

Spanning Tree Bounds for Grid Graphs

Kristopher Tapp

Submitted: Jun 12, 2023; Accepted: Jan 5, 2024; Published: Jan 26, 2024

© The author. Released under the CC BY license (International 4.0).

Abstract

Among subgraphs with a fixed number of vertices of the regular square lattice, we prove inequalities that essentially say that those with smaller boundaries have larger numbers of spanning trees and vice-versa. As an application, we relate two commonly used measurements of the compactness of district maps.

Mathematics Subject Classifications: 05C05, 05C81, 05C90, 05C70, 91F10, 60J20, 60C05

1 Introduction

For a finite connected graph G , let $\tau(G)$ denote its number of spanning trees. The study of this measurement goes back to Kirchoff's Matrix-Tree Theorem, which equates it with the product of the non-zero eigenvalues of the Laplacian of G [11].

Let $\mathcal{L}(\mathbb{Z}^2)$ denote the regular square lattice, which has vertex set \mathbb{Z}^2 and edges between pairs of vertices of Euclidean distance 1. We are interested here in *grid graphs*, by which we mean finite connected subgraphs of $\mathcal{L}(\mathbb{Z}^2)$. The *bulk limit* of $\mathcal{L}(\mathbb{Z}^2)$ is known to equal $\frac{4C}{\pi}$, where C is Catalan's constant. This means that

$$\lim_{k \rightarrow \infty} \frac{\ln(\tau(G(k)))}{|V(G(k))|} = \frac{4C}{\pi} \approx 1.166243, \quad (1)$$

where $G(1) \subset G(2) \subset \dots$ is any nested sequence of grid graphs (satisfying certain weak hypotheses) whose union equals $\mathcal{L}(\mathbb{Z}^2)$. In this paper, $V(G)$ and $E(G)$ will denote the vertex set and edge set of a graph G . For this theorem and analogous results for other lattices, see [3], [17], [18], [19], [21] and references therein.

There is a good intuition that, among grid graphs with a fixed number of vertices (or even among more general classes of graphs), the ones with higher numbers of spanning trees should have smaller boundaries and vice-versa. Asymptotic evidence for this intuition is found in [10].

Further evidence comes from recent work on the mathematics of redistricting. For the reversible version of the spanning-tree-based MCMC algorithm by which ensembles of

Saint Joseph's University (ktapp@sju.edu).

maps are commonly generated, the stationary distribution is known to assign a probability to each map that is proportional to the map's *spanning tree score* (which means the product of the numbers of spanning trees of its districts); see [5], [2], [6]. Thus, the algorithm prefers maps whose districts have larger numbers of spanning trees. Data from large ensembles of maps indicate a strong negative correlation between a map's spanning tree score and its number of cut edges (which is a discrete measurement of the total size of the district boundaries) [5]. Thus, the algorithm seems to prefer maps whose districts have small boundaries, and one purpose of this paper is to more rigorously understand this behavior.

The following important result, due to Russell Lyons, says that the bulk limit is an upper bound:

Theorem 1 (Lyons). *If G is a grid graph, then*

$$\ln(\tau(G)) < \frac{4C}{\pi} \cdot |V(G)|.$$

For completeness, we'll include Lyons' unpublished proof of this theorem in the next section. In terms of the base

$$\mathfrak{b} = \exp(4C/\pi) \approx 3.2099,$$

Lyons' theorem can be re-phrased as:

$$\tau(G) < \mathfrak{b}^{|V(G)|}. \tag{2}$$

Our main result is related to this, and is easiest to state for the following natural class of grid graphs.

Definition 2. A grid graph G is called *simple* if it is comprised of all of the vertices and edges that are on and interior to a cycle α in $\mathcal{L}(\mathbb{Z}^2)$. In this case, the set of vertices of α is called the *boundary* of G , denoted ∂G . The *area* of G , denoted $\text{Area}(G)$, means the area of the interior of α , or equivalently the number of faces of G .

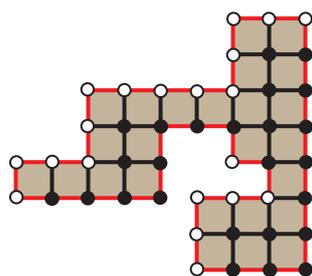


Figure 1: A simple grid graph. The white vertices lie in its *top-left boundary*.

Figure 1 illustrates a simple grid graph G . Its *bounding cycle* α , colored red, can be considered as a piecewise-linear path in \mathbb{R}^2 whose length equals $|\partial G|$. The white vertices comprise the *top-left boundary* of G , defined as:

Definition 3. The *top-left boundary* of a simple grid graph G , denoted $\hat{\partial}G$, is the set of all $v \in \partial G$ such that the face of $\mathcal{L}(\mathbb{Z}^2)$ whose bottom-right corner is v is not a face of G .

Our main result is the following theorem, which forces graphs with larger boundaries to have smaller numbers of spanning trees and vice-versa:

Theorem 4. *If G is a simple grid graph, then*

$$\mathfrak{b}^m \leq \tau(G) \leq 4^m,$$

where $m = \text{Area}(G) = |V(G)| - \frac{1}{2}|\partial G| - 1 = |V(G)| - |\hat{\partial}G|$.

The fact that $\text{Area}(G) = |V(G)| - \frac{1}{2}|\partial G| - 1$ follows from Pick's Theorem [15], while the fact that $\text{Area}(G) = |V(G)| - |\hat{\partial}G|$ will be proven in Section 4.

We'll show that the lower bound of Theorem 4 more generally makes sense and is true for *all* grid graphs, but that the upper bound is only true of *simple* grid graphs.

The example of a square with side length 1 (which contains four vertices) demonstrates that the upper bound of Theorem 4 is sharp. However when $\frac{|\partial G|}{|V(G)|}$ becomes small, this upper bound becomes worse than Theorem 1. To improve this situation, we give a much stronger upper bound in Section 6.

This paper is organized as follows. Section 2 contains Lyons' unpublished proof of Theorem 1. Section 3 explains the main idea of this paper with an illuminating example. Section 4 derives basic properties of the top-left boundary of a grid graph. Sections 5 and 6 respectively prove the lower and upper bound of Theorem 4 plus generalizations and improvements.

