

A Near-Linear Approximation Scheme for Multicuts of Embedded Graphs with a Fixed Number of Terminals*

Vincent Cohen-Addad[†] Éric Colin de Verdière[‡] Arnaud de Mesmay[§]

April 10, 2017

Abstract

For an undirected edge-weighted graph G and a set R of pairs of vertices called pairs of *terminals*, a multicut is a set of edges such that removing these edges from G disconnects each pair in R . We provide an algorithm computing a $(1 + \varepsilon)$ -approximation of the minimum multicut of a graph G in time $(g + t)^{O(g+t)^3} \cdot (1/\varepsilon)^{O(g+t)} \cdot n \log n$, where g is the genus of G and t is the number of terminals.

This result is tight in several aspects, as the minimum multicut problem is both APX-hard and W[1]-hard (parameterized by the number of terminals), even on planar graphs (equivalently, when $g = 0$).

In order to achieve this, our article leverages on a novel characterization of a minimum multicut as a family of Steiner trees in the universal cover of a surface on which G is embedded. The algorithm heavily relies on topological techniques, and in particular on the use of homotopical tools and computations in covering spaces, which we blend with classic ideas stemming from approximation schemes for planar graphs and low-dimensional geometric inputs.

*This work is supported by the French ANR Blanc projects ANR-12-BS02-005 (RDAM) and ANR-16-CE40-0009-01 (GATO).

[†]University of Copenhagen, Denmark. Work partly done while the author was at Département d'Informatique, École normale supérieure, Paris. Email: vincent.cohen.addad@ens-lyon.org

[‡]CNRS, LIGM, Université Paris-Est Marne-la-Vallée, France. Work partly done while the author was at Département d'Informatique, École normale supérieure, Paris. Email: eric.colindeverdiere@u-pem.fr.

[§]CNRS, Gipsa-lab, France. Email: arnaud.de-mesmay@gipsa-lab.fr

1 Introduction

Cuts and flows are fundamental objects in combinatorial optimization and have generated a very deep theory. One classical problem is MULTICUT: Let $G = (V, E)$ be an undirected graph, let T be a subset of vertices of G , called *terminals*, and let R be a set of unordered pairs of vertices in T , called *terminal pairs*. A subset E' of E is a *multicut* (with respect to (T, R)) if for every terminal pair $\{t_1, t_2\} \in R$, the vertices t_1 and t_2 lie in different connected components of the graph $(V, E \setminus E')$. Assuming G is positively edge-weighted, and given G, T , and R , the MULTICUT problem asks for a multicut of minimum weight.

The case $|R| = 1$ is the famous minimum cut problem and admits a polynomial-time algorithm. The case $|R| = 2$ can also be solved in polynomial time [36], but the case where R is arbitrary is much harder, and in fact even NP-hard to approximate within any constant factor assuming the Unique Game Conjecture [13]. Thus, to circumvent this hardness, a fruitful approach is to focus on specific inputs that characterize the instances that are relevant in practice.

The study of cut and flow problems originates from practical problems involving road or railway networks [49]. Since these can usually be modeled by embedded graphs, there exists an important literature towards the efficient computation of minimum cuts (or maximal flows) on planar graphs [23, 32, 33, 40, 48] (among other references) or, more generally, graphs of bounded genus [6, 8, 12, 10, 24, 26]. However, MULTICUT remains APX-hard for planar graphs [30], and so to obtain an exact or $(1 + \varepsilon)$ -approximation algorithm it is needed to consider parameterized algorithms.

A very natural parameter is the number $t = |T|$ of terminals; the case $t = 2$, corresponding to the minimum cut, has inspired a huge literature, and the problem is already NP-hard for $t = 3$ in general [19]. Unfortunately, MULTICUT is W[1]-hard when parameterized by the number of terminals, even for planar graphs [45].

Therefore, for planar graphs, the best possible result for MULTICUT parameterized by the number of terminals is a $(1 + \varepsilon)$ -approximation. We provide the first algorithm achieving such an approximation guarantee, for planar and more generally for surface-embedded graphs. Moreover, the dependence in the size of the graph is near-linear and the dependence in ε is polynomial:

Theorem 1.1. *Let G be an undirected, positively edge-weighted graph embeddable on a surface of genus g , orientable or not. Let n be the number of vertices and edges of G . Let T be a set of t terminals and R be a set of unordered pairs of T . Then for every $\varepsilon > 0$, we can compute a $(1 + \varepsilon)$ -approximation of the minimum multicut of G with respect to (T, R) in time $f(\varepsilon, g, t) \cdot n \log n$, where $f(\varepsilon, g, t) = (g + t)^{O(g+t)^3} \cdot (1/\varepsilon)^{O(g+t)}$.*

As hinted at before, we argue that our result is, from various points of view, the best possible. Indeed, through its multiple parameters, it finds a very intricate middle ground amidst the flurry of hardness results that have been established for the MULTICUT problem. More precisely:

1. **Dependence on the genus.** For general graphs, MULTICUT is known to be hard to approximate, even for fixed values of t : Dahlhaus et al. [19, Theorem 5] have proved that it is APX-hard for any fixed $t \geq 3$. Therefore (unless P=NP), the exponential dependence on the genus of the graph cannot be improved to a polynomial one, no matter the dependence on the number of terminals, as this would yield a $(1 + \varepsilon)$ -approximation for general graphs (thus contradicting Dahlhaus et al.).
2. **Dependence on the number of terminals.** MULTICUT is known to be hard to approximate on planar graphs: Garg et al. [30] have proved that it is APX-hard on unweighted stars. Therefore, the exponential dependence on the number of terminals cannot be improved to

a polynomial one, no matter the dependence on the genus, as this would yield a PTAS for planar graphs.

3. **Approximation factor.** MULTICUT is hard to solve exactly for planar graphs, even from the point of view of parameterized complexity: Marx [45] has shown that it is W[1]-hard for planar graphs when parameterized by the number of terminals. Actually, he has proved a lower bound of $n^{\Omega(\sqrt{t})}$ assuming the Exponential Time Hypothesis. Therefore, there is no exact FPT algorithm parameterized by t unless W[1]=FPT (no matter the dependence in the genus).
4. **Dependence on the approximation factor.** The previous observation can be refined in the following way, using a parameterized version of the classical argument showing that “integer-valued” strongly NP-complete problems do not admit an FPTAS. Since the reduction of Marx [45] also applies to unweighted planar graphs, it also precludes an algorithm with running time $f(1/\varepsilon)h(t)p(n)$ for some computable function h and polynomials f and p , which are even allowed to depend arbitrarily on g . indeed, such an algorithm would yield an *exact* FPT algorithm for MULTICUT on unweighted planar graphs by setting $\varepsilon = 1/\Omega(n)$, since a solution of an unweighted MULTICUT instance has value $O(n)$.

Our result lies at the interface of parameterized complexity and approximation algorithms—we refer to the article of Marx [44] for a survey of the results of this type. Within the language of that survey, our algorithm is a fixed parameter tractable approximation scheme (fpt-AS) with respect to its parameters (the genus, the number of terminals, and the approximation factor). Furthermore, we highlight that the dependence on $1/\varepsilon$ is polynomial here. In particular, our algorithm is a near-linear Fully Polynomial Time Approximation Scheme (FPTAS) if the number of terminals and the genus are constant. This is a rare occurrence for approximation schemes: Indeed, as mentioned above, in the unparameterized world, such an approximation scheme cannot exist for “integer-valued” strongly NP-complete problems. Thus, we can not rely directly on the successful techniques designed for obtaining PTASes for planar and embedded graphs, like for example “brick decompositions” (*e.g.*, [3, 4, 5, 7]), as they usually lead to algorithms that have exponential dependence in ε . Hence to obtain a polynomial dependence in ε , we develop a new, specific toolset.

Related work. In addition to the hardness proofs mentioned above, there has also been a great body of work towards exact or approximate algorithms for the MULTICUT problem. The most relevant to this article is a recent algorithm [17] by the second author of this paper which solves the exact version of MULTICUT in time $O((g+t)^{O(g+t)}n^{O(\sqrt{g^2+gt})})$. There also exist constant-factor approximations to solve MULTICUT on graphs of bounded genus (or excluding some fixed minor) using techniques based on padded decompositions (see *e.g.*, [1, 28, 41]), and Garg et al. [29] have provided a $O(\log t)$ -approximation for general graphs. Regarding parameterized algorithms, there has been a recent flurry of results, culminating into two simultaneous proofs [9, 46] that unweighted MULTICUT is FPT when parameterized by the *cost of the solution*. We emphasize that this result is not comparable to ours due to the difference in the choice of parameters.

A special instance of this problem that has been more studied is the MULTIWAY CUT problem (also known as MULTITERMINAL CUT), where the set R of pairs of terminals comprises all the pairs of vertices in the set T of terminals. All the hardness results mentioned above for MULTICUT were actually proved for MULTIWAY CUT, with one exception: The planar version of MULTIWAY CUT is not APX-hard, and indeed Bateni et al. [3] have provided a PTAS for this planar problem. Regarding exact algorithms, Klein and Marx [42] have shown how to solve planar MULTIWAY CUT exactly in time $2^{O(t)} \cdot n^{O(\sqrt{t})}$, coming very close to the lower bound of $n^{\Omega(\sqrt{t})}$ mentioned above. Chekuri and Madan [14] have recently investigated the complexity of approximating MULTICUT

with looser constraints on the terminal pairs, in both the undirected and the directed setting.

Outline of the algorithm. The previous works on cuts and flows in embedded graphs have revealed a strong relationship between these a priori purely combinatorial constructs and some *topological* aspects of the underlying surface. For example, for embedded graphs, *homology* is a key tool in the study of minimum cuts [10, 26]. Similarly, the starting point of our algorithm is the observation that the MULTICUT problem for surface-embedded graphs is inherently a topological problem. In the same way that minimal cycles are dual to minimal cuts in planar graphs, the edges dual to a minimum multicut form in some sense the “shortest dual graph” which separates topologically the required sets of terminals. While this topological point of view already underlies the early algorithms for planar MULTIWAY CUT [19], it was only recently made precise by the second author of this article [17], based on the framework of *cross-metric surfaces* [18]. He also proved that the topology of such an optimal *multicut dual* can be quantified appropriately by bounding its number of intersections with a certain *cut graph* of the surface (equivalently in this paper, a *system of arcs*). The algorithm of [17] for solving MULTICUT exactly guesses the topology (i.e., the *homotopy class*) of the edges of the optimal multicut together with the position of the vertices.

In order to obtain a fixed-parameter dependence on the genus and the number of terminals, we mix the approach described above with techniques stemming from the design of approximation schemes in planar graphs and low-dimensional geometric inputs. Our main algorithm computes a subgraph (called *skeleton*) of controlled length (w.r.t. OPT) that intersects a “structured” near-optimal solution (i.e., a multicut of cost at most $(1 + \varepsilon)\text{OPT}$). In a nutshell, the main structural property of the near-optimal solution is that it is made (roughly) acyclic after cutting along the intersections with the skeleton. The crux of the paper is to show how to compute the skeleton efficiently and to prove that such a structured near-optimal solution exists. Both these steps are rather involved technically, and they rely extensively on the use of *annular covers* of the surface. We defer their overview to Section 3.4, once we have introduced the key notion of a *good multicut dual* and the relevant topological concepts.

This structural result allows us to make use of the “portal approach” that has been developed for obtaining approximation schemes for embedded graphs. We show that the structured near-optimal solution can be modified slightly to intersect the skeleton at a bounded-size set of vertices called *portals*. Then, the rest of the algorithm consists in “guessing” both the topology of the near-optimal solution and the portals that intersect with it, and computing a family of Steiner trees linking these portals in some portion of the universal cover of the surface (which encodes the guessed topology). These Steiner trees can be computed using the fixed-parameter tractable algorithm of Dreyfus and Wagner [21, 27].

