



Maximum Cut Parameterized by Crossing Number

*Markus Chimani*¹ *Christine Dahn*² *Martina Juhnke-Kubitzke*¹
*Nils M. Kriege*³ *Petra Mutzel*² *Alexander Nover*¹

¹School of Mathematics/Computer Science, University Osnabrück, Germany

²Institute for Computer Science, University of Bonn, Germany

³Department of Computer Science, TU Dortmund University, Germany

Abstract

Given an edge-weighted graph G on n nodes, the NP-hard MAX-CUT problem asks for a node bipartition such that the sum of edge weights joining the different partitions is maximized. We propose a fixed-parameter tractable algorithm parameterized by the number k of crossings in a given drawing of G . Our algorithm achieves a running time of $\mathcal{O}(2^k \cdot p(n+k))$, where p is the polynomial running time for planar MAX-CUT. The only previously known similar algorithm [8] is restricted to embedded 1-planar graphs (i.e., at most one crossing per edge) and its dependency on k is of order 3^k . Finally, combining this with the fact that crossing number is fixed-parameter tractable with respect to itself, we see that MAX-CUT is fixed-parameter tractable with respect to the crossing number, even without a given drawing. Moreover, the results naturally carry over to the minor-monotone-version of crossing number.

Submitted: August 2019	Reviewed: December 2019	Revised: January 2020	Accepted: February 2020
	Final: February 2020	Published: March 2020	
Article type: Regular paper		Communicated by: G. Liotta	

E-mail addresses: markus.chimani@uni-osnabrueck.de (Markus Chimani) christine.dahn@cs.uni-bonn.de (Christine Dahn) juhnke-kubitzke@uni-osnabrueck.de (Martina Juhnke-Kubitzke) nils.kriege@cs.tu-dortmund.de (Nils M. Kriege) petra.mutzel@cs.uni-bonn.de (Petra Mutzel) anover@uni-osnabrueck.de (Alexander Nover)

1 Introduction

Cut problems in graphs are a well-established class of problems attracting interest since the beginning of modern algorithmic research. Given an edge-weighted undirected graph, the MAX-CUT problem asks for a node partition into two sets, such that the sum of the weights of the edges between the partition sets is maximized. The problem is getting increasing attention in the literature due to its applicability to various scenarios: these range from ℓ^1 -embeddability [11], to the layout of electronic circuits [3, 9], to solving Ising spin glass models, which are of high interest in physics [1]. Besides the theoretical merits, such models need to be solved in adiabatic quantum computation [32]. Furthermore, quadratic unconstrained binary optimization (QUBO) problems can be solved via MAX-CUT. Many combinatorial optimization problems can be stated in the form of QUBO such as multicommodity-flow problems, maximum clique, vertex cover, scheduling, and many others. Also see [10, 11] for a more in-depth overview on applications.

The MAX-CUT problem has been shown to be NP-hard for general graphs [24]. Papadimitriou and Yannakakis [36] have shown that the MAX-CUT problem is even APX-hard, i.e., there does not exist a polynomial-time approximation scheme unless P=NP. Goemans and Williamson proposed a randomized constant-factor approximation algorithm [17], which has been derandomized by Mahajan and Ramesh [31], achieving a ratio of 0.87856. Several special cases of the problem allow polynomial algorithms: If the weights of all edges are negative we obtain a MIN-CUT problem, which can be solved, e.g., via network flow. Other special cases are, e.g., graphs without long odd cycles [20] or weakly bipartite graphs [19]. The case of planar input graphs is of particular interest. Orlova and Dorfman [34] and Hadlock [21] have shown how to solve MAX-CUT in polynomial time for unweighted planar graphs. Those algorithms can be extended to weighted planar graphs; the currently fastest algorithms for the weighted case have been suggested by Shih et al. [38] and by Liers and Pardella [30], and achieve a running time of $p(n) = \mathcal{O}(n^{3/2} \log n)$ on planar graphs with n nodes. Barahona has shown that the planarity condition can be relaxed to graphs not contractible to K_5 [2].

Similarly, it has been shown that MAX-CUT can be solved in polynomial time if the graph can be embedded on a surface of constant genus g [14, 15]. However, the edge-weights have to be restricted to integers whose absolute values are bounded by a polynomial in the input. The central idea of this algorithm is to write the generating function of cuts as a linear combination of 4^g Pfaffians. This algorithm is in fact highly non-trivial to realize: In order to obtain an implementable algorithm, [15] reports on a scheme, which depends on the existence of sufficiently many prime numbers within a given interval, which cannot be guaranteed in general.

A graph is *1-planar* if it allows a drawing where each edge is involved in at most one crossing. A *1-plane* graph is such a graph, together with an embedding realizing this property. The MAX-CUT problem on 1-plane graphs with k edge crossings has recently been shown to be *fixed-parameter tractable* (FPT)

with parameter k [8]. More precisely, it was shown that such instances can be solved in $\mathcal{O}(3^k \cdot p(n))$ time, where $p(n)$ is the running time of a polynomial-time MAX-CUT algorithm on planar graph with n nodes, e.g., [30,38]. There are no restrictions on the edge weights.

Our contribution. In this paper, we improve on the latter result in several ways: Firstly, we drop the requirement of 1-planarity, i.e., we consider graphs that can be drawn with at most k crossings (even if multiple such crossings lie on the same edge). We therefore handle the case of the well-established notion of the graph's *crossing number*. Secondly, we reduce the runtime dependency on k from 3^k to 2^k . Finally, unlike the previous result, our approach can be extended to an FPT algorithm which does not even require a crossing-realizing drawing as an input; however, this increases the running time and requires a deep algorithm from the literature as a black box [18,26]. Interestingly, we achieve these results by a *simpler* approach (compared to [8]). Comparing our algorithm with [15], we have no restrictions on the edge weights. Even in the restricted scenario, our algorithm is faster for graphs whose crossing number is at most twice its genus. Furthermore, we require only easily-implementable data structures and subalgorithms (if we are given a crossing-realizing drawing), compared to advanced methods from algebra.

