A Unified FPT Framework for Crossing Number Problems

Éric Colin de Verdière
LIGM, CNRS, Univ Gustave Eiffel, F-77454 Marne-la-Vallée, France

Petr Hliněný
Masaryk University, Faculty of Informatics, Brno, Czech republic

Abstract

The basic (and traditional) *crossing number problem* is to determine the minimum number of crossings in a topological drawing of an input graph in the plane. We develop a unified framework that smoothly captures many generalized crossing number problems, and that yields fixed-parameter tractable (FPT) algorithms for them not only in the plane but also on surfaces.

Our framework takes the following form. We fix a surface $\mathcal{S}$, an integer r, and a map κ from the set of topological drawings of graphs in $\mathcal{S}$ to $\mathbb{Z}_+ \cup \{\infty\}$, satisfying some natural monotonicity conditions, but essentially describing the allowed drawings and how we want to count the crossings in them. Then deciding whether an input graph G has an allowed drawing D on $\mathcal{S}$ with $\kappa(D) \leq r$ can be done in time quadratic in the size of G (and exponential in other parameters). More generally, we may take as input an edge-colored graph, and distinguish crossings by the colors of the involved edges; and we may allow to perform a bounded number of edge removals and vertex splits to G before drawing it. The proof is a reduction to the embeddability of a graph on a two-dimensional simplicial complex.

This framework implies, in a unified way, quadratic FPT algorithms for many topological crossing number variants established in the graph drawing community. Some of these variants already had previously published FPT algorithms, mostly relying on Courcelle's metatheorem, but for many of those, we obtain an algorithm with a better runtime. Moreover, our framework extends, at no cost, to these crossing number variants in any fixed surface.

2012 ACM Subject Classification Theory of computation → Computational geometry; Theory of computation → Fixed parameter tractability; Mathematics of computing → Graph theory

Keywords and phrases computational geometry; fixed-parameter tractability; graph drawing; graph embedding; crossing number; two-dimensional simplicial complex; surface

Funding *Petr Hliněný*: Part of this work was done while this author was invited professor at LIGM, Marne-la-Vallée, supported by Paris-Est Sup.

Acknowledgements We thank Marcus Schaefer for helpful discussions on the many non-traditional crossing number variants, and reviewers for their many constructive suggestions.

1 Introduction

The (traditional) *crossing number* problem, minimizing the number of pairwise edge crossings in topological drawings of an input graph in the plane, is a long-standing and central task in graph drawing and visualization that comes in many established flavors; see the extensive dynamic survey by Schaefer [45]. In this paper, we give a framework that smoothly captures many of the variants of crossing numbers, and in a unified way provide efficient algorithms for them when parameterized by the solution value (i.e., the respective crossing number).

Many flavors of crossing numbers. There is currently a surge of crossing number variants. A motivation for such variations is that, in practical drawing applications, not every crossing or crossing pattern of edges may be "equal" to other ones. One may, e.g., want to avoid mutual crossings of important edges. Or, to allow crossings only within specific parts of a

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

graph, and not between unrelated parts. Or, to exclude crossings of edges of some type. Or, to allow only certain "comprehensible" crossing patterns in order to help visualize the graph. This is the main focus of the recent research direction called "beyond planarity" [15].

Paraphrasing Schaefer [45, Chapter 2], a crossing number variant minimizes, over all *allowed drawings* D of an input graph G in some specified *host surface*, an *objective function* related to the crossings in D. The *allowed drawings* may, e.g., restrict the crossings on one edge or in a local pattern of crossing edges, or prescribe or forbid certain local properties of the drawing (also depending on types or colors of edges). The *host surface* is usually the plane, but can be any fixed surface, orientable or not. The *objective function* is often the number of crossings between edges, but other possibilities include, e.g., the number of edges involved in crossings, or the number of pairs of edges that cross an odd number of times.

Existing algorithms for the traditional version. One usually considers the decision version of the (traditional) crossing number problem: Given an input graph G of size n, and an integer r, does G have a crossing number at most r? This problem is already NP-hard in the plane, as proved by Garey and Johnson in 1983 [18], and even in very specific cases [8, 24, 26]. Moreover, the problem is APX-hard [7], and the best known polynomial-time approximation algorithm [10] gives an approximation factor that is subpolynomial in n only for bounded-degree graphs.

Thus, the main current focus is on algorithms that are fixed-parameter tractable (FPT) in r — the runtime has the form $f(r)\cdot\mathsf{poly}(n)$, where f is a computable function of the *parameter* r. For fixed r, Grohe [20] has described an $O(n^2)$-time algorithm, and Kawarabayashi and Reed [32] have announced an $O(n)$-time algorithm. Both rely on Courcelle's metatheorem [14], which automatically entails a huge dependence in the parameter r.

Two recent approaches avoid resorting to Courcelle's theorem. First, Colin de Verdière, Magnard, and Mohar [13] have studied the problem of embedding a graph in a *two-dimensional simplicial complex* (2-complex for short), a topological space obtained from a surface by adding isolated edges and identifying vertices. Among other motivations for this problem, they observe [13, Introduction] that the crossing number problem reduces to the embeddability of the input graph in a certain 2-complex depending only on r. Then Colin de Verdière and Magnard [12] have shown that the embeddability of a graph of size n in a 2-complex of size C can be tested in $2^{\mathsf{poly}(C)}\cdot n^2$ time. Hence, this results in a $2^{\mathsf{poly}(r)}\cdot n^2$-time algorithm for the crossing number problem, which, moreover, extends to any fixed surface.

Second, even more recently, Lokshtanov, Panolan, Saurabh, Sharma, Xue, and Zehavi [35] have given a $2^{\mathcal{O}(r\log r)}\cdot n$-time algorithm for the traditional crossing number problem in the plane, by a reduction to bounded treewidth and dynamic programming, all in time linear in n.

Existing algorithms for other variants. For other flavors of crossing numbers, which are (typically) also NP-hard, the literature on FPT algorithms is scarce; besides Pelsmajer, Schaefer, and Štefankovič [41] giving FPT algorithms for the odd and pair crossing numbers, we are only aware of two very recent papers, which bring a general approach to multiple crossing number flavors and which we now detail. (The various flavors of crossing numbers are defined in Section 5.)

First, Münch and Rutter [37], extending Grohe's approach [20], provide a framework for quadratic FPT algorithms for the crossing number of several types of beyond-planar drawings of graphs in the plane, characterized by forbidden combinatorial crossing patterns. This includes the crossing number of k-planar, k-quasi-planar, min-k-planar, fan-crossing, and fan-crossing free drawings of a graph for any constant k.

Second, Hamm, Klute, and Parada [22], in a very recent preprint, have announced a generalization of the previous framework [37], also handling forbidden topological crossing patterns in drawing types in the plane and, more importantly, bringing the possibility of handling predrawn parts of the input graph (as already known for the traditional crossing number [21]). The dependence in the size n of the input graph is an unspecified polynomial.

The above results all rely on Courcelle's theorem. Apart from that, we are only aware of sporadic results for isolated variants. Kawarabayashi and Reed [32] and Jansen, Lokshtanov, and Saurabh [30] both claim, without details, a linear FPT algorithm to compute the skewness of a graph, and Nöllenburg, Sorge, Terziadis, Villedieu, Wu, and Wulms [39] give a non-uniform FPT algorithm for the splitting number in surfaces. All mentioned papers except Nöllenburg et al. are restricted to the plane.

Our contributions. Our general framework takes the following form. We fix a surface $\mathcal{S}$ and an objective function κ from the set $\mathcal{D}$ of (possibly edge-colored) topological drawings of graphs on $\mathcal{S}$ to the set $\mathbb{Z}_+ \cup \{\infty\}$; the allowed drawings D are those that satisfy $\kappa(D) < \infty$. The map κ has to satisfy some natural monotonicity conditions, defined later. Given an integer r, we prove that deciding whether an input graph G has a drawing D on $\mathcal{S}$ with $\kappa(D) \leq r$ can be done in quadratic time in the size n of G, and exponential in a polynomial of the other parameters, namely the genus of $\mathcal{S}$, the maximum size of a "fully crossing" drawing $D^\times$ satisfying $\kappa(D^\times) \leq r$, and the integer r. This remains true if one allows a bounded number of (possibly color-restricted) edge removals and vertex splits to G before drawing it; these numbers are also parameters. The proof is a reduction to the already mentioned problem of deciding the embeddability of a graph on a 2-complex [12]; while the general strategy is intuitive, many subtle technical details need to be overcome to make it all work.

We deduce FPT algorithms with runtime quadratic in the size n of the input graph, and exponential in a polynomial of the other parameters, for many established crossing number variants; we now survey some of them. First, we can restrict ourselves to many established drawing styles, such as k-planar, k-quasi-planar, min-k-planar, fan-crossing, weakly and strongly fan-planar, k-gap, and fixed-rotation drawings (k being an additional fixed parameter). For the k-gap and fixed rotation cases, no FPT algorithms were known. Second, we may assign colors to the edges of the input graph and count the crossings differently depending on the colors of the edges involved, leading to the first FPT algorithms on the joint crossing number on surfaces and its generalizations. Third, assuming suitable additional properties, we can handle problems that do not count the crossings, but rather the number of edges, or pairs of edges, involved in crossings, such as the edge, pair, and odd crossing numbers (no FPT algorithm was known for the edge crossing number). And fourth, since we allow for prior edge removals and vertex splits, our framework encompasses, e.g., the skewness, for which FPT algorithms were only sketched [30, 32] in the plane, and the splitting number, which was only known to admit a nonuniform FPT algorithm [39].

Last but not least, all these four aspects that define flavors of crossing numbers can be combined arbitrarily, and our main results are valid not only in the plane, but for arbitrary surfaces. As a side note, we allow non-orientable surfaces in our framework, homeomorphic to a sphere with a number of "crosscaps". Placing k crosscaps in the plane amounts to choosing k specific points in the plane that can be traversed by an arbitrary number of pairwise crossing edges of the planar drawing "for free", without accounting for the induced crossings at these points, which seems relevant from a graph drawing perspective. Another small but relevant note is that we allow surfaces with boundary, which is needed for some subtle variations of the fan-planar crossing number.

For most flavors of crossing number, the arguments are direct, though they really differ according to the flavor; they essentially boil down to checking some technical conditions (see Definition 3.1), and to giving an algorithm (not necessarily polynomial) to check a given drawing of bounded size and to count its crossings (see Theorem 3.2). However, for a couple of crossing number flavors (in particular, when fixing the rotation scheme), additional ingredients are needed.

Comparison with the state of the art. Our main contribution is a convenient, versatile, and unified framework capturing a very general class of crossing number variants while giving FPT algorithms. In particular, as listed above, we get new FPT results for several established variants and their generalizations (to surfaces, or by combining the features of multiple variants). It is conceivable that some of the variants newly covered by our framework could be proved FPT using other techniques, perhaps the (very powerful) metatheorem of Courcelle, as already used in [20–22, 37]. However, such developments are extremely delicate even in isolated variants and lead to hard-to-describe algorithms. Indeed, first, checking a given drawing and counting its crossings must be formalized in the MSO_2 logic of graphs, which usually needs heavy tricks tailored to the specific case; second, there remains the specific and often highly nontrivial treewidth reduction step to be done. So far, such approaches have been successful only in the case of the plane.

Moreover, all approaches using Courcelle's theorem inherently come with a high computational cost. Indeed, they result in algorithms with a huge multi-level exponential dependence in the parameter r; for instance, it is an exponential tower of height at least four in the case of [20] and three in the case of [32], see the discussion in [35]. In contrast, our algorithms are singly exponential in a polynomial in r. Also, while Courcelle's theorem itself runs in linear time in the size n of the input graph, the necessary treewidth reduction step requires quadratic time, if not more, in the known approaches for the non-traditional crossing numbers. Our algorithms have a quadratic dependence in n, and the only bottleneck to a linear-time dependence is the quadratic runtime of the current best embeddability algorithm [12].

The frameworks used in the two very recent works mentioned above [22, 37], which both use treewidth reduction and Courcelle's theorem, can handle some of the drawing styles that we encompass, but different counting functions, edge removals and vertex splits, and nonplanar surfaces are out of their reach. However, it should not be difficult to extend their framework to handle edge-colored graphs, thanks to the MSO_2 logic being able to handle arbitrary sets of edges. On the other hand, the recent breakthrough of Lokshtanov et al. [35] for the traditional crossing number problem may fuel hope for more $2^{\mathcal{O}(r \log r)} \cdot n$-time algorithms, but tweaking all ingredients of that highly technical paper seems hard, even for isolated crossing number variants, and moreover this approach is inapplicable to surfaces other than the plane, due to the use of several inherently planar techniques (3-connectivity arguments and dependence on [30]).

Organization of the paper. After the preliminaries (Section 2), we describe our general framework and state our main result (Section 3), which is then proved (Section 4). We then list the numerous applications to crossing number variants (Section 5).

2 Preliminaries

Graphs and surfaces. In this paper, graphs are finite and undirected, but not necessarily simple unless specifically noted. The *size* of a graph G is the number of vertices plus the number of edges of G. A *vertex split* of a graph G at vertex v creates a new vertex v' and replaces some of the edges incident to v by making them incident to v' instead of v.

We follow [2, 36] for surface topology. A *surface* $\mathcal{S}$ is a topological space obtained from finitely many disjoint solid, two-dimensional triangles by identifying some of their edges in pairs. Surfaces are not necessarily connected. The *boundary* of $\mathcal{S}$ is the closure of the union of the unidentified edges. The *(Euler) genus* of a connected surface is twice its number of "handles", if it is orientable, and is its number of "crosscaps", otherwise. Up to homeomorphism, each *connected* surface $\mathcal{S}$ is specified by whether it is orientable or not, by its (Euler) genus, and by its number of boundary components. In this paper, by 'genus' we always mean the Euler genus. The *topological size* of $\mathcal{S}$ equals $s := d + b + g$, where d is the number of connected components of $\mathcal{S}$, g is the sum of the (Euler) genera of its connected components, and b is the number of boundary components. The plane is not a surface according to our definition, but in this paper it can always be substituted with a disk, which is (homeomorphic to) a surface. A *self-homeomorphism* of $\mathcal{S}$ is just a homeomorphism from $\mathcal{S}$ to $\mathcal{S}$.

