

ARTICLE

Expected number of faces in a random embedding of any graph is at most linear

Jesse Campion Loth and Bojan Mohar[†]

Department of Mathematics, Simon Fraser University, Burnaby, BC, V5A 1S6, Canada

Corresponding author: Jesse Campion Loth; Email: jcampion@sfu.ca

(Received 16 February 2022; revised 10 January 2023; accepted 22 March 2023; first published online 24 April 2023)

Abstract

A random two-cell embedding of a given graph G is obtained by choosing a random local rotation around every vertex. We analyse the expected number of faces of such an embedding, which is equivalent to studying its average genus. In 1991, Stahl [5] proved that the expected number of faces in a random embedding of an arbitrary graph of order n is at most $n \log(n)$. While there are many families of graphs whose expected number of faces is $\Theta(n)$, none are known where the expected number would be super-linear. This led the authors of [1] to conjecture that there is a linear upper bound. In this note we confirm their conjecture by proving that for any n -vertex multigraph, the expected number of faces in a random two-cell embedding is at most $2n \log(2\mu)$, where μ is the maximum edge-multiplicity. This bound is best possible up to a constant factor.

Keywords: random embeddings; average genus; random topological graph theory

2020 MSC Codes: 05C10

1. Introduction

By an *embedding* of a graph G we mean a two-cell embedding of G in some orientable closed surface, and we consider two embeddings of G as being the same (or *equivalent*) if there is a homeomorphism between the corresponding surfaces that induces the identity isomorphism on G . Equivalent embeddings are considered the same, and when we speak about all embeddings of a graph, we mean all equivalence classes. It is well known that the equivalence classes of all embeddings are in bijective correspondence with *rotation systems*, which are defined as the collection of local rotations at the vertices of the graph, where by a *local rotation at v* we mean a cyclic ordering of the half-edges, (we call these darts), incident with v . We refer to [4] for more details.

It is a classical problem to study the minimum genus and maximum genus of a graph across all of its embeddings, see [2, 4, 8]. Considering the set of all two-cell embeddings of a graph is also a viable topic. An outline of various applications of graph embeddings can be found in [3]. In this work, we consider the problem of the *average genus* across all the different embeddings of a fixed graph. By Euler's formula, this is equivalent to studying the average number of faces across all embeddings of a graph. It will be more convenient to state our results in terms of the number of faces, as it better illustrates our bounds. Formally, we consider the uniform distribution across

Bojan Mohar was supported in part by the NSERC Discovery Grant R611450 (Canada) and by the Research Project N1-0218 of ARRS (Slovenia).

[†]On leave from IMFM, Jadranska 19, 1000 Ljubljana.

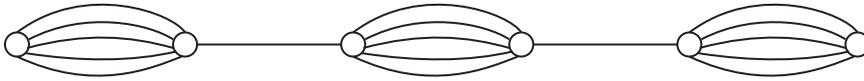


Figure 1. A chain of dipoles joined by cut edges gives a tight example for the main result of the paper.

all embeddings of a fixed graph using rotation systems, and study $\mathbb{E}[F]$ where F is the random variable denoting the number of faces in a random embedding of the graph.

This field of study was termed *random topological graph theory* by White [9]. Stahl [5] gave an upper bound on $\mathbb{E}[F]$ by proving that $\mathbb{E}[F] \leq n \log n$ for any simple graph on n vertices. It was shown in [1] that there are many examples of graphs with $\mathbb{E}[F] = \Theta(n)$: suppose a graph has maximum vertex degree d and a set \mathcal{C} of cycles, all of length at most ℓ . Then it is shown that

$$\mathbb{E}[F] \geq \frac{2|\mathcal{C}|}{(d - 1)^\ell}.$$

In particular, if the graph has small vertex-degrees and a set of $\Theta(n)$ short cycles, then we have at least linearly many expected faces. There are many examples of graphs of bounded degree and with linearly many short cycles. They all have $\mathbb{E}[F] = \Theta(n)$.

