In Claire Mathieu, editor, *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4-6, 2009*, pages 728–737. SIAM, 2009. URL: <http://dl.acm.org/citation.cfm?id=1496770.1496850>.
- 8 C. Bazgan, L. Brankovic, K. Casel, H. Fernau, K. Jansen, K.-M. Klein, M. Lampis, M. Liedloff, J. Monnot, and V. T. Paschos. The many facets of upper domination. *Theoretical Computer Science*, 717:2–25, 2018.
- 9 Rémy Belmonte, Eun Jung Kim, Michael Lampis, Valia Mitsou, and Yota Otachi. Grundy distinguishes treewidth from pathwidth. In Fabrizio Grandoni, Grzegorz Herman, and Peter Sanders, editors, *28th Annual European Symposium on Algorithms, ESA 2020, September 7-9, 2020, Pisa, Italy (Virtual Conference)*, volume 173 of *LIPICs*, pages 14:1–14:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ESA.2020.14.
- 10 Rémy Belmonte, Michael Lampis, and Valia Mitsou. Parameterized (approximate) defective coloring. *SIAM J. Discret. Math.*, 34(2):1084–1106, 2020. doi:10.1137/18M1223666.
- 11 Anand Bhalgat, Tanmoy Chakraborty, and Sanjeev Khanna. Approximating pure Nash equilibrium in cut, party affiliation, and satisfiability games. In David C. Parkes, Chrysanthos Dellarocas, and Moshe Tennenholtz, editors, *Proceedings 11th ACM Conference on Electronic Commerce (EC-2010), Cambridge, Massachusetts, USA, June 7-11, 2010*, pages 73–82. ACM, 2010. doi:10.1145/1807342.1807353.

- 12 Ali Bibak, Charles Carlson, and Karthekeyan Chandrasekaran. Improving the smoothed complexity of FLIP for max cut problems. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 897–916. SIAM, 2019. doi:10.1137/1.9781611975482.55.
- 13 Vittorio Bilò and Marios Mavronicolas. The complexity of computational problems about Nash equilibria in symmetric win-lose games. *CoRR*, abs/1907.10468, 2019. arXiv:1907.10468.
- 14 Hans L. Bodlaender and Torben Hagerup. Parallel algorithms with optimal speedup for bounded treewidth. *SIAM J. Comput.*, 27(6):1725–1746, 1998.
- 15 É. Bonnet, M. Lampis, and V. T. Paschos. Time-approximation trade-offs for inapproximable problems. *Journal of Computer and System Sciences*, 92:171–180, 2018.
- 16 Mark Braverman, Young Kun-Ko, and Omri Weinstein. Approximating the best Nash equilibrium in  $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In Piotr Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 970–982. SIAM, 2015. doi:10.1137/1.9781611973730.66.
- 17 Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. Approximate pure Nash equilibria in weighted congestion games: Existence, efficient computation, and structure. *ACM Trans. Economics and Comput.*, 3(1):2:1–2:32, 2015. doi:10.1145/2614687.
- 18 Xi Chen, Chenghao Guo, Emmanouil-Vasileios Vlatakis-Gkaragkounis, Mihalis Yannakakis, and Xinzhi Zhang. Smoothed complexity of local max-cut and binary max-CSP. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing, STOC 2020, Chicago, IL, USA, June 22-26, 2020*, pages 1052–1065. ACM, 2020. doi:10.1145/3357713.3384325.
- 19 George Christodoulou, Vahab S. Mirrokni, and Anastasios Sidiropoulos. Convergence and approximation in potential games. *Theor. Comput. Sci.*, 438:13–27, 2012. doi:10.1016/j.tcs.2012.02.033.
- 20 Vincent Conitzer and Tuomas Sandholm. New complexity results about Nash equilibria. *Games Econ. Behav.*, 63(2):621–641, 2008. doi:10.1016/j.geb.2008.02.015.
- 21 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.