(Topological) drawings of graphs. In this paper, any drawing of any graph in any topological space maps distinct vertices to distinct points, and has finitely many intersection points, each avoiding the images of the vertices; however, an intersection point may involve two or more pieces of edges. In more formal terms, and introducing more terminology, the previous sentence means the following. A *curve* c in a topological space $\mathcal{X}$ is a continuous map from the closed interval $[0, 1]$ into $\mathcal{X}$. The *relative interior* of c is the image, under c, of the open interval $(0, 1)$. In a *drawing* D of a graph G in a topological space $\mathcal{X}$, vertices are represented by points of $\mathcal{X}$ and edges by curves in $\mathcal{X}$ such that the *endpoints* $c(0)$ and $c(1)$ of a curve c are the images of the end vertices of the corresponding edge. Moreover, we assume that distinct vertices are mapped to distinct points and that the relative interiors of the edges avoid the images of the vertices. Given a point $x \in \mathcal{X}$, its *multiplicity* in D is the number of pairs (c, t) such that $c(t) = x$, where c is a curve representing an edge of G in D, and $t \in (0, 1)$. An *intersection point* (shortly an *intersection*) of D is a point with multiplicity at least two. We only consider drawings D with finitely many intersection points.

When considering a drawing D of a graph G on a *surface* $\mathcal{S}$, we *always* implicitly assume that G has no isolated vertex and that the image of D avoids the boundary of $\mathcal{S}$; these conditions are actually benign in a graph drawing context. In D, each edge e is subdivided into (finitely many) *pieces* by the intersection points along e. A point of multiplicity two corresponds to a *crossing* if no local perturbation of the curves can remove this intersection, or to a *tangency*, otherwise. D is *normal* [44] if every intersection point has multiplicity two and is a crossing (not a tangency). Moreover, D is *simple* if it is normal, no edge self-crosses, no two adjacent edges cross, and no two edges cross more than once.[1]

The (traditional) *crossing number* of a graph G in a surface $\mathcal{S}$ is the least number of crossings over all normal drawings of G in $\mathcal{S}$. Many variations exist [45], see Section 5.

Storing drawings on surfaces. We need to represent drawings of graphs on a surface $\mathcal{S}$ combinatorially, up to self-homeomorphisms of $\mathcal{S}$. Cellular embeddings can be represented

[1] Note that many authors do not consider a drawing D *simple* if D contains two parallel edges (even uncrossed), as two parallel edges necessarily intersect at their two endpoints. Likewise, uncrossed loops are usually not allowed in simple drawings. We do not make these additional restrictions here for technical reasons connected to Definition 3.1; see also Proposition 5.2.

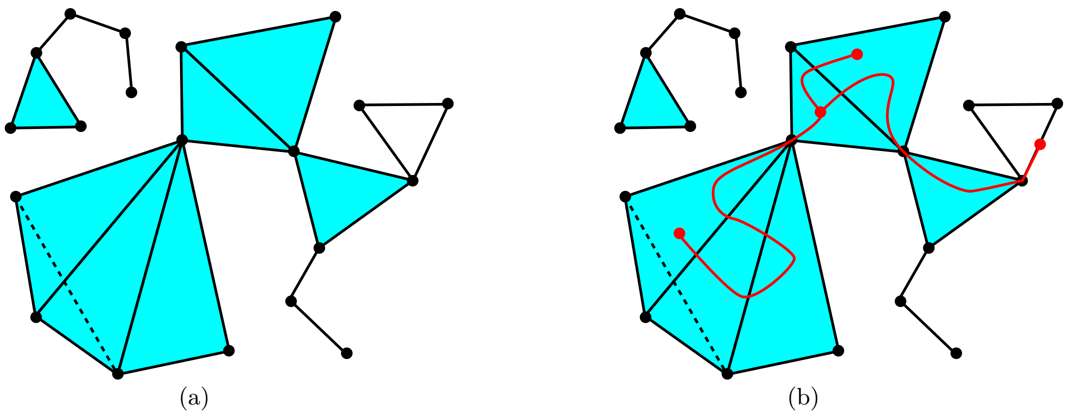

Figure 1 (a): An example of a two-dimensional simplicial complex, or 2-complex for short. (b): An embedding of some graph with four vertices and three edges on the same 2-complex.

by combinatorial maps [16], but we need to allow crossings, and moreover faces need not be disks. We can achieve this, ultimately relying on combinatorial maps. Here are the details.

We represent drawings on surfaces as subsets of the vertex-edge graph of some triangulation of $\mathcal{S}$. (An alternative possibility would be through *extended combinatorial maps* [11, Section 2.2], but this appears to be unnecessarily complicated for our purposes.)

Specifically, a *triangulation* of a surface $\mathcal{S}$ is a graph T embedded on $\mathcal{S}$ such that each face of T (each connected component of the complement of the image of T) is homeomorphic to a disk bounded by three edges, and all three incident vertices and edges are pairwise distinct. In particular, a triangulation is a 2-complex. Let T' be the graph made of the vertices and edges of T that lie in the interior of $\mathcal{S}$. Let $W = (w_1, \dots, w_k)$ be a set of k walks in T', such that each edge of T' is used at most once by the union of these walks. This set W naturally gives rise to a drawing D of a graph G on $\mathcal{S}$: The graph G has k edges, each drawn as $w_1, \dots, w_k$ on $\mathcal{S}$; the vertices of G are the at most $2k$ endpoints of the walks in W. We say that the pair (T, W) is a *representation* of D.

A representation of a colored drawing is a representation (T, W) of the corresponding uncolored drawing, together with the data of a color for each walk in W. Its *size* is the size of a data structure to represent it, which is $\mathcal{O}(tc)$, where t is the number of triangles and c is the number of colors. Conversely (as proved in Lemma 4.3), every drawing D on $\mathcal{S}$ can be represented by such a pair (T, W) in which T is made of $\mathcal{O}(s + u)$ triangles, where s is the topological size of $\mathcal{S}$, and the intersection points of D subdivide the edges of D into u pieces in total.

2-complexes. In this paper, a 2-*complex* $\mathcal{C}$ (or two-dimensional simplicial complex) is a topological space obtained from a simple graph (without loops or multiple edges) by attaching solid, two-dimensional triangles to some of its cycles of length three; see Figure 1.[2] The *simplices* of $\mathcal{C}$ are its vertices, edges, and triangles. We can easily represent 2-complexes algorithmically, by storing the vertices, edges, and triangles, and the incidences between them. In general, many such representations correspond to the same topological space, and the choice of the representation only impacts the complexity analysis of our algorithms. The

[2] Our definition of 2-complex slightly departs from the standard one; it is the same as a geometric simplicial complex of dimension at most two, realized in some ambient space of dimension large enough.

class of 2-complexes is quite general; it contains all graphs, all surfaces, all k-books (although in this paper, we only need 2-complexes in which every edge is incident to at most two triangles), and any space obtained from a surface by identifying finitely many finite subsets of points and by adding finitely many edges between any two points.

An *isolated edge* of a 2-complex $\mathcal{C}$ is an edge incident to no triangle. The *surface part* of $\mathcal{C}$ is the union of all its triangles, together with their incident vertices and edges. A *singular point* of $\mathcal{C}$ is a point that has no open neighborhood homeomorphic to an open disk, a closed half-disk, or an open segment.

Our above definition of drawings in topological spaces applies to drawings in 2-complexes. In particular, vertices of G may lie anywhere on $\mathcal{C}$, and edges of G as curves may traverse several vertices, edges, and triangles of $\mathcal{C}$. An *embedding* of G into $\mathcal{C}$ is a drawing without any intersection point.

Embeddability of graphs on 2-complexes. The *embeddability problem* takes as input a graph G and a 2-complex $\mathcal{C}$, and the task is to decide whether G has an embedding in $\mathcal{C}$. This problem is fixed-parameter tractable in the size of the input 2-complex:

▶ **Theorem 2.1** (Colin de Verdière and Magnard [12, arXiv version, Theorem 1.1][3]). *One can solve the embeddability problem (of graphs in 2-complexes) in $2^{\mathsf{poly}(C)} \cdot n^2$ time, where C is the number of simplices of the input 2-complex and n is the size of the input graph.*

3 The framework: description and results

We consider *colored graphs*, in which each *edge* is labeled with a positive integer. A *colored drawing* in a surface $\mathcal{S}$ is a drawing of a colored graph G in $\mathcal{S}$ in which each edge of the drawing inherits the color of the corresponding edge of G. Throughout this paper, let $\mathcal{S}$ be a surface, and let $\mathcal{D}$ be the set of colored drawings of graphs on $\mathcal{S}$.

The general formulation of our result is based on the following definition, in which the function κ of a drawing should be thought of as a generalization of the (traditional) crossing number — "counting the crossings", or indicating (with value ∞) that "a drawing is invalid". More precisely, κ is a quality function that does not increase upon changes that add no new crossings, and λ bounds the number of pieces into which edges are split in valid drawings:

▶ **Definition 3.1** (Crossing-counting pair). *Let $\kappa : \mathcal{D} \to \mathbb{Z}_+ \cup \{\infty\}$ be a function, and let $\lambda : \mathbb{Z}_+ \to \mathbb{Z}_+$ be a non-decreasing function. We say that (κ, λ) is a* crossing-counting pair *if:*

1. *For any $D, D' \in \mathcal{D}$, we have $\kappa(D') \le \kappa(D)$ whenever D' results from D by one of the following operations:*
   - **(a)** *the removal of a vertex or an edge;*
   - **(b)** *the addition of a vertex whose image is disjoint from the image of D (and from the boundary of $\mathcal{S}$);*
   - **(c)** *the addition of an edge, of an arbitrary color, connecting two vertices in D, and otherwise disjoint from the image of D (and from the boundary of $\mathcal{S}$);*
   - **(d)** *a self-homeomorphism of $\mathcal{S}$.*
2. *For any drawing $D^\times \in \mathcal{D}$ such that each edge carries at least one intersection point in $D^\times$ (either a self-intersection or an intersection with another edge, but common endpoints do not count) and $\kappa(D^\times) < \infty$, the number of pieces of $D^\times$ is at most $\lambda(\kappa(D^\times))$.*

[3] The conference proceedings version of [12] gives a slightly worse bound on the running time, cubic in n. The theorems that we state take into account the improvement in the latest arXiv version.

3. *One can compute* $\lambda(i)$ *in* $2^{\mathcal{O}(i)}$ *time.*

Note that Definition 3.1, Item 1, implies that κ is not affected by any of the operations (b), (c), and (d), since such operations can be reversed by an operation of type (a) or (d).

As an important special case, if $\kappa(D)$ is the number of crossings in D, or is ∞ if D is not a normal drawing (i.e., has an intersection of multiplicity at least three or a tangency intersection), and $\lambda(i) = 4i$, then (κ, λ) satisfies the conditions of Definition 3.1 (indeed, note that each piece of an edge of $D^\times$ is incident to at least one crossing point, and each crossing point involves at most four pieces). This motivates the term 'crossing-counting pair'.

We now introduce the following problem (the reader may wish to focus on the simplified case of $p = q_1 = \ldots = q_c = 0$, which implies $G = G'$):

(κ, λ)-CROSSING NUMBER (FOR THE CLASS $\mathcal{D}$ OF c-COLORED DRAWINGS, IN SURFACE $\mathcal{S}$)
Input: A colored graph G with colors in $\{1, \ldots, c\}$ for some integer c, with at most n vertices and edges, a non-negative integer r, and non-negative integers p and $q_1, q_2, \ldots, q_c$.
Question: Is there a colored drawing $D' \in \mathcal{D}$ in $\mathcal{S}$, with $\kappa(D') \leq r$, of a colored graph G', such that G' is obtained from G by performing at most p successive vertex splits and by removing, for $i = 1, \ldots, c$, at most q_i edges of color i?

Our main result is that, assuming κ is computable, the (κ, λ)-CROSSING NUMBER problem is fixed-parameter tractable in the topology of the surface, the number c of colors, the number $p + \sum_i q_i$ of allowed vertex splits and edge removals, and the value of r:

▶ **Theorem 3.2.** *Let $\mathcal{S}$ be a surface of topological size s, let $\mathcal{D}$ be the set of colored drawings of graphs in $\mathcal{S}$, and let (κ, λ) be a crossing-counting pair.*

Let δ be a non-decreasing function such that, given a representation of size i of a (colored) drawing $D \in \mathcal{D}$, one can compute $\kappa(D)$ in time $\delta(i)$.

Then the (κ, λ)-CROSSING NUMBER problem, on input instance $(G, r, p, q_1, \ldots, q_c)$, can be solved in time $2^{\mathsf{poly}(s+c+p+q+\lambda(r))} \cdot \delta(\mathcal{O}(s + \lambda(r))) \cdot n^2$, where n is the size of G and $q = \sum_i q_i$.

In all our applications, the factor $\delta(\cdot)$ in the runtime will be at most exponential in the parameters, and can thus be ignored. Moreover, for positive instances, we can compute an actual representation of the corresponding drawing; see Theorem 6.1 in Section 6.

4 Proof of Theorem 3.2

Recall that $\mathcal{D}$ denotes the set of drawings of graphs on the surface $\mathcal{S}$. Let (κ, λ) be a crossing-counting pair for drawings in $\mathcal{D}$.

4.1 Cutting a surface along a drawing

Let D be a drawing of a graph G on $\mathcal{S}$. We first need to define what we mean by *cutting* $\mathcal{S}$ along (the relative interior of the edges of) D (see Figure 2(a–b)). Intuitively, we cut $\mathcal{S}$ along the relative interior of edges of D, resulting in a topological space that is actually homeomorphic to a 2-complex. Note that each point of $\mathcal{S}$ that is an endpoint of some edge of D still corresponds to a single point on the 2-complex obtained by cutting along D. More formally, recall from Section 2 that D is represented by a pair (T, W) where T is a triangulation of $\mathcal{S}$ and W is a set of walks in the triangulation and in the interior of the surface, such that each edge of T is used at most once. The triangulation T is obtained from a collection of initially disjoint triangles by gluing some vertices and edges together. Cutting $\mathcal{S}$ along D is best described by starting with the same collection of disjoint triangles and specifying a *subset* of these gluing operations:

- whenever two directed edges e and e' of triangles are identified to a single edge in T, and that edge is not used by any walk in W, we identify e and e'. Note that this process automatically identifies the source endpoints of e and e', and similarly for their target endpoints;
- whenever two vertices of the resulting triangulation correspond to the same vertex of T, and that vertex is an endpoint of at least one walk in W, we identify these two vertices.

Since we only perform a subset of the identifications of vertices and edges in T, the resulting space is naturally a 2-complex.

4.2 Overview of the reduction

The proof of Theorem 3.2 is a parameterized Turing reduction to the embeddability problem on 2-complexes. Consider an instance $(G, r, p, q_1, \dots, q_c)$ of the (κ, λ)-Crossing Number problem. We define an uncolored graph $G_2 = G_2(G, r, p, q_1, \dots, q_c)$ and a set of 2-complexes $\Gamma = \Gamma(r, p, q_1, \dots, q_c)$ such that our (κ, λ)-Crossing Number instance is positive if and only if G_2 embeds in at least one of the 2-complexes in Γ (see a sketch in Figure 2). The set Γ is built by branching over a (small enough) set of properties for the hypothetical drawing of a graph G' (obtained from G by edge removals and vertex splits as in the problem definition) or, equivalently, by *guessing* some properties of that drawing.