The general idea of our algorithm is to recursively get rid of each crossing, each time resulting in two new subinstances. We end up with a set of up to 2^k planar graphs, each of which can be solved using a known polynomial-time MAX-CUT algorithm for planar graphs. The maximum over all these subinstances then yields a maximum cut in the original instance.

Finally, we consider parameterizing the problem by the *minor* crossing number (see below for details). This measure is always at most the graph's crossing number. While the exponential dependency on the respective parameter is identical, the running time only slightly increases in its polynomial part.

Independent work. Kobayashi et al. [28] independently and simultaneously obtained another fixed-parameter tractable algorithm for MAX-CUT parameterized by crossing number achieving the same running time. However, while we can always stay in the realm of maximum cuts when solving subinstances, they have to consider maximum weighted b -factor problems. Their preprint was uploaded to arXiv shortly after ours [6,27].

Organization of the paper. Section 2 recapitulates the basic definitions for cuts and crossings. In Section 3, we present our new algorithm and prove its correctness and running time. Section 4 extends the results to the minor crossing number case. We end with a conclusion and open problems in Section 5.

2 Preliminaries

Throughout this paper we consider undirected edge-weighted graphs. The input for our MAX-CUT problem is a graph $G = (V, E, c)$, where $c_e \in \mathbb{R}$ denotes the (positive or negative) weight of edge e , for each edge $e \in E$. A partition of the nodes V into two sets $S \subseteq V$ and $\bar{S} = V \setminus S$ defines the *cut* $\delta(S) = \{uv \in E \mid u \in S \Leftrightarrow v \notin S\}$. The *value* $c(\delta(S)) = \sum_{e \in \delta(S)} c_e$ of a cut is the sum of all edge weights of the edges in the cut. Given G , the MAX-CUT problem asks for a cut with highest value. Since a graph can have multiple cuts of equal value, only the value of a maximum cut is unique, not the cut itself.

A *non-degenerated drawing* of a graph in the plane is a map of its nodes to distinct points in \mathbb{R}^2 , and a map of its edges to curves connecting the respective endpoints, not including the points of any other node. Any point mapped to the plane either corresponds to a graph node, or is contained in at most two edge curves. A shared non-endpoint between two curves is called a *crossing*.

A graph is *planar* if it admits a drawing without any crossings. It is well known that planarity can be tested in linear time [23]. For non-planar graphs it is natural to ask for a drawing with as few crossings as possible. The smallest such number is the *crossing number* $\text{cr}(G)$ of G . Not only is it NP-hard to compute $\text{cr}(G)$ [16], but even the so called *realizability* problem turns out to be NP-hard [29]: Given a graph G and a set X of edge pairs, is there a drawing \mathcal{D} of G such that X contains an edge pair if and only if the pair's two edge curves cross in \mathcal{D} ? The key problem in testing realizability is that it is hard to figure out whether there exist *orderings* of the crossings along the respective edges that allow the above properties.

Therefore, sometimes more restricted crossing variants are considered. For example, *1-planar* graphs admit drawings where every edge is involved in at most one crossing. Not all graphs can be drawn in such a way, since 1-planar graphs can have at most $4|V| - 8$ edges; also, the 1-planar number of crossings is in general larger than $\text{cr}(G)$ [7, 35].

For a general drawing (not necessarily 1-plane), we typically encode its crossings as a *crossing configuration* \mathcal{X} . Therein, we not only store the pairs of edges that cross, but for each edge also the order of the crossings as they occur along its curve. The feasibility of a crossing configuration can be tested in time linear in $|V| + |\mathcal{X}|$ by replacing crossings with dummy nodes of degree 4, testing planarity, and checking the cyclic order around dummy nodes.¹ Although we will not require this fact in the following, this also allows us to efficiently deduce a drawing that respects \mathcal{X} . It is well understood that we can restrict ourselves to *good drawings* when considering the (traditional) crossing number of graphs: adjacent edges never cross and no edge pair crosses more than once.

¹In general, a specified edge pair may not really cross but merely “touch”; this is trivial to detect after testing planarity by checking the cyclic order of the edges around the dummy node. Given such a “flawed” configuration, we trivially obtain one with less crossings by removing such crossing pairs from \mathcal{X}

3 Algorithm

Our main idea for computing the maximum cut in an embedded weighted graph is to eliminate its crossings one by one. In the end, we use a MAX-CUT algorithm for planar graphs. We first introduce a slight variant of MAX-CUT:

Definition 1 (Partially-Fixed Maximum Cut, PF-MAX-CUT) *Given an edge weighted graph $G = (V, E, c)$ and a set of fixed edges $F \subseteq E$, find a cut of maximum value that contains all elements of F .*

A cut is *feasible* if it contains F . A PF-MAX-CUT instance is *infeasible* if it does not allow a feasible cut. It is easy to see that an instance is infeasible if and only if F contains a cycle of odd length. We denote a maximum objective value by $\text{MaxCut}_{\text{pf}}(G, F)$, and let $\text{MaxCut}_{\text{pf}}(G, F) = -\infty$ for infeasible instances.

