

Crossing number of toroidal graphs

János Pach* and Géza Tóth†

Rényi Institute, Hungarian Academy of Sciences

Abstract

It is shown that if a graph of n vertices can be drawn on the torus without edge crossings and the maximum degree of its vertices is at most d , then its planar crossing number cannot exceed cdn , where c is a constant. This bound, conjectured by Brass, cannot be improved, apart from the value of the constant. We strengthen and generalize this result to the case when the graph has a crossing-free drawing on an orientable surface of higher genus and there is no restriction on the degrees of the vertices.

1 Introduction

Let S_g be the compact orientable surface with no boundary, of genus g . Given a simple graph G , a *drawing* of G on S_g is a representation of G such that the vertices of G are represented by points of S_g and the edges are represented by simple (i.e., non-selfintersecting) continuous arcs in S_g , connecting the corresponding point pairs and not passing through any other vertex. The *crossing number* of G on S_g , $\text{cr}_g(G)$, is defined as the minimum number of edge crossings over all drawings of G in S_g . For $\text{cr}_0(G)$, the “usual” planar crossing number, we simply write $\text{cr}(G)$.

Let G be a graph of n vertices and e edges, and suppose that it can be drawn on the torus without crossing, that is, G satisfies $\text{cr}_1(G) = 0$. How large can $\text{cr}(G)$ be? Clearly, we have $\text{cr}(G) < \binom{e}{2}$, and this order of magnitude can be attained, as shown by the following example. Take five vertices and connect any pair of them by $\frac{e}{20}$ vertex-disjoint paths of lengths two. In any drawing of this graph in the plane, every subdivision of K_5 gives rise to a crossing. Therefore, the number of crossings must be at least $\frac{e^2}{400}$.

Peter Brass suggested that this estimate can be substantially improved if we impose an upper bound on the degree of the vertices. More precisely, we have

*Supported by NSF grant CCR-00-98246 and grants from NSA, PSC-CUNY, Hungarian Research Foundation, and BSF

†Supported by OTKA-T-038397 and OTKA-T-046246.

Theorem 1. *Let G be a graph of n vertices with maximum degree d , and suppose that G has a crossing-free drawing on the torus. Then we have $\text{cr}(G) \leq cdn$, where c is a constant.*

For $d \geq 3$, the bound in Theorem 1 cannot be improved, apart from the value of the constant c . Consider the following example. Let $d \geq 4$, $G = C_k \times C_k$, where $k = \sqrt{n/d}$ is a large integer and C_k denotes a cycle of length k . Obviously, this graph can be drawn on the torus without crossings. On the other hand, by a result of Salazar and Ugalde [?], its planar crossing number is larger than $(\frac{4}{5} - \varepsilon)k^2$, for any $\varepsilon > 0$, provided that k is large enough. Substitute every edge e of G by $\lfloor \frac{d}{4} \rfloor$ new vertices, each connected to both endpoints of e . The resulting graph G' has at most n vertices, each of degree at most d . It can be drawn on the torus with no crossing, and its planar crossing number is at least

$$\left(\frac{4}{5} - \varepsilon\right) k^2 \times \left\lfloor \frac{d}{4} \right\rfloor^2 > \frac{1}{100} nd.$$

To see this, it is enough to observe that there is an optimal drawing of G' in the plane with the property that any two paths of length two connecting the same pair of vertices cross precisely the same edges. The same construction can be slightly modified to show that $\text{cr}(G)$ can also grow linearly in n if the maximum degree d is equal to three.

Theorem 1 can be generalized as follows.

Theorem 2. *Let G be a graph of n vertices of maximum degree d that has a crossing-free drawing on S_g , the orientable surface of genus g . Then we have $\text{cr}(G) \leq c_{d,g}n$, where $c_{d,g}$ is a constant depending on d and g .*

We can drop the condition on the maximum degree and obtain an even more general statement.

Theorem 3. *Let G be a graph of n vertices with degrees d_1, d_2, \dots, d_n , and suppose that G has a crossing-free drawing on S_g . Then we have*

$$\text{cr}(G) \leq c_g \sum_{i=1}^n d_i^2,$$

where c_g is a constant depending on g .