Assume that G' has a colored drawing D' in $\mathcal{S}$ with $\kappa(D') \le r$. Let $D^\times$ be the subdrawing of the subgraph of G' made of the edges that are involved in at least one intersection in D'; the definition of a crossing-counting pair implies that $\kappa(D^\times) \le r$, so $D^\times$ is made of at most $\lambda(\kappa(D^\times)) \le \lambda(r)$ pieces. Viewing $D^\times$ as an abstract drawing, without its correspondence with the vertices and edges of G', this implies that we can enumerate all such colored drawings up to self-homeomorphism of $\mathcal{S}$ in time $(s + c + \lambda(r))^{\mathcal{O}(s+\lambda(r))}$. In other words, we can *guess* the appropriate colored drawing $D^\times$.

Subsequently, we cut $\mathcal{S}$ along the relative interior of the edges of $D^\times$ (Figure 2(b)) and add isolated edges connecting the endpoints of the edges in $D^\times$ (Figure 2(c)), obtaining a 2-complex $\mathcal{C}_2$ in which G' embeds. If G' is allowed to result from G by edge removals and vertex splits, we modify the 2-complex appropriately, by guessing where the endpoints of the additional isolated edges must be inserted (Figure 2(d)) and which points of the 2-complex must be identified back (doing an inverse of the vertex split, Figure 2(e)). Note that these guessed points may be picked from the surface part of $\mathcal{C}_2$, as well as among the vertices of the edges of $D^\times$ in $\mathcal{C}_2$. We obtain a 2-complex $\mathcal{C}_4$ in which G embeds. Each isolated edge of $\mathcal{C}_4$ naturally bears the color of the edge of G it carries.

It remains to ensure that, conversely, the embeddability of G into one of the resulting 2-complexes implies a positive (κ, λ)-Crossing Number instance. For this purpose, we do the following. First (Figure 3(b)), as a minor technical detail, we apply a preprocessing step — attaching a 4-clique to each vertex of G, to ensure that all vertices of G have degree at least three. Let G_1 be this new graph. Second, and more importantly (Figure 2(f) and Figure 3), we need to encode the colors of the edges. We turn each edge of G_1 and each isolated edge of the 2-complex $\mathcal{C}_4$ (in which G_1 embeds, see above) into a *necklace* encoding the color, such that a necklace of a certain encoding type in the resulting graph can only use a necklace of the same type in the resulting 2-complex. This results in an (uncolored) graph G_2 embedded in a 2-complex $\mathcal{C}_5$. Finally, the set Γ of 2-complexes in which we try to embed G_2 is made of all 2-complexes $\mathcal{C}_5$, over all possible choices (guesses) described above.

In more detail, we define a *necklace of thickness h and beads of size k* to be the graph obtained from a path of length three by (1) replacing each edge with h parallel edges, and (2) attaching k loops to each of the two internal vertices of the path (Figure 3(c)). The

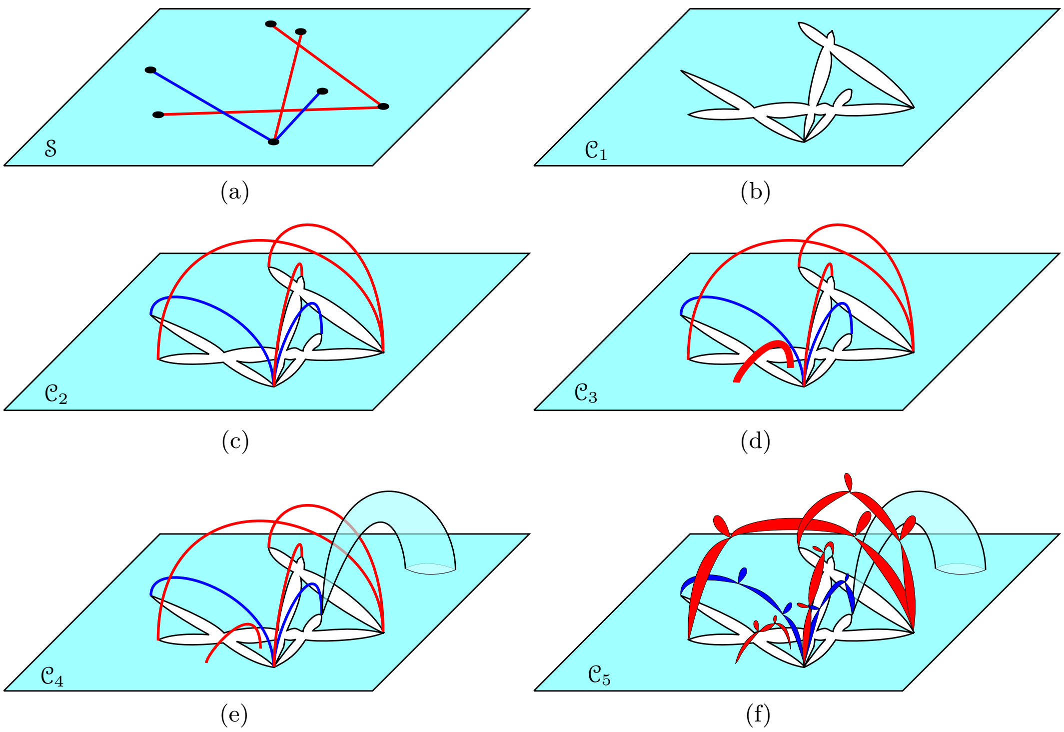

Figure 2 The construction of a 2-complex $\mathcal{C}_5$ in our reduction. (a): The surface $\mathcal{S}$ and the colored drawing $D^\times$. (b): $\mathcal{C}_1$ is obtained by cutting $\mathcal{S}$ along $D^\times$. Note that each vertex of G (in particular, the one in the bottom of the image) on $\mathcal{S}$ corresponds to a single point in $\mathcal{C}_1$, while, in this particular example, each intersection point in D of multiplicity i corresponds to $4i$ points in $\mathcal{C}_1$. (c) $\mathcal{C}_2$ is obtained by adding isolated edges corresponding to the edges of $D^\times$; we define the color of each isolated edge to be that of the corresponding edge in $D^\times$. (d) $\mathcal{C}_3$ is obtained by adding q_i isolated edges of color i for each color i. Here color 1 is represented in red, color 2 is represented in blue, and $q_1 = 1$, $q_2 = 0$. (e) $\mathcal{C}_4$ is obtained by iteratively identifying p pairs of points of $\mathcal{C}_3$ (here $p = 1$). (f) $\mathcal{C}_5$ is obtained by replacing each isolated edge of color i with a necklace (represented figuratively here) of thickness $p + 2q + 2\lambda(r) + i$ and beads of size $c - i + 2$.

intuition is that, by the minimum-degree condition on G_1, the internal vertices of necklaces of our 2-complex can essentially only be used by internal vertices of the necklaces of G_2. So, if edges of color i are encoded with necklaces of, say, thickness i and beads of size $c - i$, we effectively prevent graph necklaces from using 2-complex necklaces of different types. (We will actually use slightly refined formulas for the thickness and the size of the beads.)

4.3 Formal description of the reduction

We now define our reduction precisely. We first describe the uncolored graph $G_2 = G_2(G, r, p, q_1, \ldots, q_c)$, which will be the common input to all embeddability instances. See Figure 3.

- Let G_1 be the colored graph obtained from G by attaching a 4-clique to each vertex v of G, that is, by adding three new vertices making a 4-clique with v. The edges of each attached 4-clique get color $c + 1$. (This is a new color, and we assume $q_{c+1} = 0$.)
- Let G_2 be the uncolored graph obtained from G_1 by replacing each edge of color i with a necklace of thickness $p + 2q + 2\lambda(r) + i$, where $q = \sum_j q_j$, and beads of size $c - i + 2$.

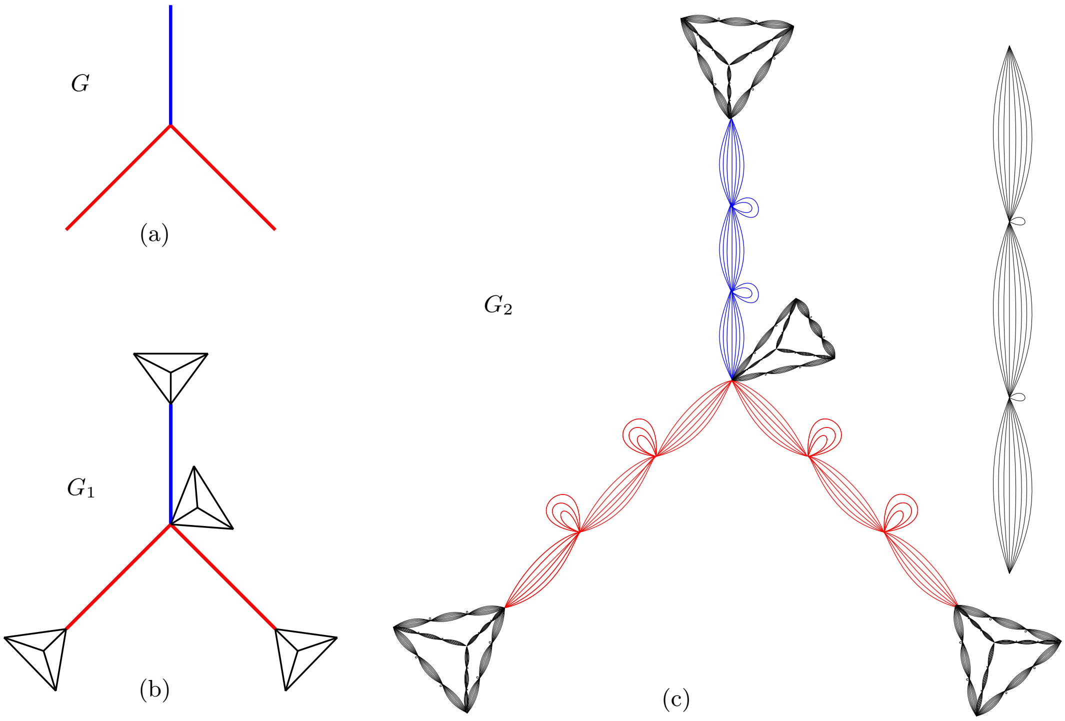

Figure 3 The construction of the graph G_2. (a): The graph G; the red edges have color 1 and the blue edge has color 2. (b): The graph G_1. The new edges, in black, have color 3. (c): The graph G_2 in the case $c = 2$, $\lambda(r) = p = q = 1$. Each edge of color i is replaced with a necklace of thickness $p + 2q + 2\lambda(r) + i = 5 + i$ and beads of size $c - i + 2 = 4 - i$. On the right, a close-up on the replacement for a black edge in G_1, which is a necklace of thickness eight and beads of size one.

Since we have a new color $c + 1$, we extend κ as follows: For a drawing D, if at least one edge of color $c+1$ is involved in an intersection, we let $\kappa(D) = \infty$; otherwise, we let $\kappa(D)$ to be equal to $\kappa(D')$, where D' is obtained from D by removing the edges with color $c+1$. It is straightforward to check that (κ, λ) is still a crossing-counting pair.

Let us now describe our set Γ of 2-complexes in which we try to embed G_2. See Figure 2. These 2-complexes depend on several successive choices. First, let $D^\times \in \mathcal{D}$ be a colored drawing on $\mathcal{S}$ such that $\kappa(D^\times) \leq r$ and such that each edge carries at least one intersection in the drawing.

- Let $\mathcal{C}_1$ be obtained from $\mathcal{S}$ by cutting $\mathcal{S}$ along the relative interior of each edge of $D^\times$.
- Let $\mathcal{C}_2$ be obtained from $\mathcal{C}_1$ by doing the following: For each edge of $D^\times$ of color i, connecting points x and y of $\mathcal{C}_1$, we add an isolated edge connecting x and y to the 2-complex, which is said to have *color* i. (This may require subdividing $\mathcal{C}_1$ to turn x and y into vertices.)

Let U be the set of points of $\mathcal{C}_2$ that correspond to a single point of the interior of $\mathcal{S}$. In detail, U is made of (1) the points of $\mathcal{C}_2$ that correspond to an endpoint of an edge in $D^\times$, and (2) the interior of the surface part of $\mathcal{C}_2$.

For each $i = 1, \dots, c$, let $(x_i^1, y_i^1), \dots, (x_i^{q_i}, y_i^{q_i})$ be pairs of points on U. (These points are not necessarily distinct. There are finitely many choices up to a self-homeomorphism of $\mathcal{S}$.)

- Let $\mathcal{C}_3$ be obtained from $\mathcal{C}_2$ by doing the following: For each $i = 1, \dots, c$ and each $j = 1, \dots, q_i$, we add to the 2-complex an isolated edge connecting x_i^j and y_i^j, which is said to have *color* i.

For each unordered partition $p = p_1 + \ldots + p_k$ of the integer p, where $p_1, \ldots, p_k$ are positive integers, and for each $i = 1, \ldots, k$, let $z_i^1, \ldots, z_i^{p_i+1}$ be $p_i + 1$ points on U (which is also a part of $\mathcal{C}_3$), such that all these $p + k$ points are pairwise distinct. (Note that the points z_i^j are not necessarily distinct from the points $x_{i'}^{j'}$ and $y_{i'}^{j'}$.)

- Let $\mathcal{C}_4$ be obtained from $\mathcal{C}_3$ by doing the following: For each $i = 1, \ldots, k$, we identify the points $z_i^1, \ldots, z_i^{p_i+1}$ to a single point z_i.
- Let $\mathcal{C}_5$ be obtained from $\mathcal{C}_4$ by replacing each isolated edge e of $\mathcal{C}_1$ of color i with a necklace of thickness $p + 2q + 2\lambda(r) + i$ and beads of size $c - i + 2$.

To conclude, $\mathcal{C}_5$ depends on the choice of the drawing $D^\times \in \mathcal{D}$, of the points x_i^j, y_i^j, of the partition of p, and of the points z_i^j. We let $\Gamma = \Gamma(r, p, q_1, \ldots, q_c)$ to be the set of 2-complexes $\mathcal{C}_5$ that can be obtained in this way, over all these choices.

4.4 Correctness of the reduction

The correctness of the reduction is established in the two following lemmas.