Throughout the proof, we use several techniques which are somewhat at odds with the setting of embedded graphs. For example, computations in some covering space, when projected back to the surface, may involve undesired crossings; and it is more natural to have a near-optimal solution take detours after the placement of portals if it is allowed to cross itself. This is why many of our arguments deal with *drawings* of graphs instead, which allow a graph to intersect itself. One of the perks of our approach is that since being a multicut dual is a homotopical property, most of our techniques dovetail naturally with this change of perspective, and the setting of cross-metric surfaces also allows for a seamless algorithmic treatment of drawings of graphs.

We also observe that the topological tools involved in the proof of our main result are essentially the same for planar and embedded graphs; neither the algorithm nor its proof of correctness get significantly simpler by considering only planar graphs.

2 Preliminaries

2.1 Graphs, surfaces, and homotopy

We recall here standard definitions on the topology of surfaces. For general background on topology, see for example Stillwell [50] or Hatcher [34]. For more specific background on the topology of surfaces in our context, see recent articles and surveys on the same topic [18, 16].

Let S be a (compact, connected) *surface* without boundary. If S is an orientable surface, its *genus* $g \geq 0$ is even, and S is (homeomorphic to) a sphere with $g/2$ handles attached. If S is a non-orientable surface, its *genus* is a positive integer g , and S is a sphere with $g \geq 1$ disjoint disks replaced by Möbius strips. A *surface with boundary* is obtained from a surface without boundary by removing a set of disks with disjoint closures.

A *path* is a continuous map from $[0, 1]$ to S . Two paths are *homotopic* if there is a continuous deformation between them on the surface that keeps the endpoints fixed. An *arc* is a path with endpoints on the boundary of the surface. A *closed curve* is a continuous map from the unit circle to S . Two closed curves are (freely) *homotopic* if there is a continuous deformation between them. A path or closed curve is *simple* if it is one-to-one.

A closed curve γ is *contractible* if it is homotopic to a constant closed curve. If γ is simple, a small neighborhood of the image of γ is either an annulus or a Möbius strip; in the former case, γ is *two-sided*, and otherwise it is *one-sided*. Of course, one-sided curves exist only if the surface is non-orientable¹.

We consider *drawings* of graphs C on the surface S , which are assumed to be in general position: There are finitely many (self-)intersection points of the drawing, each involving exactly two edges, and these two edges actually cross at that point, in their relative interior. (A similar definition holds for closed curves.) We frequently abuse terminology and identify the abstract graph C with its drawing on S . An *embedding* of the graph C is an injective (“crossing-free”) drawing. An embedding is *cellular* if its faces are disks.

Let Γ be a set of closed curves in general position on a surface S . A *monogon* is a disk on S whose interior does not meet Γ and whose boundary is formed by a subpath of a curve in Γ . A *bigon* on S is a disk on S whose interior does not meet Γ and whose boundary is formed by the concatenation of two subpaths of curves in γ (possibly the same curve) such that these two subpaths are disjoint except at their endpoints.

If C is drawn on S with crossings, there is an ambiguity when talking about *cycles* of C . We adopt the convention that a *cycle* in C is a cycle in C in the graph-theoretical sense (a closed walk, not reduced to a single vertex, without vertex repetition), or by abuse of language the image of that cycle. Thus a cycle in C is a closed curve on S (which may cross itself).

The input of our algorithm is a graph that can be embedded on a surface, orientable or not, with genus g . There is an algorithm that takes as input a graph G with n vertices and edges, and a surface S specified by its genus g and by whether it is orientable, and decides in $2^{O(g)}n$ whether G embeds on S . Moreover, in the affirmative, it computes a cellular embedding of G on a surface S' with the same orientability as S and with genus at most g [39]. We can afford using this algorithm and assume w.l.o.g. that G is cellularly embedded on a surface of genus (at most) g . Actually, it is useful to consider the surface with boundary obtained by removing small disks around each terminal; this operation transforms every terminal vertex v of G into $\deg(v)$ one-degree vertices on the boundary of S . Henceforth, S is a surface with genus g and t boundary components. Each

¹The reader not familiar with topology may safely skip all considerations involving non-orientable surfaces and one-sided curves. If the input graph is planar, then one can even restrict oneself to a surface that is a sphere with disjoint disks removed.

boundary component of S corresponds to a terminal, and we can thus view the set of terminal pairs R as a set of pairs of boundary components to be separated.

2.2 Cross-metric surfaces, systems of arcs, and topology

A *cross-metric surface* (S', G') is a surface S' (possibly with boundary) equipped with a positively edge-weighted graph G' cellularly embedded on S' [18, Section 1.2]. Curves (paths and closed curves) on (S', G') are assumed to be in general position with G' . The *length* of a curve c , sometimes denoted by $|c|$, is the sum of the weights of the edges of G' crossed by the curve, counted with multiplicity. This provides a discrete notion of metric for S' . In this paper, the cross-metric surface we frequently use is (S, G) , as defined in the previous paragraph. When we use terms such as “shortest” or “distance”, this implicitly refers to this notion of distance.

As in previous papers [17, Proposition 4.1] (see also, e.g., Chambers et al. [11]), we will use an algorithm to compute a *greedy system of arcs* K :

Proposition 2.1. *In $O(n \log n + (g + t)n)$ time, we can compute a set K of $O(g + t)$ disjoint, simple arcs on S in general position with respect to G , such that $S \setminus K$ is an open disk. Moreover, each arc is the concatenation of two shortest paths on S , and is a shortest homotopic path on S .*

We note that the presentation in [17] is slightly different because the construction takes place in the surface without boundary, but the result is the same. When we write that the arc is the concatenation of two shortest paths uv and vw , here it could be that the point v lies on an edge of G , not taken into account for computing the lengths of paths uv and vw . Also, the fact that each edge is a shortest homotopic path is folklore and follows, e.g., from the analysis in [15] and the fact that two homotopic paths are homologous.

As an illustrating special case for the above proposition, assume that S has genus zero (a sphere with disjoint open disks removed). If we contract each boundary component to a point, then K corresponds to a tree spanning the terminals. Thus, the set K can be computed using, e.g., a minimum spanning tree algorithm, slightly adapted to the cross-metric setting and to accommodate for the boundary components instead of the terminals.

Henceforth, all graphs are drawn in general position with respect to $K \cup G$, and the *complexity* of a graph drawn on S is the sum of its number of vertices, edges, and the number of intersections with G (this coincides with the sum of the lengths of the edges if G is unweighted).

Let C be a graph embedded on S and in general position with K . Since K cuts S into a disk, we can represent C by its image in the disk. More precisely, we define the *topology* of C to be the information of the combinatorial map of the overlay of C and K .

Finally, one notation: given a closed curve γ , we denote by γ^o a shortest closed curve in its homotopy class.

2.3 Covering spaces and shortest homotopic curves

Our algorithm reuses arguments used to compute shortest homotopic paths and closed curves on the surface with boundary S [18, Sections 1 and 6]. Since the setting we work on is both simpler than in that article (because S has boundaries and we do not need the best possible complexity with respect to g and t), but we extend a bit these techniques (to deal with the non-orientable surfaces), we find it worthwhile to summarize and explain these tools in a way that is tailored to our purpose; see Appendix A. The reader unfamiliar with covering spaces will also find the relevant definitions there.

3 Multicut duals and overview

3.1 Multicut duals

A *multicut dual* C is a graph drawn (possibly with crossings) on S such that every path in G connecting two boundaries of S corresponding to a terminal pair $\{t_1, t_2\} \in R$ meets the image of C . Thus, the set of edges crossed by C forms a multicut. Conversely, any multicut corresponds to a multicut dual of the same cost; thus, it suffices to compute a shortest multicut dual (see also [17, Proposition 3.1]). Throughout this paper, we focus on the problem of computing a shortest multicut dual.

Lemma 3.1 ([17, Section 5, Propositions 5.1 and 6.1]). *Some shortest multicut dual is an embedded graph with $O(g+t)$ vertices and edges in general position with respect to $K \cup G$, with $O(g+t)$ crossings with each edge of K . Moreover:*

- *the topology of that multicut dual belongs to a set of $(g+t)^{O(g+t)}$ topologies of multicut duals, which we can enumerate within the same time bound;*
- *each graph in such a topology is inclusionwise minimal, in the sense that removing any vertex or edge would violate the fact that it is a multicut dual.*

In particular, by the last condition, a graph with such a topology has no vertex of degree zero or one, and each face contains at least one terminal. A *candidate topology* is a topology of a subgraph of one of those output by this lemma that also has no vertex of degree zero or one. Obviously, every face of a graph with a candidate topology also contains at least one terminal, and one can also enumerate all the candidate topologies in time $(g+t)^{O(g+t)}$.

We let C_{OPT} be a shortest multicut dual satisfying the conditions of Lemma 3.1, and let OPT be its length.

3.2 Multicut dual property

We now introduce a sufficient condition for a graph drawn on S to be a multicut dual, which will be used throughout this paper. Let C be a graph drawn on S . We say that C is an *even graph* if every vertex has even degree. We will work with *homology* over the field \mathbb{Z}_2 , which is an equivalence relation between even graphs that is coarser than homotopy; see, e.g., [10, 16].

We remark that the property of being a multicut dual is not preserved by homological transformations, contrary to, e.g., the dual of a minimum cut [10]. For example, a shortest multicut dual needs not be an even graph; moreover, a pair of terminals is allowed to be separated by a positive even number of edges. However, homology can still be used in this realm, by also looking at subgraphs of the multicut dual: We say that C has the *multicut dual property* if, for every even subgraph C'_{OPT} of C_{OPT} , there is a subgraph C' of C such that the homology classes of C' and C'_{OPT} in the surface S are equal.

Lemma 3.2. *Let C be a graph drawn on S . Assume that C has the multicut dual property. Then C is a multicut dual.*

Proof. (In this proof, we are actually not using the fact that C_{OPT} is optimal, only that it is a multicut dual that is embedded.) Let $(t_1, t_2) \in R$ be a terminal pair. Since C_{OPT} is a multicut dual, every path connecting t_1 and t_2 on S meets C_{OPT} . Let C'_{OPT} be an inclusionwise minimal subgraph of C_{OPT} that separates t_1 from t_2 . This subgraph has exactly two faces, and every edge is incident to each of these two faces; thus, it is an even subgraph of C'_{OPT} , and there is a subgraph C' of C

homologous to C'_{OPT} in S , i.e., $C' = C'_{\text{OPT}} + U$, where U is homologically trivial. Now, any path connecting t_1 and t_2 in G crosses U an even number of times (because homology is considered in S), and it crosses C'_{OPT} an odd number of times (because C'_{OPT} has two faces, each containing one of t_1 and t_2). Therefore, each such path crosses C' an odd number of times, i.e., at least once. Since this applies for any terminal pair, this shows that C is a multicut dual. \square

3.3 Good multicut duals

We now define good multicut duals, which are a central concept of this paper. A graph C drawn in S is a *good multicut dual* if, as an abstract graph, it is the disjoint union of subgraphs (C_0, C_1, \dots, C_k) (whose images on S may overlap), satisfying the following conditions.

1. [**multicut dual property**] C has the multicut dual property.
2. [**structure**] C_1, \dots, C_k , viewed as abstract graphs, are cycles. (C_0 is arbitrary, and in particular may be non-connected.)
3. [**length**] The length of C is at most $(1 + O(\varepsilon))\text{OPT}$.
4. [**minimality**] Each vertex of C_0 has degree at least two.
5. [**number of cycles**] $k = O(g + t)$.
6. [**complexity**] C_0 has $O(g + t)$ vertices and edges.
7. [**crossings with K**] For $0 \leq i \leq k$, the number of crossings between C_i and each edge of K is $O(g + t)$.

Additionally, we say that a good multicut dual C is *embedded* if the following condition holds.

8. [**embedding**] C_0 is embedded on S , and each face of C_0 contains at least one terminal.