However, there are no known examples where $\mathbb{E}[F]$ was super-linear, and the following conjecture was proposed by Halasz, Masařík, Šámal, and the authors of this note.

Conjecture 1 ([1]). *For every simple graph of order n , the expected number of faces when selecting an orientable embedding of G uniformly at random is $O(n)$.*

In fact, a more general conjecture from [1] allowing for multiple edges of arbitrarily large multiplicity μ will be treated.

Conjecture 2 ([1]). *For every n -vertex multigraph G with maximum edge-multiplicity $\mu \geq 2$, the expected number of faces when selecting an orientable embedding of G uniformly at random is $O(n \log(\mu))$.*

We first give some examples which show that this more general conjectured bound is tight. We define a *dipole* as the graph with 2 vertices joined by μ edges. Stahl first showed [6] that for the dipole on μ edges, $\mathbb{E}[F] \leq H_{\mu-1} + 1$ where

$$H_\mu = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{\mu}$$

is the harmonic number. It was later shown [1] using Stanley’s generating function [7] that $\mathbb{E}[F] = H_{\mu-1} + \lceil \frac{\mu}{2} \rceil^{-1}$. Since $H_\mu \sim \log(\mu) + \gamma$, where γ is the Euler–Mascheroni constant, this gives a tight example for $n = 2$. In fact Stanley’s generating function may be used (see [1]) to show that for any graph with one central vertex incident to all of the μ edges, $\mathbb{E}[F] \leq H_\mu + \frac{3}{\mu}$.

A tight example, up to a constant factor, where n and μ may both tend to infinity is obtained by attaching a series of dipoles via cut edges as shown in Figure 1. More precisely, consider $n/2$ dipoles, each with μ parallel edges, joined by cut edges. Each separate dipole contains an average of at least H_μ faces, and joining these dipoles by cut-edges removes $n/2 - 1$ faces. Therefore $\mathbb{E}[F] \geq \frac{1}{2}n(H_\mu - 1)$.

Our main result confirms Conjectures 1 and 2. In fact we prove a more general bound in Theorem 7, which allows for different edge multiplicities. The following result which implies both conjectures, is a simple corollary of it.

Theorem 3. *Let G be a graph on n vertices with maximum edge-multiplicity μ , and let F be the random variable for the number of faces in a random embedding of G . Then we have:*

$$\mathbb{E}[F] \leq n(H_{2\mu} + 1).$$

In the case when G is a simple graph, we are able to show the better bound of $\frac{\pi^2}{6}n$ in Theorem 8. We are unaware of any examples of simple graphs which come close to the constant in this upper bound. A chain of triangles connected by cut edges gives an example of a graph for which $\mathbb{E}[F] = \frac{1}{3}n + 1$. We conjecture that this is the optimal bound.

Conjecture 4. *For any simple graph G of order n , $\mathbb{E}[F] \leq \frac{1}{3}n + 1$.*

2. Random embeddings

Fix a graph G on n vertices with $V(G) = [n] := \{1, 2, \dots, n\}$. Let μ_{ij} be the number of edges between vertices i and j , where we could have $i = j$. Let μ_i be the maximum of the multiplicities of edges incident with vertex i , where we count loops twice. That is, let $\mu_i = \max(2\mu_{i,i}, \max(\mu_{i,j} : j \neq i))$. Let E_i be the set of darts incident with vertex i , let $N(i)$ be the set of vertices adjacent to i , and let $d_i = |E_i|$ be the degree of the vertex. We start by outlining the random process leading to a random embedding of G that we will use to prove the main result of this paper.

Random Process A.