Finally in Section 7 we apply our main theorem to relate two different measurements of compactness that are commonly used in the mathematical redistricting literature: a map's cut edge count and its spanning tree score. Empirical evidence suggests a very strong negative correlation between these two measurements, and our results partially account for this correlation. Independent work by Procaccia and Tucker-Folz related these two measurements for general planar graphs [16]; in the case of grid graphs, our results are complimentary to theirs.

2 The bulk limit is an upper bound

We thank Russell Lyons for sharing with us the following proof. For brevity, in this section we assume knowledge of the vocabulary and results of [13] and [14].

Proof of Theorem 1. Choose a leftmost vertex x of G and a rightmost vertex y of G . For every integer $k \in \mathbb{Z}$, let G_k be a copy of G with corresponding vertices named x_k and y_k . For every integer $n > 0$, let H_n denote the connected graph formed from all of the copies G_k with $-n \leq k \leq n$, with the copies connected together by adding an edge between y_k and x_{k+1} for each $-n \leq k < n$. Notice that each H_n is isomorphic to a grid graph; that is, the construction can be embedded in $\mathcal{L}(\mathbb{Z}^2)$ as exemplified in Figure 2.

We have $\tau(H_n) = \tau(G)^{2n+1}$ and $|V(H_n)| = (2n + 1)|V(G)|$. Therefore,

$$\frac{\ln(\tau(G))}{|V(G)|} = \frac{\ln(\tau(H_n))}{|V(H_n)|}.$$

By [13, Theorem 3.2], the limit of the latter quantity is the tree entropy of the random rooted infinite graph H_∞ formed similarly from *all* copies G_k and rooted at a uniformly random vertex of G_0 . Clearly H_∞ is stochastically dominated by the entire square lattice $\mathcal{L}(\mathbb{Z}^2)$, whence the tree entropy of H_∞ is strictly less than that of $\mathcal{L}(\mathbb{Z}^2)$ by [14, Theorem 3.2]. The latter is $4C/\pi$, which proves the claimed upper bound. \square

Notice that this proof generalizes to yield the analogous result for any lattice (in any dimension) with a transitive group of translation symmetries.

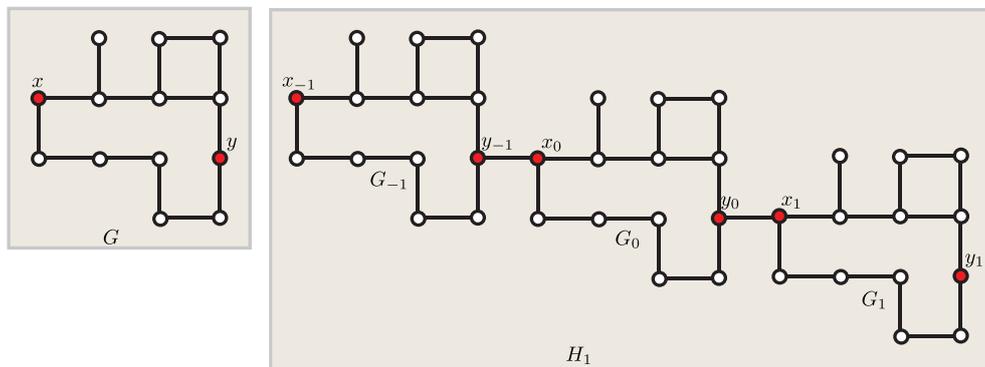


Figure 2: H_n is a grid graph.

3 Setup and example

In this section, let G be a grid graph. Our main technique involves building G by adding one vertex at a time in the words-on-a-page order (starting with the top row ordered left-to-right and ending with the bottom row ordered left-to-right), and studying the multiplicative factor by which the spanning tree count grows with each added vertex.

More precisely, let $v \in V(G)$. Let H'_v (respectively H_v) denote the subgraph of G induced by all vertices prior to v (respectively prior to and including v) with respect to the words-on-a-page ordering of $V(G)$. We will study the multiplicative growth factor:

$$m_v = \frac{\tau(H_v)}{\tau(H'_v)}.$$

Note that H_v and H'_v are not necessarily connected, even for the simple grid graph in Figure 1. To allow for the possibility of disconnected graphs, the meaning of τ here must be slightly generalized as follows. If H is a (possibly disconnected) graph, let $\mathcal{T}(H)$ denote the set of ways to choose one spanning tree from each of its connected components, and

let $\tau(H) = |\mathcal{T}(H)|$, which equals the product of the numbers of spanning trees on the connected components. Our convention here is that a component containing just a single vertex is counted as having one spanning tree. We additionally use the convention that $m_v = 1$ if v is the first vertex. With these definitions, we can recover $\tau(G)$ as:

$$\tau(G) = \prod_{v \in V(G)} m_v. \tag{3}$$

It is useful to regard $v \mapsto m_v$ as a real-valued function on $V(G)$, which we call the *multiplier function*. In fact, the primary technical goal of this paper is to understand its behavior on general grid graphs. For this, it is helpful to first gain intuition from examples.

Example 5. Figure 3 illustrates heatmaps for the multiplier function on two simple grid graphs. The left grid graph, which we call S , is the 12-by-12 square. The right grid graph, which we call D , is the diamond inside the 15-by-15 square. These examples were chosen to have similar numbers of vertices: $|V(S)| = 144$, while $|V(D)| = 141$. Each small square represents a vertex. The graphs' edges don't need to be displayed because adjacency is visually obvious, so the small squares are drawn large enough to bump into their neighbors forming a grid. The color of each small square represents the value of the multiplier function on the corresponding vertex.

For each vertex v of S or D , the underlying data shows that either $m_v = 1$ or $m_v \in (\mathfrak{b}, 4]$. The set of vertices with multiplier 1 (colored black with a white cross) is exactly the *top-left boundary*. The square's top-left boundary has 23 vertices, while the diamond's has 29. The square has more spanning trees: $\ln(\tau(S)) \approx 146.15$, $\ln(\tau(D)) \approx 136.19$. The square and the diamond both have the property that the size of the top-left boundary is one more than half the size of the boundary.