Figure 2 pictures two good multicut duals, one consisting only of the “graphical” part C_0 on the left, and one also using a cycle C_1 on the right: in both cases the three stars are separated. The rationale between the a priori strange splitting in this definition into a graph C_0 and a family of cycles C_1, \dots, C_k will be explained in the next subsection. We remark:

- Condition 1, together with Lemma 3.2, implies that a good multicut dual is a multicut dual.
- By the definition of C_{OPT} , $C = \{C_0\} = \{C_{\text{OPT}}\}$ is a good multicut dual that is embedded.
- This definition uses the $O(\cdot)$ notation. Each time the notion of good multicut dual is used, the constants hidden in the $O(\cdot)$ notation are universal, independent from the input of the algorithm, but the constants will vary across statements. More precisely, in some of the following sections, the main proposition (Propositions 6.1, 7.1, and 8.1) involves some good multicut dual, and the constants in the $O(\cdot)$ notation implicitly increase each time. Since the number of increases is bounded, there is no danger in using this notation.
- In the rest of the paper, we focus on computing a good multicut dual. Assuming we computed one, by Condition 3, we can obtain a $(1 + O(\varepsilon))$ -approximation of OPT , by taking the edges in G that are crossed by the good multicut dual.

3.4 Overview of the proof

Our main strategy is to compute a set of portals which cut a near-optimal multicut dual into a family of trees with a controlled number of possible topologies and a controlled number of locations for their leaves. One can then guess these topologies and these locations for the leaves. For each such guess, we determine the shortest trees satisfying these properties by computing Steiner trees in the universal cover. Assembling together these trees provides a graph, one for each guess; the graph corresponding to the correct guess must be a near-optimal multicut dual.

The definition of good multicut duals given in the previous subsection encapsulates exactly the structure of the near-optimal solutions we are interested in. Throughout the proof we will work with such a good multicut dual $C = (C_0, C_1, \dots, C_k)$, starting with $C = (C_0) = (C_{\text{OPT}})$, and iteratively modify it while preserving its good multicut dual structure. The goal of Sections 6 to 8 is to (1) show the existence of such a “well-behaved” near-optimal solution and (2) show how to compute a set of portals cutting this near-optimal solution into a family of trees. (Sections 4 and 5 introduce preliminary tools.)

For technical reasons, we first need to single out a small number of *long* cycles and show that there exists a good multicut dual with no long cycle (Section 6). The main idea is to modify the optimal solution by replacing each long cycle by the shortest homotopic closed curve. The algorithmic outcome of this step is Proposition 6.1 which can be summarized as follows. Since the shortest homotopic closed curves can be “guessed” (i.e.: enumerated) and computed efficiently, we only need to find portals cutting the *short* cycles of the near-optimal solution. This is the purpose of Section 7.

Section 7 is the most technical part of the paper. It describes an algorithm which finds a subgraph of bounded total length with respect to OPT that cuts the graphical part C_0 of a good multicut dual into a family of trees. We call such a subgraph a *skeleton*. The main step of our algorithm is the BUILDALLSKELETA algorithm. It computes a family of skeleta of controlled length such that one of them intersects every cycle of the graphical part C_0 of a good multicut dual. This algorithm proceeds as follows. It enumerates the possible topologies of C_0 and estimates the length of its cycles. For each such combination of cycle lengths and topologies, the algorithm runs the BUILDONESKELETON algorithm. If the combination of cycle lengths and topologies corresponds to the one of a good multicut dual, the BUILDONESKELETON algorithm outputs a graph cutting each cycle of the graphical part C_0 of this good multicut dual. The reader can glance at Figures 3 and 4 to get an intuition of how this works. This algorithm heavily relies on annular covers in both the computations and the analysis, and ultimately on Klein’s multiple source shortest path algorithm [43] to do the computations in near-linear time.

The naive application of the above approach requires to enumerate all cycles of the graphical part C_0 of a good multicut dual. This yields a correct algorithm, but with a running time doubly-exponential in $g + t$. To obtain a better dependence, we show that we can restrict our attention to a much smaller family of $O(g + t)$ cycles, called an *exhaustive family* (Section 5), which bears strong similarities with a *pants decomposition*. The proof uses topological arguments quite different from the rest of the paper.

The outcome of Section 7 is a family of skeleta, one of which cuts a good multicut dual into a family of trees. Thus, the next goal is to show that one can place evenly a set of portals along any skeleton and show that, for the right skeleton, there exists a good multicut dual that is cut into a family of trees by a subset of these portals. This is achieved in Section 8 and summarized in Proposition 8.1. This is done by making detours of small lengths. In order to make such detours, working with drawings of graphs instead of embeddings is very convenient. Even though using portals is a classical step for obtaining PTASs dating back to Arora [2] and Mitchell [47], some

additional subtleties appear in our context.

Finally, in order to prove Theorem 1.1, it suffices to compute the trees linking the portals and the cycles C_1, \dots, C_k . The cycles can be directly inferred from the guessed topology using the algorithm for shortest homotopic closed curves of the second author and Erickson [18]. For each tree, the computation is more involved. We can guess which portals it has as leaves, and also its topology (essentially, the homotopy class of each path in the tree). We then show that we can compute the shortest tree satisfying such constraints: If there were no topological constraint, the problem would boil down to computing a Steiner tree on the surface, and we would use the algorithm of Dreyfus and Wagner [21, 27], which is fixed-parameter tractable in the number of terminals². Taking into account the topology of the tree requires using the same algorithm in the universal cover of the surface. The next section provides more details (see Corollary 4.3).

4 Homotopy types

Recall that, starting with a shortest multicut dual C_{OPT} , we need to prove that it can be transformed into another near-optimal multicut dual satisfying additional structural properties. In order to do this, we need to define some operations that transform a multicut dual into another one. Actually, it is more convenient to work on graphs that satisfy the multicut dual property.

Let C and D be two graphs drawn on S . We say that C and D have the same *homotopy type* if one can be obtained from the other via a sequence of homotopies of the graph (vertices and edges move continuously), edge contractions, and edge expansions (the reverse of edge contractions—a vertex v is replaced with two new vertices w and w' , connected by a new edge, and each edge incident to v is made incident to either w or w'). The following lemma is intuitively clear:

Lemma 4.1. *Let C and D be graphs drawn on S with the same homotopy type. If C satisfies the multicut dual property, then so does D .*

Proof. The proof is easy as far as edge contractions and expansions are concerned, and for the homotopies it follows, for example, from the fact that homotopic graphs are homologous and from Lemma 3.2. \square

Furthermore, in the special case of trees in the plane, we have the following easy lemma:

Lemma 4.2. *Let T and T' be two trees drawn in the plane, such that the leaves of T and T' are at the same positions. Then there is a sequence of edge contractions, edge expansions, and homotopies that transform T into T' without moving the leaves.*

Proof. In T , contract all the edges not incident with a leaf, obtaining a star. Do the same for T' . Transform one star into the other with a homotopy. \square

These easy lemmata provide us with the corollary below, showing that a shortest multicut dual can be obtained as a family of Steiner trees in the universal cover. It will not be used directly in this paper since we will be working with approximate solutions, but it provides a good intuition as to the idea of our algorithm. For a graph D drawn on S and a set P of points on S avoiding the vertices and the self-crossings of D , the graph $D \setminus P$ is the one obtained by cutting D along all the points of P on an edge of D . More precisely, for each point $p \in E(D)$ we do the following. Let u and v be two points close to p on the edge e of D containing p , so that the points u , p , and v

²Note that the planar Steiner Tree PTAS of Borradaile, Klein and Mathieu [7] is exponential in $1/\varepsilon$, and thus too slow for our algorithm.

appear in this order. Subdivide edge e twice by inserting vertices u and v on e , and remove edge uv . Finally, push u and v towards the place where p was, without changing the graph, only its drawing on S (thus the graph has two distinct vertices that overlap).

Note that, if T is a tree drawn on S , it can be lifted to the universal cover \tilde{S} of S as follows. Choose an arbitrary root r of T , and a lift \tilde{r} of r in \tilde{S} . For each leaf ℓ of T , lift the *unique* path from r to ℓ in T to \tilde{S} (starting at \tilde{r}). The result is a tree drawn on \tilde{S} with the same number of vertices and edges as T (but possibly less self-crossings).

Corollary 4.3. *Let C_{OPT} be a shortest multicut dual drawn on S , and P a set of points on the edges of C_{OPT} (but not on self-crossing points) such that $C_{OPT} \setminus P$ is a forest $F = (T_1, \dots, T_m)$. For each $i \in \{1, \dots, m\}$, let \tilde{T}_i be a lift of T_i in the universal cover of S , and let \tilde{P}_i denote the set of lifts of points of P that are on \tilde{T}_i . Then the graph consisting of the projections on S of the Steiner trees with terminals \tilde{P}_i is a shortest multicut dual on S .*

Proof. For each $i \in \{1, \dots, m\}$, the Steiner tree with terminals \tilde{P}_i is shorter than \tilde{T}_i since it is a Steiner tree linking the same terminals, and it has the same homotopy type as \tilde{T}_i by Lemma 4.2. Thus, the graph consisting of the projections of the \tilde{T}_i s on S has the same homotopy type as C_{OPT} and is a multicut dual by Lemma 4.1. By construction, it cannot be longer than C_{OPT} , so it is a shortest multicut dual. \square

5 Exhaustive families

We say that two closed curves γ and δ are *essentially crossing* on S if for each choice of closed curves γ' and δ' homotopic to γ and δ respectively in S , the closed curves γ' and δ' intersect. Recall that a *cycle* in a graph drawn (or embedded) on S is a cycle in the graph-theoretical sense (a closed walk, not reduced to a single vertex, without vertex repetition). A family Γ of cycles in a graph D embedded on S forms an *exhaustive family* if, for each cycle δ in D , either $\delta \in \Gamma$ or there exists a cycle γ in Γ such that γ and δ are essentially crossing in S .

Our algorithm to build the skeleta (Section 7) requires computing an exhaustive family, which for the sake of efficiency should be small. A trivial exhaustive family of D is the entire family of cycles in D . Using this exhaustive family in Section 7 results in a $(1 + \varepsilon)$ -approximation of MULTICUT in time $f(\varepsilon, g, t) \cdot n \log n$, where $f(\varepsilon, g, t) = (\log(\frac{g+t}{\varepsilon})/\varepsilon)^{(g+t)^{O(g+t)}}$. The purpose of this section is to give an algorithm to compute efficiently a smaller exhaustive family, resulting in a singly-exponential f as desired (Theorem 1.1). The main result of this section is the following:

Proposition 5.1. *Let D be a graph embedded on S that has a candidate topology. We can compute an exhaustive family Γ made of $O(g+t)$ cycles of D in time $(g+t)^{O(g+t)}$. Moreover, this computation and the topology of the cycles of Γ only depend on the topology of D .*

Some remarks: Of course, a given graph D has (in general) several exhaustive families; the algorithm computes one such exhaustive family, specified by a set of cycles in (the abstract graph D). The last part of the proposition states that, to perform this computation, the exact knowledge of D is unnecessary; only the data of the topology of D is actually used by the algorithm. Of course, if we know the topology of D and the cycles of Γ (viewed as cycles in the graph D), we can infer the topology of the cycles in D (viewed as closed curves in S). Also, recall that a candidate topology is the topology of a subgraph of one of the graphs output by Lemma 3.1 without vertices of degree zero or one; the assumption that D has a candidate topology is only used to bound the number of cycles in Γ and the complexity of the algorithm.

The remaining part of this section is devoted to the proof of Proposition 5.1. It uses arguments that are independent from the rest of the paper. Thus, the reader might wish to jump to the next section in a first reading, admitting Proposition 5.1.

5.1 The algorithm

We say that two distinct cycles γ and δ in the embedded graph D **cross** at one connected component p of the intersection of the images of γ and δ (a path, possibly reduced to a single vertex) if, after contracting p in both γ and δ into a degree-four vertex v , the pieces of γ and δ alternate at v . The algorithm for Proposition 5.1 is greedy: Start with an empty set Γ and iteratively add a cycle to Γ if it does not cross any of the existing ones.