▶ **Lemma 4.1.** *Assume that there is a drawing D' in $\mathcal{D}$, with $\kappa(D') \leq r$, of a colored graph G' obtained from G by performing at most p successive vertex splits and by removing, for $i = 1, \ldots, c$, at most q_i edges of color i. Then $G_2(G, r, p, q_1, \ldots, q_c)$ embeds in some 2-complex in $\Gamma(r, p, q_1, \ldots, q_c)$.*

Proof. Let $D^\times$ be the colored subdrawing of D' made of the edges that carry at least one intersection in D', together with their endpoints. Since (κ, λ) is a crossing-counting pair, we have $\kappa(D^\times) \leq r$. Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be defined as above, for that specific choice of $D^\times$. By construction, G' embeds into $\mathcal{C}_2$ in such a way that each isolated edge of $\mathcal{C}_2$ of color i coincides with an edge of G' of color i.

It follows that, for an appropriate choice of the partition of p and of the vertices x_i^j, y_i^j, z_i^j, the graph G embeds into $\mathcal{C}_4$ (again, we reuse the notations introduced in the description of the reduction), in such a way that each isolated edge of $\mathcal{C}_4$ of color i coincides with an edge of G of color i or is not used by the embedding.

Note that, by the construction, in that embedding the vertices of G are mapped to the surface part of $\mathcal{C}_4$, and thus we can attach a 4-clique to each vertex of G while preserving the fact that we have an embedding. Thus, the graph G_1 embeds into $\mathcal{C}_4$ too, in such a way that each isolated edge of $\mathcal{C}_4$ of color i coincides with an edge of G of color i or is not used at all.

Finally, replacing each isolated edge of $\mathcal{C}_4$ of color i, and similarly each edge of G_1 of color i, with a copy of a necklace of thickness $p + 2q + 2\lambda(r) + i$, where $q = \sum_j q_j$, and beads of size $c - i + 2$, implies that G_2 embeds into $\mathcal{C}_5$. Indeed, each edge of G_1 that coincided with an isolated edge of $\mathcal{C}_4$ is replaced with a necklace that can be embedded in the corresponding necklace of $\mathcal{C}_5$; and each edge of G_1 that belonged to the surface part of $\mathcal{C}_4$ is replaced with a necklace that can be drawn without crossings in a neighborhood of the edge of G_1. This concludes the proof. ◀

Conversely:

▶ **Lemma 4.2.** *Assume that the graph $G_2(G, r, p, q_1, \ldots, q_c)$ embeds in some 2-complex in $\Gamma(r, p, q_1, \ldots, q_c)$. Then, there exists a drawing D' in $\mathcal{D}$, with $\kappa(D') \leq r$, of a colored graph G' obtained from G by performing at most p successive vertex splits and by removing, for $i = 1, \ldots, c$, at most q_i edges of color i.*

Proof. Let $G_2 := G_2(G, r, p, q_1, \ldots, q_c)$, and let $\mathcal{C}_5$ be some 2-complex in $\Gamma(r, p, q_1, \ldots, q_c)$ in which G_2 can be embedded. The 2-complex $\mathcal{C}_5$ is defined by the choice of a drawing $D^\times$

(such that $\kappa(D^\times) \leq r$, and each edge is involved in at least one intersection), of pairs (x_i^j, y_i^j), of a partition of p, and of the points z_i^j. We also freely use the notations $G_1, \mathcal{C}_1, \ldots, \mathcal{C}_5$ used in the description of the reduction.

First, in a series of claims, we prove that we can assume that each necklace of $\mathcal{C}_5$ is either not used at all by G_2, or used by a necklace of G_2 of the same type.

- Let v be a vertex of G. We claim that, in the embedding of G_2 in $\mathcal{C}_5$, v cannot be mapped to the interior of a necklace N of $\mathcal{C}_5$. Indeed, otherwise, since v has degree at least three in G_2, it is mapped to some interior vertex of N (i.e., not to an interior point of an edge of N). Moreover, v, as a vertex in G_2, has at least three distinct neighbors, each of degree at least three, and thus mapped either to a vertex of N, or to the complement of N. But, starting from the location of v in N, there do not exist three disjoint paths going to different vertices of N or to the complement of N. This contradiction proves the claim.
- Assume now that some vertex v of G_2 is mapped to the interior of a necklace N of $\mathcal{C}_5$, of thickness $p + 2q + 2\lambda(r) + i$ (where, again, $q = \sum_j q_j$), and beads of size $c - i + 2$, for some j. We claim that N actually contains a necklace of G_2 of the same type. Indeed, by the previous paragraph, v is not a vertex of G, and is thus an interior vertex of a necklace of G_2. Moreover, v has degree at least three, so it is mapped to an interior vertex x of N. This vertex x of N is connected to each of its two distinct neighbors by $p+2q+2\lambda(r)+i$ parallel edges, and is incident to $c - i + 2$ loops. Moreover, v is part of a necklace of thickness $p+2q+2\lambda(r)+j$ and beads of size $c-j+2$, for some j. Considerations similar to those of the previous paragraph imply that $i = j$. Arguing similarly with the neighbors of v that are not vertices of G yields the result.
- We claim that without loss of generality, if the endpoints of an edge e of G_2 avoid the interior of all necklaces of $\mathcal{C}_5$, then e avoids all necklaces of $\mathcal{C}_5$. To prove this, there are two cases:
  - Assume that e is a loop (part of a "bead"). Let v be the vertex of G_2 incident to e. Since v has degree at least three, v is mapped to the surface part of $\mathcal{C}_5$. So e can be redrawn to a small neighborhood of v, while preserving the fact that we have an embedding of G_2 into $\mathcal{C}_5$.
  - Assume that e is not a loop. First we remark that $\mathcal{C}_5$ has at most $p + 2q + 2\lambda(r)$ singular points, excluding the singular points in the interior of a necklace. Then, we note that e is part of a batch of at least $p + 2q + 2\lambda(r) + 1$ parallel edges, at least one of which, say e', has its relative interior that avoids all of the at most $p + 2q + 2\lambda(r)$ singular points described above. We can thus reroute e parallel to e', avoiding all necklaces, while preserving the fact that we have an embedding of G_2 into $\mathcal{C}_5$.

As a summary of the above considerations, we now have an embedding of G_2 into $\mathcal{C}_5$ in such a way that each necklace of $\mathcal{C}_5$ is either not used, or used by a necklace of G_2 of the same type. This implies that G_1 embeds in $\mathcal{C}_4$ in such a way that each isolated edge of $\mathcal{C}_4$ of color i coincides with an edge of G_1 of color i, or is not used at all. The same holds for G, since G is a subgraph of G_1.

Recall that $\mathcal{C}_4$ is obtained by identifying points of $\mathcal{C}_3$ into k points $z_1, \ldots, z_k$. Let S be the set of indices i such that z_i contains a point in the relative interior of an edge of G. We do the following:

- While, for some $i \in S$, z_i contains a point of the relative interior of an edge $e = uv$ of G, we do the following. First, we perform a vertex split of one of its vertices, say v, by replacing uv with uv', where v' is a new vertex. Second, we shrink edge uv' without moving u, so that in the new embedding of G, edge uv' does not contain any singular

point of $\mathcal{C}_4$ in its relative interior. The number of vertex splits performed so far is at most the size of S, and after this step, the image of G avoids the points z_i, $i \in S$.
- Now, for each $i = 1, \ldots, k$, we split z_i into p_i+1 points in order to obtain the 2-complex $\mathcal{C}_3$. In passing, if vertex z_i is used by a vertex v of G, we perform at most p_i splits of G at v in order to preserve the property that we have an embedding of some graph.

We now have an embedding, into $\mathcal{C}_3$, of a graph G'' obtained from G by at most $|S|+\sum_{i\in S} p_i \leq p$ vertex splits. Moreover, as above, each isolated edge of $\mathcal{C}_3$ of color i coincides with an edge of G'' of color i, or is not used at all.

If we remove the edges embedded in $\mathcal{C}_3 \setminus \mathcal{C}_2$, we obtain an embedding, in $\mathcal{C}_2$, of a graph G' obtained from G by at most p vertex splits and by removing, for each i, at most q_i edges of color i. Moreover, again, each isolated edge of $\mathcal{C}_2$ of color i coincides with an edge of G' of color i, or is not used at all.

Finally, we obtain a drawing D' of G' into $\mathcal{S}$ that results from $D^\times$ by removing vertices and edges, and by adding vertices and noncrossing edges. Thus, $\kappa(D') \leq \kappa(D^\times) \leq r$. ◀

4.5 End of proof

The two preceding lemmas show that a (κ, λ)-Crossing Number instance $(G, r, p, q_1, \ldots, q_c)$ on surface $\mathcal{S}$ is positive if and only if $G_2(G, r, p, q_1, \ldots, q_c)$ embeds in at least one of the 2-complexes in $\Gamma(r, p, q_1, \ldots, q_c)$. To conclude the proof, we show that the reduction described above turns into an actual algorithm, and analyze its complexity. We first state and prove the result announced in Section 2 on the size of a representation of a drawing:

▶ **Lemma 4.3.** *Let D be a drawing of a graph on a surface $\mathcal{S}$. Then D is represented by some pair (T, W) in which T has at most $96(s + u)$ triangles, where s is the topological size of $\mathcal{S}$, and the intersection points of D subdivide the edges of D into u pieces in total.*

Proof. We view D as a graph G embedded in the interior of $\mathcal{S}$. The vertices of G are the endpoints of the edges of D and the intersection points of D; the edges of G are the pieces of the edges of D. A *weak triangulation* is a graph T embedded on $\mathcal{S}$ such that each face of T (each connected component of the complement of the image of T) is homeomorphic to a disk bounded by three edges, but this time some incident vertices and/or edges may coincide; in other words, a weak triangulation is not necessarily a 2-complex. We gradually augment G to a weak triangulation of $\mathcal{S}$. Recall that the topological size of $\mathcal{S}$ equals $s = d + g + b$, where d is the number of connected components, g is the genus, and b is the number of connected components of $\mathcal{S}$. In a first case, we assume that $\mathcal{S}$ is connected and intersects D.

We first cover each of the b boundary components of $\mathcal{S}$ with a loop. Now, consider a face f of this new graph (still denoted G). It has at least one boundary component, and each boundary component contains at least one vertex. Recall that G has no isolated vertex; let $\ell' \geq 1$ be the number of edges on the boundary of f. We distinguish several cases:
- If f is homeomorphic to a disk and $\ell' = 1$, then we insert a new vertex in the interior of the disk, and connect it to the unique vertex on the boundary of f, turning f into $\ell' = 1$ triangle;
- if f is homeomorphic to a disk and $\ell' = 2$, then we insert a new vertex in the interior of the disk, and connect it to each of the two vertices on the boundary of f, turning f into $\ell' = 2$ triangles;
- otherwise, let $g' \geq 0$ and $b' \geq 1$ be the genus and number of boundary components of f. We now summarize a standard argument to triangulate f without adding vertices: One can add g' loops to decrease the genus of f to zero, and there are still b' boundary components, now with $\ell' + 2g'$ edges in total. Adding $b' - 1$ edges, we connect all the

boundary components of f into a single one; f is now a disk with $\ell' + 2g' + 2b' - 2$ sides, which is at least three (otherwise f would be of one of the first two types). In turn, this disk can be (weakly) triangulated into $\ell' + 2g' + 2b' - 4$ triangles.

We now have a weak triangulation of f with at most $\ell' + 2g' + 2b'$ triangles. Initially, G has at most $2u$ faces, and moreover, $\sum_f g' \leq g$, $\sum_f b' \leq b + 2u$, and $\sum_f \ell' \leq 2u + b$. So, to conclude the first case where $\mathcal{S}$ is connected and intersects D, we obtain that G is weakly triangulated using at most $2u + b + 2g + 2b + 4u = 6u + 2g + 3b$ triangles.

In the second case, if $\mathcal{S}$ is still connected but does not intersect D, we just insert a single vertex and a loop in G and apply the above process, weakly triangulating $\mathcal{S}$ with at most $6 + 2g + 3b$ triangles.

So, if $\mathcal{S}$ is connected, we can weakly triangulate it using at most $6 + 2g + 3b + 6u$ triangles. To turn this weak triangulation into a triangulation, it suffices to apply barycentric subdivisions twice; the number of triangles is multiplied by 16.

Finally, we obtain the result by summing up over all connected components. ◀

Recall that we consider drawings up to homeomorphism of $\mathcal{S}$. The next lemma explains how to enumerate (representations of) the drawings $D^\times$.

▶ **Lemma 4.4.** *Let S be the set of colored drawings $D^\times$ on $\mathcal{S}$ such that $\kappa(D^\times) \leq r$ and each edge of $D^\times$ carries at least one intersection in $D^\times$. In $(s+c+\lambda(r))^{\mathcal{O}(s+\lambda(r))} \cdot \delta(96(s+\lambda(r)))$ time, we can enumerate a set of $(s+c+\lambda(r))^{\mathcal{O}(s+\lambda(r))}$ representations of drawings in S, such that each drawing in S is represented at least once, and each representation has size $\mathcal{O}((s+\lambda(r))c)$.*

Proof. Since (κ, λ) is a crossing-counting pair, any drawing $D^\times$ such that $\kappa(D^\times) \leq r$ has at most $\lambda(r)$ pieces; so $D^\times$ can be represented by a pair (T, W) such that T has at most $96(s+\lambda(r))$ triangles by Lemma 4.3.

We first compute $\lambda(r)$ in $2^{\mathcal{O}(\lambda(r))}$. Then, starting with at most $96(s+\lambda(r))$ triangles, we identify some of their sides in pairs, obtaining a triangulation T; there are $(s+\lambda(r))^{\mathcal{O}(s+\lambda(r))}$ possibilities. For each such choice, we compute the genus, orientability, and number of boundary components of each connected component of the resulting surface using standard methods (in particular the Euler characteristic), and discard the surface if it is not homeomorphic to $\mathcal{S}$, all in time polynomial in $s + \lambda(r)$.

Then, and we choose an arbitrary set of edge-disjoint walks W on each of the remaining triangulations; there are $(s+\lambda(r))^{\mathcal{O}(s+\lambda(r))}$ possibilities. We also assign a color to each of these walks, adding a factor of $\mathcal{O}(c^{\lambda(r)})$ in the number of choices. We discard the choices in which some walk has no intersection with any other walk.

Finally, for each of the $(s+c+\lambda(r))^{\mathcal{O}(s+\lambda(r))}$ resulting representations of colored drawings $D^\times$, we compute $\kappa(D^\times)$ in $\delta(96(s+\lambda(r)))$ time, and discard those with $\kappa(D^\times) > r$. ◀

Proof of Theorem 3.2. The algorithm tests the embeddability of $G_2 = G_2(G, r, p, q_1, \ldots, q_c)$ into each 2-complex in $\Gamma = \Gamma(r, p, q_1, \ldots, q_c)$, and returns "yes" if and only if at least one embeddability test succeeded. The correctness of the reduction follows from Lemmas 4.1 and 4.2. There remains to provide some details on the algorithm and the runtime.

