Approximating the Crossing Number of Graphs Embeddable in Any Orientable Surface

Petr Hliněný[1] and **Markus Chimani**[2]

[1] Faculty of Informatics, Masaryk University, Brno,
Czech Republic
hlineny@fi.muni.cz

[2] Faculty of Computer Science, TU Dortmund,
Germany
markus.chimani@tu-dortmund.de

July 5, 2009

Abstract. The crossing number of a graph is the least number of pairwise edge crossings in a drawing of the graph in the plane. We provide an $O(n \log n)$ time constant factor approximation algorithm for the crossing number of a graph of bounded maximum degree which is "densely enough" embeddable in any fixed orientable surface.

Our approach combines some known tools with a powerful new lower bound on the crossing number of an embedded graph. This result extends previous results that gave such approximations in particular cases of projective, toroidal or apex graphs; it is a qualitative improvement over previously published algorithms that constructed low-crossing-number drawings of embeddable graphs without giving any approximation guarantees. No constant factor approximation algorithms for the crossing number problem over comparably rich classes of graphs are known to date.

Keywords: crossing number, approximation algorithm, surface embedding.

1 Introduction

The *crossing number* $cr(G)$ of a graph G is the minimum number of pairwise edge crossings in a drawing of G in the plane. Formally, a *drawing* of a graph G in some surface Σ is a mapping of its vertex set $V(G)$ into distinct points in Σ. Edges are mapped into curves in Σ between the images of their endvertices and must not contain the image of any non-incident vertex. To resolve ambiguity, we consider drawings of graphs such that no three edges intersect in a common point which is not a vertex. Then a *crossing* is an intersection point of two edges that is not a vertex. We refer to Section 2 for further definitions.

The crossing number arises in several research fields: E.g., it is a natural problem in graph drawing and diagramming applications, and can be used in VLSI design to estimate the required chip area [2]. It is of further interest in algorithm design, since graphs with small crossing number or genus can be regarded to be "similar" to planar graphs and thus potentially allow more efficient algorithms for various graph problems.

Although the crossing number has been studied for over 60 years, see [27] for an extensive bibliography, surprisingly little is known about many of its central properties. Even the crossing number of complete and complete bipartite graphs can only be conjectured [28, 15, 16, 25], even though they were the first questions asked in this context [26].

Algorithmically, we know that the problem of computing the crossing number is NP-hard [11]. This holds true even for graphs of bounded degree; in fact, even for graphs with maximum degree 3 [17]. On the other hand, it has been shown that the problem is fixed parameter tractable: We can test whether a graph has a crossing number at most k in linear time, when considering k fixed [13, 21]. While these approaches do currently not allow practical algorithms, there exist linear programming based exact algorithms that are promising for "real-world" graphs arising in graph drawing applications [8]. Yet, computing exact crossing numbers is in general extremely difficult and one usually has to resort to heuristics, see, e.g., [1, 14].

In this paper, we are especially interested in crossing number approximations. It is unknown whether the problem allows efficient approximations at all, even for bounded degree graphs which are the focus for virtually all approximation approaches. The best known polynomial algorithm for the crossing number of general graphs with bounded degree approximates the quantity $|V(G)|+cr(G)$, not directly $cr(G)$, within a factor of $\log^3 |V(G)|$ [10]. Polynomial constant factor approximations of $cr(G)$ are known only for much more restricted graph classes with bounded degree, in particular for *apex* [6, 7] and *near-planar* [18, 5] graphs on the one hand, and for *projective* [12] and *toroidal* [19] graphs on the other. An apex graphs is a graph where one can remove one specific vertex, in order to obtain a planar graph. Note that the complexity status for computing the crossing number of apex graphs, as well as for the subclass of near-planar graphs (where the removal of a single edge suffices) is still open.

In this paper, we are going to extend these latter results to graphs G embeddable (i.e., drawable without any crossings) in an orientable surface of arbitrary (fixed) genus g. This is by far the yet richest graph class allowing a fixed factor approximation. Unfortunately, as we shall see in the following, the proving techniques used for toroidal or projective graphs cannot be directly extended to higher surfaces.

Consider a graph G embedded into some surface Σ. Böröczky, Pach and Tóth [3], later improved by Djidjev and Vrt'o [9], presented an algorithm to compute a drawing of G in the plane with few crossings, using the embedding in Σ as a starting point. The algorithm is based on a simple idea: One iteratively "cuts and opens" handles of Σ, temporarily removing the affected edges from G. By greedily trying to remove the fewest number of edges and "cheaply re-inserting" them at the end, one obtains a drawing of the graph in the plane. The above papers were able to obtain upper bounds for the number of crossings generated by this approach. However, in order to get an approximation guarantee for such an algorithm, one also has to provide a matching lower bound on $cr(G)$, and this task requires a much more careful consideration of the cutting process. Generally, obtaining such lower bounds has always been the hardest task for all known crossing number approximations.

Algorithmically, our approximation procedure (Algorithm 3.1) is surprisingly simple, following the above idea of iteratively cutting "cheap" handles. Its analysis (Theorem 3.2) already gives a qualitative improvement over the previously published results in terms of the upper bound. Yet, proving a matching lower bound is by far harder and constitutes the main new contribution of this paper (Theorem 4.2). To this ends, we have to introduce a set of novel and sophisticated tools.

The next section will introduce some crucial concepts used throughout the paper. In Section 3 we show the actual approximation algorithm, analyse its running time, and estimate the number of crossings it generates. Section 4 then presents the central theoretical results necessary to obtain a constant factor approximation guarantee. The proofs for the required lower bounds are discussed in Section 5. Overall we obtain:

Theorem 1.1. *Let G be a graph embeddable in an orientable surface of genus $g \geq 1$ with nonseparating dual edge-width at least $2^{g+2}\Delta$ where Δ is the maximum degree of G. The presented Algorithm 3.1 computes a drawing of G in the plane with at most $3 \cdot 2^{3g+2} \cdot \Delta^2 \cdot cr(G)$ crossings. Hence this is a constant factor approximation algorithm of the crossing number $cr(G)$ for bounded degree Δ and bounded genus g. Its running time is $O(n \log n)$ where $n = |V(G)| + |E(G)|$ ($n = |V(G)|$ when Δ is bounded).*

In this result we need the technical restriction that something called "*nonseparating dual edge-width*" is large enough. While it is formally defined in the following section, one can think of it as the requirement that the graph is embedded "densely enough" in the surface. This restriction is necessary in the lower bound part of the proof, since even a planar graph (i.e. of crossing number 0) can be embedded (non-densely, though) in higher surfaces. Finally, Section 6 will discuss the dependency on this and other parameters in the overall algorithm, and sketch possible extensions.

2 Definitions and Tools

Our terminology is based on Mohar–Thomassen [24]. Specifically, we deal with unoriented multigraphs by default; so when speaking about a *graph*, we allow multiple edges or loops. The vertex set of a graph G is denoted by $V(G)$, the edge set by $E(G)$, and the maximum degree by $\Delta(G)$. We denote by $len(Q)$ the *length* (number of edges) of a path or a cycle Q. We call a graph H a *theta* graph if H is formed by three pairwise internally disjoint paths with common ends. Most of the time we shall deal with graphs that are *embedded* in some surface Σ, i.e., drawn on Σ without edge crossings. If $b(G)$ (as a point set in Σ) is an embedding of G in Σ, then the arc-connected components of $\Sigma \setminus b(G)$ are called the *faces* of the embedding.

We need some basic notions of classical topology. A *surface* is a compact 2-manifold without boundary. A closed curve on a surface is called a *loop*. Two loops α, β on a surface Σ are *freely homotopic* if α can be continuously transformed to β on Σ. A loop α on Σ is *contractible* if α is freely homotopic to a constant curve (α can be continuously deformed to a single point), and it is *separating* if $\Sigma \setminus \alpha$ is not arc-connected. A loop α on Σ is *one-sided* if $\Sigma \setminus \alpha$ has a connected boundary, and α is *two-sided* otherwise. A surface with no one-sided loops is *orientable*.

By the surface classification theorem, all orientable surfaces are homeomorphic to some $\mathcal{S}_g$—a sphere with g added "handles" where g is the *genus* of the surface. Notice that if α is an arbitrary nonseparating loop on $\mathcal{S}_g$, then α always "cuts one handle" of $\mathcal{S}_g$, up to homeomorphism. In particular, the simplest nonorientable surface is the *projective plane*, and the orientable surface of genus $g = 1$ (next to the sphere) is *torus*.