- (1) Start with the vertices $1, 2, \dots, n$, with d_i unlabelled darts coming out of vertex i for each i and fix a cyclic order of the darts around each vertex. Over the course of the random process we will pair darts to form the edges of G , decreasing the number of unlabelled darts at each vertex. When we pair a dart at vertex i and a dart at vertex j , thus forming the edge ij , we label the two darts and say that we have *processed* the edge ij . We write D_i for the number of unlabelled darts coming out of vertex i at some step in the process, and write μ_{ij} for the number of unprocessed edges between i and j at some step.
- (2) Repeat the following process: Pick one of Option A or Option B to use at this step. Option A: Pick an edge between i and j which hasn't yet been chosen to process. At vertices i and j choose one of the D_i and one of the D_j unlabelled darts (respectively) uniformly at random and then join them together to make an edge. Option B: Pick a dart at some vertex i . Choose one of the unprocessed edges ij incident with i uniformly at random. At vertex j , choose one of the D_j unlabelled darts uniformly at random and join it with the chosen dart at i to make an edge. In either option, we decrease each of D_i, D_j, μ_{ij} by one. Note that in the case of a loop ($i = j$), D_i is decreased by 2.
- (3) After all edges have been processed, the initial cyclic orders of darts around each vertex define a rotation system and hence an embedding of G .

By choosing the darts at step (2) of Random Process A in all possible ways, each embedding of G is obtained the same number of times, and each outcome has the same probability. This shows that the process always gives an embedding of G that is selected uniformly at random from the set of all embeddings. Let us also mention that the order in which the edges are processed, and whether we choose Option A or B, is not important. These can be chosen deterministically or randomly at each step.

Observation 5. *No matter whether we choose Option A or Option B at any step, and no matter which edge we choose when we use Option A or which dart we choose when we use Option B, at the end of the Random Process A, each embedding of G is obtained with the same probability. \square*

At each step during Process A, we have a *partial rotation* for which we can define (*partial*) *faces*. The *partial facial walk* around a partial face starts with an unlabelled dart, then it follows the already processed edges (maybe none) using the local rotation at vertices until we come to another unlabelled dart that is the end of this partial facial walk. See Figure 2, where the partial facial walks starting at c and a (respectively) are outlined with thick lines. Each dart is the beginning dart

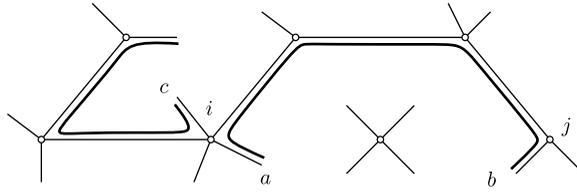


Figure 2. Partial rotation after processing five edges. Unlabelled darts are shown as short halfedges, whose local rotation is as given at the beginning, but it is not yet decided which of these will correspond to particular edges of G . Two of the partial facial walks are shown by a thick tracing line. If we are processing the edge ij and choose darts a and b to be paired, a partial face will be closed. If we choose c and b instead, the two partial faces will be merged into a larger partial face, which will not be closed.

of a partial facial walk and is also the ending dart of some partial facial walk. We call a partial walk which starts and ends on opposite sides of the same unlabelled dart a *bad partial facial walk*, and the corresponding face a *bad partial face*. This special type of partial face will be of special significance. In addition to the partial facial walks, we have facial walks that use only already processed edges. These will be unchanged for the rest of the process and will be facial walks of the final embedding of G . We say that such a completed facial walk is *closed* and is no longer considered to be a partial facial walk. Each closed facial walk became closed when we processed the last of its edges during the process. When we process the edge ij , there could be several pairs of a dart at i and a dart at j whose pairing will close a face. If there are k such pairs, we say that k faces can be closed while processing that edge.

We now discuss a special order in which we will process the edges while generating random embeddings. A particular processing of edges in Process A will be termed as the *greedy process*. This process is as follows:

1. If the partial rotation has some bad partial faces, then pick a dart in some bad partial face to process and carry out Option B. We prioritise picking a dart in a bad partial face which was also in a bad partial face at the previous step of Random Process A.
2. If the partial rotation has no bad partial faces, then take Option A and choose the next edge ij so that the number of faces that can be closed by processing this edge divided by μ_{ij} is minimum possible.

An important property that we have when using the greedy process is that whenever an edge ij is processed, there is a bound on the number of faces that can be closed by processing it.