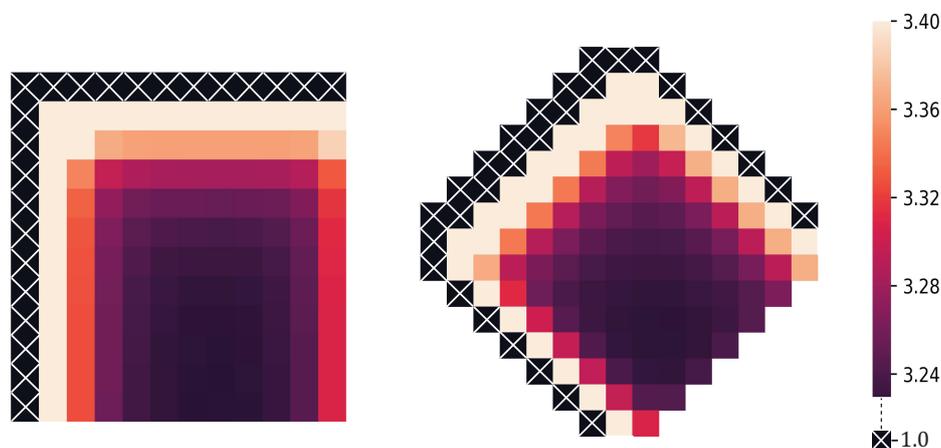


Figure 3: Heatmaps for the multiplier function on a square and a diamond.

The remainder of this paper will demonstrate that each key feature of the previous examples carries over to all grid graphs or at least all simple grid graphs.

4 The top-left boundary

In this section, we study the *top-left boundary* and prove that it behaves like the set of vertices colored black with a white cross in the examples of the previous section. We begin by generalizing Definition 3 to (not necessarily simple) grid graphs.

Definition 6. Let G be a grid graph. For each $v \in V(G)$, let \square_v denote the subgraph of $\mathcal{L}(\mathbb{Z}^2)$ comprised of the vertices and edges of the 1-by-1 square whose bottom-right corner is v . The *top-left boundary* of G is:

$$\hat{\partial}G = \{v \in V(G) \mid \square_v \text{ is not a subgraph of } G\}.$$

Lemma 7. *If G is simple and $v \in \hat{\partial}G$, then $m_v = 1$.*

Proof. Denote the coordinates of v as $v = (x, y)$. Denote the relevant neighbors of v as $a = (x, y + 1)$, $w = (x - 1, y + 1)$ and $b = (x - 1, y)$. The following three cases are straightforward:

- If $\overline{va} \notin E(G)$ and $\overline{vb} \notin E(G)$, then $\{v\}$ is a connected component of H_v , so $m_v = 1$.
- If $\overline{va} \in E(G)$ and $\overline{vb} \notin E(G)$, then $|\mathcal{T}(H'_v)| = |\mathcal{T}(H_v)|$ because “adding the edge \overline{va} ” is a bijection between $\mathcal{T}(H'_v)$ and $\mathcal{T}(H_v)$, so $m_v = 1$.
- If $\overline{va} \notin E(G)$ and $\overline{vb} \in E(G)$, then $m_v = 1$ by a similar argument.

Next assume that $\overline{va} \in E(G)$ and $\overline{vb} \in E(G)$, which is the only remaining case. We claim that a and b must lie in different connected components of H'_v . Indeed, if there were a path between a and b in H'_v , then adding \overline{va} and \overline{vb} to this path would yield a loop in H_v that encloses or contains w . But since G is simple, it contains all edges inside of any closed loop in it, so $\overline{aw}, \overline{bw} \in E(G)$, contradicting the hypothesis that $v \in \hat{\partial}G$.

In summary, a and b lie in different connected components of H'_v , but they are connected through v in H_v . Therefore, “adding \overline{va} and \overline{vb} ” is a bijection between $\mathcal{T}(H'_v)$ and $\mathcal{T}(H_v)$, so $m_v = 1$. \square

Figure 4 exhibits counterexamples to Lemma 7 when G is not simple. Each graph has the property that all of its vertices lie in its top-left boundary, but yet its red-colored vertices have multipliers larger than 1. In fact, the red vertex of the left graph has multiplier 16; we’ll soon see that this is much larger than the multiplier of *any* vertex of a simple graph.

The decision to consider the top-left boundary (rather than the top-right, bottom-left, or bottom-right) is somewhat arbitrary, but the size of $\hat{\partial}G$ is unaffected by this decision because the following proposition provides a canonical interpretation of $|\hat{\partial}G|$. The proposition also establishes the equivalence of the three expressions for m in Theorem 4.

Proposition 8. *If G is a simple grid graph, then $\hat{\partial}G \subset \partial G$ and*

$$\text{Area}(G) = |V(G)| - \frac{1}{2}|\partial G| - 1 = |V(G)| - |\hat{\partial}G|.$$

In particular, this implies that $|\partial G|$ is even and $|\hat{\partial}G| = \frac{1}{2}|\partial G| + 1$.

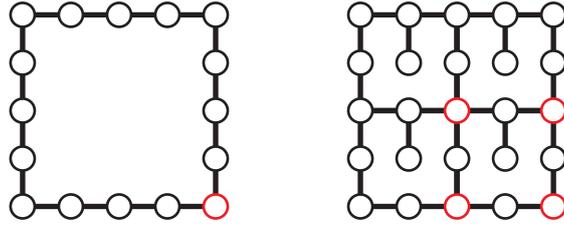


Figure 4: Counterexamples to Lemma 7 when G is not simple.

Proof. The claim that $\hat{\partial}G \subset \partial G$ is straightforward. As mentioned in the introduction, Pick's Theorem says that $\text{Area}(G) = |V(G)| - \frac{1}{2}|\partial G| - 1$. Moreover, $\text{Area}(G) = |V(G)| - |\hat{\partial}G|$ because the set of faces of G correspond one-to-one with $V(G) - \hat{\partial}G$ by matching each face with its bottom-right corner. \square

Alternatively, the previous proposition can be proven without using Pick's Theorem by establishing that $|\hat{\partial}G| = \frac{1}{2}|\partial G| + 1$ via a straightforward inductive argument.

5 A lower bound on $\tau(G)$

The goal of this section is to prove the lower bound in Theorem 4. In fact, we will prove the following generalization to (not necessarily simple) grid graphs:

Theorem 9. *If G is a grid graph, then*

$$\tau(G) \geq \mathfrak{b}^m,$$

where $m = |V(G)| - |\hat{\partial}G|$.

Theorem 9 is an immediate consequence of the following:

Proposition 10. *Let G be a grid graph and $v \in V(G)$. If $v \notin \hat{\partial}G$, then $m_v \geq \mathfrak{b}$.*

For the remainder of this section, we assume that G is a grid graph, we fix a vertex $v \in V(G)$ and we assume that $v \notin \hat{\partial}G$, with the goal of proving that $m_v \geq \mathfrak{b}$.