5.2 Some preliminary lemmas

We say that a curve δ is **minimally self-crossing** if there is no curve δ' homotopic to δ with less self-crossings. Similarly, two curves δ_1 and δ_2 are **minimally crossing** if there is no pair of curves δ'_1 and δ'_2 , homotopic respectively to δ_1 and δ_2 , such that δ'_1 and δ'_2 cross less than δ_1 and δ_2 . We first introduce auxiliary results. The following lemma relates monogons and bigons to the number of crossings of closed curves.

Lemma 5.2 (Hass and Scott [31, Lemma 3.1]). *Let γ and δ be simple closed curves on a surface. Assume that γ and δ are not minimally crossing. Then they form a bigon.*

The following lemma bounds the number of “interesting” disjoint closed curves on a surface.

Lemma 5.3 (Juvan et al. [38, Lemma 3.2]). *Let Γ be a family of pairwise disjoint simple closed curves on a surface with genus g and b boundary components. Assume that the curves in Γ are pairwise non-homotopic. Then Γ has $O(g + b)$ elements.*

Let N be a tubular neighborhood of D . The following lemma shows that one can minimize the number of self-crossings and crossings of a whole family of curves.

Lemma 5.4 (de Graaf and Schrijver [20]). *Let $(\delta_1, \dots, \delta_k)$ be a family of closed curves in N . Then there exist closed curves $(\delta'_1, \dots, \delta'_k)$ in general position in N such that:*

- for each i , the curves δ_i and δ'_i are homotopic;
- every curve δ'_i is minimally self-crossing, and
- every pair of curves δ'_i and δ'_j is minimally crossing.

5.3 Proof of Proposition 5.1

Proof of Proposition 5.1. It seems simpler to present this proof in the combinatorial setting, rather than the cross-metric one. Let Γ be the inclusionwise maximal set of cycles in D computed by the greedy algorithm. Recall that no two cycles in Γ cross. Computing Γ takes $(g+t)^{O(g+t)}$ time. Indeed, there are $(g+t)^{O(g+t)}$ cycles in D (because D has $O(g+t)$ vertices and edges, and every cycle can be represented by a permutation of a subset of edges), which we can enumerate in $(g+t)^{O(g+t)}$ time. The algorithm can easily determine whether two cycles in D cross in $O(g+t)$ time. It is clear that this algorithm only needs the topology of D , and that the topology of the resulting cycles also only depends on the topology of D .




Figure 1: A picture of the situation in P , in the proof of Proposition 5.1.

We now prove that Γ is made of $O(g+t)$ cycles. No pair of cycles in Γ is essentially crossing in S , which implies that they can all be made simple and disjoint by homotopies on N , by Lemma 5.4. These considerations lead to a family Γ' of pairwise disjoint simple closed curves in N . Moreover, no two distinct cycles in D are homotopic in D , so the cycles in Γ are pairwise non-homotopic in N . Since N has genus at most g , and at most t boundary components because every face of D contains a terminal (because D is a subgraph of a graph with a candidate topology), Lemma 5.3 implies that Γ has $O(g+t)$ cycles.

There remains to prove that Γ is an exhaustive family. Let δ be a cycle in D but not in Γ . By construction, there exists $\gamma \in \Gamma$ such that γ and δ cross.

Let us first prove that, in N , the cycles γ and δ are essentially crossing. Since N is a regular neighborhood of D , this is intuitively obvious, but the proof is somewhat heavy-handed. Let p be a connected component of the intersection of the images of γ and δ where these cycles cross. Let P be a disk neighborhood of p in N as shown in Figure 1, so that the boundary of P is made of pieces of boundaries of N and of arcs in N , four of which intersect γ or δ , which we denote by a_γ^1 , a_δ^1 , a_γ^2 , and a_δ^2 (the subscript denoting which cycle crosses the arc); since γ and δ cross, without loss of generality the arcs appear in that order when walking along the boundary of P . Let \tilde{P} be a lift of P in the universal cover \tilde{N} of N , let \tilde{a}_γ^1 , \tilde{a}_δ^1 , \tilde{a}_γ^2 , and \tilde{a}_δ^2 the corresponding lifts of the arcs, and let $\tilde{\gamma}$ and $\tilde{\delta}$ be lifts of γ and δ entering \tilde{P} . Each of these lifts of the arcs separates the universal cover of D into two connected components, and is crossed exactly once by exactly one of $\tilde{\gamma}$ and $\tilde{\delta}$. Any homotopy in N between γ and another closed curve γ' lifts to a homotopy between $\tilde{\gamma}$ and another lift $\tilde{\gamma}'$, which still has to cross \tilde{a}_γ^1 and \tilde{a}_γ^2 , and similarly if δ' is homotopic to δ in N we obtain a lift $\tilde{\delta}'$ that crosses \tilde{a}_δ^1 and \tilde{a}_δ^2 . This implies that $\tilde{\gamma}'$ and $\tilde{\delta}'$ cross, and thus γ and δ are essentially crossing in N .

Now, let γ'' and δ'' be curves in N , homotopic to γ and δ in N , and minimally crossing given these constraints. By the preceding paragraph, γ'' and δ'' have to cross, and we can choose γ'' and δ'' to be simple, again by Lemma 5.4. They form no bigon in N . They also cannot form a bigon in S : Indeed, no such bigon could be entirely in N , so any bigon has to contain at least one connected component of $S \setminus N$. Since a candidate topology is a subtopology of a multicut dual, each such component contains at least one boundary component of S , contradicting the definition of a bigon. This implies, by Lemma 5.2, that γ'' and δ'' are minimally crossing in S . In other words, γ and δ are essentially crossing in S , as desired. \square




Figure 2: Stars denote terminals. The edge $e \subseteq \gamma$ is removed from C_0 , and the cycle γ^o is added to C ; the cycle γ^o may cross C_0 .

6 Near-optimal solution with no long cycle

We say that a cycle γ of a graph D drawn on (S, G) is *long* if $|\gamma^o| = O(|\gamma|\varepsilon/(g+t))$ (for some universal, well-chosen constant hidden in the $O(\cdot)$ notation).

Proposition 6.1. *There exists a good multicut dual C that is embedded, has no long cycle, and such that C_0 has a candidate topology.*

Proof. Initially, let $C := \{C_0\} := \{C_{\text{OPT}}\}$. (Note that C_{OPT} is not necessarily connected.) We have that C is a good multicut dual that is embedded. Whenever C_0 contains a long cycle γ , we do the following: We remove one of the edges of γ from C_0 , iteratively prune vertices of degree one (by removing such vertices and their incident edges), and finally add to C (as a separate component) the cycle γ^o . We refer to Figure 2 for a picture in the planar case, the surface-embedded case bearing no difference.

Each iteration removes at least one edge in C_0 , which initially is equal to the good multicut dual C_{OPT} and therefore has $O(g+t)$ edges. Hence, there are $O(g+t)$ iterations. The new graph consists of the remaining graph C_0 and of the set of cycles that were added at each step, C_1, \dots, C_k . We now argue that this final graph $C = \{C_0, \dots, C_k\}$ is a good multicut dual that is embedded.

By construction, C_0 has no long cycle. Since C_0 is a subgraph of C_{OPT} , it has a candidate topology. Conditions 2 and 4 are obviously satisfied by construction, and since C_0 is a subgraph of C_{OPT} , it has a candidate topology. Conditions 6 and 8 are satisfied because they are satisfied for C_{OPT} , and because C_0 is a subgraph of C_{OPT} . Condition 5 holds because the number of iterations is $O(g+t)$ and each iteration adds one cycle.

Let us prove that Condition 3 holds. At each iteration, the length of C increases by at most the length of γ^o , which satisfies

$$|\gamma^o| = O(|\gamma|\varepsilon/(g+t)) = O(\text{OPT} \cdot \varepsilon/(g+t)).$$

This implies that after the $O(g+t)$ iterations, the length of the graph has increased by at most $O(\varepsilon \text{OPT})$, as desired.

Let us prove that Condition 7 holds. Recall that γ^o , being a shortest homotopic closed curve, crosses arcs of K minimally, and therefore the number of crossings between γ^o and K does not exceed the number of crossings between γ and K . In particular, the number of intersections of each

cycle C_i with each edge of K is at most the number of intersections of C_{OPT} with each edge of K , that is, $O(g+t)$. The same holds for C_0 since it is a subgraph of C_{OPT} .

Finally, let us prove that Condition 1 is satisfied. Let C'_{OPT} be an even subgraph of C_{OPT} . We need to prove that C'_{OPT} has the same homology class as some even subgraph of C . Initially, $C = \{C_{\text{OPT}}\}$, so this property was trivially true at the beginning. It thus suffices to prove that this property is maintained when performing an iteration that removes one edge e that belongs to a cycle γ and adds γ^o to C . Before this iteration, we know that C'_{OPT} has the same homology class as some subgraph C' of C . If C' does not contain the chosen edge e , this is trivially true after the replacement. Otherwise, as a chain, $C' = ((C' \setminus e) + (\gamma \setminus e)) + \gamma$, which is homologous to $((C' \setminus e) + (\gamma \setminus e)) + \gamma^o$, an even subgraph of the new C . □

7 Building skeleta

In this section, we prove the following result.

Proposition 7.1. *In $\left(\frac{g+t}{\varepsilon}\right)^{O(g+t)} \cdot n \log n$ time, the BUILDALLSKELETA algorithm builds a family of $\left(\frac{g+t}{\varepsilon}\right)^{O(g+t)}$ graphs, called **skeleta**, drawn on S , each with $O(g+t)$ vertices and edges and $O((g+t)^3 n)$ complexity, such that at least one of them, denoted by $Sk(C_0)$, satisfies the following conditions:*

- $Sk(C_0)$ has length $O((g+t)OPT)$;
- There exists a good multicut dual $C' = \{C'_0, C'_1, \dots, C'_k\}$ (not necessarily embedded) such that $Sk(C_0)$ intersects every cycle of C'_0 .

In this proposition and the next one, the data structure used to represent a skeleton and its drawing is its abstract graph, and for each of its edges e , the ordered sequence of the edges of G crossed by e in the drawing, together with the orientation of each crossing. The bound on the complexity ensures that this encoding has a controlled size.

This data structure does not encode the precise location of the skeleton inside each face of G but it is sufficient for the purpose of the algorithm. More precisely, for the proofs in this section, we rely on the precise location of the skeleta, but in the actual algorithms (see Section 8), this data structure suffices.

The proof of Proposition 7.1 relies on the exhaustive families introduced in Section 5; the proof also extensively uses covering spaces and algorithms to compute shortest homotopic curves, described in Appendix A.

7.1 Definition of a skeleton

We will define the algorithm we use, called BUILDONESKELETON, for computing a single skeleton, and will prove the following key proposition.

Proposition 7.2. *Let D be graph embedded on S that has a candidate topology and that has no long cycle. In $(g+t)^{O(g+t)} n \log n$ time, the BUILDONESKELETON algorithm computes a skeleton $Sk(D)$ drawn on S that satisfies the following properties:*

- $Sk(D)$ has $O(g+t)$ vertices and edges;
- $Sk(D)$ has length $O((g+t)|D|)$;

- $Sk(D)$ has complexity $O((g+t)^3n)$.
- there exists a graph D' drawn on S such that:
 - D' has the same homotopy type as D ,
 - $|D'| \leq (1 + O(\varepsilon))|D|$,
 - $Sk(D)$ intersects every cycle of D' ,
 - the number of crossings between D' and each edge of K is $O(g+t)$,
 - each vertex of D' has degree at least two.

Moreover, the construction does not depend on D , but only on one out of $(\frac{g+t}{\varepsilon})^{O(g+t)}$ possible choices of parameters.