We will consider only connected *cellular embeddings* in surfaces, i.e. embeddings in which every face is homeomorphic to an open disc. A cellular embedding of G on Σ is, up to homeomorphism, uniquely determined by G and a *rotation system* of G (a system of cyclic permutations of edges around each vertex), see [24, Section 4.1]. This is also the graph data structure we shall use in our algorithm. Moreover, a cellular embedding of G *determines the surface* Σ, and so we do not have to explicitly mention Σ with respect to G.

We denote by G^* the *topological dual* of an embedded graph G; the vertices of the graph G^* are the faces of G and the edges of G^* are the edge-adjacent pairs of faces of G. There is a natural one-to-one correspondence between the edges of G and the edges of G^*, and so, for arbitrary $F \subseteq E(G)$, we denote by F^* the corresponding subset of edges of $E(G^*)$. Furthermore, in the rotation-system representation of an embedded graph G, it is easy to enumerate the faces of G and hence the topological dual can be computed in linear time.

For our crossing-number problem, the input is an abstract graph G (though assumed to be embeddable on some Σ), but our algorithm will work with an actual embedding of G on Σ (although generally not unique, we can use any). The first step hence is to find such an embedding:

Theorem 2.1 (Mohar [23]). *For every surface Σ there is a linear time algorithm which, for a given graph G, either finds an embedding of G on Σ or returns a subgraph of G that is a subdivision of a "minimal obstacle" for Σ.*

A crucial ingredient of our approach is a measure of "dual density of an embedding". Assuming an embedded graph G, the shortest length of a cycle in G that forms a noncontractible (nonseparating) loop in the embedding, is called *edge-width* $ew(G)$ (*nonseparating edge-width* $ewn(G)$, respectively). Note that the edge-width of a given embedding is efficiently computable [24, Section 4.3]. Faster recent algorithms appeared, e.g., in [22].

Theorem 2.2 (Kutz [22]). *Given a graph G embedded in an orientable surface, there is an algorithm running in time $O(n \log n)$, where $n = |V(G)| + |E(G)|$, that computes the nonseparating edge-width $k = ewn(G)$ and finds a length-k nonseparating cycle in the embedding G.*

For a cycle C in a graph H, we call a path $P \subset H$ a *C-ear* if the ends r, s of P belong to C, but the rest of P is disjoint from C. We allow $r = s$, i.e., a C-ear can also be a cycle. If H is embedded in an orientable surface, then the embedding of any cycle $C \subset H$ is a two-sided loop. A C-ear P is a *C-switching ear* if the first and the last edges of P (wrt. r, s) are embedded on the opposite sides of C. Measuring the edge-width and the shortest switching-ear length in the dual graph of input G will give us, later on, the key estimate of the crossing number of G—see Theorem 4.2.

Since we will frequently deal with dual graphs in our arguments, we introduce several conventions to assist readers' understanding of the paper. When we add an adjective *dual* to a graph term, we mean this term in the topological dual of the (currently considered) graph. We will denote the faces of an embedded graph G in lowercase and treat them as vertices of its dual G^*, and we will use small Greek letters to refer to subgraphs (cycles or paths) of G^*. When there is no danger of confusion, we will not formally distinguish between a graph and its embedding. In particular, if $\alpha \subseteq G^*$ is a dual cycle, then α also refers to the loop on the surface determined by the embedding G. Finally, we will denote by $ewd(G) = ewn(G^*)$ the nonseparating edge-width of the dual G^* of G.

We finish this section with two simple and useful technical claims.

Lemma 2.3 (cf. [19, Lemma 3.1]). *If ϱ is a nonseparating dual cycle in a nonplanar embedded graph H of length $len(\varrho) = ewd(H)$, then all dual ϱ-switching ears in H have length at least $\frac{1}{2}\, ewd(H)$.*

Proof. Seeking a contradiction, we suppose that there is a ϱ-switching ear σ of length $< \frac{1}{2}\, ewd(H)$. The ends of σ on ϱ determine two dual subpaths $\varrho_1, \varrho_2 \subseteq \varrho$ (with the same ends as σ). Then both $\sigma \cup \varrho_1$ and $\sigma \cup \varrho_2$ are nonseparating loops (as witnessed by the other of ϱ_1, ϱ_2), and $len(\varrho_1) \leq \frac{1}{2}\, len(\varrho)$ up to symmetry. Hence $len(\sigma \cup \varrho_1) \leq len(\sigma) + \frac{1}{2}\, len(\varrho) < len(\varrho) = ewd(H)$, a contradiction. ∎

For the second claim, we need to introduce a tool—cutting a surface embedding of a graph G along a two-sided loop γ: Intuitively, this operation should remove all edges of G intersected by γ (we assume γ avoids vertices) and add two new faces (to "cover up" the handle cut by γ). Formally, assume an embedded graph G represented by its rotation system, and a dual cycle $\gamma \subseteq G^*$. Notice that γ, as a surface loop, intersects exactly those edges of G belonging to $E^*(\gamma)$, i.e. the edges corresponding to $E(\gamma)$ in duality. We say that an embedded graph H results by *cutting G along γ*, denoted by $H = G/\gamma$, if $V(H) = V(G)$, $E(H) = E(G) \setminus E^*(\gamma)$, and the rotations of edges around the vertices of H are the same as those of G restricted to $E(H)$.

Notice that the faces of $H = G/\gamma$ are the same as those of G, except that the faces in $V(\gamma)$ vanish and two new faces c_1, c_2, called the *γ-cut faces*, are created. For each edge

$f \in E^*(\gamma)$, exactly one endvertex will become incident with c_1 and the other endvertex with c_2 in H.

Lemma 2.4. *Let H be a graph embedded in an orientable surface of genus ≥ 2, and ϱ be a nonseparating dual cycle in H of length $\mathrm{ewd}(H)$. If $H_0 = H/\varrho$ is obtained by cutting the embedding H along ϱ, then $\mathrm{ewd}(H_0) \geq \frac{1}{2}\,\mathrm{ewd}(H)$.*

3 Drawing Algorithm (the Upper Bound)

Recall that we represent embedded graphs via a rotation system. The topological dual of a graph is easily computable in this representation. We refer to the cyclic permutation of edges incident with a vertex v in embedding H as to the *H-rotation around v*.

Algorithm 3.1 (Drawing a surface-embeddable graph in the plane). Given is a nonplanar graph G embeddable in the orientable surface $\mathcal{S}_g$ of genus g.

I) We construct an embedding G_1 of G in $\mathcal{S}_g$ using Theorem 2.1.

II) For $i = 1, 2, \ldots, g$; we use Theorem 2.2 to compute, in the dual graph G_i^*, a nonseparating dual cycle γ_i of length $c_i = ewd(G_i)$.
We construct an embedding $G_{i+1} = G_i/\gamma_i$ by cutting G_i along γ_i. Notice that G_{i+1} is a spanning subgraph of G_i and G_{i+1} has genus $g - i$.

III) Now, G_{g+1} is a planar embedding. For any edge $e \in F = E(G_1) \setminus E(G_{g+1})$ with ends v_1, v_2, let r_j^e $(j = 1, 2)$ be the face incident with v_j in G_{g+1} such that, if f_1, f_2 are the two consecutive edges of r_j^e at v_j, then e is between f_1 and f_2 in the G_1-rotation around v_j. We compute $R = \{(r_1^e, r_2^e) : e \in F\}$.
For every $(r_1, r_2) \in R$ we compute, using breadth-first search, a shortest dual path $\pi(r_1, r_2)$ between r_1 and r_2 in G_{g+1}^*. This can be done such that no two distinct paths $\pi(r_1, r_2)$, $\pi(r_1', r_2')$ intersect more than once.

IV) Within G_{g+1}, we draw every edge $e \in F$ "along" the dual path $\pi = \pi(r_1^e, r_2^e)$ while crossing the $len(\pi)$ edges of G_{g+1} that are dual to $E(\pi)$.
We output the resulting drawing $\tilde{G}$ isomorphic to input G.

Theorem 3.2. *Assume a graph G is embedded in the orientable surface $\mathcal{S}_g$ of genus g. Let $G = G_1, G_2, \ldots, G_{g+1}$ be the embedded graphs constructed in the iterations of Algorithm 3.1 where, for $i = 1, \ldots, g$, the graph $G_{i+1} = G_i/\gamma_i$ has been obtained by cutting G_i along a nonseparating dual cycle γ_i of length c_i. Let ℓ_i be the length of a shortest dual path in G_{i+1}^* between the two γ_i-cut faces.*

a) The planar drawing $\tilde{G}$ of the graph G produced by Algorithm 3.1 has at most

$$3 \cdot \left(2^{g+1} - 2 - g\right) \cdot \max\{c_i \ell_i : i = 1, 2, \ldots, g\} \quad \textit{crossings.}$$

b) Algorithm 3.1 runs in time $O(n \log n)$ where $n = |V(G)| + |E(G)|$.