- 22 Argyrios Deligkas, John Fearnley, and Rahul Savani. Inapproximability results for constrained approximate Nash equilibria. *Inf. Comput.*, 262(Part):40–56, 2018. doi:10.1016/j.ic.2018.06.001.
- 23 Louis Dublois, Tesshu Hanaka, Mehdi Khosravian Ghadikolaei, Michael Lampis, and Nikolaos Melissinos. (in)approximability of maximum minimal FVS. In *ISAAC*, volume 181 of *LIPICs*, pages 3:1–3:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- 24 Louis Dublois, Michael Lampis, and Vangelis Th. Paschos. Upper dominating set: Tight algorithms for pathwidth and sub-exponential approximation. *CoRR*, abs/2101.07550, 2021. arXiv:2101.07550.
- 25 Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Nash equilibria in graphical games on trees revisited. In Joan Feigenbaum, John C.-I. Chuang, and David M. Pennock, editors, *Proceedings 7th ACM Conference on Electronic Commerce (EC-2006), Ann Arbor, Michigan, USA, June 11-15, 2006*, pages 100–109. ACM, 2006. doi:10.1145/1134707.1134719.
- 26 Edith Elkind, Leslie Ann Goldberg, and Paul W. Goldberg. Computing good nash equilibria in graphical games. In Jeffrey K. MacKie-Mason, David C. Parkes, and Paul Resnick, editors, *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*, pages 162–171. ACM, 2007. doi:10.1145/1250910.1250935.



- 27 Robert Elsässer and Tobias Tscheuschner. Settling the complexity of local max-cut (almost) completely. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *Automata, Languages and Programming - 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I*, volume 6755 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2011. doi:10.1007/978-3-642-22006-7\_15.
- 28 Hiroshi Eto, Tesshu Hanaka, Yasuaki Kobayashi, and Yusuke Kobayashi. Parameterized Algorithms for Maximum Cut with Connectivity Constraints. In *IPEC 2019*, pages 13:1–13:15, 2019.
- 29 Michael Etscheid and Heiko Röglin. Smoothed analysis of local search for the maximum-cut problem. *ACM Trans. Algorithms*, 13(2):25:1–25:12, 2017. doi:10.1145/3011870.
- 30 Alex Fabrikant, Christos H. Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing, Chicago, IL, USA, June 13-16, 2004*, pages 604–612. ACM, 2004. doi:10.1145/1007352.1007445.
- 31 Dimitris Fotakis, Vardis Kandiros, Thanasis Lianas, Nikos Mouzakis, Panagiotis Patsilinakos, and Stratis Skoulakis. Node-max-cut and the complexity of equilibrium in linear weighted congestion games. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPICs*, pages 50:1–50:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.ICALP.2020.50.
- 32 Dimitris Fotakis, Spyros C. Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. *Theor. Comput. Sci.*, 410(36):3305–3326, 2009. doi:10.1016/j.tcs.2008.01.004.
- 33 F. Furini, I. Ljubić, and M. Sinnl. An effective dynamic programming algorithm for the minimum-cost maximal knapsack packing problem. *European Journal of Operational Research*, 262(2):438–448, 2017.
- 34 Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- 35 Laurent Gourvès and Jérôme Monnot. On strong equilibria in the max cut game. In Stefano Leonardi, editor, *Internet and Network Economics, 5th International Workshop, WINE 2009, Rome, Italy, December 14-18, 2009. Proceedings*, volume 5929 of *Lecture Notes in Computer Science*, pages 608–615. Springer, 2009. doi:10.1007/978-3-642-10841-9\_62.
- 36 Laurent Gourvès, Jérôme Monnot, and Aris Pagourtzis. The lazy bureaucrat problem with common arrivals and deadlines: Approximation and mechanism design. In *FCT*, volume 8070 of *Lecture Notes in Computer Science*, pages 171–182. Springer, 2013.
- 37 Gianluigi Greco and Francesco Scarcello. On the complexity of constrained Nash equilibria in graphical games. *Theor. Comput. Sci.*, 410(38-40):3901–3924, 2009. doi:10.1016/j.tcs.2009.05.030.