Note that we can apply Proposition 7.2 with $D = C_0$. Hence, we define the BUILDALLSKELETA to be the algorithm that, for every possible choice of parameters, computes the corresponding skeleton using the BUILDONESKELETON algorithm and returns the family consisting of all these skeleta. Assuming this Proposition 7.2, we can directly prove Proposition 7.1:

Proof of Proposition 7.1, assuming Proposition 7.2. The proof is immediate and only involves definition-chasing: we apply Proposition 7.2 for all the possible choices of parameters, which can be done in $(\frac{g+t}{\varepsilon})^{O(g+t)} \cdot n \log n$ time. All the resulting skeleta satisfy the bounds on the number of vertices and edges and the complexity. Out of these, for the parameters corresponding to the C_0 of Proposition 6.1, we obtain $Sk(C_0)$, which has length $O((g+t)|C_0|) = O((g+t)OPT)$, and we denote by C'_0 the graph D' of Proposition 7.2. We claim that $C' = (C'_0, C_1, \dots, C_k)$ is a good multicut dual. Since C'_0 has the same homotopy type as C_0 , Lemma 4.1 proves that Condition 1 is fulfilled. Conditions 2, 3, 4, 5, 6 and 7 follow directly by construction and by the properties of C'_0 , as does the property that $Sk(C_0)$ intersects every cycle of C'_0 . \square

The algorithm BUILDONESKELETON takes as input a set of *parameters*, which we now describe. The first of these is the topology of D . By Proposition 5.1, an exhaustive family Γ can be computed for D just with the knowledge of its topology, and it consists of $O(g+t)$ cycles. There is one additional parameter per two-sided cycle in Γ , which corresponds to an estimate of its length. More precisely, fix a two-sided cycle $\gamma \in \Gamma$. Since D has no long cycle, we have $|\gamma| \leq |\gamma^0|(g+t)/\varepsilon$, hence $|\gamma|$ lies within one of the following *ranges*:

$$[|\gamma^0|, (1+\varepsilon)|\gamma^0|), \quad [(1+\varepsilon)|\gamma^0|, (1+\varepsilon)^2|\gamma^0|), \quad \dots, \quad [(1+\varepsilon)^\ell|\gamma^0|, (1+\varepsilon)^{\ell+1}|\gamma^0|),$$

for ℓ satisfying $(1+\varepsilon)^{\ell+1} = \Theta((g+t)/\varepsilon)$, and thus $\ell = \Theta(\log(\frac{g+t}{\varepsilon})/\varepsilon)$. The parameter associated to the two-sided cycle γ indicates in which range $|\gamma|$ lies.

We now describe the BUILDONESKELETON algorithm for constructing the *skeleton* $Sk(D)$, computed as follows. First, we compute the exhaustive family Γ associated to D . We now look at each cycle $\gamma \in \Gamma$ in turn, and add some vertices and edges in $Sk(D)$, as follows.

If γ is one-sided, we just add it to the skeleton.

Otherwise, γ is two-sided; let \hat{S}_γ be the annular cover of S corresponding to γ , which is (topologically) an annulus; let $\hat{\gamma}$ be a cyclic lift of γ in this annulus. Let \hat{p} be an arbitrary lift of a path of K that connects the two boundaries of \hat{S}_γ ; there exists such a lift because otherwise γ would be contractible. For every face f of \hat{S}_γ crossed by \hat{p} , let $\hat{\gamma}_f$ be a shortest non-contractible closed curve in the annulus \hat{S}_γ that goes through f . Among all the closed curves $\hat{\gamma}_f$ of length within or lower than the range of γ , let $\hat{\gamma}^1$ and $\hat{\gamma}^2$ be ones corresponding to faces f which are the leftmost




Figure 3: The construction of the skeleton, as represented on the surface S . The cycle γ is in thick lines. The closed curves γ^1 and γ^2 are chosen so that, among the closed curves with the same homotopy class as γ and within the same range or lower, they cross p as left or as right as possible, respectively. Left: In this case, γ has incident edges only on one side, so only γ^1 and γ^2 are added to the skeleton. The same construction would be used if γ were incident to no other edge of C_0 . Right: In this case, γ has incident edges on both of its sides, and thus not only γ^1 and γ^2 are added to the skeleton, but also a shortest path (denoted by q here) connecting γ^1 and γ^2 in the region bounded by γ^1 and γ^2 .

and rightmost along \hat{p} (see Figure 3). Without loss of generality, these closed curves can be chosen to be simple and disjoint. Finally, let \hat{A}_γ be the annulus bounded by $\hat{\gamma}^1$ and $\hat{\gamma}^2$, and let \hat{D} be the lift of D containing $\hat{\gamma}$. We add the projections γ^1 and γ^2 of the curves $\hat{\gamma}^1, \hat{\gamma}^2$ to the skeleton $Sk(D)$ (see Figure 3, left). Moreover, if on each side of the two-sided cycle $\hat{\gamma}$ there is at least one edge of \hat{D} incident to $\hat{\gamma}$, we add the projection of the shortest path in \hat{A}_γ between $\hat{\gamma}^1$ and $\hat{\gamma}^2$ to Sk (see Figure 3, right).

Note that this construction only depends on the topology of D and the additional parameters: Indeed, the topology of D suffices to compute the exhaustive family Γ and the topology of its cycles. Furthermore, it specifies the length of γ^o and whether a cycle γ of Γ is one-sided or two-sided. If γ is two-sided, the knowledge of its topology also allows us to compute its annular cover \hat{S}_γ . Then, to define and compute the cycles $\hat{\gamma}^1$ and $\hat{\gamma}^2$, one only needs to additionally know the range of γ . The topology of D also indicates whether a shortest path between $\hat{\gamma}^1$ and $\hat{\gamma}^2$ is added in the skeleton.

7.2 Bounding the number of skeleta

In line with the description of the BUILDONESKELETON algorithm, a candidate topology and a choice of a range for every cycle of an exhaustive family associated to that topology is called a *choice of parameters*. Out of all the possible skeleta computed with all the possible choices of parameters, one will have the choice of parameters corresponding to D , i.e., the correct topology and the correct choice of range for each cycle of D , and it will be $Sk(D)$.

Lemma 7.3. *The number of possible choices of parameters is $\left(\frac{g+t}{\varepsilon}\right)^{O(g+t)}$.*

Proof. By Lemma 3.1, there are $(g+t)^{O(g+t)}$ candidate topologies, and to a given topology corresponds an exhaustive family Γ . Moreover, the number of ranges for every cycle is $\ell = O(\log(\frac{g+t}{\varepsilon})/\varepsilon)$, and the number of cycles in Γ is $O(g+t)$. Thus, the total number of skeleta is $\left(\frac{g+t}{\varepsilon}\right)^{O(g+t)}$, as desired. \square

7.3 Computing a skeleton

Lemma 7.4. *The BUILDONESKELETON algorithm runs in time $(g+t)^{O(g+t)} \cdot n \log n$ and outputs a skeleton of complexity $O((g+t)^3 n)$.*

Proof. We first compute the exhaustive family Γ corresponding to the graph D , using Proposition 5.1, in $(g+t)^{O(g+t)}$ time. It is made of $O(g+t)$ cycles, each of them crossing K $O(g+t)$ times. (We do not need to compute these cycles exactly on S , but need only to keep the information of the edges of K crossed, together with the orientation of the crossings.) For each such cycle γ , we compute the curve $\hat{\gamma}$, as well as $\hat{\gamma}^1$ and $\hat{\gamma}^2$, and possibly the shortest path inbetween. Computing the shortest homotopic closed curve γ^ρ can be done in $O((g+t)jn \log n)$ time, where j is the number of crossings between γ and K , as explained in Lemmas A.2 and A.3. By Lemma 3.1, for a candidate topology we have $j = O((g+t)^2)$.

Observe that, using the same technique as in the proofs of these lemmas (which rely on Klein's multiple shortest path algorithm [43]), we can also compute, in the same amount of time, the length of all the closed curves $\hat{\gamma}_f$, and thus also the closed curves $\hat{\gamma}^1$ and $\hat{\gamma}^2$. Finally, if needed, the algorithm adds the shortest path between $\hat{\gamma}^1$ and $\hat{\gamma}^2$.

Since the skeleton is made of $O(g+t)$ projections of shortest paths in a relevant region of size $O((g+t)^2 n)$, the bound on the complexity follows. \square

7.4 Properties of the skeleton $Sk(D)$

For an embedded graph D having a candidate topology and no long cycle, recall that $Sk(D)$ is the skeleton computed with the choice of parameters corresponding to D . To prove Proposition 7.1, it remains to prove that $Sk(D)$ has the desired properties.

Lemma 7.5. *The skeleton $Sk(D)$ has length $O((g+t)|D|)$.*

Proof. Let γ be a cycle of Γ . For $i = 1, 2$, we have $|\gamma^i| = O(|\gamma|)$ by the choice of the ranges. Moreover, if we added to $Sk(D)$ the projection of a shortest path connecting $\hat{\gamma}^1$ to $\hat{\gamma}^2$, it is because $\hat{\gamma}$ had some edges leaving it on each on its sides. In that case, since D has a candidate topology, it has no degree one vertex and no face without terminal, and since the annulus \hat{A}_γ bounded by $\hat{\gamma}^1$ and $\hat{\gamma}^2$ contains no terminal, it must be that \hat{D} connects $\hat{\gamma}^1$ to $\hat{\gamma}^2$. So the shortest path between $\hat{\gamma}^1$ and $\hat{\gamma}^2$ has length at most that of D (see Figure 3, right). Thus, each time we consider a cycle in the exhaustive family Γ , the length of the skeleton increases by $O(|D|)$. The desired result follows since Γ has $O(g+t)$ cycles. \square

Lemma 7.6. *There is a graph D' drawn on S such that:*

- D' has the same homotopy type as D ,
- $|D'| \leq (1 + O(\varepsilon))|D|$,
- $Sk(D)$ intersects every cycle of D' ,
- the number of crossings between D' and each edge of K is $O(g+t)$.

To prove this lemma, we will need the following lemma of Hass and Scott.

Lemma 7.7 (Hass and Scott [31, Theorem 2.4]). *Let γ be a closed curve in general position on a surface with non-empty boundary, not simple but homotopic to a simple curve. Then γ forms a monogon or a bigon.*




Figure 4: In this simple case, we modify the multicut dual by pushing it to γ^2 , increasing its length by a factor of at most $1 + \varepsilon$.

Proof of Lemma 7.6. We start by providing an overview of the proof. We will inductively build D' from D by pushing cycles of D whenever $Sk(D)$ does not meet them. If D has a cycle γ that is disjoint from $Sk(D)$, then γ must belong to the exhaustive family Γ , and the shortest path between $\hat{\gamma}^1$ and $\hat{\gamma}^2$ has not been added. So we can push $\hat{\gamma}$, together with the attached parts in \hat{A}_γ , in the direction of these attached parts until it hits $\hat{\gamma}^2$. This does not increase too much the length, because γ^2 is not much longer than γ , and preserves the fact that one has a multicut dual.

Here are the details. In an intermediate step, some parts of D' and $Sk(D)$ will overlap, so that these two graphs will *not* be in general position. Initially, let $D' := D$. Our construction is iterative, modifying D' as long as some cycle in D' does not intersect $Sk(D)$. We maintain the following invariant: the part of D' outside $Sk(D)$ is included in D .

So, assume that there exists a cycle γ in D' that is disjoint from $Sk(D)$. By the invariant above, γ appears at the same place in D , and is thus embedded. We claim that γ belongs to the exhaustive family Γ corresponding to D . Indeed, otherwise, by the definition of an exhaustive family, γ and some cycle $\gamma' \in \Gamma$ would be essentially crossing; but, by construction, the skeleton $Sk(D)$ contains some closed curves homotopic to γ' , which would therefore cross γ , which is a contradiction. Thus $\gamma \in \Gamma$. Moreover, $Sk(D)$ contains a closed curve homotopic to γ , which γ would cross if it were one-sided; so γ is two-sided.