To simplify the presentation and to emphasize the main idea of the proof, in Section 2 first we settle the simplest (planar) case (Theorem 1). In Section 3, we reduce Theorem 3 to a similar upper bound on the crossing number of G in S_{g-1} (Theorem 3.1). This latter result is established in Section 4.

2 The simplest case: Proof of Theorem 1

We can assume that $d \geq 3$. It is sufficient to prove that $\text{cr}(G) \leq cd(n-1)$ holds for any *two-connected* graph G satisfying the conditions. Indeed, if G is disconnected or has a cut vertex, then it can be obtained as the union of two graphs G_1 and G_2 with n_1 and n_2 vertices that have at most one vertex in common, so that we have $n_1 + n_2 = n$ or $n + 1$. Arguing for G_1 and G_2 separately, we obtain by induction that

$$\text{cr}(G) = \text{cr}(G_1) + \text{cr}(G_2) \leq cd(n_1 - 1) + cd(n_2 - 1) \leq cd(n - 1),$$

as required.

Let G be a two-connected graph with maximum degree d and $\text{cr}_1(G) = 0$. Fix a crossing-free drawing of G on the torus. We can assume that the boundary of each face is connected. Indeed, if one of the faces contains a cycle not contractible within the face, then cutting the torus along this cycle we do not damage any edge of G . Therefore, G is a planar graph and there is nothing to prove.

If our drawing is not a triangulation, then by adding $O(n)$ extra vertices and edges we can turn it into one so that the maximum degree of the vertices increases by at most a factor of three. We have to apply the following easy observation.

Lemma 2.1. *Let G be a two-connected graph with n vertices of degree at most d ($d \geq 3$). Suppose that G has a crossing-free drawing on the orientable surface of genus g such that the boundary of each face is connected. Any such drawing can be extended to a triangulation of the surface with at most $19n + 36(g - 1)$ vertices of maximum degree at most $3d$.*

Proof. First consider a cycle $f = x_1x_2 \dots x_{n(f)}$ bounding a single face in the drawing of G . Note that some vertices $x_i \in V(G)$ and even some edges may appear along this cycle several times. Take a simple closed curve $\gamma_0 = p_1p_2 \dots p_{n(f)}$ inside the face, running very close to f and passing through the (new) points p_i in this cyclic order. In the ring between f and γ_0 , connect each vertex x_i to p_i and p_{i+1} (where $p_{n(f)+1} := p_1$).

Divide γ_0 into $m_0 := \lceil \frac{n(f)}{d-1} \rceil$ connected pieces, each consisting of at most d vertices, such that the last vertex of each piece π_i is the first vertex of π_{i+1} , where $1 \leq i \leq m_0$ and $\pi_{m_0+1} := \pi_1$. Place a simple closed curve $\gamma_1 = q_1q_2 \dots q_{m_0}$ in the interior of γ_0 . In the ring between γ_0 and γ_1 , connect each q_i to all points in π_i . (If $m_0 = 1$ or 2 , then γ_1 degenerates into a point or a single edge.) If γ_1 has more than three vertices, repeat the same procedure for γ_1 in the place of γ_0 , and continue as long as the interior of the face is not completely triangulated. We added

$$n(f) + m_0 + m_1 + \dots < n(f) + n(f) + \frac{n(f)}{2} + \frac{n(f)}{4} + \dots < 3n(f)$$

new vertices, and their maximum degree is at most $d + 4$. The degree of every original vertex of f increased by at most twice the number of times it appeared in f .

If we triangulate every face of G in the above manner, the resulting drawing G' defines a triangulation of the surface with fewer than $n + \sum_f 3n(f) \leq n + 6|E(G)|$ vertices, each of degree at most $d' := 3d$. By Euler's formula, we have $n + 6|E(G)| \leq n + 18(n - 2 + 2g)$, as required. \square

In the sequel, slightly abusing the notation, we write G for the triangulation G' and d for its maximum degree d' .