Observe that (as for MAX-CUT) we do not need to consider a given crossing configuration \mathcal{X} as part of the problem description (see Corollary 7). However, since having \mathcal{X} allows for simplifications and a better running time, we will for now assume that we are given the graph together with a crossing configuration \mathcal{X} .² We will explain later how to remove this assumption in Corollary 7.

Given any edge vw with weight c_{vw} in a PF-MAX-CUT instance, we define the operation to *bisubdivide vw at v* as follows: Subdivide vw twice, i.e., replace vw by a path of length 3 with two new degree-2 nodes. We denote the new node incident to v or w by \bar{v} or \bar{w} , respectively. We consider the notation $\bar{\cdot}$ an operand.³

The edges $v\bar{v}$ and $\bar{v}\bar{w}$ have weight 0, $\bar{w}w$ retains the weight c_{vw} . Furthermore, we add $v\bar{v}, \bar{v}\bar{w}$ to F , and if $vw \in F$, we replace it in F by $\bar{w}w$. Clearly, both $v\bar{v}, \bar{v}\bar{w}$ will be in any feasible cut; node \bar{w} will always lie in the same partition set as v , and \bar{v} in the other (cf. Figure 1b). Most importantly this gives:

Lemma 2 *The feasible cuts in an original PF-MAX-CUT instance $\langle G, F \rangle$ are in 1-to-1 correspondence to feasible cuts of equal value in a bisubdivided instance $\langle G', F' \rangle$.*

Proof: Let vw be the edge in G that is bisubdivided at v to obtain $\langle G', F' \rangle$. By construction, we know that both edges $v\bar{v}, \bar{v}\bar{w}$ have cost 0 and are in F' , and thus in any F' -feasible cut. Consequently, in any F' -feasible cut, v and \bar{w} will lie in a common partition set. Let $S' \subset V(G')$ be a node subset that induces some feasible (with respect to F') cut in G' . Then, the node set $S = S' \setminus \{\bar{v}, \bar{w}\}$ induces a feasible (with respect to F) cut in G . Cut $\delta(S)$ contains edge vw if and only if $\delta(S')$ contains $\bar{w}w$. Since both these edges have identical cost, the total costs of both cuts are equal.

²As noted above, we may assume that no such \mathcal{X} ever specifies “touching points”; we can reduce such configurations whenever our algorithm retrieves a new crossing configuration.

³Observe that per recursion step, we will bisubdivide at most one edge per incident node (recall that adjacent edges never cross in good drawings). Thus, the above simple notation is unambiguous. In the graphs of the subproblems, see below, we may assume the nodes to be named afresh, and thus we may again perform bisubdivisions without creating notational ambiguity.

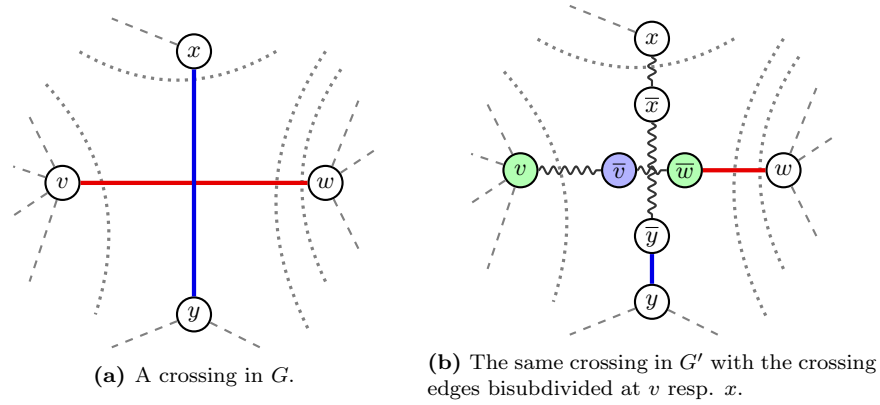


Figure 1: The situation at a crossing between vw and xy in G . In G' , the two edges of the crossing are bisubdivided at v and x , respectively, and the zig-zag edges are added to the set of fixed edges F' . As an example, the node coloring at v, \bar{v}, w gives a partition of these nodes that is forced by the respective newly added edges in F' . (Dashed and dotted edges show examples of other edges in G , resp. G' .)

Inversely, let $S \subset V(G)$ be a node subset that induces some feasible (with respect to F) cut in G . Then, consider the cut in G' induced by $S' = S \cup \{s\}$, where $s = \bar{w}$ if $v \in S$, and $s = \bar{v}$ otherwise. Both fixed edges $v\bar{v}, \bar{v}\bar{w}$ are in $\delta(S')$ and the cut is thus feasible. Again, $\delta(S')$ contains edge $\bar{w}w$ if and only if $\delta(S)$ contains vw , and both cut values are thus equal. \square

When we *identify* two nodes a, b in a graph with one another, they become a common entity that is incident to all of their former neighbors. We will only identify nodes that are neither adjacent nor share neighbors.

When identifying nodes in G of some PF-MAX-CUT instance $\langle G, F \rangle$, the set F is retained, subject to replacing the edges formerly incident to a or b with their new counterparts.

We are now ready to describe our algorithm. We are given a MAX-CUT instance $G = (V, E, c)$, together with some crossing configuration \mathcal{X} with k crossings. Let $F = \emptyset$ be the set of fixed edges and consider $\langle G, F \rangle$ as a PF-MAX-CUT instance. From $\langle G, F, \mathcal{X} \rangle$, we pick a crossed edge vw , and derive two new triplets $T_i = \langle G_i, F_i, \mathcal{X}_i \rangle$, for $i \in \{v, w\}$. Both derived crossing configurations \mathcal{X}_i attain at most $k - 1$ crossings and we can call our algorithm recursively on T_v and T_w . As a base case, the derived graphs become planar and (after a preprocessing to deal with the fixed edges) we apply an efficient MAX-CUT algorithm for planar graphs. The solutions of $\langle G_i, F_i \rangle$, for $i \in \{v, w\}$, yield a solution of $\langle G, F \rangle$. Observe, however, that $\langle G_i, F_i \rangle$ may become infeasible.