Let $\hat{\gamma}$ be the lift of γ in the annular cover \hat{S}_γ . Since $\hat{\gamma}^1$ and $\hat{\gamma}^2$ belong to $Sk(D)$, they are not crossed by $\hat{\gamma}$; by the choice of these closed curves, $\hat{\gamma}$ is entirely contained in the annulus \hat{A}_γ . Let \hat{D}'_γ be the intersection of the interior of \hat{A}_γ with the connected component of the lift of D' that contains $\hat{\gamma}$. Since γ is not crossed by $Sk(D)$, the shortest path between $\hat{\gamma}^1$ and $\hat{\gamma}^2$ was not added in the construction of $Sk(D)$. This implies that \hat{D}'_γ cannot touch both $\hat{\gamma}^1$ and $\hat{\gamma}^2$; for example, let us assume that it does not touch $\hat{\gamma}^1$. We note that, by our invariant, the projection D'_γ of \hat{D}'_γ to S is also part of D , no face of which is a disk because D has a candidate topology. Thus, since D has no degree zero or degree one vertex, \hat{D}'_γ consists of the cycle $\hat{\gamma}$ together with some disjoint trees attached to it and with leaves on $\hat{\gamma}^2$.

The naïve strategy now is to push \hat{D}'_γ onto $\hat{\gamma}^2$, or more precisely to change D' by removing \hat{D}'_γ and by adding $\hat{\gamma}^2$, as in Figure 4. However, there is a possible pitfall here, as two distinct points of \hat{D}'_γ may project to the same point of S (see Figure 5). We therefore use a more complicated argument. Let γ^2 be the projection of $\hat{\gamma}^2$ on S . If γ^2 self-intersects, and since it is homotopic to the simple cycle γ , it contains a monogon or a bigon (Lemma 7.7). We flip that monogon or bigon, so that the new curve has the same image as γ^2 except in a small neighborhood of some self-crossing points and has strictly less self-crossing points. We iterate the construction, obtaining a simple closed curve γ'^2 homotopic to γ^2 and with essentially the same image. Now, γ and γ'^2 are simple,




Figure 5: Left: An example where γ^2 is non-simple (although its lift $\hat{\gamma}^2$ always is). In that example, some points on the thin horizontal edge lifts to two distinct points in \hat{A}_γ . We iteratively flip the monogons and bigons of γ^2 to make it simple (middle), and can then replace γ with the new γ^2 (right).

pairwise disjoint, non-contractible, homotopic in S ; furthermore, γ cannot be contractible, since otherwise it would bound a disk [22, Theorem 1.7]. This implies that they bound an annulus A in S [22, Lemma 2.4].

We claim that the edges of D incident to γ leave γ on the side of A . Indeed, the pieces of \hat{D}'_γ belong to the annulus bounded by $\hat{\gamma}$ and $\hat{\gamma}^2$. When we flip a monogon or bigon on γ^2 , this changes $\hat{\gamma}^2$ to another closed curve that is disjoint from $\hat{\gamma}$, and thus has to be on the same side of $\hat{\gamma}$ in the annular cover (because flipping a monogon or bigon does not move the whole curve). Thus, at the end of the process, the lift of the annulus A is on the same side of $\hat{\gamma}$ as was the annulus bounded by $\hat{\gamma}$ and $\hat{\gamma}^2$.

We can now safely replace the part of D that belongs to A by the single cycle γ'^2 . For the same reason as above, the part of D that belongs to A consists of the cycle γ together with disjoint trees with leaves on γ^2 , so this operation can be represented as a sequence of edge compressions, edge expansions, and homotopies for D . This increases the length of D' by a factor of at most $1 + \varepsilon$, because γ^2 (or γ'^2) is at most a factor $1 + \varepsilon$ longer than γ . Finally, we slightly deform the part of D' that overlaps γ'^2 onto γ^2 , so that D' really overlaps the skeleton; this does not affect the length of D' , and is possible because the image of γ^2 is obtained from that of γ'^2 by a small perturbation that makes pieces of γ'^2 touch.

We repeat this operation as long as there is a cycle γ in D' that is not crossed by $Sk(D)$. As proved above, such a cycle is necessarily part of Γ , and after the iteration it is pushed onto $Sk(D)$, so this process eventually stops and the number of iterations is at most the number of cycles in Γ .

Finally, we slightly perturb D' so that D' and $Sk(D)$ are in general position to remove the overlaps. We do so by ensuring that each edge of D' intersecting $Sk(D)$ before the perturbation still intersects it after the perturbation. This preserves the fact that $Sk(D)$ intersects every cycle in D' . This concludes the definition of D' , which by construction has the same homotopy type as D .

At each of the $O(g + t)$ iterations above, the length of D' increases by at most $(1 + \varepsilon)$ times the length of the cycle γ considered in that iteration. But the cycles γ considered are edge-disjoint, because once a cycle γ is considered, it is contained in $Sk(D)$. So, after all the iterations, the length of D' has increased by a factor of at most $(1 + \varepsilon)$. Hence the length of D' is at most $(1 + O(\varepsilon))$ times the length of D , as desired.

Finally, let us bound the number of crossings between D' and K . It suffices to prove that each lift \hat{q} of an edge of K in \hat{S}_γ is not crossed more by $\hat{\gamma}^2$ than by $\hat{\gamma}$. This follows from the fact that edges of K are shortest homotopic paths and from the tie-breaking rule for shortest paths outlined in Section 2.2. \square

Proof of Proposition 7.2. The result now immediately follows from Lemmas 7.3, 7.4, 7.5, and 7.6.

(It is clear, by construction, that each skeleton has $O(g+t)$ vertices and edges.) \square

8 Near-optimal solution made of trees and cycles

Proposition 8.1. *There exists an algorithm that computes in $((g+t)/\varepsilon)^{O(g+t)} \cdot n \log n$ time, a family of $k = ((g+t)/\varepsilon)^{O(g+t)}$ **portal sets** $\{\mathcal{P}_1, \dots, \mathcal{P}_k\}$ such that each set consists of $O((g+t)^2/\varepsilon)$ points on S , called **portals**, and one of the portal sets, denoted by \mathcal{P} , satisfies the following condition: There exists a good multicut dual C'' such that $C'' \setminus \mathcal{P}$ is a forest.*

Proof. We describe how to construct the sets $\mathcal{P}_1, \dots, \mathcal{P}_k$. There is a set \mathcal{P}_i for each skeleton obtained by Proposition 7.1. Fix a skeleton Σ and let e be an edge of Σ . Locate portals on e in a way that each point on e is at distance $O(\varepsilon|\Sigma|/(g+t)^2)$ of some portal, where as usual the $O(\cdot)$ notation hides a universal constant. This requires at most $O(|e|(g+t)^2/(\varepsilon|\Sigma|))$ portals. Doing this for each of the edges of Σ produces $O((g+t)^2/\varepsilon)$ portals; these constitute the set \mathcal{P}_i corresponding to Σ . Doing this for each of the $((g+t)/\varepsilon)^{O(g+t)}$ skeleta yields a family of sets $\{\mathcal{P}_1, \dots, \mathcal{P}_k\}$. Recall that the algorithm of Section 7 encodes each skeleton by its abstract graph and, for each of its edges e , by the ordered sequence of the edges of G crossed by e in the drawing, together with the orientation of each crossing (see Section 7). Thus, the algorithm has enough information to define the family of sets $\{\mathcal{P}_1, \dots, \mathcal{P}_k\}$.

Now, following Proposition 7.1, let $Sk(C_0)$ be a skeleton of length $O((g+t)\text{OPT})$ and C' be the corresponding good multicut dual, such that $C'_0 \setminus Sk(C_0)$ is a forest. This skeleton corresponds to one of the sets of portals, and we denote it by \mathcal{P} .

Consider the following subgraph C''_0 obtained by modifying C'_0 as follows: For every edge e of C'_0 crossed by $Sk(C_0)$, we insert, on e , a *single* detour to the closest portal in \mathcal{P} via a shortest path. Let C''_0 be the resulting graph. By construction, each detour has length $O(\varepsilon|Sk(C_0)|/(g+t)^2)$. We now prove that $C'' = (C''_0, C''_1, \dots, C''_k)$ is a good multicut dual.

By construction, Conditions 2, 4, 5 and 6 are satisfied. The number of detours made is $O(g+t)$, and each of them has length $O(\varepsilon|Sk(C_0)|/(g+t)^2) = O(\varepsilon\text{OPT}/(g+t))$, so the length of C''_0 is at most $(1+\varepsilon)\text{OPT}$, and Condition 3 is fulfilled. Since C''_0 crosses each edge of K at most $O(g+t)$ times and the number of detours is $O(g+t)$, each consisting of two shortest paths, and every edge in K is the concatenation of two shortest paths, the number of crossings between C''_0 and each edge of K is $O(g+t)$ which gives Condition 7. Finally, C''_0 is obtained from C'_0 by a homotopy of the edges, so Condition 1 is also fulfilled. Therefore C'' is a good multicut dual.

For each skeleton in the skeleta, the portals can be placed greedily by “walking” along the skeleton, in time linear in the complexity of the skeleton, which is $O((g+t)^3n)$. This yields an overall complexity of $((g+t)/\varepsilon)^{O(g+t)} \cdot n \log n$. \square

9 The algorithm

The final step of our algorithm is to compute separately, for each set of portals \mathcal{P}_i and for each possible position of a good multicut dual with respect to \mathcal{P}_i and K , the optimal layout of this multicut dual. By Proposition 8.1, there exists a good multicut dual C'' such that C''_0 is cut into a forest by a set of portals \mathcal{P} . Then, finding C''_0 is achieved by (1) guessing the topology of the trees of the forest and (2) for each of these, computing a Steiner tree corresponding to this topology in some portion of the universal cover of S . The cycles C''_i for $i \geq 1$ are found by computing shortest homotopic closed curves.

For a given set of portals, we can compute a Steiner tree efficiently by using the following theorem, which follows from the classical algorithm of Dreyfus and Wagner [21], sped up using the

techniques of Erickson, Monma and Veinott [27] (this is not explicitly stated in that article, but it is folklore that the method applies to this problem, see, e.g., Vygen [51, Introduction]):

Theorem 9.1. *Given a planar graph $G = (V, E)$ with n vertices and a set $T \subseteq V$ of t terminals, one can compute the optimal Steiner tree on (G, T) in time $2^{O(t)} n \log n$.*

Similarly to the computation of shortest homotopic paths described in Section A.1, the appropriate setting to compute the Steiner trees in our algorithm is a relevant region of the universal cover. A tree T drawn on S linking a set of portals P can be lifted in the universal cover of S to a tree \tilde{T} linking some lifts of the portals \tilde{P} . Computing a Steiner tree on the points \tilde{P} and projecting it back to S yields the shortest tree with the same homotopy type as T linking the portals P . Indeed, after choosing a basepoint for the lifts, every tree with the same homotopy type lifts to a tree that connects the same lifts, by the definition of a universal cover. Furthermore, it is enough to compute a Steiner tree in the relevant region of \tilde{S} , which is the portion of \tilde{S} , cut by the edges of K , visited by \tilde{T} . As for shortest homotopic paths, this follows from the fact that arcs in K are shortest homotopic arcs.

For a set of portals cutting C_0'' into a forest $F = (T_1, \dots, T_m)$, the **layout** of C_0'' is the family of homotopy types of these trees. They can be lifted to the universal cover, and (after fixing a basepoint) by Lemma 4.2, specifying a layout is the same as specifying the sets of lifts of portals corresponding to each tree. It is a **good layout** if the forest F has $O(g+t)$ leaves and if it crosses each edge of K at most $O(g+t)$ times. A relevant region of the universal cover for *all* the possible good layouts has size $(g+t)^{O((g+t)^2)} n$: since there are $O((g+t)^2)$ crossings between the edges of C_0'' and the edges of K , one can compute it by cutting S along K , yielding a disk D , and pasting new copies of D at its boundaries in an arborescent way up to depth $O((g+t)^2)$. Since there are $O(g+t)$ edges in K , the corresponding relevant region has size $(g+t)^{O((g+t)^2)} n$.