Denote the coordinates of v as $v = (x, y)$ and denote the top and left neighbors of v as $a = (x, y + 1)$ and $b = (x - 1, y)$. Since $v \notin \hat{\partial}G$, we know that $\overline{va}, \overline{vb} \in E(H_v)$.

Lemma 11. *Let P_v denote the probability that a uniformly random member of $\mathcal{T}(H_v)$ contains both \overline{va} and \overline{vb} . Then $0 < P_v < 1$ and*

$$m_v = \frac{2}{1 - P_v}.$$

Proof. Partition the members of $\mathcal{T}(H_v)$ into three sets, $\mathcal{T}(H_v) = T_1 \cup T_2 \cup T_3$, according to whether they:

(T_1) Contain \overline{va} but not \overline{vb} .

(T_2) Contain \overline{vb} but not \overline{va} .

(T_3) Contain \overline{va} and \overline{vb} .

Since \square_v is a subgraph of H_v , it is straightforward to see that all three sets in this partition are nonempty. For example, T_3 is nonempty because a member of T_3 can be obtained from any member of $\mathcal{T}(H'_v)$ by adding \overline{va} and \overline{vb} and removing any other edge of the resulting cycle that this creates.

Furthermore, $|\mathcal{T}(H'_v)| = |T_1|$ because “adding the edge \overline{va} ” is a bijection between these sets. Similarly $|\mathcal{T}(H'_v)| = |T_2|$. Thus, $m_v = \frac{\tau(H_v)}{\tau(H'_v)} = \frac{2\tau(H'_v)+|T_3|}{\tau(H'_v)}$. Solving $P_v = \frac{|T_3|}{2\tau(H'_v)+|T_3|}$ for $|T_3|$ and substituting completes the proof. \square

Lemma 12. *Let E_v denote the probability that a simple random walk on H_v starting at v “escapes to b ,” which means that it reaches b before returning to v . We have:*

$$m_v = \frac{2E_v}{2E_v - 1}.$$

Proof. Define P_v as in Lemma 11. We can better understand P_v via the Aldous-Broder algorithm for generating a uniformly random spanning tree of a connected graph [1], [4] (Wilson’s algorithm from [20] would also work here). Their algorithm works as follows. Start at any vertex and do a simple random walk. Each time a vertex is first encountered, mark the edge from which it was encountered. When all vertices have been encountered, the set of marked edges is a uniformly random spanning tree.

We apply the Aldous-Broder algorithm as follows. Let $\mathcal{W}(a)$ denote a simple random walk starting at a on the connected component of H_v that contains a . Denote this connected component as H_v^0 , and note that it also contains v and b because $v \notin \hat{\partial}G$.

It is straightforward to see that P_v equals the probability that, in the walk $\mathcal{W}(a)$, the vertex b is first encountered along the edge \overline{vb} . In fact, this is the only way in which the set of marked edges will end up containing both \overline{va} and \overline{vb} .

Next let $\mathcal{W}(v)$ denote a simple random walk starting at v on H_v^0 . Here is a review of the definitions of P_v and E_v together with a new definition of Q_v :

- P_v = the probability in $\mathcal{W}(a)$ that b is first encountered along \overline{vb} .
- Q_v = the probability in $\mathcal{W}(a)$ of reaching v before reaching b .
- E_v = the probability in $\mathcal{W}(v)$ of reaching b before returning to v .

Since E_v is the probability of escaping to b on the first step plus the probability of escaping after more than one step, we have:

$$E_v = \frac{1}{2} + \frac{1}{2}(1 - Q_v) \tag{4}$$

In particular $E_v > \frac{1}{2}$ because $Q_v \neq 1$, which follows from the fact that v is not in the top-left boundary.

It remains to relate P_v and Q_v . For this, let $P(k)$ denote the probability in $\mathcal{W}(a)$ that b is first encountered along \overline{vb} immediately following the walk's k^{th} visit to v . We have:

$$P_v = \sum_{k \geq 1} P(k) = \sum_{k \geq 1} \left(\frac{Q_v}{2}\right)^k = \frac{Q_v}{2 - Q_v} \quad (5)$$

Combining Equations 4 and 5 with Lemma 11 yields the following expressions for the multiplier:

$$m_v = \frac{2}{1 - P_v} = \frac{2 - Q_v}{1 - Q_v} = \frac{2E_v}{2E_v - 1}. \quad \square$$

The problem is now reduced to understanding the escape probability E_v . A standard trick in the literature is to bound escape probabilities using Rayleigh's Monotonicity Laws, whose intuition comes from the long studied connection between random walks and electrical circuits. We recommend [7] for an elementary introduction to this connection and to Rayleigh's Laws. We'll require the following special case:

Proposition 13 (Rayleigh's Monotonicity Law). *Let \tilde{H} be a connected graph, let H be a subgraph of \tilde{H} , and let $v_0, b_0 \in V(H)$ be distinct vertices. Assume that H contains all edges in \tilde{H} incident to v_0 . Let \tilde{E} (respectively E) denote the probability that a simple random walk on \tilde{H} (respectively on H) starting at v_0 "escapes to b_0 ," which means it reaches b_0 before returning to v_0 . Then $E \leq \tilde{E}$.*

Thus, there is a greater probability of escape on the larger graph than on the smaller subgraph. In our application of Rayleigh's Law, the smaller graph will be H_v , while the larger will be the infinite subgraph, \mathcal{U} , of $\mathcal{L}(\mathbb{Z}^2)$ whose vertex set is:

$$V(\mathcal{U}) = \{(x, y) \in \mathbb{Z}^2 \mid y \geq 1 \text{ or } (y = 0 \text{ and } x \leq 0)\}. \quad (6)$$

We can think of $V(\mathcal{U})$ as the set of points of \mathbb{Z}^2 prior to (and including) the origin $\mathbf{0} = (0, 0)$ in the words-on-a-page sense. After applying a translation for notational convenience, we can assume that v is positioned at the origin; that is, we can assume that $v = \mathbf{0} = (0, 0)$, $a = (0, 1)$, and $b = (-1, 0)$. With this understanding, H_v is a subgraph of \mathcal{U} .

Lemma 14. *Let $E(\infty)$ denote the probability that a simple random walk on \mathcal{U} starting at $v = (0, 0)$ escapes to $b = (-1, 0)$. Then $E(\infty) = \frac{\mathfrak{b}}{2(\mathfrak{b}-1)}$.*

We postpone the proof of Lemma 14 until the end of the next section. For now, we will use the lemma to finish off Proposition 10 and hence also Theorem 9.