Let us describe this recursion step formally (cf. also Figures 1 and 2). We define the *crossing split* operation that, given a triplet $\langle G, F, \mathcal{X} \rangle$, yields the two

triplets T_v and T_w : Let $\langle G = (V, E, c), F \rangle$ be a PF-MAX-CUT instance and \mathcal{X} a crossing configuration of G . Consider a crossing $\chi \in \mathcal{X}$ with crossing edges vw and xy . For $j \in \{v, w, x, y\}$, let Y_j be the *ordered sets of crossings* in \mathcal{X} between j and χ (cf. the dotted edges in Figure 1: e.g., the crossings between the two dotted edges and $\bar{w}w$ in Figure 1b are in Y_w as they are between χ and w in Figure 1a). Let the intermediate instance $\langle G', F' \rangle$ be obtained from $\langle G, F \rangle$ by bisubdividing vw at v and bisubdividing xy at x . For $i \in \{v, w\}$, let $\langle G_i, F_i \rangle$ be the PF-MAX-CUT instance obtained from $\langle G', F' \rangle$ by identifying \bar{x} with \bar{i} (see Figures 2c and 2d). Intuitively, the two graphs obtained by the identifications represent the two possibilities whether x is on the same side of the cut as v or not. We obtain a corresponding crossing configuration \mathcal{X}_i from \mathcal{X} by removing χ and placing the crossings Y_j (retaining their order) on the edge $j\bar{j}$, for all $j \in \{v, w, x, y\}$. The triplets $T_v = \langle G_v, F_v, \mathcal{X}_v \rangle$ and $T_w = \langle G_w, F_w, \mathcal{X}_w \rangle$ are the results of the crossing split operation with respect to $\langle \chi, vw, xy \rangle$.

Lemma 3 *Let $\langle G = (V, E, c), F \rangle$ be a PF-MAX-CUT instance and \mathcal{X} a crossing configuration of G with k crossings. Let $\chi \in \mathcal{X}$ be any crossing with some crossing edges vw and xy , and consider the crossing split operation with respect to $\langle \chi, vw, xy \rangle$. For $i \in \{v, w\}$, let $\langle G_i, F_i, \mathcal{X}_i \rangle$ be the resulting triplets. Then we have:*

1. *for $i \in \{v, w\}$, \mathcal{X}_i is a feasible crossing configuration for G_i with at most $k - 1$ crossings; and*
2. $\text{MaxCut}_{\text{pf}}(G, F) = \max_{i \in \{v, w\}} \{ \text{MaxCut}_{\text{pf}}(G_i, F_i) \}$.

Proof: Consider any drawing \mathcal{D} of G realizing \mathcal{X} . By routing the new paths $(v\bar{v}, \bar{v}\bar{w}, \bar{w}w$ resp. $x\bar{x}, \bar{x}\bar{y}, \bar{y}y)$ along the curves of their original edges $(vw$ resp. $xy)$ we obtain a drawing \mathcal{D}' of G' from \mathcal{D} . Thereby, for $j \in \{v, w, x, y\}$, we place the new nodes \bar{j} in a close neighborhood of χ on the curve segment between j and χ , so that $\bar{x}\bar{y}$ is only crossed by $\bar{v}\bar{w}$ and vice versa. Note that the number of crossings in \mathcal{D}' is equal to that of \mathcal{D} , since all crossings in Y_j in \mathcal{D} are transferred to the edge $j\bar{j}$ in \mathcal{D}' , for all $j \in \{v, w, x, y\}$, and the original crossing χ between xy and vw in \mathcal{D} has a counterpart χ' in \mathcal{D}' between the edges $\bar{v}\bar{w}$ and $\bar{x}\bar{y}$. Since the edges $\bar{v}\bar{w}$ and $\bar{x}\bar{y}$ are crossing free except for χ' , we can follow (in a close neighborhood) the curves of $\bar{v}\bar{w}$ from any of its end points up to χ' , and onwards from there along the curve of $\bar{v}\bar{w}$ to any of its end points. Since these routes are crossings-free, we call them *free routes*. When we now identify \bar{x} with \bar{v} , we can locally redraw our drawing such that χ' vanishes and no other crossings arise, see Figure 2c. Observe that \bar{x} has precisely two neighbors: \bar{y} and x . The identification is thus such that we may remove \bar{x} and insert edges $\bar{y}\bar{v}$ and $x\bar{v}$ instead. The former can trivially be drawn without any crossings along the free route between \bar{y} and \bar{v} . The curve for the latter edge is the concatenation of the former curve of $x\bar{x}$ and the free route between \bar{x} and \bar{v} . The number of crossings along the edge $x\bar{x}$ (with now $\bar{x} = \bar{v}$) does thus not change. We can perform the analogous redrawing when identifying \bar{x} with \bar{w} , see Figure 2d. This establishes claim (1).