Similarly, for a closed curve γ drawn on S , its **layout** is its homotopy class, and the layout of a closed curve is **good** if the shortest homotopic closed curve in this layout crosses each edge of K at most $O(g+t)$ times.

We now have all the tools to describe our algorithm.

MAIN ALGORITHM

1. Compute the family of skeleta using the BUILDALLSKELETA algorithm of Section 7.
2. For each skeleton, compute a set of portals \mathcal{P}_i following Section 8.
3. Compute a relevant region \bar{S} of the universal cover of S of size $(g+t)^{O((g+t)^2)} n$.
4. For every set of portals \mathcal{P}_i and for every good layout of a graph drawn on S with respect to a subset of $O(g+t)$ portals of \mathcal{P}_i , compute the Steiner trees connecting the lifts of portals specified by this layout in \bar{S} . Once projected onto S , this gives a family of graphs drawn on S .
5. For every choice of good layouts of j closed curves, for every integer $0 \leq j \leq O(g+t)$, compute shortest homotopic closed curves corresponding to these layouts. All the possible combinations of these additional closed curves with the previous family of graphs yield the candidate solutions.
6. Output the shortest of the candidate solutions that is a multicut dual.

Proof of Theorem 1.1. Running time. The computations in Sections 7 and 8 have already been shown to take time $((g+t)/\varepsilon)^{O((g+t))} n \log n$. We argue that the number of good layouts of graphs is

$1/\varepsilon^{O(g+t)} \cdot (g+t)^{O((g+t)^3)}$. Indeed, since a set of portals \mathcal{P}_i has size $O((g+t)^2/\varepsilon)$, they lift in the relevant region of the universal cover to a set of lifts of portals of size $O((g+t)^2/\varepsilon) \cdot (g+t)^{O(g+t)^2} = (g+t)^{O(g+t)^2}/\varepsilon$. In order to specify a good layout, we first choose $O(g+t)$ portals out of this set, for which there are $\left((g+t)^{O(g+t)^2}/\varepsilon\right)^{O(g+t)} = 1/\varepsilon^{O(g+t)} \cdot (g+t)^{O((g+t)^3)}$ possible choices, then we choose the number $m = O(g+t)$ of trees and pick for each of them a subset of $O(g+t)$ portals out of these; this multiplies the number of possible choices by $(g+t)^{O(g+t)}$. Thus the number of good layouts of graphs is $1/\varepsilon^{O(g+t)} \cdot (g+t)^{O((g+t)^3)}$.

A good layout of a closed curve is characterized by the sequence of crossings with K , and thus there are $(g+t)^{O(g+t)^2}$ good layouts of closed curves. This gives at most $(g+t)^{O(g+t)^3}$ choices for the set of additional closed curves in the main algorithm.

Now, for each good layout of a graph, we compute the family of Steiner trees corresponding to it, i.e., the Steiner trees with the chosen portals as a set of $O(g+t)$ terminals in the relevant region of the universal cover. By Theorem 9.1, such a Steiner tree can be computed in time $2^{O(g+t)}(g+t)^{O(g+t)^2}n \log n = (g+t)^{O(g+t)^2}n \log n$. Similarly, for each good layout of a closed curve we can compute the shortest homotopic closed curve in the relevant region of its annular cover, which is done using the tools exposed in Section A.2. Since the relevant region of the annular cover has size $O((g+t)n)$, these computations are negligible compared to the other ones.

Hence, the algorithm has an overall running time of $1/\varepsilon^{O(g+t)}(g+t)^{O((g+t)^3)} \cdot n \log n$.

Correctness and approximation guarantee. By Proposition 8, there is a good multicut dual C'' and a set of portals \mathcal{P} such that $C''_0 \setminus \mathcal{P}$ is a forest. Since in order to make a graph acyclic, one only needs to cut each edge at most once, one can assume that only a subset of $O(g+t)$ portals of \mathcal{P} are used. From this and the properties of being a good multicut dual, we can infer that C'' forms a good layout with respect to \mathcal{P} . For this choice of portals and layout, our algorithm will output a graph C''' made of a set of cycles C''_i , $i \geq 1$, and a graph C'''_0 (possibly non-connected) such that:

- C'''_0 has the same homotopy type as C''_0 , since the trees it consists of have the same homotopy type as the ones of C''_0 .
- C'''_0 is no longer than C''_0 , since C'''_0 is made of Steiner trees that are by definition no longer than the corresponding trees in C''_0 .
- Each C'''_i is homotopic to C''_i , since they have the same layout, and each C'''_i is shorter than C''_i by construction.

Therefore, by Lemma 4.1, subgraphs of C''' have all the homology classes of subgraphs of C'' , and thus by Lemma 3.2, C''' is a multicut dual. Hence the graph output by our algorithm is a multicut dual of length at most the length of C'' , which concludes the proof since C'' is a good multicut dual and thus a near-optimal solution. □

References

- [1] I. ABRAHAM, C. GAVOILLE, A. GUPTA, O. NEIMAN, AND K. TALWAR, *Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs*, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing, STOC '14, ACM, 2014, pp. 79–88.

- [2] S. ARORA, *Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems*, Journal of the ACM, 45 (1998), pp. 753–782.
- [3] M. BATENI, M. HAJIAGHAYI, P. N. KLEIN, AND C. MATHIEU, *A polynomial-time approximation scheme for planar multiway cut*, in Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, SIAM, 2012, pp. 639–655.
- [4] M. BATENI, M. HAJIAGHAYI, AND D. MARX, *Approximation schemes for steiner forest on planar graphs and graphs of bounded treewidth*, Journal of the ACM (JACM), 58 (2011), p. 21.
- [5] G. BORRADAILE, E. D. DEMAINE, AND S. TAZARI, *Polynomial-time approximation schemes for subset-connectivity problems in bounded-genus graphs*, in Proceedings of the 26th Annual Symposium on Theoretical Aspects of Computer Science (STACS), 2009, pp. 171–182.
- [6] G. BORRADAILE, D. EPPSTEIN, A. NAYYERI, AND C. WULFF-NILSEN, *All-Pairs Minimum Cuts in Near-Linear Time for Surface-Embedded Graphs*, in 32nd International Symposium on Computational Geometry (SoCG 2016), S. Fekete and A. Lubiw, eds., vol. 51 of Leibniz International Proceedings in Informatics (LIPIcs), Dagstuhl, Germany, 2016, Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, pp. 22:1–22:16.
- [7] G. BORRADAILE, P. KLEIN, AND C. MATHIEU, *An $O(n \log n)$ approximation scheme for Steiner tree in planar graphs*, ACM Transactions on Algorithms (TALG), 5 (2009), p. 31.
- [8] G. BORRADAILE AND PH. KLEIN, *An $O(n \log n)$ algorithm for maximum st-flow in a directed planar graph*, Journal of the ACM, 56 (2009).
- [9] N. BOUSQUET, J. DALIGAULT, AND S. THOMASSÉ, *Multicut is FPT*, in Proceedings of the forty-third annual ACM symposium on Theory of computing, ACM, 2011, pp. 459–468.
- [10] E. CHAMBERS, J. ERICKSON, AND A. NAYYERI, *Minimum cuts and shortest homologous cycles*, in Proceedings of the 25th Annual Symposium on Computational Geometry (SOCG), ACM, 2009, pp. 377–385.
- [11] E. W. CHAMBERS, É. COLIN DE VERDIÈRE, J. ERICKSON, F. LAZARUS, AND K. WHITTLESEY, *Splitting (complicated) surfaces is hard*, Computational Geometry: Theory and Applications, 41 (2008), pp. 94–110.
- [12] E. W. CHAMBERS, J. ERICKSON, AND A. NAYYERI, *Homology flows, cohomology cuts*, SIAM Journal on Computing, 41 (2012), pp. 1605–1634.
- [13] S. CHAWLA, R. KRAUTHGAMER, R. KUMAR, Y. RABANI, AND D. SIVAKUMAR, *On the hardness of approximating multicut and sparsest-cut*, Computational Complexity, 15 (2006), pp. 94–114.
- [14] C. CHEKURI AND V. MADAN, *Approximating multicut and the demand graph*, in Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, P. N. Klein, ed., 2017, pp. 855–874.
- [15] É. COLIN DE VERDIÈRE, *Shortest cut graph of a surface with prescribed vertex set*, in Proceedings of the 18th European Symposium on Algorithms (ESA), part 2, no. 6347 in Lecture Notes in Computer Science, 2010, pp. 100–111.

- [16] ———, *Computational topology of graphs on surfaces*, in Handbook of Discrete and Computational Geometry, J. E. Goodman, J. O'Rourke, and C. Toth, eds., CRC Press LLC, third ed., 2017, ch. 23. To appear.
- [17] ———, *Multicuts in planar and bounded-genus graphs with bounded number of terminals*, Algorithmica, (2017). In press, available at <http://link.springer.com/article/10.1007/s00453-016-0258-0>.
- [18] É. COLIN DE VERDIÈRE AND J. ERICKSON, *Tightening nonsimple paths and cycles on surfaces*, SIAM Journal on Computing, 39 (2010), pp. 3784–3813.
- [19] E. DAHLHAUS, D. S. JOHNSON, C. H. PAPADIMITRIOU, P. D. SEYMOUR, AND M. YANNAKAKIS, *The complexity of multiterminal cuts*, SIAM Journal on Computing, 23 (1994), pp. 864–894.
- [20] M. DE GRAAF AND A. SCHRIJVER, *Making curves minimally crossing by Reidemeister moves*, Journal of Combinatorial Theory, Series B, 70 (1997), pp. 134–156.
- [21] S. E. DREYFUS AND R. A. WAGNER, *The Steiner problem in graphs*, Networks, 1 (1971), pp. 195–207.
- [22] D. B. A. EPSTEIN, *Curves on 2-manifolds and isotopies*, Acta Mathematica, 115 (1966), pp. 83–107.
- [23] J. ERICKSON, *Maximum flows and parametric shortest paths in planar graphs*, in Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, pp. 794–804.
- [24] J. ERICKSON, K. FOX, AND A. NAYYERI, *Global minimum cuts in surface embedded graphs*, in Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2012, pp. 1309–1318.
- [25] J. ERICKSON AND S. HAR-PELED, *Optimally cutting a surface into a disk*, Discrete & Computational Geometry, 31 (2004), pp. 37–59.
- [26] J. ERICKSON AND A. NAYYERI, *Minimum cuts and shortest non-separating cycles via homology covers*, in Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011, pp. 1166–1176.
- [27] R. E. ERICKSON, C. L. MONMA, AND A. F. VEINOTT, JR, *Send-and-split method for minimum-concave-cost network flows*, Mathematics of Operations Research, 12 (1987), pp. 634–664.
- [28] J. FAKCHAROENPHOL AND K. TALWAR, *An improved decomposition theorem for graphs excluding a fixed minor*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Springer, 2003, pp. 36–46.
- [29] N. GARG, V. V. VAZIRANI, AND M. YANNAKAKIS, *Approximate max-flow min-(multi) cut theorems and their applications*, SIAM Journal on Computing, 25 (1996), pp. 235–251.
- [30] N. GARG, V. V. VAZIRANI, AND M. YANNAKAKIS, *Primal-dual approximation algorithms for integral flow and multicut in trees*, Algorithmica, 18 (1997), pp. 3–20.
- [31] J. HASS AND P. SCOTT, *Intersections of curves on surfaces*, Israel Journal of Mathematics, 51 (1985), pp. 90–120.