G_2 has size $\mathcal{O}((c+p+q+\lambda(r))n)$ and can be computed in $\mathcal{O}((c+p+q+\lambda(r))n)$ time.

Using Lemma 4.4, we can compute a set of $(s+\lambda(r))^{\mathcal{O}(s+c+\lambda(r))}$ representations of the drawings $D^\times$ appearing in the definition of the set of 2-complexes Γ in $\mathcal{O}(s+c+\lambda(r))^{\mathcal{O}(s+\lambda(r))} \cdot \delta(96(s+\lambda(r)))$ time, and each representation has size $\mathcal{O}((s+\lambda(r))c)$. Given such a representation, we can compute $\mathcal{C}_1$ and then $\mathcal{C}_2$ in polynomial time, each of size $\mathcal{O}(s+\lambda(r))$ (ignoring the colors on the isolated edges).

The number of partitions of the integer p is $2^{\mathcal{O}(p)}$. We select the points x_i^j, y_i^j, and z_i^j on U. There are at most $\lambda(r)$ points in U that correspond to an endpoint of an edge in $D^\times$; moreover $D^\times$ has at most $s+\lambda(r)$ faces in $\mathcal{S}$, and any two points in the same face are equivalent (up to a self-homeomorphism of $\mathcal{S}$) as far as the choice of the point is concerned. There are at most $2p+2q$ points, and thus at most $(s+2\lambda(r)+2p+2q)^{2p+2q}$ choices in total. Thus, for a given $D^\times$, there are $(s+p+q+\lambda(r))^{\mathcal{O}(p+q)}$ possible choices of 2-complexes $\mathcal{C}_5$.

To compute the 2-complex corresponding to a given choice, if necessary, we first refine the 2-complex $\mathcal{C}_2$ by subdividing triangles so that the points x_i^j, y_i^j, and z_i^j become vertices of the new 2-complex; this increases its size by $\mathcal{O}(p+q)$. The 2-complex $\mathcal{C}_3$ is obtained from $\mathcal{C}_2$ by adding q isolated edges, which takes time $\mathcal{O}(q)$ since the endpoints of these edges are now vertices of the 2-complex. Similarly, $\mathcal{C}_4$ is obtained by performing p identifications of vertices, which also takes time $\mathcal{O}(p)$. Finally, $\mathcal{C}_5$ is obtained from $\mathcal{C}_4$ by replacing the isolated edges with necklaces, which increases its size by $\mathcal{O}(c+p+q+\lambda(r))\lambda(r)$.

So there are $(s+c+p+q+\lambda(r))^{\mathcal{O}(s+c+p+q+\lambda(r))}$ 2-complexes in Γ, each of size $\mathcal{O}(s+(c+p+q+\lambda(r))\lambda(r))$, and the computation of Γ takes $\mathcal{O}((s+c+p+q+\lambda(r))^{\mathcal{O}(s+c+p+q+\lambda(r))}\cdot\delta(96(s+\lambda(r))))$ time.

Finally, we test the embeddability of G_2 into each 2-complex in Γ using Theorem 2.1, which takes $2^{\mathsf{poly}(s+c+p+q+\lambda(r)}\cdot\delta(96(s+\lambda(r)))\cdot n^2$ time. ◀

5 Applications: FPT algorithms for diverse crossing number variants

In this section, we demonstrate the wide applicability of our framework to crossing number problems. First, as a toy example, we reprove that we can compute the traditional crossing number in quadratic FPT time on surfaces [12]:

▶ **Proposition 5.1.** *Let $r\geq 0$ be an integer, and $\mathcal{S}$ a surface with topological size s. Deciding whether an input graph G of size n has crossing number at most r in $\mathcal{S}$ can be solved in time $2^{\mathsf{poly}(s+r)}\cdot n^2$.*

Proof. We apply our framework with a single color ($c=1$), no vertex splits or edge removals ($p=q_1=0$), and in the case where $\mathcal{D}$ is the set of drawings of graphs on $\mathcal{S}$. For any drawing D, let $\kappa(D)$ be the number of crossings in D, or ∞ if D is not a normal drawing. Moreover, let $\lambda(i)=4i$. Since $\kappa(D)<\infty$ only if D is normal, (κ,λ) is a crossing-counting pair, as noted in Section 3. Theorem 3.2 implies the result. ◀

We now turn to the promised applications. We survey, in order, specific drawing styles, colored crossing number problems, various counting methods, and problems where prior edge removals and vertex splits are allowed. Throughout this section, again, let $\mathcal{S}$ be a surface and $\mathcal{D}$ the set of colored drawings of graphs on $\mathcal{S}$. In our reductions to the (κ,λ)-Crossing Number problem, we always set $p=q_1=\ldots=q_c=0$ except in Section 5.4. In nearly all cases, only normal drawings of graphs are considered. Note that in all cases, for positive instances, we can compute a corresponding drawing using Theorem 6.1.

5.1 Variations on the traditional crossing number

k-planar drawings. Introduced by Pach and Tóth [40], *k-planar drawings* (for $k\geq 0$) are normal drawings in the plane such that every edge is involved in at most k crossings. For any drawing D, let $\kappa(D)$ be the number of crossings in D, if D is a normal k-planar drawing,

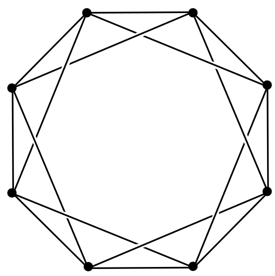

Figure 4 A drawing of some graph that is a 1-gap drawing, as certified by the casing.

and $\kappa(D) = \infty$ otherwise. The *k-planar crossing number* is then the smallest value of $\kappa(D)$ over all normal drawings D of G. A graph may have no k-planar drawing at all, and then its k-planar crossing number is ∞ by convention. Already for $k = 1$, deciding the existence of a 1-planar drawing is NP-hard even for very restricted inputs [8, 19]. This definition naturally extends to an arbitrary surface $\mathcal{S}$, which we call the *k-surface crossing number* in $\mathcal{S}$.

k-quasi-planar and min-k-planar drawings. There are similar concepts of *k-quasi-planar* [1] (no k edges pairwise cross) and of *min-k-planar* [6] (for every two edges with a common crossing, one carries at most k crossings) normal drawings in the plane, which give rise to the corresponding crossing number flavors in the plane, and we may likewise generalize them to the *k-quasi-surface* and *min-k-surface crossing numbers* in a surface $\mathcal{S}$.

Many authors restrict the problems to simple graphs and require the admissible drawings to be simple; this restriction defines the *simple-drawing k-surface* (*k-quasi-surface*, *min-k-surface*) crossing number(s). These simple-drawing variants can be very different from the non-simple ones, see, e.g., [44, Chapter 7] and [27], but we will show that they also smoothly fit into our framework.

k-gap drawings. Another known concept is that of *k-gap crossing number* [3, 46], minimizing the number of crossings over *k-gap drawings* in the plane (Figure 4), namely, normal drawings D of a graph G admitting a mapping from each crossing in D to one of the two involved edges such that at most k crossings are mapped to each edge of G. This problem is NP-hard for $k = 1$ [3]. Again, the k-gap crossing number is naturally generalized to any surface $\mathcal{S}$.

Applying our framework. Given $r \geq 0$, deciding whether a graph G has k-planar (k-quasi-planar, or min-k-planar) crossing number at most r can be done in quadratic FPT-time [37] using Courcelle's theorem. The parameterized complexity of these three problems in other surfaces, and of the k-gap crossing problem altogether, has not been studied prior to us. Even in the known planar cases, we largely improve the dependence in r in the runtime:

▶ **Proposition 5.2.** *Let $k \geq 1$ and $r \geq 0$ be integers, $\mathcal{S}$ a surface with topological size s, and let $\Pi \in \{$'k-surface', 'k-quasi-surface', 'min-k-surface', 'k-gap'$\}$. Deciding whether an input graph G of size n has Π-crossing number at most r in $\mathcal{S}$ can be solved in time $2^{\mathsf{poly}(s+r)} \cdot n^2$. The same conclusion holds for the simple-drawing Π-crossing number of a simple input graph G.*

Proof. The proof relies on Theorem 3.2 with a single color ($c = 1$) and no splits or removals ($p = q_1 = 0$). For each Π, we must give an appropriate crossing-counting pair (κ, λ) such that κ is efficiently computable. Let $\kappa = \kappa_\Pi$ map a drawing D to ∞ if it is not a normal

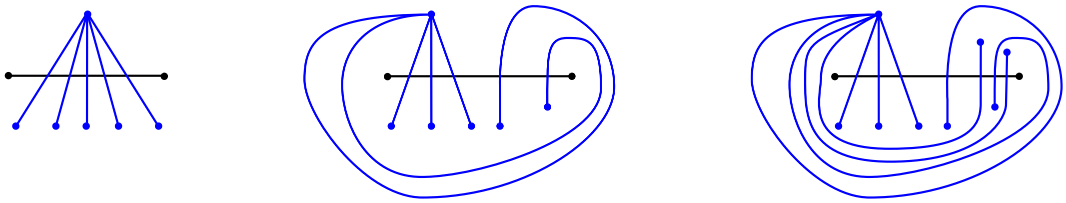

(a) (b) (c)

Figure 5 (a) A strongly fan-planar drawing. (b) A weakly fan-planar drawing that is not strongly fan-planar. (c) A fan-crossing drawing that is not weakly fan-planar.

drawing or if it violates the property Π, and to the number of crossings of D otherwise. In particular, if k is greater than the number of crossings in D, then D always satisfies Π.

Given a drawing D, one may in time polynomial in the size of a representation of D verify that D is a normal drawing satisfying Π, and so compute $\kappa_\Pi(D)$. This is an easy routine in all cases of Π except when $\Pi =$ 'k-gap'; in the latter case we additionally employ a standard maximum flow algorithm to decide the existence of a mapping from the crossings to their incident edges respecting "capacities" of edges to accept at most k crossings each.

Since we only consider normal drawings, $\lambda(i) = 4i$ satisfies the assumption of Definition 3.1. Likewise, it is easy to see that κ_Π is non-increasing under the operations listed in Definition 3.1. Hence, our conclusion follows by an application of Theorem 3.2 to (κ_Π, λ) and the graph G.

Next, we consider the simple-drawing Π-crossing number, assuming G is simple. We now, additionally, set $\kappa_\Pi(D) = \infty$ if D is a non-simple drawing. For this case, it is important that our definition of a simple drawing does not immediately exclude drawings of non-simple graphs; while this is irrelevant for our simple graph G, fulfillment of Definition 3.1, Item 1(c), depends on this technical twist. The rest of the argument is identical. ◀

Fan-crossing and fan-planar drawings. A normal drawing is *fan-crossing* [31] if all edges crossing the same edge are incident to a common vertex. One can restrict this notion to *strongly* or *weakly fan-planar* drawings, which not necessarily lead to the same crossing number problems, see Figure 5 and Cheong, Förster, Katheder, Pfister, and Schlipf [9]. The *fan-crossing*, *strongly fan-planar*, and *weakly fan-planar crossing numbers* are the minimum number of crossings over all strongly fan-planar, weakly fan-planar, and fan-crossing drawings in the plane, respectively.

Formally, a fan-crossing drawing is called *weakly fan-planar* if, whenever edges f_1 and f_2 with a common end v both cross an edge e, the (sub)arcs of both f_1, f_2 from v to the crossing with e lie on the same side of e; and it is called *strongly fan-planar* if, moreover, in the described situation of f_1, f_2 both crossing e, the union $e \cup f_1 \cup f_2$ in the drawing does not enclose both ends of e (in the plane). We refer to Cheong et al. [9] for a closer discussion.

FPT algorithms, all using Courcelle's theorem, exist for the fan-crossing crossing number, by Münch and Rutter [37], and for the weakly fan-planar crossing number, by Hamm, Klute, and Parada [22], and it seems that with further adjustments they ([22]) could also handle the strongly fan-planar version. We solve all three variants, again with a better runtime. For the strongly fan-planar version, we use the fact that our framework allows for surfaces with boundary, and thus can easily capture the "infinite face" of the drawing; indeed, strong fan-planarity is invariant under self-homeomorphisms of the plane, but not of the sphere (e.g., the drawings in Figure 5(a) and (b) are homeomorphic when seen on the sphere).

▶ **Proposition 5.3.** *Let $r \geq 0$ be an integer. The problem to decide whether an input graph G of size n has strongly fan-planar, weakly fan-planar, or fan-crossing crossing number at most r in the plane can be solved in time $2^{\mathsf{poly}(r)} \cdot n^2$.*

Proof. The proof is very similar to that of Proposition 5.2, specialized to the plane. We only need to be able to recognize strongly fan-planar, weakly fan-planar, or fan-crossing drawings in polynomial time, which is not difficult directly from their definitions. ◀

Extensions of Proposition 5.3 to arbitrary surfaces are obviously possible.

Drawings that may not be normal. To illustrate the applicability of our framework to non-normal drawings, we introduce the following example. For an integer $k \geq 2$, we define the *k-intersecting crossing number* of a graph G to be the minimum value r such that there is a drawing of G with at most r intersection points in which the multiplicity of every intersection point is at most k. For every graph G, the traditional crossing number of G is at most $\binom{k}{2}$ times its k-intersecting crossing number, and this bound can be tight. We prove:

▶ **Proposition 5.4.** *Let $k \geq 2$ and $r \geq 0$ be integers, $\mathcal{S}$ a surface with topological size s. The problem to decide whether an input graph G of size n has k-intersecting crossing number at most r in $\mathcal{S}$ can be solved in time $2^{\mathsf{poly}(s+k+r)} \cdot n^2$.*

Proof. We apply Theorem 3.2 in a setup very similar to the proof of Proposition 5.2; we only need a different definition of the crossing-counting pair (κ, λ). Let $\kappa(D) = \kappa_k(D)$ be equal to ∞ if D has an intersection point of multiplicity greater than k, and to the number of intersection points of D otherwise. Setting $\lambda(i) = \lambda_k(i) = 2ki$, it is easy to check that the pair (κ_k, λ_k) fulfills Definition 3.1. Hence, our conclusion follows again by applying Theorem 3.2 to (κ_k, λ_k) and graph G. ◀

Note that in this problem, it does matter that we allow tangential intersections of edges in our framework (however, if a restricted version of k-intersecting crossing number forbidding tangential intersections was considered, our framework would handle it smoothly as well.).

