Lines: 665

The Euclidean MST-ratio for Bi-colored Lattices

Sebastiano Cultrera di Montesano
ISTA (Institute of Science and Technology Austria), Klosterneuburg, Austria

Ondřej Draganov
ISTA (Institute of Science and Technology Austria), Klosterneuburg, Austria

Herbert Edelsbrunner
ISTA (Institute of Science and Technology Austria), Klosterneuburg, Austria

Morteza Saghafian
ISTA (Institute of Science and Technology Austria), Klosterneuburg, Austria

Abstract

Given a finite set, $A \subseteq \mathbb{R}^2$, and a subset, $B \subseteq A$, the *MST-ratio* is the combined length of the minimum spanning trees of B and $A \setminus B$ divided by the length of the minimum spanning tree of A. The question of the supremum, over all sets A, of the maximum, over all subsets B, is related to the Steiner ratio, and we prove this sup-max is between 2.154 and 2.427. Restricting ourselves to 2-dimensional lattices, we prove that the sup-max is 2.0, while the inf-max is 1.25. By some margin the most difficult of these results is the upper bound for the inf-max, which we prove by showing that the hexagonal lattice cannot have MST-ratio larger than 1.25.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases Minimum spanning trees, Steiner ratio, lattices, partitions.

Funding This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant no. 788183, from the Wittgenstein Prize, Austrian Science Fund (FWF), grant no. Z 342-N31, and from the DFG Collaborative Research Center TRR 109, 'Discretization in Geometry and Dynamics', Austrian Science Fund (FWF), grant no. I 02979-N35.

1 Introduction

The recent development of measuring the interaction between two or more sets of points with methods from topological data analysis motivates the discrete geometric question about minimum spanning trees studied in this paper; see [2, 8] for background in this general area. We refer to the measured interaction as *mingling*, deliberately choosing an ambiguous term while leaving the concrete meaning to the geometric and algebraic constructions described in [6]. As explained in the appendix of the current paper, one of these measurements can be expressed in elementary terms:

▶ **Definition.** *Given a finite set, $A \subseteq \mathbb{R}^2$, we write* MST(A) *for the* (Euclidean) minimum spanning tree *of the complete graph on A, with edge weights equal to the distances between the points. For $B \subseteq A$, the* MST-ratio *of A and B is the combined length of the minimum spanning trees of B and $A \setminus B$, divided by the length of the minimum spanning tree of A:*

$$\mu(A, B) = \frac{|\mathrm{MST}(B)| + |\mathrm{MST}(A \setminus B)|}{|\mathrm{MST}(A)|}. \tag{1}$$

To make use of this measure for statistical or other purposes, we ought to know how small and how large the ratio can get (the extremal question), and how it behaves for random data. A first result in the latter direction can be found in [7], who prove that for points chosen uniformly at random in the unit square, the expected MST-ratio for a random partition into two subsets is at least $\sqrt{2} - \varepsilon$, for any $\varepsilon > 0$.

Given any set, A, the minimum MST-ratio is achieved by removing the longest edge from MST(A) and letting B and $A \setminus B$ be the vertices of the resulting two trees, so it is less than 1.0. More interestingly, the maximum MST-ratio is related to the *Steiner ratio* of the Euclidean plane [9, 10], and we exploit this connection to prove that the supremum is between 2.154 and 2.427 (Theorem 2.1 in Section 2). The infimum of the maximum is again less interesting: allowing ourselves to pick points arbitrarily close to each other, this infimum can be seen to be arbitrarily close to 1.0.

This motivates us to study the MST-ratio for a restricted class of sets, and our choice are the (Euclidean) lattices, which are well studied objects with many applications in mathematics and beyond; see e.g. [12]. Taking a sequence of progressively larger but finite portions of such a lattice, we have well defined minimum spanning trees and can study the asymptotic behavior of the MST-ratio. Our main result is that the maximum MST-ratio of the hexagonal lattice is 1.25 (Theorem 4.2 in Section 4). Observe that this is an upper bound on the infimum, over all lattices, of the maximum MST-ratio. We complement this with a matching lower bound (Claim 3.5 in Section 3), and with matching lower and upper bounds for the supremum maximum MST-ratio, which we establish is 2.0 (Claims 3.2 and 3.4 in Section 3).

2 The Maximum MST-ratio

The main question we ask to what extent two minimum spanning trees can be longer than a single minimum spanning tree of the same points; see the definition of the MST-ratio of a set $A \subseteq \mathbb{R}^2$ and a subset $B \subseteq A$ in the introduction. We are interested in the maximum MST-ratio, over all subsets $B \subseteq A$, and in the supremum and infimum of this maximum, over all finite sets $A \subseteq \mathbb{R}^2$.