Two nodes v and x can either be on the same side of a cut, or they are on opposite sides. Therefore, we create two new subproblems in which v and x are in the same partition set or not, respectively. In G_v (where we identify \bar{x} with \bar{v}), we have a path of two edges between v and x (namely $v\bar{v}$ and $\bar{v}x$), both of which are in F_v . Thus, v and x have to be in the same partition set, see Figure 2c. Conversely, in G_w (where we identify \bar{x} with \bar{w}), we have a path of three edges between v and x (namely $v\bar{v}$, $\bar{v}\bar{w}$, and $\bar{w}x$), all of which are in F_w . Thus, v and x have to be in different partition sets, see Figure 2d. We can see that the respective constructions do not induce any further restrictions on the set of cuts. In particular, both derived instances still allow any partition choice between w and x , between w and y , and between x and y . Overall, every feasible cut in $\langle G', F' \rangle$ can be realized either in $\langle G_v, F_v \rangle$ or in $\langle G_w, F_w \rangle$.

If we know the maximum cut in instance $\langle G_v, F_v \rangle$ and the maximum cut in instance $\langle G_w, F_w \rangle$, we can pick the larger of these two cuts and transfer it back to $\langle G', F' \rangle$. By applying Lemma 2 twice (once for the bisubdivision of vw at v and once for the bisubdivision of xy at x), the maximum cut in $\langle G', F' \rangle$ induces a maximum cut in $\langle G, F \rangle$ of the same value. Claim (2) follows. \square

If we are in a base case – the considered graph is planar – we can use an efficient MAX-CUT algorithm for planar graphs:

Lemma 4 *Consider a PF-MAX-CUT instance $\langle G = (V, E, c), F \rangle$ with a planar graph G . Let $p(|V|)$ be a polynomial upper bound on the running time of a MAX-CUT algorithm on the planar graph G . We can compute an optimal solution to $\langle G, F \rangle$ – or decide that the instance is infeasible – in $\mathcal{O}(p(|V|))$.*

Proof: We transform the PF-MAX-CUT instance into a traditional MAX-CUT instance by attaching a large weight to the edges in F . Namely, we add M to the weight of each edge $f \in F$, where $M = 2 \cdot \sum_{e \in E} |c_e|$. The omission of a single edge of F from the solution cut (even if picking all other edges of positive weight) will already result in a worse objective value than picking all of F and all edges of negative weight. The instance is infeasible if and only if the computed cut does not contain all of F ; this can also be deduced purely by checking whether the objective value is at least $M \cdot |F| + \sum_{e \in E: c_e < 0} c_e$. \square

We proved our lemma above for a general case (by adding M to the weight of each edge in F), but in fact we only require a slightly weaker version, since in our algorithm $c_f = 0$ for all $f \in F$. Thus it suffices to set $c_f = M$ instead of adding M to c_f . Using any of the currently fastest MAX-CUT algorithms for planar graphs [30, 38] leads to $\mathcal{O}(|V|^{3/2} \log |V|)$ time in the above lemma. We could speed-up infeasibility detection by checking whether F contains a cycle of odd length prior to the transformation; while this only requires $\mathcal{O}(|V|)$ time via depth-first search, the overall asymptotic runtime for the lemma's claim does of course not improve.

Theorem 5 *Let $G = (V, E, c)$ be an edge-weighted graph and \mathcal{X} a crossing configuration of G with k crossings. Let $p(n)$ be a polynomial upper bound on*

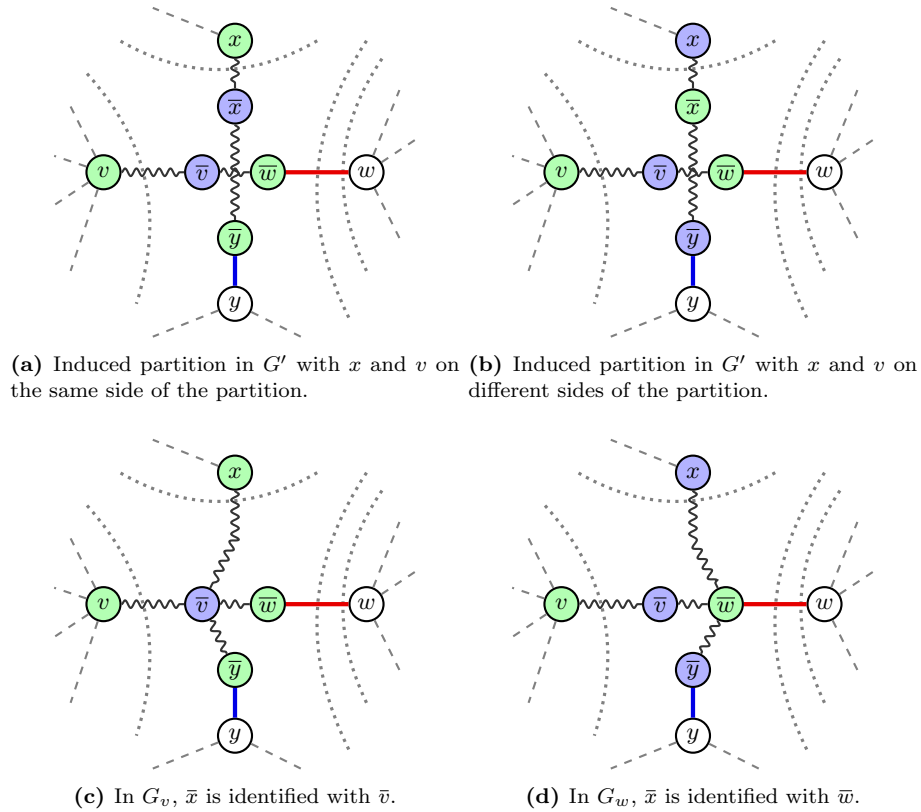


Figure 2: An illustration of the two cases where v and x are either on the same side of the partition (a/c) or on opposite sides (b/d). In the two graphs G_v and G_w , the crossing was removed while retaining the partition property. The node coloring gives a partition of the nodes that is induced by the newly added edges in F' , resp. F_v or F_w . (Dashed and dotted edges show examples of other edges in G' , resp. G_v or G_w .)

the running time of a MAX-CUT algorithm on planar graphs with n nodes. We can compute a maximum cut in G in $\mathcal{O}(2^k \cdot p(|V| + k))$ time.