Proof. (a) Let $F_i = E(G_i) \setminus E(G_{i+1})$ be the set of edges cut by γ_i at step i. It is easy to prove by induction that, for $k \in \{1, \dots, g\}$ and any edge $e \in F_k$, $len\left(\pi(r_1^e, r_2^e)\right) \le \ell_k + \ell_{k+1} + \cdots + \ell_g$.

So, in step IV of the algorithm, every edge $e \in F_k$ is routed across the plane graph G_{g+1} at cost of $len\left(\pi(r_1^e, r_2^e)\right) \le \sum_{j=k}^{g} \ell_j$ crossings. Given that $|F_k| = c_k$, $2\ell_k \ge c_k$ by Lemma 2.3, and also counting all potential crossings between edges of F_k and of $F_k \cup F_{k+1} \cup \cdots \cup F_g$, we get—over all choices $k \in \{1, \dots, g\}$—the following upper bound on the total number of crossings in the drawing $\tilde{G}$:

$$\sum_{k=1}^{g} c_k \cdot \left(\sum_{j=k}^{g} (c_j + \ell_j) \right) \le \sum_{k=1}^{g} c_k \cdot \left(\sum_{j=k}^{g} 3\ell_j \right) = 3 \sum_{j=1}^{g} \ell_j \cdot \left(\sum_{i=1}^{j} c_i \right)$$

By inductive application of Lemma 2.4, it is $c_i \le 2^{j-i} c_j$ for all $1 \le i < j \le g$, and so we continue using $M = \max\{c_i \ell_i : i = 1, 2, \dots, g\}$.

$$3 \sum_{j=1}^{g} \ell_j \cdot \left(\sum_{i=1}^{j} c_i \right) \le 3 \sum_{j=1}^{g} \ell_j c_j (2^{j-1} + \cdots + 2^1 + 2^0) = 3 \sum_{j=1}^{g} c_j \ell_j (2^j - 1) \le$$
$$\le 3M \cdot (2^1 + 2^2 + \cdots + 2^g - g) = 3(2^{g+1} - 2 - g) \cdot M \tag{1}$$

(b) Step I of Algorithm 3.1 takes time $O(n)$, and the g iterations (g is a constant) in step II take $O(n \log n)$ each. The set R in step III can be computed in time $O(n)$, and since it is easy to prove that $|R| = O(2^g)$, all paths $\pi(r_1, r_2)$ are computed in time $O(n)$ again assuming constant g. Finally, step IV takes time $O\big(n + cr(\tilde{G})\big)$ which is $O(n + M)$ for constant g by (1). However, $M = O\big(cr(G)\big)$ by Theorem 4.1, and $cr(G) = O(n)$ in this case for constant g, e.g. by [3]. Therefore, also step IV is finished in time $O(n)$. ∎

4 Approximation Guarantee (the Lower Bound)

In order to prove that Algorithm 3.1 approximates the optimum crossing number of the input graph G, we have to provide a lower bound on $cr(G)$ that "matches" Theorem 3.2. The involved proof of this lower bound in Theorem 4.1 presents the main new mathematical contribution of our paper.

Theorem 4.1. *Assume the notation of Theorem 3.2. If $ewd(G) \ge 2^{g+2}\Delta$, then*

$$cr(G) \ge 2^{-2g-1} \cdot \Delta^{-2} \cdot \max\{c_i \ell_i : i = 1, 2, \dots, g\}\,.$$

The first observation regarding Theorem 4.1 is that we can safely assume $c_1 \ell_1 = \max\{c_i \ell_i : i = 1, \dots, g\}$ for its proof: If $\max\{c_i \ell_i : i = 1, \dots, g\} = c_j \ell_j$ for $1 < j \le g$, then the embedding $G_j \subset G_1 = G$ (see Theorem 3.2 for the notation) is on a surface of genus $g' = g + 1 - j$ and satisfies $ewd(G_j) \ge 2^{1-j}\, ewd(G_1) \ge 2^{g'+2}\Delta$ using Lemma 2.4. Hence it is enough to prove an analogous statement for $G' = G_j$ and g' instead of G and g. Therefore, we can restate Theorem 4.1 as equivalent Theorem 4.2 below which is formulated independently of Algorithm 3.1.

Theorem 4.2. *Let G be a graph embedded in the orientable surface of genus $g \geq 1$ with nonseparating dual edge-width $c = \mathrm{ewd}(G) \geq 2^{g+2}\Delta(G)$, and let γ be any nonseparating dual cycle in G of length c. If the shortest γ-switching ear in G^* has length ℓ, then the crossing number of G satisfies*

$$cr(G) \geq 2^{-2g-1} \cdot \Delta(G)^{-2} \cdot c\ell \,. \tag{2}$$

The proof of Theorem 4.2 is, however, not straightforward and needs a prior introduction of several new technical terms and claims (Section 5). To motivate these terms (and assist readers' understanding), we first provide an informal outline of our proof ideas. We remark in advance that the coming arguments deal with an actual embedding of the graph G, while such an embedding may not be unique; our proofs can then work with any such embedding.

Proving a lower bound on the crossing number of a graph is quite a difficult task in general. In our previous [19, Theorem 3.3], i.e. in the toroidal case ($g = 1$) of Theorem 4.2, we have found a $\max\left(\left\lfloor \frac{2}{3}\frac{c}{\lfloor \Delta/2 \rfloor}\right\rfloor, \left\lceil \frac{\ell}{\lfloor \Delta/2 \rfloor}\right\rceil\right) \times \left\lfloor \frac{2}{3}\frac{c}{\lfloor \Delta/2 \rfloor}\right\rfloor$ *toroidal grid* (a Cartesian product of two cycles) minor in G, and then used known lower bounds [20] on its crossing number to derive our conclusions (cf. Lemma 5.1).

An extension from this base toroidal case ($g = 1$) to higher surfaces may seem straightforward at a first glance; we should, perhaps, continue cutting the "extra" surface handles in the embedding G while preserving γ and (at least approximately) the parameters c and ℓ, until we get to the toroidal case. Though this could have been considered as a process similar to Algorithm 3.1, it is a fundamentally different task due to the different objectives.

Some deep theoretical problems associated with such a cutting process are, for instance, that cutting a handle of G can drastically decrease the dual nonseparating edge-width on one hand, or turn the loop γ into a separating one on the other hand (hence leaving us with no usable toroidal grid minor). These problems cannot be easily overcome if we need to preserve bounded maximum degree of the graphs. It moreover seems that neither known results on "planarizing cycles", nor the homotopy-related tools from [4], lead to an alternative solution. That is why Theorem 4.2 is actually much harder than the toroidal case in [19].

To resolve the mentioned problems, we introduce a new parameter that is more general than our "$c\ell$ product" from Theorem 4.2. We consider two dual cycles α and β in embedded G: α and β are in a *one-leap* position, adj. *one-leaping*, if the intersection $\alpha \cap \beta$ has exactly one connected component π (a dual path or vertex) such that α and β meet transversally in π (intuitively, they "cross each other" on π). Notice that α, β are then both nonseparating loops, and that $\alpha \cap \beta$ may contain other components in which the cycles meet non-transversally. We define the *stretch* of a non-planar embedded graph G, denoted by $\mathit{stretch}(G)$, as the smallest possible value of $\mathit{len}(\alpha) \cdot \mathit{len}(\beta)$ over all pairs of dual cycles α, β in G in a one-leap position.

The stretch parameter is easier to work with in proofs than "$c\ell$" since stretch is not tied to a particular pair of loops. It can be easily shown that always $\mathit{stretch}(G) \leq 2c\ell$, and that $c\ell \leq \mathit{stretch}(G)$ if G embeds in the torus.

It is relatively straightforward to give a lower bound on $cr(G)$ for any fixed genus g in terms of $stretch(G)$, using successive cuts along shortest dual nonseparating cycles (Lemmas 5.2 and 5.1). However, apart from the toroidal case, it can easily happen that $stretch(G) \ll c\ell$. To overcome this complication (and to show that the stretch eventually "becomes" $\Omega(c\ell)$ during the cutting process), we trace in our graphs a pair of rather artificial objects (see in Lemma 5.4, a,b) which initially correspond to γ and its switching ear in G^*; later on they "vanish" whenever the stretch of the cut-subgraph of G becomes large enough. So, this approach finally leads to the claimed lower bound (2).