Lemma 6. *During Random Process A, if the partial rotation at the start of the step has no bad partial faces then there is always an unprocessed edge ij , for which there are $2\mu_{ij}$ faces that can be closed by processing this edge. Processing this edge either closes zero, one or two of the $2\mu_{ij}$ possible closeable faces.*

Also, any partial rotation appearing at some step of the greedy version of Random Process A has at most 2 bad partial faces. Each bad partial face appears in at most two consecutive steps of the greedy process.

Proof. For the first claim, notice that the rotation system of the unlabelled darts and edges at each step is fixed. Each step of the process will join up two of the unlabelled darts into an edge. After $|E|$ steps, we will end up with an embedding of G . Recall that we defined a partial face as a face which has not yet been completed (closed) during the previous steps of the process. The walk along a partial face starts with an unlabelled dart at some vertex. It will then alternate along edges and vertices until it eventually reaches an unlabelled dart (possibly the same one we started with), with which the partial walk ends.

At the start of the random process we have $2|E|$ partial faces, where each partial face consists of two darts that are consecutive around the vertex in the local rotation. At each step we join together

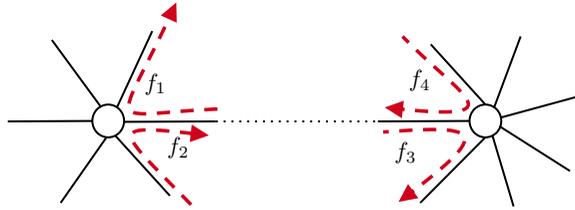


Figure 3. The situation when we have chosen two darts, and are replacing them with an edge. The partial facial walks f_1, f_4 will merge, and the partial facial walks f_2, f_3 will merge. This may also add one or two closed faces.

two darts to make an edge. We claim that this always reduces the number of partial faces by two, and possibly creates one or two closed faces. Indeed each of these darts we are joining to make an edge is the start and end of a partial face: write f_1, f_2 for the partial faces starting and ending respectively at one of the darts, and f_3, f_4 for the partial faces starting and ending respectively at the other dart as shown in Figure 3. Note that f_1 and f_3 start with different darts, so they cannot be equal. Similarly, we have $f_2 \neq f_4$. There are a couple of cases:

- f_1, f_2, f_3, f_4 are all distinct. Then joining the two darts into an edge joins f_1 and f_4 into a partial face, and f_2 and f_3 into a partial face.
- $f_1 = f_4$ and $f_2 \neq f_3$, then we close the partial face $f_1 = f_4$ into one completed closed face, and join f_2 and f_3 into a partial face. The case where $f_1 \neq f_4$ and $f_2 = f_3$ is the same.
- $f_1 = f_4$ and $f_2 = f_3$, then we close both of these partial faces into two closed faces.

This covers all the cases. Notice that in all of the above cases we reduce the number of partial faces by two, proving the claim.

This means that after k edges have been processed, there are $|E| - k$ remaining unprocessed edges and $2|E| - 2k$ partial faces. Each partial face starts with a dart at some vertex i , and ends with a dart at some vertex j , where we may have $i = j$. Each unprocessed edge is also associated to a pair of vertices ij , noting that $\sum \mu_{ij} = |E| - k$. By the pigeonhole principle there is at least one pair i, j (where we could have $i = j$) with $\mu_{ij} \geq 1$ unprocessed edges and at most $2\mu_{ij}$ partial faces associated to it. Hence we can always choose an edge such that processing it has at most $2\mu_{ij}$ different faces that could be closed by processing it.

For the second claim, we use induction. Initially we may assume G has no vertices of degree one, as these will not affect the final number of faces in an embedding of G , so we have no bad partial faces. Then, suppose that the partial rotation we have at the start of a step has at most two bad partial faces.

Case 1: It has no bad partial faces. Then since processing the edge affects at most two faces, we can add at most 2 bad partial faces.