Proof: As described above, we solve the instance by considering the PF-MAX-CUT instance $\langle G, F = \emptyset \rangle$ together with \mathcal{X} . Thus the triplet $\langle G, F, \mathcal{X} \rangle$ forms the initial input of our recursive algorithm \mathcal{R} .

Algorithm \mathcal{R} proceeds as follows on a given triplet: If the triplet's graph is planar, we solve $\langle G, F \rangle$ via Lemma 4. Otherwise, we use Lemma 3 to obtain two new input triples T_v, T_w , for each of which we call \mathcal{R} recursively. Their returned solutions (i.p., their solution values) induce the optimum solution for the current input triplet. However, while the number of crossings decreases by (at least) one per recursion step, the graph's size increases by three nodes.

The runtime complexity follows from the fact that we consider two choices per crossing in the given \mathcal{X} , and thus construct 2^k graphs. For each such graph,

which has $|V| + 3k$ nodes, we run the planar MAX-CUT algorithm. \square

Above, we trivially have $k \in \mathcal{O}(|V|^4)$ and thus $|V| + k \in \mathcal{O}(\text{poly}(|V|))$.

Corollary 6 *The above algorithm is an FPT algorithm with parameter k , provided that a crossing configuration \mathcal{X} with k crossings is part of the input. Moreover, the attained running time is polynomial for any $k \in \mathcal{O}(\log |V|)$. Using the currently fastest MAX-CUT algorithm for planar graphs [30, 38], our algorithm yields a running time of $\mathcal{O}(2^k \cdot (|V| + k)^{3/2} \log(|V| + k))$.*

Quite sophisticated results by Grohe [18] and Karabayashi and Reed [26] show that the problem to compute the crossing number of a graph is in FPT (even in linear time) with respect to its natural parameterization: Given a graph G and a number $k \in \mathbb{N}$, we can answer the question “ $\text{cr}(G) \leq k$?” in time $\mathcal{O}(f(k) \cdot n)$. In case of a yes-instance, we obtain a corresponding crossing configuration \mathcal{X} as a witness. The computable function $f(k)$ is purely dependent on k . However, the dependency $f(k)$ is double exponential, and the algorithm far from being practical. Still, these results formally allow us to get rid of the requirement that \mathcal{X} is part of the input:

Corollary 7 *Given an edge-weighted undirected graph G . Computing a maximum cut in G is FPT with parameter $\text{cr}(G)$.*

4 Minor Crossing Number

We say G is a *minor* of H , denoted by $G \preceq H$, if G can be obtained from H by deletion and contraction of edges. The *minor crossing number* of G is given by $\text{mcr}(G) = \min\{\text{cr}(H) \mid G \preceq H\}$. A *realization* of $\text{mcr}(G)$ is a pair (H, \mathcal{X}) with $G \preceq H$ and \mathcal{X} being a crossing configuration of H with $\text{mcr}(G)$ crossings. It is easy to see that for graphs G' of maximum degree 3 we have $\text{cr}(G') = \text{mcr}(G')$. Similarly, any graph G allows a realizing graph H ($\text{cr}(H) = \text{mcr}(G)$) of maximum degree 3 where vertices of G are replaced by disjoint cubic trees.

By definition we always have $\text{mcr}(G) \leq \text{cr}(G)$; as such $\text{mcr}(G)$ can be a stronger FPT-parameter. Also, in contrast to crossing number, the minor crossing number is monotone with respect to graph minors, i.e., the family $\{G \mid \text{mcr}(G) \leq k\}$ is minor closed. Thus, by [37], we can (theoretically) check whether $\text{mcr}(G) \leq k$ in $\mathcal{O}(|V(G)|^3)$ time for fixed $k \in \mathbb{N}$.

Given a connected graph G with $\text{mcr}(G) = k$, we can obtain a graph H from G realizing $\text{mcr}(G)$ in polynomial time as follows: Choose a node v of degree at least 4. Try different pairs of neighbors $w_1, w_2 \in N(v)$ until finding the first with $\text{mcr}(\tilde{G}) \leq k$, where \tilde{G} is obtained from G by splitting v into two nodes v_1 and v_2 with $N(v_1) = \{v_2, w_1, w_2\}$, $N(v_2) = (N(v) \cup \{v_1\}) \setminus \{w_1, w_2\}$ ⁴. We call the edge v_1v_2 a *split edge*. Iterating this for each high degree node, yields a graph H of maximum degree 3 realizing $\text{mcr}(G) = \text{cr}(H)$. Note that H has at most $\mathcal{O}(|E(G)|)$ nodes.

⁴Observe that in general this splitting operation may increase mcr ; we search for a split




Figure 3: Visualization of the split operation to obtain an mcr -realization. Left: part of a graph G with $\text{cr}(G) > \text{mcr}(G)$. Right: part of \tilde{G} after splitting v five times. Bold green lines denote *split edges*.

Let $N = -3 \cdot \sum_{e \in E(G)} |c_e|$. Attaching the weight N to each split edge, we can make sure that these edges are not in any maximum cut of H . Clearly, the cuts in H not containing any split edges are in one-to-one correspondence with cuts in G . Using Theorem 5, we obtain an algorithm computing a maximum cut on G parameterized by the $\text{mcr}(G)$. Similarly to Corollary 7 we do not require an explicit realization as part of the input (using the above construction method for H).