- [32] R. HASSIN, *Maximum flow in (s, t) -planar networks*, Information Processing Letters, 13 (1981), p. 107.
- [33] R. HASSIN AND D. B. JOHNSON, *An $O(n \log^2 n)$ algorithm for maximum flow in undirected planar networks*, SIAM Journal on Computing, 14 (1985), pp. 612–624.
- [34] A. HATCHER, *Algebraic topology*, Cambridge University Press, 2002. Available at <http://www.math.cornell.edu/~hatcher/>.
- [35] M. R. HENZINGER, PH. KLEIN, S. RAO, AND S. SUBRAMANIAN, *Faster shortest-path algorithms for planar graphs*, Journal of Computer and System Sciences, 55 (1997), pp. 3–23.
- [36] T. C. HU, *Multi-commodity network flows*, Operations Research, 11 (1963), pp. 344–360.
- [37] G. F. ITALIANO, Y. NUSSBAUM, P. SANKOWSKI, AND CH. WULFF-NILSEN, *Improved algorithms for Min Cut and Max Flow in undirected planar graphs*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC), 2011, pp. 313–322.
- [38] M. JUVAN, A. MALNIČ, AND B. MOHAR, *Systems of curves on surfaces*, Journal of Combinatorial Theory, Series B, 68 (1996), pp. 7–22.
- [39] K.-I. KAWARABAYASHI, B. MOHAR, AND B. REED, *A simpler linear time algorithm for embedding graphs into an arbitrary surface and the genus of graphs of bounded tree-width*, in Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2008, pp. 771–780.
- [40] S. KHULLER AND J. S. NAOR, *Flow in planar graphs: a survey of recent results*, in Planar graphs, W. T. Trotter, ed., vol. 9 of DIMACS Series in Discrete Mathematics and Theoretical Computer Science, AMS, 1993, pp. 59–84.
- [41] P. KLEIN, S. A. PLOTKIN, AND S. RAO, *Excluded minors, network decomposition, and multicommodity flow*, in Proceedings of the twenty-fifth annual ACM symposium on Theory of computing, ACM, 1993, pp. 682–690.
- [42] P. N. KLEIN AND D. MARX, *Solving PLANAR k -TERMINAL CUT in $O(n^{c\sqrt{k}})$ time*, in Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP), 2012, pp. 569–580.
- [43] PH. N. KLEIN, *Multiple-source shortest paths in planar graphs*, in Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2005, pp. 146–155.
- [44] D. MARX, *Parameterized complexity and approximation algorithms*, The Computer Journal, 51 (2008), pp. 60–78.
- [45] —, *A tight lower bound for planar multiway cut with fixed number of terminals*, in International Colloquium on Automata, Languages, and Programming, Springer, 2012, pp. 677–688.
- [46] D. MARX AND I. RAZGON, *Fixed-parameter tractability of multicut parameterized by the size of the cutset*, SIAM Journal on Computing, 43 (2014), pp. 355–388.
- [47] J. S. B. MITCHELL, *Guillotine subdivisions approximate polygonal subdivisions: a simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems*, SIAM Journal on Computing, 28 (1999), pp. 1298–1309.

- [48] J. H. REIF, *Minimum $s - t$ cut of a planar undirected network in $O(n \log^2(n))$ time*, SIAM Journal on Computing, 12 (1983), pp. 71–81.
- [49] A. SCHRIJVER, *On the history of combinatorial optimization (till 1960)*, Handbooks in Operations Research and Management Science, 12 (2005), pp. 1–68.
- [50] J. STILLWELL, *Classical topology and combinatorial group theory*, Springer-Verlag, New York, second ed., 1993.
- [51] J. VYGEN, *Faster algorithm for optimum Steiner trees*, Information Processing Letters, 111 (2011), pp. 1075–1079.

A Covering spaces and shortest homotopic curves

In this appendix, we introduce the covering spaces that are used throughout this article, and summarize the most important points of the algorithms to compute shortest homotopic paths and closed curves [18, Sections 1 and 6], eluding many technicalities from that paper that are irrelevant to our problem.

A map $\pi : S' \rightarrow S$ between two surfaces is called a **covering map** if it is a local homeomorphism, or more precisely if each point $x \in S$ lies in an open neighborhood U such that (1) $\pi^{-1}(U)$ is a countable union of disjoint open sets $U_1 \cup U_2 \cup \dots$ and (2) for each i , the restriction $\pi|_{U_i} : U_i \rightarrow U$ is a homeomorphism. We say that S' (more formally, it should be (S', π)) is a **covering space** of S . The definition implies that every path p in S can be **lifted** to a path p' in S' (the opposite operation of **projecting** with the map π), and that moreover this can be done in a unique way if the starting point x' of p' is prescribed (such that $\pi(x')$ equals the initial point of p).

An important type of covering space is the **universal cover** of S , denoted by \tilde{S} , in which every closed curve is contractible. The universal cover is unique; a closed curve in S is contractible if and only if its lifts in \tilde{S} are closed curves.

A.1 Universal cover and shortest homotopic paths

Let p be a path on S . By definition of the universal cover, to compute a shortest path homotopic to p , it suffices to (1) compute a lift \tilde{p} of p in the universal cover \tilde{S} , which is naturally a (non-compact) cross-metric surface, (2) compute a shortest path between the endpoints of \tilde{p} , and (3) return the projection of the resulting path to S . However, the universal cover is infinite, and one needs to restrict oneself to a finite portion of the universal cover, its **relevant region** with respect to \tilde{p} . Let us explain in more detail.

Let D be the disk obtained by cutting S along K ; this is naturally a cross-metric surface defined by the image of G . The universal cover \tilde{S} can be built by gluing together infinitely many copies of D along lifts of the arcs in K ; the **relevant region** of \tilde{S} with respect to \tilde{p} is defined to be the set of lifts of D visited by \tilde{p} . The claim is that a shortest path in \tilde{S} connecting the endpoints of \tilde{p} stays in the relevant region, and this follows from the fact that, since the arcs in K are shortest homotopic arcs, they lift to shortest paths, and each of these shortest paths separate the universal cover.

We can easily build the relevant region of the universal cover incrementally: Initially, it is the single copy of D containing the source of \tilde{p} . We walk along \tilde{p} in the relevant region, and each time we are about to exit the relevant region via a lift \tilde{a} of an arc of K , we attach a copy of D across \tilde{a} , preserving the local homeomorphism condition. This results in a set of copies of D attached together in a tree-like fashion.




Figure 6: Proof of Lemma A.1: The shortest non-contractible closed curve γ in the Möbius strip M corresponds, in this representation, to pairwise disjoint simple arcs connecting the left and right sides of the quadrangle. If γ contains more than one arc, as is depicted, then a_1 must be identified to a_1' and similarly a_2 must be identified with a_2' . This implies that γ crosses a exactly twice, which in turn implies that it is contractible, a contradiction.

From the above discussion, it also follows that a shortest path \tilde{p}' with the same endpoints as \tilde{p} crosses each lift of an arc in K at most once, and crosses only lifts of arcs of K that are already crossed by \tilde{p} . This is because each lift of K is a shortest path; to avoid \tilde{p}' crossing twice the same lift of K , we can use the number of crossings with K as a tie-breaking measure for the notion of length. In particular, the projection p' of p crosses each arc of K at most as many times as p does.

Finally, since computing shortest paths in a planar graph takes linear time [35], and since the relevant region has complexity $O((g+t)kn)$ where k is the number of crossings between p and K , computing a shortest homotopic path can be done in time $O((g+t)kn)$.

A.2 Annular covers and shortest homotopic closed curves

The same strategy can be used for computing a shortest homotopic closed curve homotopic to a non-contractible closed curve γ in S , but now we have to use the *annular cover* of S with respect to γ , denoted by \hat{S}_γ ; this is a covering space in which every simple closed curve is either contractible or homotopic to a lift of γ or to its reverse. If γ is two-sided, the annular cover is homeomorphic to an annulus with some points of the boundary removed. Otherwise, γ is one-sided, and the annular cover is homeomorphic to a Möbius strip with some points of the boundary removed. The shortest homotopic closed curve in the annular cover crosses every lift of an arc in K at most once, and crosses only arcs that are crossed by the lift $\hat{\gamma}$ of γ that is a closed curve (so, as above with paths, a shortest homotopic closed curve crosses each arc at most as many times as the original closed curve does). In the annulus case, the proof is as before, and in the case where the annular cover is a Möbius strip, the proof rests on the following lemma.

Lemma A.1. *Let M be a cross-metric surface that is a Möbius strip. Let a be a non-separating arc that is as short as possible in its homotopy class. Then some shortest non-contractible closed curve in M crosses a exactly once.*

Proof. Let γ be a shortest non-contractible closed curve in M , crossing a as few times as possible; γ is simple and must cross a at least once (for otherwise it would be contractible in M). We now prove that it crosses a exactly once; see Figure 6. Cut M along a , obtaining a quadrangle Q with four sides, two opposite sides (say the left and right sides) corresponding to the sides of a . Because a is a shortest homotopic arc, and because the number of crossings with a is minimal, the image of γ in Q is made of simple disjoint arcs connecting the opposite sides of Q corresponding to a . Assume that there are at least two arcs. Because the arcs are simple and disjoint, the left endpoint of the topmost arc, a_1 , must be identified to the right endpoint of the bottommost arc, a_1' , and

similarly the left endpoint of the bottommost arc, a_2 , must be identified to the right endpoint of the topmost arc, a'_2 , which implies that there are exactly two arcs, implying that γ crosses a exactly twice. This is impossible since a closed curve is non-contractible if and only if it crosses γ an odd number of times. \square

A variation of the construction described in Section A.1 allows to build the *relevant region* of the annular cover \hat{S}_γ , which is topologically an annulus (in the two-sided case) or a Möbius strip (in the one-sided case) made of finitely many copies of D . In essence, we first start at an arbitrary point of γ and build the relevant region of the closed path starting at that point. This is a topological disk obtained by gluing together copies of D in a tree-like fashion. We only keep the copies of D that are between the starting and ending points of the lift, discarding the other copies, obtaining a set of copies of D glued together along a path. Finally, we identify the initial and final copies of D in this path, leading to an annulus or a Möbius strip.

Algorithmically, we need to compute a shortest non-contractible closed curve in an annulus or a Möbius strip of complexity $O((g+t)kn)$, where k is the number of crossings between γ and K . We can do this in $O((g+t)kn \log((g+t)kn))$ time. We describe a possible algorithm for the case of the annulus, the ideas of which will be reused in Section 7 (the result is actually well-known, because it reduces to a minimum cut computation in the dual, which can even be done in $O((g+t)kn \log \log((g+t)kn))$ time [37]).

Lemma A.2. *Let A be a cross-metric annulus with complexity m . In $O(m \log m)$ time, we can compute a shortest non-contractible closed curve in A .*

Proof. Compute a shortest path a between the two boundaries of A . It is easy to see that a desired shortest non-contractible closed curve crosses a exactly once. Cut A along a , obtaining a quadrangle Q with four sides, two opposite sides (say the left and right sides) corresponding to the sides of a .

The shortest non-contractible closed curve in A corresponds to a shortest path between matching pairs of points in Q . We can compute the distance between all the matching pairs in $O(m \log m)$ time using the planar multiple-source shortest path algorithm by Klein [43]; the shortest non-contractible closed curve in A connects the matching pair with the smallest distance. \square

In the case of the Möbius strip, we expect this to be known in some circles, but to our knowledge the result has not appeared anywhere:

Lemma A.3. *Let M be a cross-metric Möbius strip with complexity m . In $O(m \log m)$ time, we can compute a shortest non-contractible closed curve in M .*

Proof. Attach a disk to the boundary component of M , obtaining a projective plane P , inheriting the cross-metric structure from M . In P , compute a shortest non-contractible loop based at some point inside the disk [25, Lemma 5.2]; that loop is simple and crosses the boundary of M exactly twice. Its trace on M is a shortest arc a among all non-separating arcs in M . By Lemma A.1, some shortest non-contractible closed curve in M crosses a exactly once, and thus corresponds to a shortest path between matching pairs of points in the disk obtained by cutting M along a , as in the annulus case; we can similarly conclude using Klein's algorithm [43]. \square