If G has no *noncontractible* cycle, i.e., no cycle represented on the torus by a closed curve not contractible to a point, then we are done, because G is a planar drawing so that $\text{cr}(G) = 0$. Otherwise, choose a noncontractible cycle C with the minimum number of vertices, fix an orientation of C , and let $k := |V(C)|$. Let E_l (and E_r) denote the set of edges not belonging to C that are incident to at least one vertex of C and in a small neighborhood of this vertex lie on the left-hand side (respectively right-hand side) of C . Note that the sets E_l and E_r are disjoint, but this fact is not necessary for the proof.

Replace C by two copies, C_r and C_l , lying on its right-hand side and left-hand side. Connect each edge of E_r (respectively E_l) to the corresponding vertex of C_r (respectively C_l). Cut the torus along C , and attach a disk to each side of the cut.

The resulting spherical (planar) drawing G_1 represents a graph, slightly different from G . To transform it into a drawing of G , we have to remove C_l and (re)connect the edges of E_l to the corresponding vertices of C_r . In what follows, we describe how to do this without creating too many crossings.

Let \hat{G}_1 denote the *dual* graph of G_1 , that is, place a vertex of \hat{G}_1 in each face of G_1 , and for any $e \in E(G_1)$ connect the two vertices assigned to the faces meeting at e by an edge $\hat{e} \in E(\hat{G}_1)$. Let r and l denote the vertices of \hat{G}_1 lying in the faces bounded by C_r and C_l .

Lemma 2.2. *In \hat{G}_1 , there are k vertex-disjoint paths between the vertices r and l .*

Proof. By Menger's theorem, the maximum number p of (internally) vertex-disjoint paths connecting r and l in \hat{G}_1 is equal to the minimum number of vertices whose deletion separates r from l . Choose p such separating vertices, and denote the corresponding triangular faces of G by f_1, \dots, f_p . The interior of the union of these faces must contain a noncontractible closed curve that does not pass through any vertex of G . Let δ be such a curve whose number of intersection points with the edges of G is minimum. Choose an orientation of δ . Let e_1, \dots, e_q denote the circular sequence of edges of G intersected by δ . By the minimality of δ , we have $q \leq p$, because the interior of each triangle f_i contains at most one maximal connected piece of δ . Let v_i be the right endpoint of e_i with respect to the orientation of δ . Notice that v_i is adjacent to or identical with v_{i+1} , for every $1 \leq i \leq q$ (where $v_{q+1} := v_1$). Therefore, the circular sequence of vertices v_1, \dots, v_q induces a cycle in G that can be continuously deformed to δ . Thus, we have a noncontractible cycle of length




Figure 1: C is the shortest noncontractible cycle

$q \leq p$ in G , which implies that k , the length of the shortest such cycle, is at most p , as required. \square

By Lemma 2.1, the graph \hat{G} has at most $2|V(G)| \leq 38n$ vertices. According to Lemma 2.2, there is a path connecting r and l in \hat{G} with fewer than $\frac{38n}{k}$ internal vertices. The corresponding faces of G_1 form a “corridor” B between C_r and C_l . Delete now the vertices of C_l from G_1 . Pull every edge in E_l through B , and connect each of them to the corresponding vertex of C_r . See Figures 1 and 2. Notice that during this procedure one can avoid creating any crossing between edges belonging to E_l .

We give an upper bound on the number of crossings in the resulting planar drawing of G . Using that $|C| = k$ and $|E_l| \leq dk$, we can conclude that by pulling each edge through the corridor B , we create at most $\frac{38n}{k}$ crossings per edge. Thus, the total number of crossings cannot exceed $dk \cdot \frac{38n}{k} = 38dn$, which completes the proof of Theorem 1. \square

3 Reducing Theorem 3 to Theorem 3.1

Given a graph G , let $n(G)$ and $\sigma(G)$ denote the number of vertices of G and the sum of the squares of their degrees.

Theorem 3 provides an upper bound for the crossing number of a graph G that can be drawn on S_g without crossing. Next we show that this bound can be deduced by repeated application of the following result. In each step, we reduce the genus of the surface by one.