Finally, Theorem 4.2 implies Theorem 4.1 and, combining the latter with Theorem 3.2, we get the main conclusion:

Corollary 4.3 (Theorem 1.1). *Let an input graph G be embeddable in the orientable surface of genus $g \geq 1$ with dual edge-width $\mathrm{ewd}(G) \geq 2^{g+2}\Delta(G)$. Then Algorithm 3.1 outputs a drawing of G in the plane with at most $3 \cdot 2^{3g+2} \cdot \Delta(G)^2 \cdot \mathrm{cr}(G)$ crossings.*

∎

5 Lower Bound Proof

In this section, we give a formal proof of Theorem 4.2. As already mentioned, the central notion of this proof is that of the *stretch* of an embedded graph. We now give three technical claims describing its properties.

Lemma 5.1 (cf. Hliněný and Salazar [19]). *Let G be a graph embedded in the torus such that $\mathrm{ewd}(G) \geq 8\Delta(G)$. Then $\mathrm{cr}(G) \geq \frac{1}{8}\Delta(G)^{-2} \cdot stretch(G)$.*

Lemma 5.2. *Let G be a graph embedded in an orientable surface of genus at least 2, and ϱ be a nonseparating dual cycle in G of length $\mathrm{ewd}(G) = \mathrm{len}(\varrho)$. Denote by $G_0 = G/\varrho$ the embedding obtained by cutting G along ϱ. Then $stretch(G_0) \geq \frac{1}{4}\, stretch(G)$.*

Lemma 5.3. *Let H be a graph embedded in an orientable surface of genus ≥ 2, and let $\alpha, \beta \subseteq H^*$ be a one-leaping pair of dual cycles gaining the stretch $\mathrm{len}(\alpha) \cdot \mathrm{len}(\beta)$ of H such that $\mathrm{len}(\alpha) \leq \mathrm{len}(\beta)$. We denote by $H_0 = H/\alpha$ the embedding obtained by cutting H along α. Then $\mathrm{ewd}(H_0) \geq \frac{1}{2}\,\mathrm{ewd}(H)$.*

The rest of the proof of Theorem 4.2 needs a further generalization of the concepts of switching and leaping. We assume a graph H cellularly embedded in a surface Σ, and choose a subgraph (not necessarily connected) $D \subseteq H$. The *H-induced* embedding $\tilde{D}$ of the graph D is determined by the system of H-rotations around vertices of D restricted to $E(D)$. Intuitively, $\tilde{D}$ is obtained from the usual subembedding of D in Σ via replacing all non-cellular faces with discs. Notice that $\tilde{D}$ has a separate surface for each connected component of D.

We consider a dual subgraph $\delta \subseteq H^*$, and its H^*-induced embedding $\tilde{\delta}$. If $\tilde{\delta}$ can be face-bicoloured, then we say that δ is *bipolar in H^**, and we associate one chosen facial bicolouring of $\tilde{\delta}$ with δ (notice that this bicolouring is not unique when δ is not connected).

We will refer to the facial colours of $\tilde{\delta}$ (white and black) as to the δ-polarities in H^* (positive and negative). More formally, let a *halfedge* be a pair $\langle e, v\rangle$ where e is an edge and v is one of the two ends of e. For $v \in V(\delta)$ and $e \notin E(\delta)$, the halfedge $\langle e, v\rangle$ ("e at v") has a *positive* (*negative*) δ-polarity if the position of e in the H^*-rotation around v is between consecutive edges of a white (black) $\tilde{\delta}$-face.

Clearly, a dual cycle in any embedding is always bipolar. On the other hand, a bipolar graph δ must be Eulerian. A δ-ear π is *δ-polarity switching* if the halfedges of π incident with the ends of π are of distinct δ-polarities. If δ is a dual cycle, then being "δ-polarity switching" is equivalent to being "δ-switching".

We now consider a bipolar dual subgraph δ in H^*, and a (closed) dual walk $\omega \subseteq H^*$. A proper subwalk μ of ω is called a *leap (of ω and δ)* if μ belongs to the intersection $\delta \cap \omega$, neither the dual edge f_0 preceding μ in ω nor the dual edge f_1 succeeding μ in ω belong to δ, and the halfedges of f_0, f_1 incident to μ are of distinct δ-polarities. We say that ω is *odd-leaping* δ if the number of all proper subwalks of ω which are leaps is odd, and ω is *even-leaping* δ otherwise. Notice that being "one-leaping" (Section 4) implies "odd-leaping" in this new sense.

The core step of our inductive approach to Theorem 4.2 is the next claim. As outlined in Section 4, the intuition behind the application of Lemma 5.4 is to suitably (by careful choice of α in the lemma) "cut down" the embedding G to a toroidal one, while "preserving γ" (actually represented by δ and δ_0 in (a,a') below), and also keeping the "switching distance" (see (c,c') below) sufficiently long. The conditions (b) and (b') in Lemma 5.4 have purely technical purpose.

Notice, for instance, that if (b) is true, then the embedding H is not planar (and so the stretch of H is well defined): A closed walk odd-leaping a bipolar planar graph δ cannot exist since plane δ equals its H^*-induced embedding $\tilde{\delta}$, which means that δ is face-bicoloured, too, and a simple parity argument then gives a contradiction. For a similar "parity reason", (b) implies that (c) a δ-polarity switching ear in H^* must exist. Moreover, as we proceed in the cutting process, nonplanarity implied by (b') guarantees that we will eventually arrive at the exceptional conclusion (d) *len*$(\beta) \geq h$ in Lemma 5.4.

Lemma 5.4. *Let a graph H be embedded in an orientable surface, and assume*

a) there is a bipolar dual subgraph δ in H^*,
b) there exists a closed walk in H^* that is odd-leaping δ, and
c) the shortest δ-polarity switching ear in H^* has length h.

Let α, β be a one-leaping pair of dual cycles in H^* such that $\mathrm{len}(\alpha) \leq \mathrm{len}(\beta)$ *and* $\mathrm{stretch}(H) = \mathrm{len}(\alpha) \cdot \mathrm{len}(\beta)$. *We denote by $H_0 = H/\alpha$ the embedded subgraph of H obtained by cutting H along α. Unless (d)* $\mathrm{len}(\beta) \geq h$, *the following hold*

a') there is a bipolar dual subgraph δ_0 ("induced" by δ) in H_0^*,
b') there exists a closed walk in H_0^* that is odd-leaping δ_0, and
c') the shortest δ_0-polarity switching ear in H_0^* has length $h_0 \geq h - \frac{1}{2}\,\mathrm{len}(\alpha)$.

We are finally ready to give the proof of our main theorem.

Proof of Theorem 4.2. Our proof applies induction based on Lemma 5.4. Notice that all the conditions (a),(b),(c) of Lemma 5.4 are satisfied by the graph $G = H$, its bipolar dual cycle $\gamma = \delta$, and by $h = \ell$. Precisely, we prove the following claim by induction on g:

(5.5) Let H_1 be a graph embedded in an orientable surface of genus g, and h_1 be an integer. Assume that $H = H_1$ either (i) satisfies the conditions (a),(b),(c) of Lemma 5.4 with $h = h_1$ and some δ, or (ii) $g \geq 1$ and $\mathit{stretch}(H_1) \geq h_1 \cdot \mathit{ewd}(H_1)$. Then there exists a subgraph $T_0 \subseteq H_1$ that embeds in the torus with $\mathit{ewd}(T_0) \geq 2^{1-g}\,\mathit{ewd}(H_1)$, and T_0 has $\mathit{stretch}(T_0) \geq 2^{2-2g} \cdot h_1 \cdot \mathit{ewd}(H_1)$.

If (5.5) is true, then the rest of the proof is easily finished. We set $H_1 = G$ and $h_1 = \ell$, and hence immediately $\mathit{stretch}(T_0) \geq 2^{2-2g} c\ell$. Recalling $c = \mathit{ewd}(G) \geq 2^{g+2}\Delta(G)$, we get $\mathit{ewd}(T_0) \geq 2^3\Delta(G)$, and so by Lemma 5.1,

$$\mathit{cr}(G) \geq \mathit{cr}(T_0) \geq \frac{1}{8} \cdot \Delta(T_0)^{-2} \cdot 2^{2-2g} \cdot c\ell \geq 2^{-2g-1} \cdot \Delta(G)^{-2} \cdot c\ell\,.$$

∎

6 Concluding Remarks

It is a natural question, whether the bounds obtained above can be improved. Comparing our bound with related results, the dependency on $\Delta(G)^2$ seems unavoidable. The exponential dependency on g, on the other hand, is much more interesting—it pops up independently in multiple places within the proof, and these occurrences seem unavoidable on a local scale, when considering each inductive step independently. Yet, it is hard to construct an example where the exponential decrease can actually be observed; it might be that a completely different approach with a global view can reduce the dependency to some $\mathit{poly}(g)$ factor, cf. also [9].