The supremum is related to the well-studied Steinter tree problem. Given a finite set, $X \subseteq \mathbb{R}^2$, the *Steiner tree* of X is the minimum spanning tree of $X \cup B$, in which $B = B(X)$ is chosen to minimize the length of this tree. The *Steiner ratio* of the Euclidean plane is the infimum of the length ratio, $|\mathrm{MST}(X \cup B)|/|\mathrm{MST}(X)|$, over all finite sets in the plane. There are sets $X \subseteq \mathbb{R}^2$ for which the ratio is only $\sqrt{3}/2 = 0.866\ldots$; take for example the vertices of an equilateral triangle as X and the barycenter of this triangle as the sole point in B. It is conjectured that $\sqrt{3}/2$ is the Steiner ratio of the Euclidean plane [9], but the current best lower bound proved in [3] is only $0.824\ldots$. We use this bound to prove upper and lower bounds for the supremum maximum MST-ratio:

▶ **Theorem 2.1.** *The supremum, over all finite $A \subseteq \mathbb{R}^2$, of the maximum, over all subsets $B \subseteq A$, of the MST-ratio satisfies* $2.154 \leq \sup_A \max_B \mu(A, B) \leq 2.427$.

Proof. We first prove the upper bound. Since B is a subset of A, the MST of A cannot be shorter than the Steiner tree of B. Similarly, the MST of A cannot be shorter than the Steiner tree of $A \setminus B$. Hence, $|\mathrm{MST}(A)| \geq 0.824\ldots \cdot |\mathrm{MST}(B)|$ and $|\mathrm{MST}(A)| \geq 0.824\ldots \cdot |\mathrm{MST}(A \setminus B)|$. It follows that the ratio satisfies

$$\mu(A, B) \leq \frac{2 \cdot [|\mathrm{MST}(B)| + |\mathrm{MST}(A \setminus B)|]}{0.824\ldots \cdot [|\mathrm{MST}(B)| + |\mathrm{MST}(A \setminus B)|]} = 2.426\ldots. \quad (2)$$

This inequality holds for every $B \subseteq A$. We second prove the lower bound for the sup-max by constructing a set A of seven points that implies the inequality. Let $B \subseteq A$ be the three vertices of an equilateral triangle with unit length edges, and let $A \setminus B$ be the vertices of another equilateral triangle with unit length edges, but this time together with the barycenter. Hence, $|\mathrm{MST}(B)| = 2$ and $|\mathrm{MST}(A \setminus B)| = \sqrt{3}$. Assuming the distance between corresponding vertices of the two equilateral triangles is less than $\varepsilon > 0$, we have

$|\mathrm{MST}(A)| < \sqrt{3} + 3\varepsilon$. This implies

$$\mu(A, B) > \frac{2 + \sqrt{3}}{\sqrt{3} + 3\varepsilon} > 2.154\ldots - 4\varepsilon. \tag{3}$$

Since we can make $\varepsilon > 0$ arbitrarily small, this implies the claimed lower bound. ◀

The example used to establish the lower bound can be extended to larger numbers of points, e.g. the following disjoint union of three lattices: B is the hexagonal lattice (to be defined shortly), and $A \setminus B$ is a slightly shifted copy of the hexagonal lattice, together with the barycenters of the triangles in every fourth row, which is a rectangular lattice with distances 1 and $\sqrt{3}$ between consecutive rows and columns.

The question about the infimum of the maximum MST-ratio turns out to be less interesting, with 1.0 as answer. To see the lower bound, set $B = A$, in which case $|\mathrm{MST}(B)| = |\mathrm{MST}(A)|$ and $|\mathrm{MST}(A \setminus B)| = 0$. The ratio is therefore equal to 1. We get the upper bound by constructing a set A of $n \geq 2$ points. It contains the origin, $n - 2$ points each at distance at most $\varepsilon > 0$ from the origin, and another point at unit distance from the origin. Call the latter point b, assume $b \in B$, and consider the case in which B contains at least one other point of A. Then

$$1 \leq \quad |\mathrm{MST}(A)| \quad \leq 1 + 2(n - 2)\varepsilon, \tag{4}$$

$$1 - \varepsilon \leq \quad |\mathrm{MST}(B)| \quad \leq 1 + 2(n - 2)\varepsilon, \tag{5}$$

$$0 \leq |\mathrm{MST}(A \setminus B)| \leq 2(n - 3)\varepsilon. \tag{6}$$

For any given $\delta > 0$, we can choose $\varepsilon > 0$ sufficiently small such that the ratio is smaller than $1 + \delta$. In the other case, in which $B = \{b\}$, we have $|\mathrm{MST}(B)| = 0$ and $|\mathrm{MST}(A \setminus B)| \leq 2(n - 2)\varepsilon$, so we can make the ratio arbitrarily small and certainly smaller than 1.0.