Proof of Proposition 10. Rayleigh's Law together with Lemma 14 gives $E_v \leq E(\infty) = \frac{\mathfrak{b}}{2(\mathfrak{b}-1)}$. Recall from the previous proof that $E_v > 1/2$. On the domain $E_v > 1/2$, the function $m_v = \frac{2E_v}{2E_v-1}$ is decreasing. Therefore $m_v \geq \frac{2E(\infty)}{2E(\infty)-1} = \mathfrak{b}$. \square

Proof of Theorem 9.

$$\tau(G) = \prod_{v \in V(G)} m_v \geq \prod_{v \in V(G) - \hat{\partial}G} m_v \geq \mathfrak{b}^m,$$

where $m = |V(G)| - |\hat{\partial}G|$. \square

6 An upper bound on $\tau(G)$

The goal of this section is to prove the upper bound of Theorem 4. Figure 4 shows that this upper bound is false in the non-simple case (with m re-expressed in terms of $\hat{\partial}G$ as in Theorem 9). This upper bound will follow immediately from Lemma 7 together with the following:

Proposition 15. *If G is a simple grid graph and $v \in V(G)$ with $v \notin \hat{\partial}G$, then $m_v \leq 4$*

We will prove this proposition (and more general upper bounds on m_v) via Rayleigh's Law by comparing H_v to a smaller subgraph constructed as follows.

For each integer $k \geq 1$, define $\tilde{\mathcal{U}}(k)$ to be the subgraph of \mathcal{U} (from Equation 6) induced by all vertices within graph-distance k from $\mathbf{0}$, and then obtain $\mathcal{U}(k)$ from $\tilde{\mathcal{U}}(k)$ by removing all vertices of degree 1 and their adjacent edges. That is,

$$\begin{aligned} V(\tilde{\mathcal{U}}(k)) &= \{p \in V(\mathcal{U}) \mid \text{dist}(\mathbf{0}, p) \leq k\}, \\ V(\mathcal{U}(k)) &= \{p \in V(\tilde{\mathcal{U}}(k)) \mid \text{degree}(p) \neq 1\}, \end{aligned}$$

where "dist" is the edge distance of the graph. The first few are shown in Figure 5.

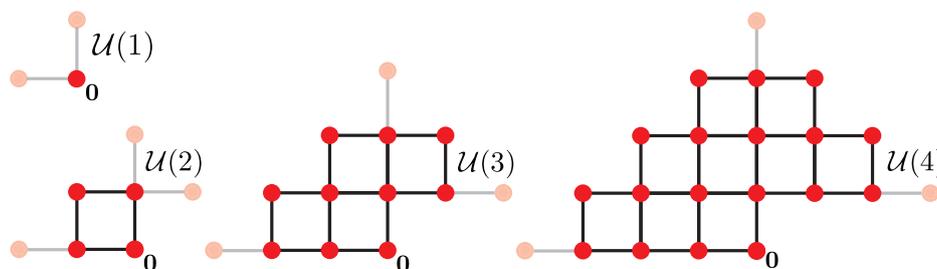


Figure 5: $\mathcal{U}(k)$ for $k \in \{1, 2, 3, 4\}$. The transparent vertices and edges belong to $\tilde{\mathcal{U}}(k)$ but not $\mathcal{U}(k)$.

Assume for the remainder of the section that G is a grid graph and $v \in V(G)$. As in the previous section, assume (after applying a translation) that $v = \mathbf{0} = (0, 0)$ so that $H_v \subset \mathcal{U}$. Define:

$$d_v = \max\{k \mid \mathcal{U}(k) \subset H_v\}. \quad (7)$$

Notice that $v \in \hat{\partial}G$ if and only if $d_v = 1$.

Lemma 16. *If $v \notin \hat{\partial}G$ (or equivalently if $d_v \geq 2$), then $m_v \leq F(d_v)$, where F is a function defined in the proof below, whose first few values are given in Table 1.*

Proof. Set $k = d_v$. Define Q_v and E_v as in the proof of Lemma 12, in which are found the relations:

$$m_v = \frac{2E_v}{2E_v - 1} = \frac{2 - Q_v}{1 - Q_v}. \quad (8)$$

Analogously define $Q(k)$ and $E(k)$ with respect to random walks in $\mathcal{U}(k)$; that is:

Table 1: Some values of F rounded to 4 decimals

k	$F(k)$
2	4
3	3.4833
4	3.3486
5	3.2936
\vdots	\vdots
12	3.2193
	\downarrow
	$\mathfrak{b} \approx 3.2099$

- $Q(k)$ is the probability that a simple random walk in $\mathcal{U}(k)$ starting at a reaches v before reaching b .
- $E(k)$ is the probability that a simple random walk in $\mathcal{U}(k)$ starting at v escapes to b .

Define

$$F(k) = \frac{2E(k)}{2E(k) - 1} = \frac{2 - Q(k)}{1 - Q(k)}. \tag{9}$$

Rayleigh's Monotonicity Law implies that $E_v \geq E(k)$ and therefore that $m_v \leq F(k)$.

To compute $F(k)$, it will suffice to compute $Q(k)$ via the method of [7, Section 1.2.6], which we briefly review here. Regard the random walk on $\mathcal{U}(k)$ starting at a as an absorbing Markov chain with absorbing states $\{v, b\}$. Index the vertices of $\mathcal{U}(k)$ with these absorbing states listed first, so the transition matrix of the Markov chain has the block form $\begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{R} & \mathbf{Q} \end{pmatrix}$. The absorption probabilities are given by $\mathbf{B} = (\mathbf{I} - \mathbf{Q})^{-1}\mathbf{R}$. Thus $Q(k)$ equals the entry of \mathbf{B} whose column corresponds to the absorbing state v and whose row corresponds to the non-absorbing state a . \square

The function $F : \{2, 3, \dots\} \rightarrow \mathbb{R}$ defined in the previous proof has the following properties.

Lemma 17. *F is non-increasing, and $\lim_{k \rightarrow \infty} F(k) = \mathfrak{b}$.*

Proof. The function $k \mapsto E(k)$ is non-decreasing by Rayleigh's Monotonicity Law because $\mathcal{U}(k) \subset \mathcal{U}(k + 1)$. Therefore Equation 9 implies that the function $k \mapsto F(k)$ is non-increasing. It follows that $\lim_{k \rightarrow \infty} F(k)$ exists.