It would be interesting to extend our approach also to non-orientable surfaces. While this looks promisingly straightforward at a first glance, everything becomes much more difficult due to the fact that a "cheapest" cut through an embedding may have three different forms: a two-sided loop cutting a handle or an anti-handle, or a one-sided loop cutting a crosscap. In particular, one has to consider projective, toroidal, and Klein-bottle grids together as the base cases, and to go through many more cases in the (already complicated) cutting process.

Another interesting topic of further research is the role of the "density" requirement $\mathit{ewd}(G) \geq 2^{g+2}\Delta$ in Theorem 4.2 or 1.1. On one hand, an assumption like this one is clearly necessary for proving a lower bound on $\mathit{cr}(G)$ of order $\Omega(c\ell)$ as in Theorem 4.2 (since planar graphs can have "non-dense" cellular embeddings in higher surfaces). On the other hand, if $\mathit{ewd}(G)$ was low (constant in the case of bounded genus and degree), then one could use a solution to the multiple-edge insertion problem, instead, to approximate $\mathit{cr}(G)$ by [7]. Unfortunately, no polynomial time (though approximation) algorithm is currently known for solving this multiple-edge insertion problem—unlike for the single edge and vertex insertion problems [6].

References

1. C. Batini, M. Talamo and R. Tamassia, *Computer Aided Layout of Entity-Relationship Diagrams*, J. of Systems and Software 4 (1984), 163–173.
2. S.N. Bhatt, F.T. Leighton, *A framework for solving VLSI graph layout problems*. J. of Computer and System Sciences 28 (1984), 300–343.
3. K. Böröczky, J. Pach, and G. Tóth, *Planar crossing numbers of graphs embeddable in another surface*. Internat. J. Found. Comput. Sci. 17 (2006), 1005–1015.
4. R. Brunet, B. Mohar, R.B. Richter, *Separating and nonseparating disjoint homotopic cycles in graph embeddings*. J. of Combinatorial Theory ser. B 66 (1996), 201–231.
5. S. Cabello and B. Mohar, *Crossing and weighted crossing number of near planar graphs*. In Graph Drawing 2008, LNCS 5417, Springer Verlag (2009), 38–49.
6. M. Chimani, C. Gutwenger, P. Mutzel, and C. Wolf, *Inserting a vertex into a planar graph*. In ACM-SIAM Symposium on Discrete Algorithms 2009; ACM Press (2009), 375–383.
7. M. Chimani, P. Hliněný, and P. Mutzel, *Vertex Insertion approximates the Crossing Number for Apex Graphs*. manuscript (2008).
   A poster in Graph Drawing 2008, LNCS 5417, Springer Verlag (2009), 432–434.
8. M. Chimani, P. Mutzel, and I. Bomze, *A New Approach to Exact Crossing Minimization*. In European Symposium on Algorithms 2008; LNCS 5193, Springer (2008), 284–296.
9. H. Djidjev and I. Vrt'o, *Planar crossing numbers of genus g graphs*. In: Proc. 33rd ICALP, Lecture Notes in Computer Science 4051, Part I, Springer–Verlag (2006), 419–430.
10. G. Even, S. Guha, and B. Schieber, *Improved approximations of crossings in graph drawing*. In ACM Symposium on Theory of Computing 2000, (2000), 296–305.
11. M.R. Garey and D.S. Johnson, *Crossing number is NP-complete*. SIAM J. Algebraic Discrete Methods 4 (1983), 312–316.
12. I. Gitler, P. Hliněný, J. Leanos, and G. Salazar, *The crossing number of a projective graph is quadratic in the face-width*. Electr. J. of Combinatorics 15 (2008), #R46.
13. M. Grohe, *Computing Crossing Numbers in Quadratic Time*. J. Comput. Syst. Sci. 68 (2004), 285–302.
14. C. Gutwenger and P. Mutzel, *An Experimental Study of Crossing Minimization Heuristics*, In Proc. GD '03; LNCS 2912, Springer (2004), 13–24.
15. R.K. Guy, *The Decline and Fall of Zarankiewicz's Theorem*, In: Proof Techniques in Graph Theory, Proc. 2nd Ann Arbor Graph Theory Conference; Academic Press (1969), 63–69.
16. R.K. Guy, *Crossing numbers of graphs*, In: Proc. Graph Theory and Applications, LNM, Springer (1972), 111–124.
17. P. Hliněný, *Crossing Number is Hard for Cubic Graphs*, J. of Combinatorial Theory, Series B 96 (2006), 455–471.
18. P. Hliněný and G. Salazar, *On the Crossing Number of Almost Planar Graphs*. In Graph Drawing 2006; LNCS 4372, Springer (2007), 162–173.
19. P. Hliněný and G. Salazar, *Approximating the Crossing Number of Toroidal Graphs*, In ISAAC 2007; LNCS 4835, Springer (2007), 148–159.

20. H.A. Juárez and G. Salazar, *Drawings of $C_m \times C_n$ with one disjoint family II.* J. of Combinatorial Theory ser. B 82 (2001), 161–165.
21. K. Kawarabayashi and B. Reed, *Computing crossing number in linear time*, In Proc. STOC '07, (2007), 382–390.
22. M. Kutz, *Computing shortest non-trivial cycles on orientable surfaces of bounded genus in almost linear time.* In: Annual Symposium on Computational Geometry 2006, ACM Press (2006), 430–438.
23. B. Mohar, *A linear time algorithm for embedding graphs in an arbitrary surface.* SIAM J. Discrete Math. 12 (1999), 6–26.
24. B. Mohar and C. Thomassen, Graphs on surfaces. Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press (2001), Baltimore MD, USA.
25. S. Pan and R.B. Richter, *The crossing number of K_{11} is* 100, J. of Graph Theory 56(2) (2007), 128–134.
26. P. Turán, *A note of welcome*, J. of Graph Theory 1 (1977), 7–9.
27. I. Vrťo, Bibliography on Crossing Numbers. Available online at `http://www.ifi.savba.sk/~imrich/crobib.ps`.
28. K. Zarankiewicz, *The solution of a certain problem on graphs of P. Turán*, Bulletin de l'Academie Polonaise des sciences, Cl. III 1 (1953), 167–168.

7 APPENDIX

Here we include additional proofs for all the claims in our paper.

7.1 Supplements for Section 2

Let H result by cutting an embedded graph G along γ. From a dual point of view, H^* results from G^* by contracting the dual cycle γ into a single vertex, and then splitting it into c_1 and c_2, the two γ-cut faces. So every dual edge in $E(H^*)$ has a naturally corresponding dual edge in $E(G^*)$, and for every dual subgraph $\sigma \subseteq H^*$ there is a unique dual graph $\hat{\sigma} \subseteq G^*$ (the *lift of* σ) induced by the edges corresponding to $E(\sigma)$ in $E(G^*)$.

Lemma 2.4. *Let H be a graph embedded in an orientable surface of genus ≥ 2, and ϱ be a nonseparating dual cycle in H of length* ewd(H)*. If $H_0 = H/\varrho$ is obtained by cutting the embedding H along ϱ, then* ewd$(H_0) \geq \frac{1}{2}$ ewd(H)*.*

Proof. We denote by r_1, r_2 the two ϱ-cut faces of H_0, in other words the new dual vertices $r_1, r_2 \in V(H_0^*) \setminus V(H^*)$. Suppose that σ is a nonseparating cycle in H_0^* of length $ewd(H_0)$. If σ avoids both r_1, r_2, then its lift $\hat{\sigma}$ in H^* is a cycle again, and so $ewd(H) \leq len(\sigma) = ewd(H_0)$. If σ hits both r_1, r_2 and $\pi \subseteq \sigma$ is one of the dual paths with the ends r_1, r_2, then the lift $\hat{\pi}$ is a ϱ-switching ear in H^* as can be seen from the definition. So $ewd(H_0) = len(\sigma) \geq len(\hat{\pi}) \geq \frac{1}{2}\, ewd(H)$ by Lemma 2.3.