We refer to Section 5.3 for more examples of different methods for counting crossings.

5.2 Crossing problems with colored edges

We now turn our attention to problems in which the colors of intersecting edges play a role. As a natural example, we choose the joint crossing number problem introduced a while ago by Negami [38].

Joint crossing number and generalizations. In the *joint crossing number* problem [38], the goal is to embed two input graphs simultaneously on the same surface while minimizing the number of crossings between them. It is NP-hard for any fixed Euler genus $g \geq 6$ [25, 28], and the complexity parameterized by the solution value has not been studied prior to our work. We actually solve the following new problem generalizing the joint crossing number:

Color-Constrained Crossing Problem (in surface $\mathcal{S}$)
Input: A colored graph G with colors in $\{1, \ldots, c\}$, and a symmetric matrix M of size $c \times c$ with nonnegative integer values.
Question: Is there a normal colored drawing D of G in $\mathcal{S}$ such that, for each pair $i, j \in \{1, \ldots, c\}$, the number of crossings involving two edges, one colored i and the other j, is at most $M_{i,j}$?

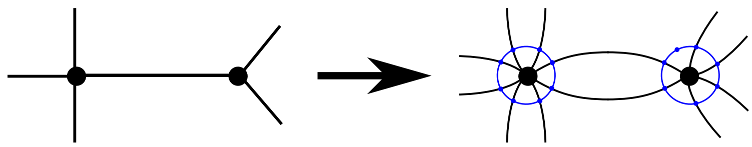

Figure 6 The hardening of a graph (fixing the rotation system in an orientable surface).

▶ **Proposition 5.5.** *Let $\mathcal{S}$ be a surface of with topological size s. The* Color-Constrained Crossing Problem *in $\mathcal{S}$ of the input graph G of size n and matrix M can be solved in time $2^{\mathsf{poly}(s+r)} \cdot n^2$, where $r = \sum_{1\le i\le j\le c} M_{i,j}$.*

Consequently, one can decide whether the joint crossing number of graphs G_1 and G_2, each embeddable in $\mathcal{S}$, is at most r in time $2^{\mathsf{poly}(s+r)} \cdot n_0^2$, where n_0 is the size of $G_1 \cup G_2$.

Proof. Without loss of generality, we assume that $c \le r+1$, since merging the colors corresponding to empty rows of M will not change the problem. Again, as in the proof of Proposition 5.2, we have to define an appropriate crossing-counting pair (κ_M, λ). Let $\kappa_M(D)$ be equal (straightforwardly) to the number of crossings in a normal drawing D, except that $\kappa_M(D) = \infty$ if D is not normal, or if for some pair $i, j \in \{1, \dots, c\}$, the number of crossings involving two edges, one colored i and the other j, is larger than $M_{i,j}$. This is easily computable in polynomial time in the representation of D. Moreover, κ_M is non-increasing under the operations listed in Definition 3.1. Letting $\lambda(i) = 4i$, Theorem 3.2 clearly applies to (κ_M, λ) and the graph G, and yields the result.

In the case of the joint crossing number problem, for $i = 1, 2$, we use color i for G_i, define G to be the disjoint union of G_1 and G_2, let $M = \begin{pmatrix} 0 & r \\ r & 0 \end{pmatrix}$, and apply the previous result. ◀

Homeomorphic joint crossing number. In the related *homeomorphic joint crossing number* problem in an orientable surface $\mathcal{S}$ without boundary, the input graphs G_1 and G_2 are cellularly embedded in $\mathcal{S}$ (every face is homeomorphic to a disk) and the solution must be composed of embeddings of G_1 and G_2 that are each homeomorphic to the input one (the rotation systems must be the same, possibly up to reversal). Given an integer $r \ge 0$, the goal is to determine whether there is a way to achieve this with at most r crossings between G_1 and G_2. The problem is also NP-hard if the Euler genus of $\mathcal{S}$ is at least 12 [28]. We prove:

▶ **Proposition 5.6.** *Let $\mathcal{S}$ be an orientable surface of topological size s without boundary, and an integer $r \ge 0$. The* homeomorphic joint crossing number *in $\mathcal{S}$ of the input graphs G_1 and G_2 can be solved in time $2^{\mathsf{poly}(s+r)} \cdot n_0^2$, where n_0 is the size of $G_1 \cup G_2$.*

The strategy is to reduce the problem to the regular joint crossing number. First, for a graph H cellularly embedded in an orientable surface $\mathcal{S}$, let the *hardening* of H be the graph $\bar{H}$ constructed as follows (Figure 6):

- For each edge $e = uv \in E(H)$, we replace e with two new internally disjoint paths P_e and Q_e of length three, called the *twin paths of* e, with the new vertices denoted in such a way that $P_e = (u, w^1_{u,e}, w^2_{v,e}, v)$ and $Q_e = (u, w^2_{u,e}, w^1_{v,e}, v)$.
- For each vertex $v \in V(H)$ with the cyclic order of adjacent edges $(e_1, e_2, \dots, e_p)$ in the given embedding in $\mathcal{S}$, we add the edges of a cycle (the *hardening cycle* of v) on $(w^1_{v,e_1}, w^2_{v,e_1}, w^1_{v,e_2}, w^2_{v,e_2}, \dots, w^1_{v,e_p}, w^2_{v,e_p})$ in this cyclic order.

In relation to the operation of hardening, we prove the following technical lemma:

▶ **Lemma 5.7.** *Let H^2 be the graph constructed from a graph H by replacing each edge with a pair of parallel edges (called a* twin edge pair*). If D is a drawing of H^2 in a surface $\mathcal{S}$ with* *at most $4k$ crossings (but no self-crossings), then there is a subdrawing $D' \subseteq D$ isomorphic to H, obtained by selecting one edge out of each twin edge pair of H^2, with at most k crossings.*

Proof. We choose $D' \subseteq D$ by selecting one out of each twin edge pair of H^2 independently at random. The probability that a crossing x of D is a crossing in D' is $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$ (both edges of x are selected), and so the expected number of crossings in D' is $\frac{1}{4} \cdot 4k = k$ and there has to be a choice of D' with at most k crossings. ◀

▶ **Lemma 5.8.** *Let G_1 and G_2 be two graphs cellularly embedded in an orientable surface $\mathcal{S}$, and let $\bar{G}_1$ and $\bar{G}_2$ denote their hardenings. The homeomorphic joint crossing number of G_1 and G_2 in $\mathcal{S}$ is at most k if and only if the joint crossing number of $\bar{G}_1$ and $\bar{G}_2$ in $\mathcal{S}$ is at most $4k$.*

Proof. Let $G := G_1 \dot\cup G_2$. The '⇒' direction is easy. Assuming D is a drawing of G with at most k crossings solving the homeomorphic joint crossing number problem of G_1 and G_2, one may draw the twin paths of each edge of G closely along this edge in D (and each old crossing hence generates 4 new crossings), and the hardening cycles of all vertices of G can be drawn there without crossings thanks to having homeomorphic subdrawings of G_1 and G_2 in D.

In the '⇐' direction, let $\bar{D}$ be a drawing witnessing a solution to the joint crossing number problem of $\bar{G}_1$ and $\bar{G}_2$ with at most $4k$ crossings. We can clearly assume that no edge self-crosses in $\bar{D}$. Restricting to only the subdrawing of $\bar{D}$ formed by the twin paths of all edges of G and suppressing degre-2 vertices, we may apply Lemma 5.7 and conclude that there is a subdrawing $D' \subseteq \bar{D}$ with at most k crossings which is isomorphic to (a subdivision of) G. It remains to verify that, for $i = 1, 2$, the subdrawing of D' representing G_i (which is an embedding by the definition of the joint crossing number problem) is homeomorphic to the given embedding of G_i.

It is well known that two cellular embeddings of the same graph in an orientable surface are homeomorphic (mirror image of the whole embedding allowed) if and only if their rotation systems are the same. For the embedding of G_i represented by D', this is ensured by embeddings of the hardening cycles of $\bar{G}_i$ within the solution $\bar{D}$. This finishes the proof. ◀

Fixing the rotation system. More generally, a natural restriction on crossing number problems is to fix the clockwise cyclic order of edges around each vertex (the *rotation system*). The traditional crossing number with a fixed rotation system remains NP-hard [43], and its parameterized complexity has not been considered so far. We solve the variant of the COLOR-CONSTRAINED CROSSING PROBLEM where the input also prescribes the cyclic permutation of the edges around each vertex of G, and give a solution in the plane. (We remark that the problem naturally generalizes to orientable surfaces, but the proof is more complicated then, so in this version we only formulate it in the plane.)

We formally introduce the problem as follows:

COLOR-CONSTRAINED CROSSING PROBLEM WITH ROTATION SYSTEM
Input: A colored graph G with colors in $\{1, \ldots, c\}$, a symmetric matrix M of size $c \times c$ with nonnegative integer values, and a system of cyclic permutations $\big(\pi_v : v \in V(G)\big)$ where the domain of π_v is the set of edges incident to v in G.
Question: Is there a normal colored drawing D of G in the *plane* such that:
- for each pair $i, j \in \{1, \ldots, c\}$, the number of crossings in D involving two edges, one colored i and the other j, is at most $M_{i,j}$, and
- for each vertex $v \in V(G)$, the clockwise rotation of edges around v in D is π_v?

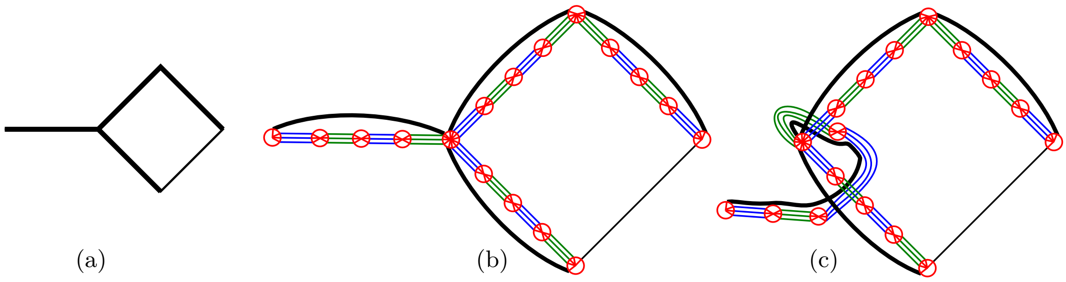

Figure 7 The construction of the graph G_1 for the proof of Theorem 5.9. (a) The graph G, with the spanning tree T in thick lines, shown with the desired rotation system. (b) The graph G_1; its original edges are black, and the new colors red, blue, and green are as pictured. Red edges cannot carry any crossing, and blue-blue and green-green crossings are forbidden. (c) A valid drawing of the graph G_1 with crossings. In general, we must allow crossings between the chains, and these are made possible with the alternating green-blue paths; forbidding blue-blue and green-green crossings forces the consistency of the rotation systems of the wheels.

▶ **Theorem 5.9.** *The* Color-Constrained Crossing Problem with Rotation System *(in the plane) of the input connected graph G of size n, matrix M, and rotation system $\big(\pi_v : v \in V(G)\big)$ can be solved in time $2^{\mathsf{poly}(r)} \cdot n^2$, where $r = \sum_{1 \le i \le j \le c} M_{i,j}$.*

The proof, not very difficult but quite technical, uses an "encoding" of the rotations and (more importantly) of their orientations in the drawings by gadgets sketched in Figure 7 to reduce to the ordinary Color-Constrained Crossing Problem.

Proof. We are going to reduce this problem to the ordinary Color-Constrained Crossing Problem (with $c+3$ colors and a suitable matrix M'), and then apply Proposition 5.5.

For our reduction, we introduce three new colors which we call red, green, and blue. Recall that a *wheel* of k spokes is the graph obtained from the cycle C_k by adding a new vertex adjacent to all vertices of the cycle (called a *rim*). We modify the given colored graph G to a colored (including the three new colors) graph G_1 in the following steps; see Figure 7 where the case of $r = 1$ is shown:

- We choose any spanning tree $T \subseteq G$ (notice that G is assumed connected).
- For every vertex $v \in V(G)$ we create a new wheel W_v of k_v spokes where $k_v = \deg_G(v) + 3\deg_T(v)$, assign the color red to all edges of W_v, and fix an orientation of the rim cycle—deemed *clockwise*. For each edge e incident with v, we choose a set R_v^e of consecutive vertices on the rim cycle of W_v, made of four vertices if $e \in T$ and of a single vertex otherwise, in such a way that these sets R_v^e are pairwise disjoint and the clockwise cyclic order of the terminals on W_v agrees with the rotation π_v. These wheels will be used to "fix the rotation" of edges incident to v up to mirror image.
- For every edge $e \in E(G)$ we create a new isolated edge e^o, called an *original edge*, which inherits the color of e from G.
- For every edge $e \in E(T)$ we create a new graph N_e, called a 3-*chain*, as follows. Start with $4r-1$ copies of the wheel of 6 spokes, with all edges red and denoted by $N_e^1, \dots, N_e^{4r-1}$. Then, for $i = 0, 1, \dots, 4r-1$, we make three new edges joining three consecutive rim vertices of N_e^i (if $i \ge 1$, otherwise we leave pendant vertices) to three consecutive rim vertices of N_e^{i+1} (if $i+1 \le 4r-1$, otherwise we leave pendant vertices) in such a way that N_e is planar. These three edges get color blue if i is even, and green otherwise. The purpose of the 3-chains will be to "transfer the orientation" from one vertex of G to another.

- Finally, we construct the graph G_1 by identifying vertices from the previous building components. For every $e = uv \in E(G) \setminus E(T)$ we identify one end of the original edge e^o with the unique vertex in R_e^u (of the wheel W_u), and the other end with the unique vertex in R_e^v (of W_v). If $e = uv \in E(T)$, we assign the four vertices of R_e^u, in clockwise order, to one endpoint of e^o and then to three pendant vertices of one side of N_e, and assign the four vertices of R_e^v, in counterclockwise order, to the other endpoint of e^o and then to the three pendant vertices of the other side of N_e, in such a way that the resulting graph $W_u \cup N_e \cup W_v$ is planar.

Then we define a matrix M' which extends M as follows; for $i =$ 'red' and any j let $M'_{i,j} = M'_{j,i} = 0$ (red edges cannot be crossed), for $i \in \{$'green','blue'$\}$ and $j \in \{1, \ldots, c\}$ let $M'_{i,j} = M'_{j,i} = 3r$, and for $i =$ 'blue' and $j =$ 'green' let $M'_{i,i} = M'_{j,j} = 0$ and $M'_{i,j} = M'_{j,i} = 9r$.