Corollary 8 (i) Let $G = (V, E, c)$ be an edge-weighted undirected graph with $\text{mcr}(G) = k$, (H, \mathcal{X}) a realization of $\text{mcr}(G)$, and $p(n)$ be a polynomial upper bound on the running time of a MAX-CUT algorithm on planar graphs with n nodes. We can compute a maximum cut in G in $\mathcal{O}(2^k \cdot p(|E(G)| + k))$ time.

(ii) Given an edge-weighted undirected graph G , computing a maximum cut in G is FPT with parameter $\text{mcr}(G)$.

5 Conclusion and Open Problems

Given a graph together with a feasible crossing configuration with k crossings, we previously only knew that MAX-CUT is polynomial time solvable if k is constant and the graph is 1-planar, i.e., each edge is involved in at most one crossing. The runtime dependency on k has been to the order of 3^k [8].

Herein, we improved on this in several ways: Firstly, we decreased the dependency on k to the order of 2^k . Secondly, we extended the applicability to any graph with (at most) k crossings: our parameter becomes the true crossing number of the graph, without any 1-planarity restriction. This shows that

(which has to exist) for which it does not increase. Since the split is an inverse minor operation, mcr can never decrease.

MAX-CUT is in FPT with respect to the graph's crossing number. Moreover, we achieve these improvements by introducing simpler ideas than those proposed for the former result, yielding an overall surprisingly simple algorithm. Compared to the result of Kobayashi et al. [28], we are able to stay within the realm of MAX-CUT. Finally, our result naturally carries over to the minor crossing number.

The *skewness* of a graph is the minimum number of edges to remove such that the graph becomes planar. The *genus* of a graph is the minimum oriented genus of a surface onto which the graph can be embedded without crossings. In FPT research, there are many algorithmic approaches that consider graphs with bounded genus g , see, e.g. [4, 5, 12, 13]. However, the obtained FPT algorithms are typically parameterized by the objective value z , or by the combined parameter (z, g) . There are much fewer results that obtain FPT algorithms parameterized purely with g . Notable examples are the graph genus problem itself [33] (where z and g coincide by definition), and the graph isomorphism problem [25] (which generalizes the linear-time algorithm for the problem on planar graphs). There are even fewer parameterized results with respect to skewness; the probably best known example is that maximum flow can be solved in the running time of planar graphs, if the graph's skewness is fixed [22]. Our above algorithm seems to be the first time that the crossing number has been proposed as an efficient non-trivial FPT parameter for any widely known problem.

Besides the weight-restricted case of [15] (briefly described in the introduction), it is unclear whether MAX-CUT could be FPT with respect to either skewness or genus. We deem this an interesting question for further research.

References

- [1] F. Barahona. On the computational complexity of Ising spin glass models. *Journal of Physics A: Mathematical and General*, 15(10):3241, 1982. doi:10.1088/0305-4470/15/10/028.
- [2] F. Barahona. The Max-Cut problem on graphs not contractible to K_5 . *Operations Research Letters*, 2(3):107–111, 1983. doi:10.1016/0167-6377(83)90016-0.
- [3] F. Barahona, M. Grötschel, M. Jünger, and G. Reinelt. An application of combinatorial optimization to statistical physics and circuit layout design. *Operations Research*, 36(3):493–513, 1988. doi:10.1287/opre.36.3.493.
- [4] H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (Meta) Kernelization. *Journal of the ACM*, 63(5):44:1–44:69, 2016. doi:10.1145/2973749.
- [5] J. Chen, I. A. Kanj, L. Perkovic, E. Sedgwick, and G. Xia. Genus characterizes the complexity of certain graph problems: Some tight results. *Journal of Computer and System Sciences*, 73(6):892–907, 2007. doi:10.1016/j.jcss.2006.11.001.
- [6] M. Chimani, C. Dahn, M. Juhnke-Kubitzke, N. M. Kriege, P. Mutzel, and A. Nover. Maximum cut parameterized by crossing number. *CoRR*, abs/1903.06061, 2019. URL: <http://arxiv.org/abs/1903.06061>.
- [7] M. Chimani, P. Kindermann, F. Montecchiani, and P. Valtr. Crossing numbers of beyond-planar graphs. In *Graph Drawing and Network Visualization - 27th International Symposium, GD 2019*, pages 78–86, 2019. doi:10.1007/978-3-030-35802-0_6.
- [8] C. Dahn, N. M. Kriege, and P. Mutzel. A fixed-parameter algorithm for the Max-Cut problem on embedded 1-planar graphs. In C. S. Iliopoulos, H. W. Leong, and W. Sung, editors, *Combinatorial Algorithms - 29th International Workshop, IWOCA 2018*, volume 10979 of *Lecture Notes in Computer Science*, pages 141–152. Springer, 2018. doi:10.1007/978-3-319-94667-2_12.
- [9] C. De Simone, M. Diehl, M. Jünger, P. Mutzel, G. Reinelt, and G. Rinaldi. Exact ground states of Ising spin glasses: New experimental results with a branch-and-cut algorithm. *Journal of Statistical Physics*, 80(1-2):487–496, 1995. doi:10.1007/BF02178370.
- [10] M. Deza and M. Laurent. Applications of cut polyhedra I. *J. Comput. Appl. Math.*, 55(2):191–216, 1994. doi:10.1016/0377-0427(94)90020-5.
- [11] M. Deza and M. Laurent. Applications of cut polyhedra II. *J. Comput. Appl. Math.*, 55(2):217–247, 1994. doi:10.1016/0377-0427(94)90021-3.