Hence it remains to consider that σ, up to symmetry, hits r_1 and avoids r_2. Then its lift $\hat{\sigma}$ is a ϱ-ear; if $\hat{\sigma}$ itself is a cycle, then we are done as above. Otherwise, $\hat{\sigma} \cup \varrho \subset H^*$ forms a theta dual subgraph, and so there are exactly three dual cycles $\gamma_1, \gamma_2, \gamma_3 \subseteq \hat{\sigma} \cup \varrho$. Loop σ is nonseparating in the embedding surface of H/ϱ, so each of $\gamma_1, \gamma_2, \gamma_3$ is nonseparating in that of H, and hence $len(\gamma_i) \geq ewd(H)$ for $i = 1, 2, 3$. Since every edge of $\hat{\sigma} \cup \varrho$ is in two of $\gamma_1, \gamma_2, \gamma_3$, it is $len(\gamma_1) + len(\gamma_2) + len(\gamma_3) = 2\, len(\varrho) + 2\, len(\hat{\sigma}) = 2\, ewd(H) + 2\, len(\hat{\sigma})$ and $len(\gamma_1) + len(\gamma_2) + len(\gamma_3) \geq 3\, ewd(H)$, from which we get $ewd(H_0) = len(\sigma) = len(\hat{\sigma}) \geq \frac{1}{2}\, ewd(H)$ again. ∎

7.2 Supplements for Section 3

Proof of Theorem 3.2. We refer to the notation from Algorithm 3.1. The proof skips an argument for the following claim: for $k \in \{1, \ldots, g\}$ and any edge $e \in F_k$, $len\left(\pi(r_1^e, r_2^e)\right) \leq \ell_k + \ell_{k+1} + \cdots + \ell_g$.

Let the ends of e be $v_1 v_2$. For the inductive argument, let $s_j^{e,i}$, $j = 1, 2$, denote the face of G_{i+1}, $k \leq i \leq g$, defined analogously to r_j^e above. By induction on i, the dual distance between $s_1^{e,i}$ and $s_2^{e,i}$ in G_{i+1}^* is $d_i(s_1^{e,i}, s_2^{e,i}) \leq \ell_k + \cdots + \ell_i$. This holds at equality for $i = k$ by the definition of ℓ_k. Considering step $i + 1$, we see that the dual distance between $s_1^{e,i+1}$ and $s_2^{e,i+1}$ may grow by at most ℓ_{i+1}—the dual distance between the two γ_{i+1}-cut faces in G_{i+2}^*. Finally, $len\left(\pi(r_1^e, r_2^e)\right) = d_g(s_1^{e,g}, s_2^{e,g})$. ∎

7.3 Supplements for Section 5

Lemma 7.2. *If δ is a nonseparating dual cycle in an embedded graph G, and ε is a δ-switching ear in G^*, then* $stretch(G) \leq len(\delta) \cdot \big(len(\varepsilon) + \frac{1}{2}\, len(\delta)\big)$.

Proof. The ends of ε partition δ into two paths $\delta_1, \delta_2 \subseteq \delta$, and, up to symmetry, it is $len(\delta_1) \leq \frac{1}{2}\, len(\delta)$. In a degenerate case, δ_1 can be a single vertex. Since δ and $\varepsilon \cup \delta_1$ are dual cycles in a one-leap position, we have $stretch(G) \leq len(\delta) \cdot (len(\varepsilon) + len(\delta_1))$. ∎

Lemma 5.1. *Let G be a graph embedded in the torus such that $ewd(G) \geq 8\Delta(G)$. Then* $cr(G) \geq \frac{1}{8}\Delta(G)^{-2} \cdot stretch(G)$.

Proof. Let γ and c, ℓ be defined as in Theorem 4.2, and λ be a γ-switching ear of length ℓ in G^*. The result [19, Corollary 3.4] claims, under the assumptions that G is toroidal and $ewd(G) \geq 8\Delta(G)$, the bound $cr(G) \geq \frac{1}{4}\Delta(G)^{-2} \cdot c\ell$. By Lemma 2.3 ($\varrho = \gamma$), $\ell \geq \frac{c}{2}$. Hence by Lemma 7.2, $stretch(G) \leq c(\ell + \frac{c}{2}) \leq 2c\ell$, and it follows $cr(G) \geq \frac{1}{8}\Delta(G)^{-2} \cdot 2c\ell \geq \frac{1}{8}\Delta(G)^{-2} \cdot stretch(G)$. ∎

Lemma 5.2. *Let G be a graph embedded in an orientable surface of genus at least 2, and ϱ be a nonseparating dual cycle in G of length $ewd(G) = len(\varrho)$. Denote by $G_0 = G/\varrho$ the embedding obtained by cutting G along ϱ. Then* $stretch(G_0) \geq \frac{1}{4}\, stretch(G)$.

Proof. We denote by r_1, r_2 the two dual vertices (the cut-faces in G_0) of the cut ϱ, and recall the notion of a lift from Section 2. To prove the statement, we are going to give an upper bound $stretch(G) \leq 4\, stretch(G_0)$. We assume that $stretch(G_0) = ab$ is gained by a pair of dual one-leaping cycles α, β in G_0^* such that $a = len(\alpha)$, $b = len(\beta)$. Using Lemma 2.4 and the fact that both α, β are nonseparating, we get

$$a, b \geq ewd(G_0) \geq \frac{1}{2}\, ewd(G) = \frac{1}{2}\, len(\varrho). \tag{3}$$

We firstly suppose that $V(\alpha \cup \beta)$ contains both the cut-faces r_1, r_2 (which are naturally treated as vertices in G_0^*). Then there exists a path $\pi \subseteq \alpha \cup \beta$ connecting r_1 to r_2 such that $len(\pi) \leq \frac{1}{2}(a + b)$. Clearly, its lift $\hat{\pi}$ is a ϱ-switching ear in G^*, and so by Lemma 7.2 and (3),

$$\begin{aligned} stretch(G) &\leq len(\varrho) \cdot (len(\hat{\pi}) + \frac{1}{2}\, len(\varrho)) \leq len(\varrho) \cdot \frac{1}{2}(a + b + len(\varrho)) \\ &\leq \frac{1}{2}(2ba + 2ab + 4ab) = 4\, stretch(G_0). \end{aligned}$$

Secondly, we suppose that, up to symmetry, $r_2 \notin V(\alpha \cup \beta)$ but possibly $r_1 \in V(\alpha \cup \beta)$. For the lift $\hat{\alpha}$ in G^* (which is a ϱ-ear in the case $r_1 \in V(\alpha)$), we define $\bar{\alpha}$ equal to $\hat{\alpha}$ if it is a cycle, and otherwise $\bar{\alpha} = \hat{\alpha} \cup \varrho_0$ where $\varrho_0 \subseteq \varrho$ is the shorter subpath with the same ends on ϱ as $\hat{\alpha}$. We define $\bar{\beta}$ analogously. With the help of a simple case-analysis, we argue that $\bar{\alpha}$, $\bar{\beta}$ forms a one-leaping pair of dual cycles in G^*, and so again

$$
\begin{aligned}
stretch(G) &\le len(\bar{\alpha}) \cdot len(\bar{\beta}) \le (a + \frac{1}{2}\, len(\varrho)) \cdot (b + \frac{1}{2}\, len(\varrho)) \\
&\le (a + a) \cdot (b + b) = 4\, stretch(G_0).
\end{aligned}
$$

∎

Lemma 5.3. *Let H be a graph embedded in an orientable surface of genus ≥ 2, and let $\alpha, \beta \subseteq H^*$ be a one-leaping pair of dual cycles gaining the stretch $len(\alpha) \cdot len(\beta)$ of H such that $len(\alpha) \le len(\beta)$. We denote by $H_0 = H/\alpha$ the embedding obtained by cutting H along α. Then $ewd(H_0) \ge \frac{1}{2}\, ewd(H)$.*

Proof. A 3-*path condition* of a property $\mathcal{P}$ says that if T is a theta graph, and two of the three cycles of T do not possess $\mathcal{P}$, then neither the third cycle does (cf. [24, Section 4.3]). Our key observation at this point is:

(7.6) Given an embedded graph G and a fixed dual cycle $\gamma \subseteq G^*$, the dual cycles in G^* satisfy the 3-path condition w.r.t. the property of being odd-leaping γ.