We apply Proposition 5.5 to the Color-Constrained Crossing Problem in the plane with the input G_1 and M', and prove (in a constructive way) that this modified problem has a solution if and only if the given instance of Color-Constrained Crossing Problem with Rotation System with the input G, M and $\big(\pi_v : v \in V(G)\big)$ has a solution.

($\Leftarrow$) First, assume that a colored drawing D of G is a solution of the latter problem. For every vertex v of D we thus know that the clockwise rotation of edges incident to v agrees with the specified π_v. Let δ_v be a closed disk centered at v in D such that δ_v is disjoint from all edges non-incident to v in D and from all intersection points of D. We replace, in the drawing D, the interior of δ_v with an embedding of the wheel W_v in δ_v such that, for each edge e of G incident to v, the vertex of W_v that is the end of the original edge e^o of G_1 coincides with a point in which the image of e in D intersects the boundary of δ_v (this is possible because D respects the cyclic order π_v at v). Doing this for all $v \in V(G)$, we turn D into a drawing D_0 of the subgraph of G_1 formed by the vertex wheels and the original edges, where D_0 has precisely the same set of (colored) crossings as D.

We now turn D_0 into a drawing D_1 of G_1 by adding images of the 3-chains N_e for every $e \in E(T)$. We draw N_e in a close neighborhood of the original edge e^o in D_0 such that the red edges stay uncrossed. For every $f \in E(G)$ such that f^o crosses e^o in D_0, we let f^o cross a selected triple of green or blue edges of N_e. Moreover, if $f \in E(T)$, we let the same triple of green (or blue) edges of N_e cross a triple of blue (respectively, green) edges of N_f. Since N_e contains (alternately) $2r$ green triples and $2r$ blue triples, we can always do the choice such that the mutually crossing triples of N_e and N_f are of the opposite colors among blue and green (even if we have all r crossings as self-crossings of the edge e^o).

Every crossing (of e^o and f^o) in D_0 thus contributes in D_1 at most 3 crossings of original edges with a blue or a green edge, and at most $3 \cdot 3 = 9$ crossings between a green and a blue edge. There is no green-green and no blue-blue crossing. Thus, the obtained drawing D_1 of G_1 is admissible with respect to the constraint matrix M' defined above.

($\Rightarrow$) We now consider a drawing D'_1 of G_1 in the plane respecting the constraints of M'. In particular, no red edge, and so none of the wheels used in the construction of G_1, is crossed. We may without loss of generality assume that each of the wheels W_v for $v \in V(G)$ in D'_1 is drawn such that the rest of D'_1 is drawn in the face of W_v bounded by the rim cycle; otherwise, we can redraw each of the rim edges closely along the uncrossed spoke edges.

Let D' be the colored drawing of G obtained from D'_1 by removing all 3-chains and contracting each wheel W_v for $v \in V(G)$ into a single vertex. Then D' immediately satisfies the crossing constraints of M. Furthermore, for every $v \in V(G)$, since W_v is uniquely embeddable in the plane, the clockwise cyclic ordering of the edges around v in D' is either π_v or its reversal. Thus, there only remains to prove that the orientation (clockwise or counterclockwise) of all the wheels in D' is the same.

To prove this, we remark that the subgraph of a 3-chain made of two consecutive wheels together with the three edges connecting them is 3-connected, and thus has a unique planar embedding. Since every such subgraph is embedded in D' (because there are no blue-blue or green-green crossings), we obtain that consecutive wheels in D' are either both oriented in the same way as in the planar embedding of N_e, or both oriented in the opposite way. Using this argument $4r$ times implies that two wheels N_u and N_v are oriented in the same way if u and v are connected by an edge of T. The result follows since T is a spanning tree of G. ◀

5.3 Non-traditional methods of counting crossings

We now look at different problems which do not count the crossings simply "one by one".

Edge crossing number. Perhaps the most natural problem that fits into this category is the edge crossing number. The *edge crossing number* of a graph G is the smallest r such that G has a normal drawing having at most r edges with at least one crossing. Computing the edge crossing number is NP-hard [4], and according to Schaefer [45, Section 3.2: Edge crossing number], no FPT algorithm parameterized by the solution size r is known even in the plane.

▶ **Proposition 5.10.** *Let $r \geq 0$ be an integer and $\mathcal{S}$ be a surface with topological size s. The problem to decide whether an input graph G of size n has edge crossing number at most r in $\mathcal{S}$ can be solved in time $2^{\mathsf{poly}(s+r)} \cdot n^2$.*

Proof. For a drawing D in $\mathcal{S}$, let $\kappa(D)$ be equal to the number ℓ of edges of D carrying an intersection point if the total number of intersection points in D is at most $\binom{\ell}{2}$, and let $\kappa(D) = \infty$ otherwise. We also set $\lambda(i) = 4\binom{i}{2}$. It is again routine to check that (κ, λ) is a crossing-counting pair and that $\kappa(D)$ is computable in polynomial time.

It is well known, see, e.g., Schaefer [44], that any drawing D with ℓ edges carrying an intersection point can be turned into a normal drawing D in which no two edges cross twice, without introducing new crossed edges, and thus having at most $\binom{\ell}{2}$ crossings. Thus, the edge crossing number of G indeed equals the minimum value of $\kappa(D)$ over all drawings D of G. Hence, Theorem 3.2 implies an algorithm running in time $2^{\mathsf{poly}(s+r)} \cdot n^2$. ◀

A natural edge-colored generalization can be solved by a straightforward combination of the ideas in the proof of Proposition 5.5 and of the previous proof.

Odd and pair crossing numbers. The *pair (resp. odd) crossing number* of a graph G is the minimum r such that there exists a normal drawing of G in which at most r pairs of edges mutually cross (resp. cross an odd number of times). While the question whether the pair crossing number coincides with the traditional crossing number is one of the biggest open problems in crossing numbers, Pelsmajer, Schaefer, and Štefankovič [42] proved that the odd crossing number can be lower than the traditional crossing number.

FPT algorithms for the pair and odd crossing numbers were given again by Pelsmajer, Schaefer, and Štefankovič [41], with an unspecified dependence on r; however, because these algorithms are based on an adaptation of Grohe's algorithm [20], they are quadratic in the size n of the input and the dependence on r is at least an exponential tower of height four. We improve over their algorithms by providing algorithms also quadratic in n but with a better dependence in the parameter r:

▶ **Proposition 5.11.** *Let $r \geq 0$ be an integer, and let $\Pi \in \{$'pair', 'odd'$\}$. The problems to decide whether an input graph G of size n has Π-crossing number at most r in the* plane *can be solved in time $2^{2^{\mathcal{O}(r)}} \cdot n^2$.*

Proof. The approach is analogous to the proof of Proposition 5.10, with the additional ingredient given by the bounds by Pelsmajer et al. [41]: if G is of pair (respectively, odd) crossing number ℓ, then there exists a normal drawing D of G achieving this optimum such that the number of crossings in D is at most $\ell 2^\ell$ (respectively, 9^ℓ).

We define the crossing-counting pair (κ, λ) such that $\kappa(D)$ is equal to the Π-crossing number ℓ of D, if D is normal and D has at most $\max(\ell 2^\ell, 9^\ell)$ crossings in total, and $\kappa(D) = \infty$ otherwise. This can clearly be computed in polynomial time from given D. We correspondingly set $\lambda(i) = 4\max(i2^i, 9^i)$, which together with κ fulfills Definition 3.1, thanks to the bounds by Pelsmajer et al. [41]. Hence, an application of Theorem 3.2 gives an algorithm running in time $2^{2^{\mathcal{O}(r)}} \cdot n^2$. ◀

We remark that the same approach would extend to surfaces, assuming some bounds analogous to those given by Pelsmajer et al. [41] for the plane.

5.4 Allowing edge removals and vertex splits before drawing

Finally, we consider crossing number variants that allow graph simplifications before drawing.

Skewness. The *skewness* of a graph G is the smallest number of edges whose removal from G leaves a planar graph. Deciding whether the skewness of G is at most q is NP-complete [34], and linear-time FPT algorithms are claimed, without details, by Kawarabayashi and Reed [32] and Jansen, Lokshtanov, and Saurabh [30]. We generalize the problem as follows.

Color-Constrained $\mathcal{S}$-Skewness with Crossings
Input: A colored graph G with colors in $\{1, \ldots, c\}$, non-negative integers $q_1, \ldots, q_c$ and r.
Question: Can we remove, for each $i = 1, \ldots, c$, at most q_i edges of color i from G, such that the resulting graph has crossing number at most r in $\mathcal{S}$?

▶ **Proposition 5.12.** *Let $\mathcal{S}$ be a surface with topological size s. The* Color-Constrained $\mathcal{S}$-Skewness with Crossings *problem for an input graph G of size n and parameter r can be solved in time $2^{\mathsf{poly}(s+q+r)} \cdot n^2$, where $q = \sum_i q_i$.*

Proof. Without loss of generality, we assume that $c \le q + 1$, since the colors i with $q_i = 0$ can be merged into one. We simply define $\kappa(D)$ as the number of crossings in D if D is a normal drawing, and $\kappa(D) = \infty$ otherwise. We also let $\lambda(i) = 4i$. Then, setting $p = 0$, we apply Theorem 3.2 to get the desired algorithm. ◀

Splitting number. The smallest integer p such that a graph obtained from the given graph G by p successive vertex splits is embeddable in $\mathcal{S}$ is called the *splitting number* of G in $\mathcal{S}$ [23,29]. This problem is NP-hard [17], already in the plane. Nöllenburg, Sorge, Terziadis, Villedieu, Wu, and Wulms [39] proved that the property of having splitting number at most p is minor-monotone in any fixed surface, and so the $\mathcal{S}$-splitting number has a nonuniform FPT algorithm parameterized by p using the theory of graph minors. We generalize and improve the latter result to a uniform FPT algorithm:

▶ **Proposition 5.13.** *Let $p, r \ge 0$ be integers and $\mathcal{S}$ be a surface with topological size s. The problem to decide whether an input graph G of size n has, after at most p vertex splits, a crossing number at most r in $\mathcal{S}$ can be solved in time $2^{\mathsf{poly}(s+p+r)} \cdot n^2$.*

Proof. We use the same crossing-counting pair (κ, λ) as in the proof of Proposition 5.12. Then, we set $q_1 = \ldots = q_c = 0$ and apply Theorem 3.2 to get the desired algorithm with respect to given p and r. ◀

6 Computing drawings for positive instances

In this section, we prove the following result.

▶ **Theorem 6.1.** *For every positive instance of the (κ, λ)-Crossing Number problem, we can compute a representation of the corresponding drawing, of size the size n of the input graph times a polynomial in $s + c + p + q + \lambda(r)$, without overhead in the runtime compared to Theorem 3.2.*

We need some preliminaries.

A graph embedding on a surface is *cellular* if all its faces are homeomorphic to open disks. Cellular graph embeddings can be manipulated through their *combinatorial maps*. Algorithmically, refined data structures such as *graph encoded maps* [33] (see also [16, Section 2]) are used.

The idea of the proof of Theorem 6.1 is the following. Colin de Verdière and Magnard [12] not only gave a decision algorithm for the problem of embedding a graph on a 2-complex (Theorem 2.1), but for positive instances their algorithm can compute a representation of an embedding [12, arXiv version, Theorem 9.1]. There remains to turn this embedding into $\mathcal{C}$ into a drawing. The following result is obtained by a careful inspection of the proof of Theorem 2.1.

▶ **Theorem 6.2** (Colin de Verdière and Magnard [12, arXiv version, Theorem 9.1])**.** *In Theorem 2.1, if G has an embedding into $\mathcal{C}$, and $\mathcal{C}$ has no edge incident to three triangles and no isolated vertex, an embedding of G can be computed without overhead in the asymptotic runtime. In detail, an embedding of a graph H is computed where:*

- *H is obtained from G by augmenting it with at most $2C$ vertices and at most $3C + 2u$ edges, and performing at most C edge subdivisions, where C is the size of $\mathcal{C}$ and u is the number of connected components of G;*
- *the images of the vertices of H cover the singular points of $\mathcal{C}$;*
- *the restriction of H to the surface part of $\mathcal{C}$ is cellular, specified by its combinatorial map, and by which point is mapped to which singular point of $\mathcal{C}$;*
- *the restriction of H to the isolated edges of $\mathcal{C}$ is specified by the sequence of vertices and edges of H appearing along each isolated edge of $\mathcal{C}$.*

Some comments on this theorem are in order. First, the requirements on the 2-complex $\mathcal{C}$ are satisfied for all 2-complexes considered in this paper. Second, it is necessary to add vertices to H in order to cover the singular points of $\mathcal{C}$, to add edges in order to make the graph cellularly embedded on the surface part of $\mathcal{C}$, and to subdivide edges in order to ensure that edges of the graph are either entirely inside, or entirely outside, the surface part of $\mathcal{C}$.

Proof of Theorem 6.1. Recall from the proof of Theorem 3.2 that if an instance of the (κ, λ)-Crossing Number problem is positive, the corresponding graph G' embeds into at least one 2-complex $\mathcal{C}_5 \in \Gamma$, where G' has size $\mathcal{O}((c + p + q + \lambda(r))n)$ and $\mathcal{C}_5$ has size $\mathcal{O}(s + (c + p + q + \lambda(r))\lambda(r))$. Thus, the size of the graph H obtained from Theorem 6.2 is that of G times a polynomial in $s + c + p + q + \lambda(r)$.

The proof basically follows the same steps as the proof of Lemma 4.2. First, up to modifying the embedding, we can assume that each necklace of $\mathcal{C}_5$ is either not used at all, or used by a necklace of H of the same type. We thus have an embedding of a subgraph H_1 of H (a supergraph of G_1, in which some edges may be subdivided) in such a way that each isolated edge of $\mathcal{C}_4$ coincides with an edge of G_1 of the same color, or is not used at all.

In a second step, we perform the vertex splits as described in the proof of Lemma 4.2, resulting in an embedding in $\mathcal{C}_3$ of a graph H'' obtained from H_1 by at most p vertex splits.

A minor detail: In the case where a singular point z_i, $i \in S$, contains a point in the relative interior of an edge $e = u$ of H, the proof of Lemma 4.2 performs a split of one of its vertices, say v, as prescribed, by replacing uv with uv' where v' is a new vertex and shrinking uv' without moving u so that in the new embedding, uv' does not contain any singular point of $\mathcal{C}_4$ in its interior. In order to keep the graph cellular on the surface part, we cannot remove the edges on the surface part of the complex during this process, but now mark them as extra edges.

In a third step, we remove the edges of the graph H'' using edges in $\mathcal{C}_3 \setminus \mathcal{C}_2$. The resulting graph H', embedded into $\mathcal{C}_2$, is a supergraph, possibly subdivided, of a graph G' obtained from G by at most p vertex splits and by removing, for each i, at most q_i edges of color i. Moreover, again, each isolated of $\mathcal{C}_2$ of color i coincides with an edge of G' of color i, or is not used at all.