To prove that this limit equals \mathfrak{b} , we first let $S(n)$ denote a square with n^2 vertices in $\mathcal{L}(\mathbb{Z}^2)$. Equation 1 implies

$$\lim_{n \rightarrow \infty} \frac{\ln(\tau(S(n)))}{|V(S(n))|} = \ln \mathfrak{b}. \tag{10}$$

Let μ_n denote the geometric mean of $\{m_v \mid v \in V(S(n))\}$. Using Equation 3, we can rewrite Equation 10 as:

$$\lim_{n \rightarrow \infty} \mu_n = \mathfrak{b}. \quad (11)$$

For $n \in \mathbb{N}$ and $v \in V(S(n))$, define:

$$\begin{aligned} d_v &= \max\{k \mid \mathcal{U}(k) \subset H_v\}, \\ D_v &= \min\{k \mid H_v \subset \mathcal{U}(k)\}. \end{aligned}$$

The first equation is just the previous definition of d_v from Equation 7. Notice that $d_v \leq D_v$. All of this is illustrated in Figure 6 for a particular vertex $v \in V(S(9))$ for which $d_v = 4$ and $D_v = 9$.

Rayleigh's Monotonicity Law applied to both the inclusion $\mathcal{U}(d_v) \subset H_v$ and the inclusion $H_v \subset \mathcal{U}(D_v)$ yields

$$E(d_v) \leq E_v \leq E(D_v) \quad (12)$$

Equations 8 and 9 convert this into bounds on the multiplier of v :

$$F(D_v) \leq m_v \leq F(d_v) \quad (13)$$

The value d_v is controlled by the distance from v to the top, left and right edges of $S(n)$. When n is large, the vast majority of the vertices of $S(n)$ have large distance to these edges and hence have large values of d_v (and thus also of D_v). More precisely, it is straightforward to show that for every $K \in \mathbb{N}$ and every $\epsilon > 0$, there exists $N \in \mathbb{N}$ such that if $n > N$, then the equation

$$d_v > K$$

is true for at least a $(1 - \epsilon)$ portion of the vertices $v \in S(n)$.

Together with Equation 13, this implies that $\lim_{n \rightarrow \infty} \mu_n = \lim_{k \rightarrow \infty} F(k)$. Combining this with Equation 11 completes the proof. \square

The value $F(2) = 4$ in Table 1 is exact (not rounded). Proposition 15 is an immediate consequence of this value.

We now use Lemma 17 to fill in a missing proof from Section 5.

Proof of Lemma 14. Let $E(k)$ denote the probability that a simple random walk starting at $v = (0, 0)$ on $\mathcal{U}(k)$ escapes to $b = (-1, 0)$. Then

$$E(\infty) = \lim_{k \rightarrow \infty} E(k) = \lim_{k \rightarrow \infty} \frac{F(k)}{2(F(k) - 1)} = \frac{\mathfrak{b}}{2(\mathfrak{b} - 1)}. \quad \square$$

Finally, we prove the upper bound of Theorem 4 as a quick consequence of Lemma 7 and Proposition 15.

Proof of upper bound of Theorem 4.

$$\tau(G) = \prod_{v \in V(G)} m_v = \prod_{v \in V(G) - \hat{\partial}G} m_v \leq 4^m,$$

where $m = |V(G)| - |\hat{\partial}G| = |V(G)| - \frac{1}{2}|\partial G| - 1$. \square

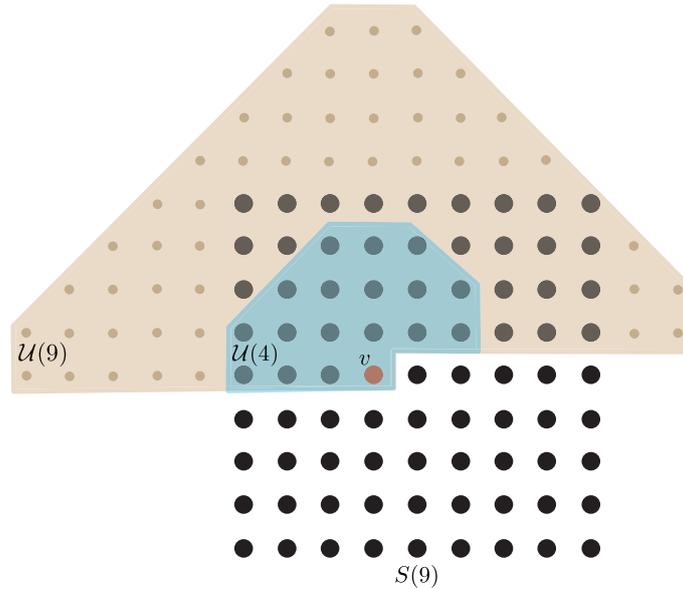


Figure 6: $\mathcal{U}(4) \subset H_v \subset \mathcal{U}(9)$ for this $v \in V(S(9))$.

An improvement on the upper bound of Theorem 4 can be obtained by considering the level sets of d . That is, for each $k \geq 1$ define:

$$G^k = \{v \in V(G) \mid d_v = k\}.$$

Notice that $G^1 = \hat{\partial}G$. Assuming that G is simple, Lemma 16 gives:

$$\ln(\tau(G)) \leq \sum_{k \geq 2} \ln(F(k)) \cdot |G^k|. \quad (14)$$

Equation 14 is stronger than the upper bound of Theorem 4, but it is not clear whether Equation 14 is necessarily stronger for all simple grid graphs than Theorem 1.

7 Application to redistricting

In this section, we apply our results to help account for an empirically observed correlation between two different compactness measurements for district maps, which we now describe.

The starting point of modern redistricting models is a graph G whose vertices represent the precincts of a state. Two vertices are connected by an edge if the corresponding precincts share a geographic boundary with non-zero length. A *map* \mathcal{M} is a partition¹ of G into subgraphs $\{G_1, \dots, G_k\}$ called *districts*, which are required to satisfy certain legal requirements.

¹More precisely, this means a partition $\{V_1, \dots, V_K\}$ of $V(G)$, with $\{G_1, \dots, G_K\}$ denoting the subgraphs of G induced by these vertex sets.