Let a dual theta graph $\tau \subseteq G^*$ be formed by three paths $\tau = \tau_1 \cup \tau_2 \cup \tau_3$ connecting dual vertices s, t in G^*. We consider a connected component μ of $\gamma \cap \tau$. If $\mu = \emptyset$ or $\mu = \gamma$, then the 3-path condition holds. Otherwise, μ is a dual path with ends m_1, m_2 in G^*. We denote by f_1, f_2 the edges in $E(\gamma) \setminus E(\mu)$ incident with m_1, m_2, respectively, and by μ^+ the union of μ and the halfedges $\langle f_1, m_1 \rangle$ and $\langle f_2, m_2 \rangle$. We show that the number q of leaps of μ^+ summed over all three cycles in τ is always even.

If $m_i \notin \{s, t\}$ for $i \in \{1, 2\}$, then contracting the edge of μ incident to m_i clearly does not change the number q. Iteratively applying this argument, we can assume that finally either (i) $m_1 = m_2$ (and possibly $m_1 \in \{s, t\}$), or (ii) $m_1 = s$, $m_2 = t$, and $\mu = \tau_1$. In case (i), μ^+ leaps either none or two of the cycles of τ in the single vertex m_1, and so $q \in \{0, 2\}$.

For $i = 1, 2, 3$, let the edge of τ_i incident to s be e_i and the one incident to t be e'_i. In case (ii) we routinely investigate, up to symmetry, two cyclic permutations (e_1, f_1, e_2, e_3) and (e_1, e_2, f_1, e_3) in the G^*-rotation around s, and all six cyclic permutations of e'_1, e'_2, e'_3, f_2 in the G^*-rotation around t. In all twelve possibilities, we get $q \in \{0, 2\}$ except the case of (e_1, e_2, f_1, e_3), (e'_1, e'_2, f_1, e'_3) when μ^+ leaps twice the cycle $\tau_2 \cup \tau_3$ and $q = 4$.

Altogether, the number of leaps of γ summed over all three cycles in τ is even. Hence the number of cycles of τ which are odd-leaping with γ is also even, and the 3-path condition follows.

Using (7.6), we observe another useful claim:

(7.7) If ψ, φ is an odd-leaping pair of dual cycles in H^*, then $stretch(H) \le len(\psi) \cdot len(\varphi)$.

We choose the odd-leaping pair ψ, φ such that $len(\psi) \cdot len(\varphi)$ is minimized. Up to symmetry, $len(\psi) \le len(\varphi)$. Since $\psi \cap \varphi \neq \emptyset$, there is a set $\mathcal{S} = \{\varphi_1, \ldots, \varphi_k\}$ of pairwise edge-disjoint ψ-ears in φ, such that $E(\varphi_1) \cup \cdots \cup E(\varphi_k) = E(\varphi) \setminus E(\psi)$. By a simple parity argument, there exists a ψ-switching ear in $\mathcal{S}$. Hence if $|\mathcal{S}| = 1$, then ψ, φ are one-leaping, and the claim is trivial.

If more than one ψ-ear in $\mathcal{S}$ is switching, then we pick, say, φ_1 as the shorter of these. It is $len(\varphi_1) \leq \frac{1}{2}\, len(\varphi)$, and so by Lemma 7.2, we get

$$stretch(H) \leq len(\psi) \cdot \left(len(\varphi_1) + \frac{1}{2}\, len(\psi) \right) \leq len(\psi) \cdot \left(\frac{1}{2}\, len(\varphi) + \frac{1}{2}\, len(\varphi) \right) .$$

Finally, $|\mathcal{S}| > 1$ and exactly one ψ-ear in $\mathcal{S}$ is switching, say φ_1 again. We pick any $\varphi_j \in \mathcal{S}$, $j > 1$, with the ends u, v on ψ, and compare the distance d between u and v on ψ with $len(\varphi_j)$. If $d > len(\varphi_j)$, then both cycles of $\psi \cup \varphi_j$ containing φ_j are shorter than $len(\psi)$, and one of them is odd-leaping with φ by (7.6). This contradicts our minimum choice of ψ. Hence $len(\varphi_j) \geq d$, and summing this inequalities over all $j = 1, \ldots, k$ we get $len(\varphi_1) \leq len(\varphi) - s$, where s is the distance between the ends of φ_1 on ψ. In the same way as in Lemma 7.2, we thus get (7.7)

$$stretch(H) \leq len(\psi) \cdot (len(\varphi_1) + s) \leq len(\psi) \cdot (len(\varphi) - s + s) .$$

Now we turn our attention to the statement of the lemma. Assume that σ is a nonseparating dual cycle in H_0^* of length $ewd(H_0)$. If its lift $\hat{\sigma}$ is a cycle again, then $ewd(H) \leq len(\hat{\sigma}) = ewd(H_0)$ since $\hat{\sigma}$ is nonseparating in H^*. Otherwise, $\hat{\sigma}$ contains an α-ear $\pi \subseteq \hat{\sigma}$ such that $\alpha \cup \pi$ is a theta graph, and we denote by $\alpha_1, \alpha_2 \subseteq \alpha$ the subpaths divided by the ends of π on α. By (7.6), exactly two of the cycles of $\alpha \cup \pi$ are odd-leaping with β—one of them is α and the other one, say, is $\alpha_1 \cup \pi$. Then $len(\alpha_1 \cup \pi) \geq len(\alpha)$ using (7.7), and so $len(\pi) \geq len(\alpha_2)$. Furthermore, $\alpha_2 \cup \pi$ is nonseparating in H^*, and we conclude

$$ewd(H) \leq len(\alpha_2 \cup \pi) \leq 2\, len(\pi) \leq 2\, len(\hat{\sigma}) = 2\, ewd(H_0) .$$

∎

Lemma 5.4. *Let a graph H be embedded in an orientable surface, and assume*

a) there is a bipolar dual subgraph δ in H^*,
b) there exists a closed walk in H^* that is odd-leaping δ, and
c) the shortest δ-polarity switching ear in H^* has length h.

Let α, β be a one-leaping pair of dual cycles in H^* such that $\mathrm{len}(\alpha) \leq \mathrm{len}(\beta)$ *and* $stretch(H) = \mathrm{len}(\alpha) \cdot \mathrm{len}(\beta)$. *We denote by $H_0 = H/\alpha$ the embedded subgraph of H obtained by cutting H along α. Unless (d)* $\mathrm{len}(\beta) \geq h$, *the following hold*

a') there is a bipolar dual subgraph δ_0 ("induced" by δ) in H_0^*,
b') there exists a closed walk in H_0^* that is odd-leaping δ_0, and
c') the shortest δ_0-polarity switching ear in H_0^* has length $h_0 \geq h - \frac{1}{2}\, len(\alpha)$.

Proof. Recall the definition of cutting an embedding H along a dual cycle α. The dual graph $(H/\alpha)^*$ is obtained from H^* by successive contractions of all the dual edges in $E(\alpha)$ into one dual vertex a, and then "splitting" a into two a_1, a_2 (giving the two α-cut faces of

H/α). This "stepwise contraction" perspective of cutting turns out to be very useful in our proof.

(a') Let ε denote the subgraph of H_0^* induced by the edges $E(\delta) \setminus E(\alpha)$. If $\alpha = \delta$, then clearly (d) $len(\beta) \geq h$, and so ε can be considered nonempty further on. We show that we can choose $\delta_0 = \varepsilon$ unless α contains a δ-polarity switching ear. If the latter happened, then it would mean (d) $len(\beta) \geq len(\alpha) \geq h$ by (c).

To avoid confusion with the operation of cutting, we denote by $G \triangleleft e$ the result of contracting the edge e in a graph G. From the definition of bipolarity, one straightforwardly derives, in our context:

(7.9) If $f \in E(H^*)$ is not a loop-edge and not a δ-polarity switching ear, then the dual graph $H^* \triangleleft f$ (after contraction of f) is embedded in the same surface as H^*, and the dual subgraph δ' induced by $E(\delta) \setminus \{f\}$ in $H^* \triangleleft f$ is bipolar again, where the δ'-polarities are naturally inherited from δ-polarities.

Since we assume that α contains no δ-polarity switching ear, we can iteratively apply (7.9) to all edges of α except some (last one) $f_1 \in E(\alpha) \setminus E(\beta)$. In this way we get an "intermediate" embedding $H_1^* = H^* \triangleleft \big(E(\alpha) \setminus \{f_1\}\big)$ such that f_1 is a nonseparating dual loop-edge in H_1^*, and bipolar $\varepsilon_1 \subseteq H_1^*$ is naturally derived from δ. Let a be the face of H_1 that is the double end of f_1, and let the H_1^*-rotation of edges around a be $e_1, \ldots, e_i, f_1, e_1', \ldots, e_j', f_1$. The last step in the construction of H_0^* (and of ε) is to remove f_1 and split a into a_1, a_2 such that the H_0^*-rotation around a_1 is $e_1, \ldots, e_i$ and around a_2 is $e_1', \ldots, e_j'$.