Case 2: It has one or two bad partial faces. Then the greedy version of Random Process A will pair a dart in a bad partial face, removing it. At most one other partial face will be affected by adding this edge, so we can add at most one new bad partial face.

In either case, the number of bad partial faces in the new partial rotation is also at most two. Also, since there is at most one dart in a bad partial face which was not processed at this step, the greedy process must process this dart at the next step. Therefore this unlabelled dart appears in a bad partial face in at most two consecutive steps of the random process. \square

An analysis of the greedy version of Random Process A gives our main result. Recall that μ_i was defined as the maximum multiplicity of edges incident with vertex i , counting loops twice.

Theorem 7. $\mathbb{E}[F] \leq n + \sum_{i=1}^n H_{\mu_i}$.

Proof. At the start of each step in our random process, we have some partial rotation. The random process then fixes one edge into the embedding to obtain a new partial rotation. Since each step fixes one edge, there are $|E|$ total steps. If the partial rotation R_k appears at the start of step k for $k = 1, \dots, |E|$, then we say that the sequence $(R_1, R_2, \dots, R_{|E|})$ occurs at this run of Random Process A. Note that no edges have been fixed in R_1 , and all but one edge has been fixed in $R_{|E|}$. However there will only be one place to put the final edge, so this determines a unique embedding. We denote by $\mathbb{P}[R_1, R_2, \dots, R_{|E|}]$ the probability that we obtain this sequence. Similarly let $\mathbb{P}[R_k = R]$ denote the probability that we obtain R at the start of step k . Let X_k denote the random variable for the number of faces closed at step k of the random process.

Now suppose the partial rotation at the start of a step is R . The random process will then add an edge to this partial rotation, which will possibly close some faces. Let $X(R)$ denote the random variable for the number of faces closed at this step of the random process. Recall that $X(R) \in \{0, 1, 2\}$, as shown in Lemma 6.

The total probability formula gives that:

$$\mathbb{E}[X_k] = \sum_R \mathbb{P}[R_k = R] \mathbb{E}[X_k | R_k = R] = \sum_R \mathbb{P}[R_k = R] \mathbb{E}[X(R)]$$

where the sum runs over all possible partial rotations.

Using linearity of expectation yields:

$$\begin{aligned} \mathbb{E}[F] &= \sum_{k=1}^{|E|} \mathbb{E}[X_k] = \sum_{k=1}^{|E|} \sum_R \mathbb{P}[R_k = R] \mathbb{E}[X(R)] \\ &= \sum_{k=1}^{|E|} \sum_R \sum_{\substack{(R_1, R_2, \dots, R_{|E|}) \\ R=R_k}} \mathbb{P}[R_1, R_2, \dots, R_{|E|}] \mathbb{E}[X(R)]. \end{aligned}$$

Switching the order of summation, and then taking the maximum element in the sum, gives the following:

$$\begin{aligned} \mathbb{E}[F] &= \sum_{(R_1, R_2, \dots, R_{|E|})} \mathbb{P}[R_1, R_2, \dots, R_{|E|}] \sum_{k=1}^{|E|} \mathbb{E}[X(R_k)] \\ &\leq \max_{(R_1, R_2, \dots, R_{|E|})} \left\{ \sum_{k=1}^{|E|} \mathbb{E}[X(R_k)] \right\}. \end{aligned}$$

Therefore we may analyse each step of the random process separately, over any fixed possible sequence of partial rotations $(R_1, R_2, \dots, R_{|E|})$. Suppose that we are at the start of step k of the random process, and we have the partial rotation R_k . Further suppose that R_k has no bad partial faces, and that we have chosen an edge $e = ij$ to process using Option A. Let us first suppose that $i \neq j$. Recall that D_i and D_j are the number of unlabelled darts at vertices i and j at this step, respectively. So, there are $D_i D_j$ choices of places to place the edge across two darts at these vertices. However in the greedy version of Process A we choose an edge to process with only $2\mu_{ij}$ partial faces that could be closed. Each partial face is closed by only one choice out of the $D_i D_j$ total placements of the edge, hence the probability that we close a face is $\frac{1}{D_i D_j}$. By Lemma 6 at most two faces may be closed by the same choice of edge placement, but in any case we have $\mathbb{E}[X(R_k)] \leq \frac{2\mu_{ij}}{D_i D_j}$. Note that $\mu_{ij} \leq \min(D_i, D_j)$. Write $D_i = D_i(R_k), D_j = D_j(R_k)$ for the values of D_i, D_j at this step. Then at the next step we have $D_i(R_{k+1}) = D_i(R_k) - 1, D_j(R_{k+1}) = D_j(R_k) - 1$.