Theorem 3.1. *Let G be a two-connected graph with $\text{cr}_g(G) = 0$. Then we have $\text{cr}_{g-1}(G) \leq c_g^* \sigma(G)$, for some constant $c_g^* \geq 1$.*

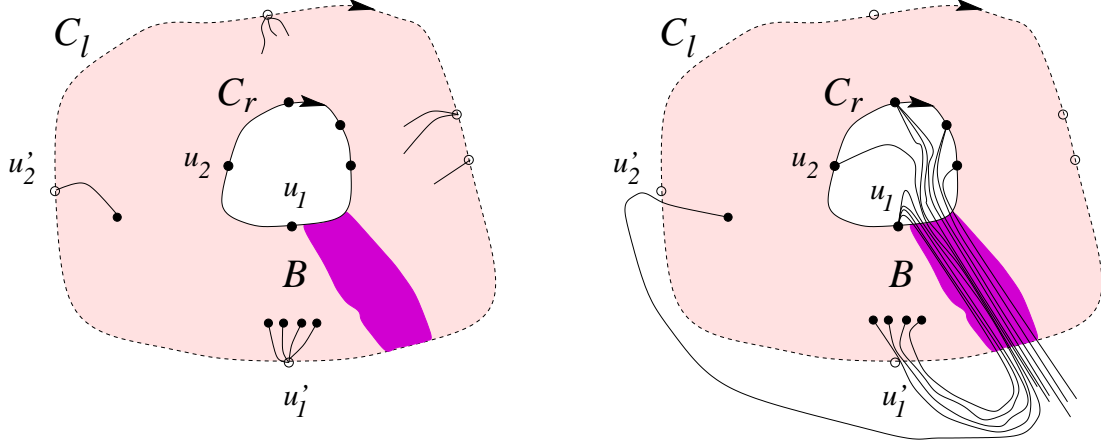


Figure 2: Pulling the edges in E_l through the corridor B .

Proof of Theorem 3 using Theorem 3.1. As in the proof of Theorem 1, we can assume that G is two-connected. Consider a crossing-free drawing of $G_0 := G$ on S_g . According to Theorem 3.1, G_0 can be drawn on S_{g-1} with at most $c\sigma(G)$ crossings. Place a new vertex at each crossing, and apply Theorem 3.1 to the resulting graph G_1 . Proceeding like this, we obtain a series of graphs G_2, G_3, \dots, G_g , drawn on $S_{g-2}, S_{g-3}, \dots, S_0$, respectively, with no crossing.

We claim that for any i , $0 \leq i \leq g$,

$$\sigma(G_i) \leq (17)^i \left(\prod_{g-i < j \leq g} c_j^* \right) \sigma(G)$$

holds. This is obviously true for $i = 0$. Let $0 < i \leq g$, and assume that the claim has already been verified for $i - 1$. Notice that, apart from the original vertices of G_{i-1} , every other vertex of G_i has degree four. Thus, applying Theorem 3.1 to the graph G_{i-1} that had a crossing-free drawing on S_{g-i+1} , we obtain

$$\begin{aligned} \sigma(G_i) &\leq \sigma(G_{i-1}) + 16\text{cr}_{g-i}(G_{i-1}) \leq \sigma(G_{i-1}) + 16c_{g-i}^* \sigma(G_{i-1}) \\ &\leq (1 + 16c_{g-i}^*)(17)^{i-1} \left(\prod_{g-i+1 < j \leq g} c_j^* \right) \sigma(G) \leq (17)^i \left(\prod_{g-i < j \leq g} c_j^* \right) \sigma(G), \end{aligned}$$

which proves the claim.

It follows from the construction that G_g is a planar graph, and we have

$$n(G_g) - n(G) < \sigma(G_g) \leq 17^g \left(\prod_{j=1}^g c_j^* \right) \sigma(G).$$

Replacing the $n(G_g) - n(G)$ “new” vertices of G_g by proper crossings, we obtain a drawing of G in the plane with at most $17^g \left(\prod_{j=1}^g c_j^* \right) \sigma(G)$ crossings. This completes the proof of Theorem 3. \square .