Many states require their congressional and legislative maps to be *compact*, which roughly means that the districts should have plump shapes without the winding tentacles commonly associates with gerrymandering. The term “compact” is defined only vaguely in most states, but there are several ways to quantify it. For example, let $\mathcal{C} = \mathcal{C}(\mathcal{M})$ denote the set of *cut edges*, which means the edges between pairs of vertices of G that belong to different districts. The value $|\mathcal{C}|$ is frequently used as a discrete measurement of the map’s overall compactness; see [8] for advantages of this measurement compared to other compactness measurements. A second common measurement of compactness is the *spanning tree score*, defined as

$$\mathcal{S}(\mathcal{M}) = \ln \left(\prod \tau(G_i) \right).$$

Figure 7 exhibits a very strong negative correlation between $|\mathcal{C}|$ and $\mathcal{S}(\mathcal{M})$ for an ensemble of 1000 random partitions of the square with 30^2 vertices into 9 districts². Our goal is to account for this negative correlation.

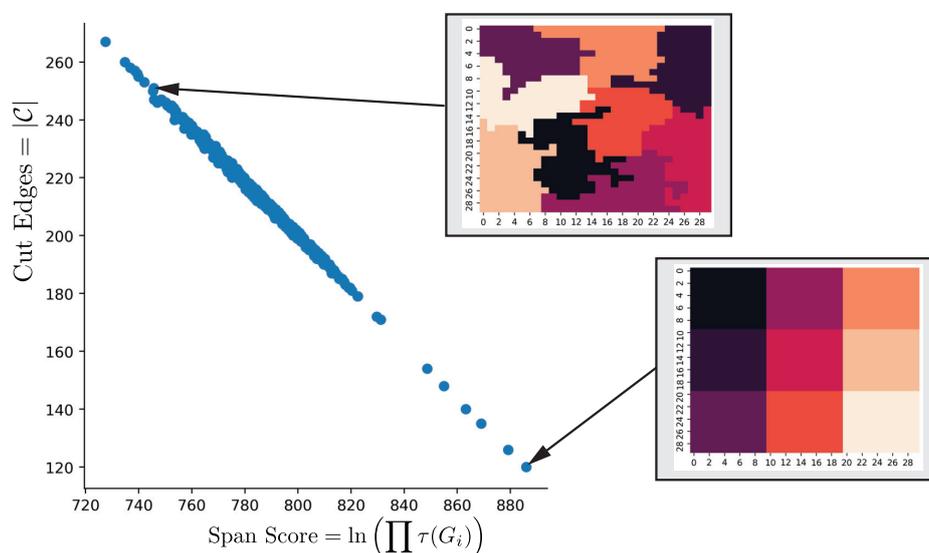


Figure 7: Spanning tree score vs. cut edge count for an ensemble of 1000 partitions of the square with 30^2 vertices into 9 districts.

Procaccia and Tucker-Folz related these two measurements for general planar graphs in [16] with bounds that depend on degree bounds in G and its dual. In the case of large grid graphs, their result say for a pair of maps \mathcal{M}_1 and \mathcal{M}_2 that if $|\mathcal{C}(\mathcal{M}_2)| \geq 7.23|\mathcal{C}(\mathcal{M}_1)|$ then $\mathcal{S}(\mathcal{M}_1) \geq \mathcal{S}(\mathcal{M}_2)$.

We assume henceforth that G and each G_i is a simple grid graph. The main result of this section is:

²This ensemble was created with the ReCom algorithm from [6] with 5% population deviation using the pictured tic-tac-toe arrangement as the initial partition.

Proposition 18.

$$|\mathcal{C}| = \underbrace{\text{Area}(G) + K - 1}_{\text{denote as } C_1} - \sum_{i=1}^K \text{Area}(G_i).$$

Notice that C_1 is a constant that does not depend on the partition. Also notice that $\text{Area}(G)$ and $\text{Area}(G_i)$ do not equal the areas of any geographic regions because G and G_i are *dual graphs*.

Proposition 18 and Theorem 4 together imply:

$$C_1 - \frac{1}{\ln(b)} \cdot \ln \left(\prod \tau(G_i) \right) \leq |\mathcal{C}| \leq C_1 - \frac{1}{\ln(4)} \cdot \ln \left(\prod \tau(G_i) \right). \quad (15)$$

Figure 8 contains the same data as Figure 7 (zoomed out in order to show the axes) with the upper and lower bounds of Equation 15 displayed as red lines. The slopes of these red lines are $-\frac{1}{\ln(b)}$ and $-\frac{1}{\ln(4)}$. Their common vertical intercept is $C_1 = 29^2 + 9 - 1 = 849$. Notice that all of the data points lie between the two red lines, even though the maps in this ensemble do not satisfy all of our hypotheses – their districts are not all simple.

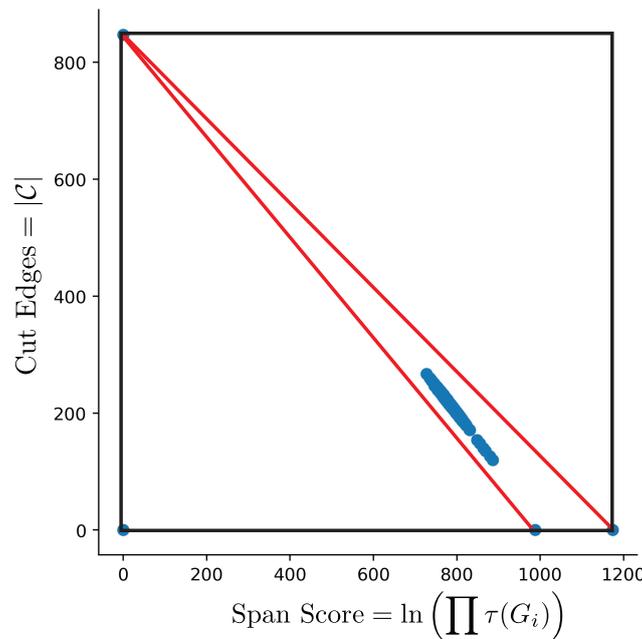


Figure 8: The red lines represent the inequalities of Equation 15.

Proof of Proposition 18. Let F denote the set of faces of G that are not faces of any of the districts and are not the unbounded external face. It will suffice to prove that $|\mathcal{C}| = |F| + K - 1$.

It is possible to select a subset $S \subset \mathcal{C}$ of size $K - 1$ such that S connects the districts into a spanning tree. More precisely, S induces a spanning tree, T_S , on the *district quotient graph* of G , which is defined to contain one vertex for each district, and to have an edge

between each pair of vertices if the corresponding pair of districts is connected by at least one cut edge. Figure 9 provides an illustration in which the districts are dark grey, the faces of F are light grey, the members of S are dashed red lines, and the members of $\mathcal{C} - S$ are solid blue lines.