If R_k has at least one bad partial face, then recall that the greedy version of Random Process A will take a dart in a bad partial face, incident with some vertex i . It will then pick an unprocessed

edge ij incident with i uniformly at random, and a dart incident with j uniformly at random to pair the dart at i with. Observe that a face is closed at this step if and only if the dart we're pairing with is also in a bad partial face. Since there is at most one other bad partial face, there is at most one choice of dart to pair with which will close a face. Suppose that this dart in the other bad partial face is incident with vertex j' , if such a dart exists. Therefore the probability we make this choice, and hence close one face, is $\mathbb{E}[X(R_k)] \leq \frac{\mu_{ij'}}{D_i D_j}$. At the next step we reduce D_i and D_j by one. Note we don't reduce D_j by one, but at the next step if $j' \neq j$, we will process the dart in a bad partial face incident with vertex j' . Therefore at the following step we will reduce D_j by one.

The case when we are processing a loop from vertex i to itself is similar. If the partial rotation has no bad partial faces, then there are $\binom{D_i}{2}$ choices of placements for the edge, and there are at most $2\mu_{ii}$ partial faces which may be closed. Therefore by the same reasoning as in the previous case, we have $\mathbb{E}[X(R_k)] \leq \frac{2\mu_{ii}}{D_i(D_i-1)/2} = \frac{4\mu_{ii}}{D_i(D_i-1)}$. At the next step we have $D_i(R_{k+1}) = D_i(R_k) - 2$. A similar reasoning holds for the case when we process a dart incident with a bad partial face using Option B.

Fix some vertex i . Over the sequence of partial rotations $(R_1, R_2, \dots, R_{|E|})$, let R_{k_1}, \dots, R_{k_c} be the partial rotations for which an edge at vertex i is processed, where c is equal to the number of edges incident with vertex i , counting each loop only once. We have that $D_i(R_{k_1}) = d_i$. If we are in option A of the random process, and the first edge processed at vertex i was not a loop, then $D_i(R_{k_2}) = d_i - 1$. If it was a loop then $D_i(R_{k_2}) = d_i - 2$. If we are in option B of the random process, then we could have $D_i(R_{k_2}) = d_i - 1$ or $D_i(R_{k_2}) = d_i$. However if $D_i(R_{k_2}) = d_i$ then necessarily $D_i(R_{k_3}) = d_i - 1$. Therefore the values of D_i decrease at each of these steps until $D_i(R_{k_c}) = 1$ or 2 , then these remaining darts are processing at this step.

Initially we have that $D_i = d_i$ and $D_j = d_j$. For $i \neq j$, $\mu_{ij} \leq \min(\mu_i, \mu_j)$, and $2\mu_{ii} \leq \mu_i$. When we process a non-loop edge using Option A, $\mathbb{E}[X(R_k)]$ is bounded by some $\frac{2\mu_{ij}}{D_i D_j}$, and D_i, D_j and μ_{ij} all decrease by one. When we process a loop edge using Option A, $\mathbb{E}[X(R_k)]$ is bounded by $\frac{4\mu_{ii}}{D_i(D_i-1)}$, D_i decreases by two and μ_{ii} decreases by one. When we process a dart using Option B, $\mathbb{E}[X(R_k)]$ is bounded by $\frac{\mu_{ij'}}{D_i D_j}$ for some j' . For some j , D_i, D_j decrease by one and μ_{ij} decreases by one. Also when $i \neq j$, μ_{ij} is the number of unprocessed edges between vertices i and j , so by definition we have that $\mu_{ij} \leq \min(D_i, D_j, \mu_i, \mu_j)$. Similarly, we have that $2\mu_{ii} \leq \min(D_i, \mu_i)$.