- 38 Tesshu Hanaka, Hans L. Bodlaender, Tom C. van der Zanden, and Hirotaka Ono. On the maximum weight minimal separator. *Theoretical Computer Science*, 796:294–308, 2019.
- 39 Elad Hazan and Robert Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.*, 40(1):79–91, 2011. doi:10.1137/090766991.
- 40 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 41 Ken Iwaide and Hiroshi Nagamochi. An improved algorithm for parameterized edge dominating set problem. *J. Graph Algorithms Appl.*, 20(1):23–58, 2016.
- 42 David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988. doi:10.1016/0022-0000(88)90046-3.

- 43 Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structural parameters, tight bounds, and approximation for  $(k, r)$ -center. *Discret. Appl. Math.*, 264:90–117, 2019. doi:10.1016/j.dam.2018.11.002.
- 44 Ioannis Katsikarelis, Michael Lampis, and Vangelis Th. Paschos. Structurally parameterized  $d$ -scattered set. *Discrete Applied Mathematics*, 2020. doi:10.1016/j.dam.2020.03.052.
- 45 Kaveh Khoshkhan, Mehdi Khosravian Ghadikolaei, Jérôme Monnot, and Florian Sikora. Weighted upper edge cover: Complexity and approximability. *J. Graph Algorithms Appl.*, 24(2):65–88, 2020.
- 46 Michael Lampis. Parameterized approximation schemes using graph widths. In Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias, editors, *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, volume 8572 of *Lecture Notes in Computer Science*, pages 775–786. Springer, 2014. doi:10.1007/978-3-662-43948-7\_64.
- 47 Michael Lampis and Valia Mitsou. Treewidth with a quantifier alternation revisited. In Daniel Lokshantov and Naomi Nishimura, editors, *12th International Symposium on Parameterized and Exact Computation, IPEC 2017, September 6-8, 2017, Vienna, Austria*, volume 89 of *LIPICs*, pages 26:1–26:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPICs.IPEC.2017.26.
- 48 Martin Loeb. Efficient maximal cubic graph cuts (extended abstract). In Javier Leach Albert, Burkhard Monien, and Mario Rodríguez-Artalejo, editors, *Automata, Languages and Programming, 18th International Colloquium, ICALP91, Madrid, Spain, July 8-12, 1991, Proceedings*, volume 510 of *Lecture Notes in Computer Science*, pages 351–362. Springer, 1991. doi:10.1007/3-540-54233-7\_147.
- 49 Lorenz Minder and Dan Vilenchik. Small clique detection and approximate Nash equilibria. In Irit Dinur, Klaus Jansen, Joseph Naor, and José D. P. Rolim, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 12th International Workshop, APPROX 2009, and 13th International Workshop, RANDOM 2009, Berkeley, CA, USA, August 21-23, 2009. Proceedings*, volume 5687 of *Lecture Notes in Computer Science*, pages 673–685. Springer, 2009. doi:10.1007/978-3-642-03685-9\_50.
- 50 Dominik Peters. Graphical hedonic games of bounded treewidth. In Dale Schuurmans and Michael P. Wellman, editors, *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 586–593. AAAI Press, 2016. URL: <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12400>.
- 51 Svatopluk Poljak. Integer linear programs and local search for max-cut. *SIAM J. Comput.*, 24(4):822–839, 1995. doi:10.1137/S0097539793245350.
- 52 Alejandro A. Schäffer and Mihalis Yannakakis. Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87, 1991. doi:10.1137/0220004.
- 53 Grant Schoenebeck and Salil P. Vadhan. The computational complexity of Nash equilibria in concisely represented games. *ACM Trans. Comput. Theory*, 4(2):4:1–4:50, 2012. doi:10.1145/2189778.2189779.
- 54 M. Zehavi. Maximum minimal vertex cover parameterized by vertex cover. *SIAM Journal on Discrete Mathematics*, 31(4):2440–2456, 2017.