- [12] J. A. Ellis, H. Fan, and M. R. Fellows. The dominating set problem is fixed parameter tractable for graphs of bounded genus. *Journal of Algorithms*, 52(2):152–168, 2004. doi:10.1016/j.jalgor.2004.02.001.
- [13] F. V. Fomin, D. Lokshtanov, V. Raman, and S. Saurabh. Bidimensionality and EPTAS. In D. Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011*, pages 748–759. SIAM, 2011. doi:10.1137/1.9781611973082.59.
- [14] A. Galluccio and M. Loeb. On the theory of Pfaffian orientations. II. T-joins, k -cuts, and duality of enumeration. *Electronic Journal of Combinatorics*, 6, 1998.
- [15] A. Galluccio, M. Loeb, and J. Vondrak. Optimization via enumeration: a new algorithm for the max cut problem. *Mathematical Programming*, 90:273–290, 2001. doi:10.1007/PL00011425.
- [16] M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 4(3):312–316, 1983. doi:10.1137/0604033.
- [17] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995. doi:10.1145/227683.227684.
- [18] M. Grohe. Computing crossing numbers in quadratic time. *Journal of Computer and System Sciences*, 68(2):285–302, 2004. doi:10.1016/j.jcss.2003.07.008.
- [19] M. Grötschel and W. R. Pulleyblank. Weakly bipartite graphs and the Max-Cut problem. *Operations Research Letters*, 1(1):23–27, 1981. doi:10.1016/0167-6377(81)90020-1.
- [20] M. Grötschel and G. L. Nemhauser. A polynomial algorithm for the Max-Cut problem on graphs without long odd cycles. *Mathematical Programming*, 29(1):28–40, 1984. doi:10.1007/BF02591727.
- [21] F. Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4(3):221–225, 1975. doi:10.1137/0204019.
- [22] J. M. Hochstein and K. Weihe. Maximum s - t -flow with k crossings in $O(k^3 n \log n)$ time. In N. Bansal, K. Pruhs, and C. Stein, editors, *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007*, pages 843–847. SIAM, 2007. doi:10.1145/1283383.1283473.
- [23] J. E. Hopcroft and R. E. Tarjan. Dividing a graph into triconnected components. *SIAM Journal on Computing*, 2(3):135–158, 1973. doi:10.1137/0202012.

- [24] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Proceedings of a symposium on the Complexity of Computer Computations, New York*, The IBM Research Symposia Series, pages 85–103. Plenum Press, New York, 1972. URL: <http://www.cs.berkeley.edu/~luca/cs172/karp.pdf>.
- [25] K. Kawarabayashi. Graph isomorphism for bounded genus graphs in linear time. *arXiv.org*, abs/1511.02460, 2015. URL: <http://arxiv.org/abs/1511.02460>.
- [26] K. Kawarabayashi and B. A. Reed. Computing crossing number in linear time. In D. S. Johnson and U. Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, STOC 2007*, pages 382–390. ACM, 2007. doi:10.1145/1250790.1250848.
- [27] Y. Kobayashi, Y. Kobayashi, S. Miyazaki, and S. Tamaki. An FPT algorithm for Max-Cut parameterized by crossing number. *CoRR*, abs/1904.05011, 2019. URL: <http://arxiv.org/abs/1904.05011>.
- [28] Y. Kobayashi, Y. Kobayashi, S. Miyazaki, and S. Tamaki. An improved fixed-parameter algorithm for Max-Cut parameterized by crossing number. In *Combinatorial Algorithms - 30th International Workshop, IWOCA 2019*, pages 327–338, 2019. doi:10.1007/978-3-030-25005-8_27.
- [29] J. Kratochvíl. String graphs. II. recognizing string graphs is NP-hard. *Journal of Combinatorial Theory, Series B*, 52(1):67 – 78, 1991. doi:10.1016/0095-8956(91)90091-W.
- [30] F. Liers and G. Pardella. Partitioning planar graphs: a fast combinatorial approach for Max-Cut. *Computational Optimization and Applications*, 51(1):323–344, 2012. doi:10.1007/s10589-010-9335-5.
- [31] S. Mahajan and H. Ramesh. Derandomizing semidefinite programming based approximation algorithms. *SIAM Journal on Computing*, 28(5):1641–1663, 1995. doi:10.1109/SFCS.1995.492473.
- [32] C. C. McGeoch. *Adiabatic Quantum Computation and Quantum Annealing: Theory and Practice*. Synthesis Lectures on Quantum Computing. Morgan & Claypool Publishers, 2014. doi:10.2200/S00585ED1V01Y201407QMC008.
- [33] B. Mohar. A linear time algorithm for embedding graphs in an arbitrary surface. *SIAM J. Discrete Math.*, 12:6–26, 1999.
- [34] G. I. Orlova and Y. G. Dorfman. Finding maximum cut in a graph. *Engineering Cybernetics*, 10(3):502–506, 1972.
- [35] J. Pach and G. Tóth. Graphs drawn with few crossings per edge. *Combinatorica*, 17(3):427–439, 1997. doi:10.1007/BF01215922.

- [36] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43(3):425 – 440, 1991. doi:10.1016/0022-0000(91)90023-X.
- [37] N. Robertson and P. Seymour. Graph minors. XIII. the disjoint paths problem. *Journal of Combinatorial Theory, Series B*, 63(1):65 – 110, 1995. doi:10.1006/jctb.1995.1006.
- [38] W.-K. Shih, S. Wu, and Y.-S. Kuo. Unifying maximum cut and minimum cut of a planar graph. *IEEE Transactions on Computers*, 39(5):694–697, 1990. doi:10.1109/12.53581.