4 Reducing the genus by one: Proof of Theorem 3.1

It remains to prove Theorem 3.1.

All noncrossing closed curves C on S_g belong to one of the following three categories:

1. C is *contractible* (to a point);
2. C is *noncontractible* and *twosided*, i.e., it separates S_g into two connected components;
3. C is *noncontractible* and *onesided*.

Let us cut the surface S_g along C , and attach a disk along each side of the cut. If C is contractible, we obtain two surfaces: one homeomorphic to S_g and the other homeomorphic to the sphere S_0 . If C is noncontractible and twosided, then we obtain two surfaces homeomorphic to S_a and S_b , for some $a, b > 0$ with $a + b = g$. Finally, if C is noncontractible and onesided, then we get only *one* surface, S_{g-1} [?].

First we need an auxiliary statement, interesting on its own right.

Theorem 4.1. *Let G be a graph with a crossing-free drawing on S_g . If G has no noncontractible onesided cycle, then G is a planar graph.*

Proof. We follow the approach of Cairns and Nikolayevsky [?], developed to handle a similar problem on generalized thrackles. Let S be a very small closed neighborhood of the union of all edges of the drawing of G on S_g . Then S is a compact connected surface whose boundary consists of a finite number of closed curves. Attaching a disk to each of these closed curves, we obtain a surface S' with no boundary. We show that S' is a sphere. To verify this claim, consider two closed curves, α' and β' , on S' . They can be continuously deformed into closed walks, α_1 and β_1 , along the edges of G . Let α and β be the corresponding closed walks along the edges of G in the original drawing on S_g . By the assumption, α divides S_g into two parts, therefore, β crosses α an even number of times. Since the original drawing of G on S_g was crossing-free, every crossing between α and β occurs at a vertex of G . Using the fact that in the new drawing of G on S' , the cyclic order

of the edges incident to a vertex is the same as the cyclic order of the corresponding edges in the original drawing, we can conclude that α_1 and β_1 cross an even number of times. It is not hard to argue that then the same was true for α' and β' . Thus, S' is a surface with no boundary in which any two closed curves cross an even number of times. This implies that S' is a sphere. Consequently, we have a crossing-free drawing of G on the sphere, that is, G is a planar graph. \square

Proof of Theorem 3.1. As in the previous section, let $\sigma(G)$ denote the the sum of the squared degrees of the vertices of G . A *grid* of size $k \times k$ is the cross product $P_k \times P_k$ of two paths of length k . The vertices of $P_k \times P_k$ with degrees less than four are said to form the *boundary* of the grid. The proof of Theorem 3.1 is based on the same idea as that of Theorem 1, but some important details have to be modified.

Suppose that G is a two-connected graph of n vertices, drawn on S_g without crossing. We can also assume that G has no crossing-free drawing on S_{g-1} , otherwise Theorem 3.1 is trivially true. In particular, it follows that every face of the drawing of G on S_g has a connected boundary.

Replace each vertex v of degree $d(v) > 4$ by a grid of size $d(v) \times d(v)$ and connect the edges incident to v to distinct vertices on the boundary of the grid, preserving their cyclic order. The resulting crossing-free drawing of G' has at most $\sigma(G)$ vertices, each of degree at most four. Every face has a connected boundary, so that we can apply Lemma 2.1 to turn G' into a triangulation G'' with at most $19\sigma(G) + 36(g - 1)$ vertices, each of degree at most twelve. Restricting G' and G'' to any grid substituting for a vertex in G , the only difference between them is that each quadrilateral face in G' is subdivided by one of its diagonals into two triangles in G'' . Color all edges along the boundaries of the grids *blue*, and all other grid and diagonal edges of G'' that lie in the interior of some grid *red*.

If G'' has no noncontractible onesided cycle, then we are done by Theorem 4.1. Otherwise, pick such a cycle C with the smallest number k of vertices. Without increasing its length too much, we can replace all red edges of C by blue edges. Indeed, the first vertex and the last vertex of any maximal red path in C must belong to the boundary of the same grid. Replace each such path by the shortest blue path connecting its first and last vertices along the boundary of the grid containing them. The resulting cycle C' is noncontractible, onesided, and its length is at most $2k$. It has no red edges, and we can assume without loss of generality that it does not intersect itself. Fix an orientation of C' .