Clearly, $\varepsilon_1 = \varepsilon$ stays bipolar in H_0^* if $a \notin V(\varepsilon_1)$, and so we assume $a \in V(\varepsilon_1)$. Let $\tilde{\varepsilon}$ denote the H_0^*-induced embedding of ε. Let e_a and e_b be the first and last element of the list $e_1, \ldots, e_i$, respectively, that are also edges of ε. Note that both ends of f_1 in the H_1^*-rotation around a are between e_b and e_a. Then, e_b, e_a appear consecutively on a unique face x of $\tilde{\varepsilon}$. We define a face x' in $\tilde{\varepsilon}$ analogously at a_2. Loosely speaking, x, x' are the dual $\tilde{\varepsilon}$-faces "inheriting" the two H_1^*-faces incident with f_1. If $f_1 \notin E(\varepsilon_1)$, then both halfedges of f_1 are of the same ε_1-polarity by our assumption on α, say positive. Hence both $\tilde{\varepsilon}$-faces x and x' will get (consistently) positive polarity, and so ε is bipolar in H_0^*.

If, on the other hand, $f_1 \in E(\varepsilon_1)$, then one of the two faces incident with f_1 in the H_1^*-induced embedding $\tilde{\varepsilon}_1$ of ε_1 is positive, say the one containing edge(s) from $e_1, \ldots, e_i$, and the other one is negative. Then the $\tilde{\varepsilon}$-face x will be (consistently) positive and x' negative. Hence again, $\varepsilon = \delta_0$ is bipolar in H_0^*.

(b') As in (a'), we assume that α contains no δ-polarity switching ear. Similarly with β; if there is a δ-polarity switching ear contained in β, then $len(\beta) \geq h$.

The following counterpart of claim (7.9), formulated for any closed dual walk ψ in H^*, is easily derived from our definition of a leap:

(7.10) Suppose $f \in E(H^*)$ is not a loop-edge and not a δ-polarity switching ear, and denote by δ', ψ' the dual subgraphs induced by $E(\delta) \setminus \{f\}$ and $E(\psi) \setminus \{f\}$ in $H^* \triangleleft f$ (after contraction of f). Then the number of leaps of δ' and ψ' in $H^* \triangleleft f$ is the same as the number of leaps of δ and ψ in H^*, with an exception when $f \in E(\psi) \setminus E(\delta)$ and both ends of f are incident with leaps of δ and ψ in H^*—in this case the two leaps vanish in $H^* \triangleleft f$.

We proceed in the same way as in (a'), and use the same notation H_1^*, f_1, a, ε_1, etc., as in (a'). Let ω be a dual closed walk in H^* odd-leaping δ, and ω_1, β_1 denote the dual closed walks in H_1^* induced by $E(\omega) \cap E(H_1^*)$ and $E(\beta) \cap E(H_1^*)$. By iterative application of (7.10) to all edges in $E(\alpha) \setminus \{f_1\}$, we get that the parity of leaping between δ and ω (δ and β) in H^* is the same as that between ε_1 and ω_1 (ε_1 and β_1) in H_1^*. Hence ω_1 is odd-leaping ε_1, and β_1 is even-leaping ε_1, since β contains no δ-polarity switching ear in H^* and so β is not odd-leaping δ.

We note that $a \in V(\beta_1)$ since α intersects β, and recall $f_1 \notin E(\beta)$. If $f_1 \in E(\omega)$, then we moreover remove f_1 from ω_1 which does not change the parity of leaping between ε_1 and ω_1. We say that the dual walk ω_1 *passes through* a in H_1^* if one edge of ω_1 is from $e_1, \dots, e_i$ and the next edge of ω_1 is among $e'_1, \dots, e'_j$, or vice versa. Every time ω_1 passes through a, we replace this pass by one iteration of the cycle β_1. The resulting closed dual walk ω_2 in H_1^* (which does not pass through a) is again odd-leaping ε_1 since β_1 is even-leaping ε_1. Then ω_0 which is induced by $E(\omega_2)$ in the graph H_0^* is a closed dual walk odd-leaping $\varepsilon = \delta_0$.

(c') Let σ be a δ_0-polarity switching ear in H_0^* of length h_0. If $V(\sigma)$ contains both α-cut faces a_1, a_2, then the lift $\hat{\nu}$ of a subpath $\nu \subseteq \sigma$ between a_1 and a_2 is a δ-polarity switching ear, and hence (c) $h \le \mathit{len}(\hat{\nu}) \le h_0$. Otherwise, the lift $\hat{\sigma}$ in H^* is an $(\alpha \cup \delta)$-ear which means that, for some subpath $\pi \subseteq \alpha$ of length at most $\frac{1}{2}\,\mathit{len}(\alpha)$ (possibly empty), $\hat{\sigma} \cup \pi$ is a δ-ear. Since σ is δ_0-polarity switching in H_0^*, and the δ_0-polarities are inherited from those of δ in H^* by (a', 7.9), we conclude that $\hat{\sigma} \cup \pi$ is a δ-polarity switching ear. Therefore, $h \le \mathit{len}(\hat{\sigma} \cup \pi) \le h_0 + \frac{1}{2}\,\mathit{len}(\alpha)$ as claimed. ∎

Proof of Theorem 4.2. It remains to finish proof of the following claim:

(5.5) Let H_1 be a graph embedded in an orientable surface of genus g, and h_1 be an integer. Assume that $H = H_1$ either (i) satisfies the conditions (a),(b),(c) of Lemma 5.4 with $h = h_1$ and some δ, or (ii) $g \ge 1$ and $\mathit{stretch}(H_1) \ge h_1 \cdot \mathit{ewd}(H_1)$. Then there exists a subgraph $T_0 \subseteq H_1$ that embeds in the torus with $\mathit{ewd}(T_0) \ge 2^{1-g}\,\mathit{ewd}(H_1)$, and T_0 has $\mathit{stretch}(T_0) \ge 2^{2-2g} \cdot h_1 \cdot \mathit{ewd}(H_1)$.

We first consider that the assumption (i) is true, and apply Lemma 5.4 to H_1 and h_1. Notice that the condition (b) implies that $g \ge 1$, and hence no explicit assumption on g is needed in this part. If the exceptional case $\mathit{len}(\beta) \ge h_1$ happens, then $\mathit{stretch}(H_1) = \mathit{len}(\alpha) \cdot \mathit{len}(\beta) \ge \mathit{ewd}(H_1) \cdot h_1$, and hence the assumption (ii) is also true. See below.

Otherwise, the new embedded graph $H_0 = H_1/\alpha$ has genus $g-1$, and H_0 satisfies the conditions of Lemma 5.4 (implying $g-1 \ge 1$ again) with

$$h_0 = h \ge h_1 - \frac{1}{2}\,\mathit{len}(\alpha) \ge h_1 - \frac{1}{2}\,\mathit{len}(\beta) \ge \frac{1}{2} h_1.$$

By inductive application of (5.5) to $H_1 = H_0$ and $h_1 = h_0$ in genus $g-1$, we get a toroidal graph $T_0 \subseteq H_0 \subseteq H_1$ that suits our needs. Using Lemma 5.3, we have

$$\mathit{ewd}(T_0) \ge 2^{1-(g-1)} \cdot \mathit{ewd}(H_0) = 2^{1-g} \cdot 2\,\mathit{ewd}(H_0) \ge 2^{1-g} \cdot \mathit{ewd}(H_1),$$

and similarly,

$$\begin{aligned} stretch(T_0) &\geq 2^{2-2(g-1)} \cdot h_0 \cdot ewd(H_0) = 2^{2-2g} \cdot 2h_0 \cdot 2ewd(H_0) \\ &\geq 2^{2-2g} \cdot h_1 \cdot ewd(H_1). \end{aligned}$$

Secondly, we consider that the assumption (ii) is true. If $g = 1$ (the base case), then we are done. Otherwise, we define $H_0 = H_1/\varrho$ as the embedded graph of genus $g-1$ obtained from H_1 by cutting along some nonseparating dual cycle ϱ of length $ewd(H_1)$. The conclusions then straightforwardly follow from Lemmas 2.4 and 5.2, and the inductive assumption (ii) for H_0 and $\frac{1}{2}h_1$. ∎