This embedding of H' in $\mathcal{C}_2$ turns into a an embedding, in $\mathcal{S}$, of a graph H obtained from a drawing D' of G' satisfying $\kappa(D') \leq r$ by the following steps:

1. for each crossing point between two edges e and e' of G, subdivide e and e' and identify the two new vertices together (the new vertex corresponds to the crossing between e and e');
2. take a supergraph of the resulting graph by subdividing edges and by adding vertices and edges, increasing the size by a factor that is at most polynomial in $s + c + p + q + \lambda(r)$.

The graph H is specified by its combinatorial map. If necessary, we push the image of G' away from the boundary of $\mathcal{S}$, by adding more vertices and edges (at most doubling the size of the graph). Then we triangulate each face of H. We obtain a representation (T, W) of the graph G' of size n times a polynomial in $s + c + p + q + \lambda(r)$. ◀

7 Conclusions

More applications. Theorem 3.2 allows for almost endless combinations for crossing number variants, in terms of drawing styles, ways to count crossings, allowed vertex splits and edge removals, all possibly taking edge colors into account, all on an arbitrary surface. This shows the versatility of our approach. Altogether, our general framework encompasses most existing crossing number variations, and even more general ones, implying, in a unified way, fixed-parameter tractable algorithms for these crossing number problems in any surface.

Limitations. Nonetheless, in our definition of a crossing-counting pair, $\kappa(D)$ must functionally bound the total number of intersection points in the drawing D, and the value of κ cannot grow upon deletion of vertices and edges. The first requirement immediately excludes, e.g., the local crossing number (the minimum r such that a given graph has an r-planar drawing, which is NP-complete to compute already for $r = 1$ [8, 19]), and the second requirement excludes some recently introduced drawing styles such as, for instance, 1^+ and 2^+-real face drawings [5]. It is conceivable that Courcelle-based approaches would be still able to handle 1^+ and 2^+-real face drawings, albeit not easily.

Another weakness, in particular compared to the very recent preprint by Hamm, Klute, and Parada [22], is our quite restricted ability to handle predrawn parts of the input graph, essentially limited to fixing *uncrossable* parts of the graph via rigid subembeddings, and to special restrictions like the fixed rotation system in Theorem 5.9. In contrast, [21] and [22] are very general in this respect and, in particular, allow crossings of fixed parts with the unfixed rest of the graph.

Remarks and possible extensions. Beyond their use for strongly fan-planar drawings, surfaces with boundaries can be useful for other problems, because they allow to pinpoint specific regions of the surface. One could actually consider a version with colored boundaries, restricting Item 1(d) of the definition of a crossing-counting pair to color-preserving self-homeomorphisms. Our arguments carry through; this only affects the running time by a $b!$ factor (see Lemma 4.4).

Finally, a possible extension of the (κ, λ)-Crossing Number problem would be to replace the surface $\mathcal{S}$ with an arbitrary 2-complex. We leave open whether such an extension is possible, and also potential applications.

References

1 Pankaj K. Agarwal, Boris Aronov, János Pach, Richard Pollack, and Micha Sharir. Quasi-planar graphs have a linear number of edges. *Comb.*, 17(1):1–9, 1997. doi:10.1007/BF01196127.

2 Mark Anthony Armstrong. *Basic topology.* Undergraduate Texts in Mathematics. Springer-Verlag, 1983.

3 Sang Won Bae, Jean-François Baffier, Jinhee Chun, Peter Eades, Kord Eickmeyer, Luca Grilli, Seok-Hee Hong, Matias Korman, Fabrizio Montecchiani, Ignaz Rutter, and Csaba D. Tóth. Gap-planar graphs. *Theor. Comput. Sci.*, 745:36–52, 2018. doi:10.1016/j.tcs.2018.05.029.

4 Martin Balko, Petr Hliněný, Tomás Masařík, Joachim Orthaber, Birgit Vogtenhuber, and Mirko H. Wagner. On the uncrossed number of graphs. In Stefan Felsner and Karsten Klein, editors, *32nd International Symposium on Graph Drawing and Network Visualization, GD 2024, September 18-20, 2024, Vienna, Austria*, volume 320 of *LIPIcs*, pages 18:1–18:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.GD.2024.18.

5 Carla Binucci, Giuseppe Di Battista, Walter Didimo, Vida Dujmovic, Seok-Hee Hong, Michael Kaufmann, Giuseppe Liotta, Pat Morin, and Alessandra Tappini. Graphs drawn with some vertices per face: Density and relationships. *IEEE Access*, 12:68828–68846, 2024. doi:10.1109/ACCESS.2024.3401078.

6 Carla Binucci, Aaron Büngener, Giuseppe Di Battista, Walter Didimo, Vida Dujmovic, Seok-Hee Hong, Michael Kaufmann, Giuseppe Liotta, Pat Morin, and Alessandra Tappini. Min-k-planar drawings of graphs. In *GD (1)*, volume 14465 of *Lecture Notes in Computer Science*, pages 39–52. Springer, 2023. doi:10.1007/978-3-031-49272-3_3.

7 Sergio Cabello. Hardness of approximation for crossing number. *Discrete Comput. Geom.*, 49(2):348–358, 2013. doi:10.1007/s00454-012-9440-6.

8 Sergio Cabello and Bojan Mohar. Adding one edge to planar graphs makes crossing number and 1-planarity hard. *SIAM J. Comput.*, 42(5):1803–1829, 2013. doi:10.1137/120872310.

9 Otfried Cheong, Henry Förster, Julia Katheder, Maximilian Pfister, and Lena Schlipf. Weakly and strongly fan-planar graphs. In Michael A. Bekos and Markus Chimani, editors, *Graph Drawing and Network Visualization - 31st International Symposium, GD 2023, Isola delle Femmine, Palermo, Italy, September 20-22, 2023, Revised Selected Papers, Part I*, volume 14465 of *Lecture Notes in Computer Science*, pages 53–68. Springer, 2023. doi:10.1007/978-3-031-49272-3_4.

10 Julia Chuzhoy and Zihan Tan. A subpolynomial approximation algorithm for graph crossing number in low-degree graphs. In *STOC*, pages 303–316. ACM, 2022. doi:10.1145/3519935.3519984.

11 Éric Colin de Verdière and Arnaud de Mesmay. Testing graph isotopy on surfaces. *Discrete & Computational Geometry*, 51(1):171–206, 2014. doi:10.1007/s00454-013-9555-4.

12 Éric Colin de Verdière and Thomas Magnard. An FPT algorithm for the embeddability of graphs into two-dimensional simplicial complexes. In *Proceedings of the 29th European Symposium on Algorithms (ESA)*, pages 32:1–32:17, 2021. Full improved version in arXiv:2107.06236v2. doi:10.4230/LIPIcs.ESA.2021.32.

13 Éric Colin de Verdière, Thomas Magnard, and Bojan Mohar. Embedding graphs into two-dimensional simplicial complexes. *Computing in Geometry and Topology*, 1(1):Article 6, 2022.

14 Bruno Courcelle. The monadic second-order logic of graphs. I. recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. doi:10.1016/0890-5401(90)90043-H.

15 Walter Didimo, Giuseppe Liotta, and Fabrizio Montecchiani. A survey on graph drawing beyond planarity. *ACM Comput. Surv.*, 52(1):4:1–4:37, 2019. doi:10.1145/3301281.

16 David Eppstein. Dynamic generators of topologically embedded graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 599–608, 2003. URL: http://dl.acm.org/citation.cfm?id=644108.644208.

17 Luérbio Faria, Celina M. H. de Figueiredo, and Candido Ferreira Xavier de Mendonça Neto. Splitting Number is NP-complete. *Discret. Appl. Math.*, 108(1-2):65–83, 2001.

18 Michael R. Garey and David S. Johnson. Crossing number is NP-complete. *SIAM J. Algebr. Discrete Methods*, 4(3):312–316, September 1983.

19 Alexander Grigoriev and Hans L. Bodlaender. Algorithms for graphs embeddable with few crossings per edge. *Algorithmica*, 49(1):1–11, 2007. doi:10.1007/S00453-007-0010-X.

20 Martin Grohe. Computing crossing numbers in quadratic time. *J. Comput. Syst. Sci.*, 68(2):285–302, 2004. doi:10.1016/j.jcss.2003.07.008.

21 Thekla Hamm and Petr Hliněný. Parameterised partially-predrawn crossing number. In Xavier Goaoc and Michael Kerber, editors, *38th International Symposium on Computational Geometry, SoCG 2022, June 7-10, 2022, Berlin, Germany*, volume 224 of *LIPIcs*, pages 46:1–46:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.SOCG.2022.46.

22 Thekla Hamm, Fabian Klute, and Irene Parada. Computing crossing numbers with topological and geometric restrictions. *CoRR*, abs/2412.13092, 2024. doi:10.48550/ARXIV.2412.13092.

23 Nora Hartsfield, Brad Jackson, and Gerhard Ringel. The splitting number of the complete graph. *Graphs and Combinatorics*, 1:311–329, 1985. doi:10.1007/BF02582960.

24 Petr Hliněný. Crossing number is hard for cubic graphs. *Journal of Comb. Theory, Ser. B*, 96(4):455–471, 2006. doi:10.1016/j.jctb.2005.09.009.

25 Petr Hliněný. Complexity of anchored crossing number and crossing number of almost planar graphs. In *MFCS*, volume 345 of *LIPIcs*, pages 59:1–59:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.MFCS.2025.59.

26 Petr Hliněný and Liana Khazaliya. Crossing number is NP-hard for constant path-width (and tree-width). In Julián Mestre and Anthony Wirth, editors, *35th International Symposium on Algorithms and Computation, ISAAC 2024, December 8-11, 2024, Sydney, Australia*, volume 322 of *LIPIcs*, pages 40:1–40:15. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.ISAAC.2024.40.

27 Petr Hliněný and Csenge Lili Ködmön. Note on min-k-planar drawings of graphs. In Stefan Felsner and Karsten Klein, editors, *32nd International Symposium on Graph Drawing and Network Visualization, GD 2024, September 18-20, 2024, Vienna, Austria*, volume 320 of *LIPIcs*, pages 8:1–8:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.GD.2024.8.

28 Petr Hliněný and Gelasio Salazar. On hardness of the joint crossing number. In *ISAAC*, volume 9472 of *Lecture Notes in Computer Science*, pages 603–613. Springer, 2015. doi:10.1007/978-3-031-49272-3_4.

29 Brad Jackson and Gerhard Ringel. Splittings of graphs on surfaces. *Proceedings of the 1st Colorado Symposium on Graph Theory*, pages 203–219, 1982.

30 Bart M. P. Jansen, Daniel Lokshtanov, and Saket Saurabh. A near-optimal planarization algorithm. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1802–1811, 2014. doi:10.1137/1.9781611973402.130.

31 Michael Kaufmann and Torsten Ueckerdt. The density of fan-planar graphs. *Electron. J. Comb.*, 29(1), 2022. doi:10.37236/10521.

32 Ken-ichi Kawarabayashi and Bruce A. Reed. Computing crossing number in linear time. In *STOC*, pages 382–390. ACM, 2007. doi:10.1145/1250790.1250848.

33 Sóstenes Lins. Graph-encoded maps. *Journal of Combinatorial Theory, Series B*, 32:171–181, 1982. URL: https://doi.org/10.1016/0095-8956(82)90033-8.

34 P.C. Liu and R.C. Geldmacher. On the deletion of nonplanar edges of a graph. In *Proceedings of the Tenth Southeastern Conference on Combinatorics, Graph Theory and Computing (Florida Atlantic Univ.,Boca Raton, Fla., 1979)*, pages 727–738, 1979.

35 Daniel Lokshtanov, Fahad Panolan, Saket Saurabh, Roohani Sharma, Jie Xue, and Meirav Zehavi. Crossing number in slightly superexponential time. In *Proceedings of the 36th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1412–1424, 2025. doi:10.1137/1.9781611978322.44.

36 Bojan Mohar and Carsten Thomassen. *Graphs on surfaces*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, 2001.

37 Miriam Münch and Ignaz Rutter. Parameterized algorithms for beyond-planar crossing numbers. In Stefan Felsner and Karsten Klein, editors, *32nd International Symposium on Graph Drawing and Network Visualization, GD 2024, September 18-20, 2024, Vienna, Austria*, volume 320 of *LIPIcs*, pages 25:1–25:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.GD.2024.25.

38 Seiya Negami. Crossing numbers of graph embedding pairs on closed surfaces. *J. Graph Theory*, 36(1):8–23, 2001. doi:10.1002/1097-0118(200101)36:1\%3C8::AID-JGT2\%3E3.0.CO;2-O.

39 Martin Nöllenburg, Manuel Sorge, Soeren Terziadis, Anaïs Villedieu, Hsiang-Yun Wu, and Jules Wulms. Planarizing graphs and their drawings by vertex splitting. In *GD*, volume 13764 of *Lecture Notes in Computer Science*, pages 232–246. Springer, 2022. doi:10.1007/978-3-031-22203-0_17.

40 János Pach and Géza Tóth. Graphs drawn with few crossings per edge. *Comb.*, 17(3):427–439, 1997. doi:10.1007/BF01215922.

41 Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Crossing numbers and parameterized complexity. In *Graph Drawing, 15th International Symposium, GD 2007, Sydney, Australia, September 24-26, 2007. Revised Papers*, volume 4875 of *Lecture Notes in Computer Science*, pages 31–36. Springer, 2007. doi:10.1007/978-3-540-77537-9_6.

42 Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Odd crossing number and crossing number are not the same. *Discret. Comput. Geom.*, 39(1-3):442–454, 2008. doi:10.1007/S00454-008-9058-X.

43 Michael J. Pelsmajer, Marcus Schaefer, and Daniel Štefankovič. Crossing numbers of graphs with rotation systems. *Algorithmica*, 60(3):679–702, 2011. doi:10.1007/S00453-009-9343-Y.

44 Marcus Schaefer. *Crossing Numbers of Graphs*. Discrete mathematics and its applications. CRC Press, Taylor & Francis Group, 2017.

45 Marcus Schaefer. The graph crossing number and its variants: A survey. *Electron. J. Comb.*, Dynamic Surveys, #DS21, 2024. Eighth Edition: May 17, 2024. doi:10.37236/2713.

46 Nathan van Beusekom, Irene Parada, and Bettina Speckmann. Crossing numbers of beyond-planar graphs revisited. *J. Graph Algorithms Appl.*, 26(1):149–170, 2022. doi:10.7155/jgaa.00586.