Let E_l (and E_r) denote the set of edges not belonging to C' that are incident to at least one vertex of C' and in a small neighborhood of this vertex lie on the left-hand side (respectively right-hand side) of C' .

Replace C' by two copies, C_r and C_l , lying on its right-hand side and left-hand side. Connect each edge of E_r and E_l to the corresponding vertex of C_r and C_l . Cut S_g along C , and attach a disk to each side of the cut. The resulting surface is S_{g-1} , and it contains

a crossing-free drawing G_1 of a graph slightly different from G'' . To obtain a drawing of G'' from G_1 , we have to remove C_l and (re)connect the edges of E_l to the corresponding vertices of C_r without creating too many crossings.

Let \hat{G}_1 be the *dual* drawing of G_1 on S_{g-1} . Let r (respectively l) be the vertex of \hat{G}_1 lying in the face bounded by C_r (respectively C_l). Color *blue* each vertex of \hat{G}_1 that corresponds to a face lying inside a grid in G'' .

Repeating the proof of Lemma 2.2, we obtain

Lemma 4.2. *In \hat{G}_1 , there are k vertex-disjoint paths between the vertices r and l . \square*

The number of cells in G_1 is equal to the number of cells in G'' plus 2. Therefore, by Euler's formula, \hat{G}_1 has at most

$$2|V(G'')| + 4(g-1) + 2 \leq 2(19\sigma(G) + 36(g-1)) + 4(g-1) + 2 < 40(\sigma(G) + 2g)$$

vertices. Thus, by Lemma 4.2, there is a path $P(rl)$ between r and l , of length at most $40(\sigma(G) + 2g)/k$. Replacing all blue vertices of $P(rl)$ by others, we obtain a new path $P'(rl)$, not much longer than $P(rl)$. First observe that r and l , the two endpoints of $P(rl)$, are not blue. Let $uv_1v_2 \dots v_jv$ be an interval along P such that all v_i 's are blue ($1 \leq i \leq j$), but u and v are not. Then the faces corresponding to u and v must be adjacent to the boundary of some grid in G_1 . These two faces are connected by two chains of faces following the outer boundary of the grid. Replace v_1, v_2, \dots, v_j by the sequence of vertices corresponding to the shorter of these two chains. Since the degree of every vertex in G_1 is at most twelve, the length of this chain is at most $12j$. Repeating this procedure for each maximal blue interval of $P(rl)$, we obtain a new path $P'(rl)$, whose length is at most $480(\sigma(G) + 2g)/k$.

The corresponding faces of G_1 form a "corridor" B between C_r and C_l . Now delete r , l , and the vertices of C_l . In the same way as in the proof of Theorem 1, "pull" all edges of E_l through B , and connect them to the corresponding vertices of C_r . This step can be carried out without creating any crossing between the edges in E_l .

Now we count the number of crossings in the resulting drawing. Since $|C'| \leq 2k$, $|E_l| \leq 20k$. Pulling them through the corridor B , we create at most $480(\sigma(G) + 2g)/k$ crossings per edge, that is, altogether at most $X := 9600(\sigma(G) + 2g)$ crossings.

Deleting the extra vertices and edges from G_1 and collapsing each grid into a vertex, we obtain a drawing of G on S_{g-1} , in which the number of crossings cannot exceed X . This concludes the proof of Theorem 3.1. \square

Acknowledgement. We are very grateful to Zoltán Szabó (Princeton) for many valuable suggestions.

References

- [CN00] G. Cairns and Y. Nikolayevsky, Bounds for generalized thrackles, *Discrete Comput. Geom.* **23** (2000), 191–206.
- [MT01] B. Mohar and C. Thomassen: Graphs on surfaces, Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, 2001.
- [SU04] G. Salazar and E. Ugalde: An improved bound for the crossing number of $C_m \times C_n$: a self-contained proof using mostly combinatorial arguments, *Graphs Combin.* **20** (2004), 247–253.