If at step k we process an edge $e = ij$, then write $D_i(e) = D_i(R_k), D_j(e) = D_j(R_k)$. For this sequence of partial rotations $(R_1, R_2, \dots, R_{|E|})$, write E_A, E_B for the set of edges processed under Options A and B respectively. Then we have:

$$\begin{aligned} \sum_{k=1}^{|E|} \mathbb{E}[X(R_k)] &\leq \sum_{\substack{(i,j) \in E_A \\ i \neq j}} \frac{2 \min(D_i(e), D_j(e), \mu_i, \mu_j)}{D_i(e)D_j(e)} + \sum_{\substack{(i,j) \in E_A \\ i=j}} \frac{2 \min(D_i(e), \mu_i)}{D_i(e)(D_i(e) - 1)} \\ &+ \sum_{(i,j) \in E_B} \frac{\min(D_i(e), D_j(e), \mu_i, \mu_j)}{D_i(e)D_j(e)}. \end{aligned}$$

We first note that for any $a, b > 0$:

$$\begin{aligned} \frac{2 \min(a, b, \mu_i, \mu_j)}{ab} &\leq \frac{\min(a, b, \mu_i, \mu_j)}{a^2} + \frac{\min(a, b, \mu_i, \mu_j)}{b^2} \\ &\leq \frac{\min(a, \mu_i)}{a^2} + \frac{\min(b, \mu_j)}{b^2}. \end{aligned}$$

This means we can rewrite the expectation as:

$$\begin{aligned} \sum_{k=1}^{|E|} \mathbb{E}[X(R_k)] &\leq \sum_{\substack{(i,j) \in E_A \\ i \neq j}} \left(\frac{\min(D_i(e), \mu_i)}{D_i(e)^2} + \frac{\min(D_j(e), \mu_j)}{D_j(e)^2} \right) \\ &+ \sum_{\substack{(i,j) \in E_A \\ i=j}} \left(\frac{\min(D_i(e), \mu_i)}{D_i(e)^2} + \frac{\min(D_i(e), \mu_i)}{(D_i(e) - 1)^2} \right) \\ &+ \frac{1}{2} \sum_{(i,j) \in E_B} \left(\frac{\min(D_i(e), \mu_i)}{D_i(e)^2} + \frac{\min(D_j(e), \mu_j)}{D_j(e)^2} \right). \end{aligned}$$

Now fix some vertex i . For each non-loop edge $e \in E_A$ incident with i we obtain a term of $\frac{\min(D_i(e), \mu_i)}{D_i(e)^2}$ in the preceding sum. For each loop in E_A incident with vertex i , we obtain a term of $\frac{\min(D_i(e), \mu_i)}{D_i(e)^2} + \frac{\min(D_i(e), \mu_i)}{(D_i(e) - 1)^2}$. Recall that when we process an edge incident with i , in Option A, we reduce D_i by one if the edge is not a loop, and by two if the edge is a loop. When we process an edge in E_B incident with i , we obtain a term of $\frac{\min(D_i(e), \mu_i)}{2D_i(e)^2}$. If we don't decrease D_i at this step, then we do reduce D_i at the following step and obtain another term of $\frac{\min(D_i(e), \mu_i)}{2D_i(e)^2}$. This means that in the whole sum, D_i appears (as some $D_i = D_i(e)$) for each of the values in $\{1, 2, \dots, d_i\}$ in at most one term of the form $\frac{\min(D_i(e), \mu_i)}{D_i(e)^2}$. Also recall that by the preceding arguments, it is enough to bound $\sum_{k=1}^{|E|} \mathbb{E}[X(R_k)]$ for an arbitrary $(R_1, R_2, \dots, R_{|E|})$ in order to bound $\mathbb{E}[F]$. Therefore we may bound the expectation as:

$$\mathbb{E}[F] \leq \sum_{i=1}^n \sum_{t=1}^{d_i} \frac{\min(t, \mu_i)}{t^2} < \sum_{i=1}^n \sum_{t \geq 1} \frac{\min(t, \mu_i)}{t^2}.$$