It will suffice to find a bijection from F to $\mathcal{C} - S$. For this, we will consider F as a graph in which a pair $f_1, f_2 \in F$ are connected by an edge if they are adjacent across a member of $\mathcal{C} - S$. Considered in this way, F is acyclic because T_S is connected. Thus, F is a union of disjoint trees. We'll call $f \in F$ an *end face* if it is adjacent across an edge of $\mathcal{C} - S$ with a face of $\mathcal{L}(\mathbb{Z}^2)$ that's not a face of G . Since T_S is acyclic, each connected component of F contains at least one end face.

Imagine following a path in F and marking the faces and cut edges crossed along the way. Since faces and cut edges alternate, we can insure we mark an equal number of each by starting with a face and ending with a cut edge. Let's call such a path a *good path*. To build a bijection of F with $\mathcal{C} - S$, it will suffice to find a finite collection of good paths that together mark all of the faces in F and all of the cut edges in $\mathcal{C} - S$. This can be achieved by repeating the following two steps until all faces have been marked:

1. Select any face of $f \in F$ that hasn't yet been marked.
2. There exists a path in F from f to an end face. Traverse this path (marking the faces and edges along the way) until either reaching this end face or reaching a previously marked face.

In Figure 9, one possible collection of good paths is illustrated in green. When this algorithm terminates, all edges of $\mathcal{C} - S$ must be marked because any unmarked edge could be added to S without creating a cycle in T_S . \square

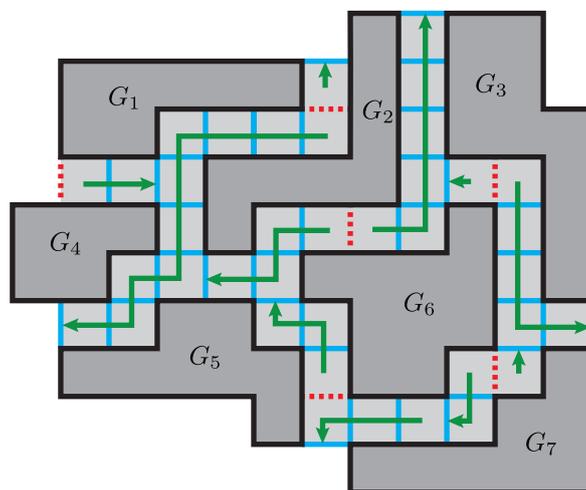


Figure 9: The green paths provide a bijection between F (the light grey faces) and $\mathcal{C} - S$ (the blue cut edges).

Acknowledgments

The author is pleased to thank Russell Lyons and the referee for valuable feedback.

References

- [1] Aldous, D., *The random walk construction of uniform spanning trees and uniform labelled trees*, SIAM J. Discrete Math. 3 (1990), 450–465.
- [2] Autry, E., Carter, D., Herschlag, G., Hunter, Z., Mattingly, J., *Metropolized Forest Recombination for Monte Carlo sampling of graph partitions*, SIAM J. Appl. Math. 83 (2023), 1366–1391.
- [3] Burton, R.M. and Pemantle, R., *Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances*, Ann. Probab. 21 (1993), 1329–1371.
- [4] Broder, A., *Generating random spanning trees*, Foundations of Computer Science (1989), pp. 442–447.
- [5] Cannon, S., Duchin, M., Rule, P., Randall, D., *Spanning tree methods for sampling graph partitions*, 2022, [arXiv:2210.01401](https://arxiv.org/abs/2210.01401).
- [6] DeFord, D., Duchin, M., Solomon, J., *Recombination: A family of Markov chains for redistricting*, Harvard Data Science Review, Issue 3.1, Winter 2021.
- [7] Doyle, P., Snell, L., *Random walks and electric networks*, Carus Math. Monogr., 22 Mathematical Association of America, Washington, DC, 1984. xiv+159 pp. ISBN:0-88385-024-9.
- [8] Duchin, M., Tenner, B., *Discrete geometry for electoral geography*, 2018, [arXiv:1808.05860v2](https://arxiv.org/abs/1808.05860v2).
- [9] Feng, L., Yu, G., Jiang, Z., Ren, L., *Sharp upper bounds for the number of spanning trees of a graph*, Appl. Anal. Discrete Math 2 (2008), 255–259.
- [10] Kenyon, R., *The asymptotic determinant of the discrete Laplacian*, Acta Math. 185 (2000), 239–286.
- [11] Kirchhoff, G., *Über die Auflösung der Gleichungen, auf welche man bei der Untersuchung der linearen Verteilung galvanischer Ströme geführt wird*, Ann. Phys. Chem. 72 (1847), 497–508.
- [12] Kostochka, A., *The number of spanning trees in graphs with a given degree sequence*, Random Structures and Algorithms, 6 (1995), 269–274.
- [13] Lyons, R., *Asymptotic enumeration of spanning trees*, Combin. Probab. Comput. 14 (2005), 491–522.
- [14] Lyons, R., *Identities and inequalities for tree entropy*, Combin. Probab. Comput. 19 (2010), 303–313.

- [15] Pick, *Geometrisches zur Zahlenlehre*, Sitzungsberichte des deutschen naturwissenschaftlich-medizinischen Vereines für Böhmen “Lotos” in Prag. (Neue Folge). 19 (1899) pp. 311–319.
- [16] Procaccia, A., Tucker-Foltz, J., *Compact redistricting plans have many spanning trees*, Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 3754–3771.
- [17] Shrock, R., Wu, F. *Spanning trees on graphs and lattices in d dimensions*, J. Phys. A, 33 (2000), 3881–3902.
- [18] Temperley, H. *The enumeration of graphs on large periodic lattices*, Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972), pages 285–294. Inst. Math. Appl., Southend, 1972.
- [19] Teufl, E., Wagner, S., *On the number of spanning trees on various lattices*, J. Phys. A., 43 (2010), 415001.
- [20] Wilson, D., *Generating random spanning trees more quickly than cover time*, Proceedings of the Twenty-eighth Annual ACM Symposium on the Theory of Computing (1996), pp. 296–303.
- [21] Wu, F., *Number of spanning trees on a lattice*. J. Phys. A, 10 (1977), L113–L115.