For each vertex i we obtain a sum of terms of the form:

$$\sum_{t \geq 1} \frac{\min(t, \mu_i)}{t^2} \leq \sum_{t=1}^{\mu_i} \frac{1}{t} + \sum_{t \geq \mu_i+1} \frac{\mu_i}{t^2} = H_{\mu_i} + \mu_i \sum_{t \geq \mu_i+1} \frac{1}{t^2}.$$

Note that we have:

$$\mu_i \sum_{t \geq \mu_i+1} \frac{1}{t^2} < \mu_i \int_{\mu_i}^{\infty} x^{-2} dx = 1.$$

Therefore the total contribution from all the vertices is bounded by:

$$\mathbb{E}[F] < \sum_{i=1}^n (H_{\mu_i} + 1) = \sum_{i=1}^n H_{\mu_i} + n.$$

This gives the required bound. □

If the graph is simple, then every μ_i is equal to 1. In this case we can obtain a slightly better upper bound.

Theorem 8. *If G is a simple graph of order n , then*

$$\mathbb{E}[F] < \frac{\pi^2}{6} n.$$

Proof. From the proof of the general case, we have the following sum as an upper bound:

$$\mathbb{E}[F] \leq \sum_{e=ij} \frac{2}{ab} \leq \sum_{e=ij} \left(\frac{1}{a^2} + \frac{1}{b^2} \right),$$

where the pairs $\{a, b\}$ of the summands exhaust the multiset

$$\Delta = \{1, 2, \dots, d_1, 1, 2, \dots, d_2, \dots, 1, 2, \dots, d_n\}.$$

Each term of $1/a^2$ appears at most n times for each $a \geq 1$, so we obtain the upper bound:

$$\mathbb{E}[F] < n \sum_{a \geq 1} \frac{1}{a^2} = \frac{\pi^2}{6} n.$$

□

Acknowledgements

The authors would like to thank Kevin Halasz, Tomáš Masařík and Robert Šámal for helpful discussions on the topic.

References

- [1] Loth, J. C., Halasz, K., Masařík, T., Mohar, B. and Šámal, R. (2022) Random 2-cell embeddings of multistars. *Proc. Am. Math. Soc.* **150**(9) 3699–3713.
- [2] Gross, J. L. and Tucker, T. W. (1987) *Topological Graph Theory*. Wiley-Interscience, New York.
- [3] Lando, S. K., Zvonkin, A. K. and Zagier, D. B. (2004) *Graphs on Surfaces and their Applications*. Springer, Berlin, Vol. 75.
- [4] Mohar, B. and Thomassen, C. (2001) *Graphs on Surfaces*. Johns Hopkins University Press, Baltimore, MD, Johns Hopkins Studies in the Mathematical Sciences.
- [5] Stahl, S. (1991) An upper bound for the average number of regions. *J. Combin. Theory Ser. B* **52**(2) 219–221.
- [6] Stahl, S. (1995) On the average genus of the random graph. *J. Graph Theory* **20**(1) 1–18.
- [7] Stanley, R. P. (2011) Two enumerative results on cycles of permutations. *Eur. J. Combin.* **32**(6) 937–943.
- [8] White, A. T. (1973) *Graphs, Groups and Surfaces*. North-Holland Publishing Co.; American Elsevier Publishing Co., Inc., Amsterdam-London; New York, North-Holland Mathematics Studies, **8**.
- [9] White, A. T. (1994) An introduction to random topological graph theory. *Combin. Probab. Comput.* **3**(4) 545–555.