3 Two-dimensional Lattices

Motivated by the triviality of the infimum maximum MST-ratio for general finite sets, we aim for a restriction that disallows extremely unbalanced distributions. There are many choices, and we opt for a maximally restricted setting in which the MST-ratio is still an interesting question. Specifically, we focus on 2-dimensional lattices.

▶ **Definition.** *The* (Euclidean) lattice *spanned by two linearly independent vectors,* $\mathbf{u}, \mathbf{v} \in \mathbb{R}^2$*, consists of all integer combinations of these vectors:* $\Lambda(\mathbf{u}, \mathbf{v}) = \{i\mathbf{u} + j\mathbf{v} \mid i, j \in \mathbb{Z}\}$.

By definition, lattices are infinite. To cope with the difficulty of constructing the minimum spanning tree of infinitely many points, we take progressively larger but finite portions of a lattice and monitor the sequence of MST-ratios. Specifically, we consider squares centered at the origin and rhombi spanned by the shortest basis of the lattice to generate such neighborhoods.

If this sequence converges, we call the limit the *MST-ratio* of the lattice. A particularly interesting lattice is the *hexagonal* or *hexagonal lattice*, which is spanned by $\mathbf{u} = (1, 0)$ and $\mathbf{v} = \frac{1}{2}(1, \sqrt{3})$; see the left panel in Figure 1. The minimum distance between its points is 1, so all edges of the MST have length 1. The two partitions illustrated in the middle and right panels of Figure 1 have MST-ratios $1.245\ldots$ and 1.25, respectively. In one way or another, we use this lattice to prove all four bounds claimed in the following theorem.

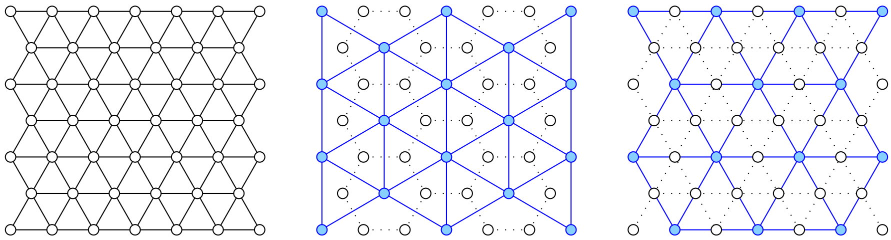

Figure 1: *Left:* a portion of the hexagonal lattice and all its shortest edges. *Middle:* a partition into one and two thirds of the points, with MST-ratio converging to $(2+\sqrt{3})/3 = 1.245\ldots$. *Right:* a partition into one and three quarters of the points, with MST-ratio converging to 1.25.

▶ **Theorem 3.1.** *The supremum and infimum, over all* 2*-dimensional lattices,* Λ*, of the maximum, over all subsets,* $B \subseteq \Lambda$*, of the MST-ratio are* $C_0 = \sup_\Lambda \max_B \mu(\Lambda, B) = 2.0$ *and* $c_0 = \inf_\Lambda \max_B \mu(\Lambda, B) = 1.25$.

Each of the subsequent subsections restates and proves one of the four bounds, except for the last subsection, which only sketches the proof strategy, with the proof presented in Section 4.

3.1 Lower Bound for Sup-Max

This subsection exhibits a lattice, and a partition of this lattice into two sets, such that the MST-ratio of progressively larger finite portions of the lattice approaches the supremum of the maximum MST-ratio claimed in Theorem 3.1 from below.

▷ Claim 3.2. $C_0 \geq 2.0$.

Proof. Let Λ be the hexagonal lattice horizontally stretched by a factor 9, and let $B \subseteq \Lambda$ be the one third of the points drawn blue in Figure 2. The (vertical) distance between neighboring points in a column of Λ is $\sqrt{3}$, and the (horizontal) distance between two neighboring columns is $\frac{9}{2}$. For each $r \geq 0$, let $\Lambda_r \subseteq \Lambda$ and $B_r \subseteq B$ be the points in

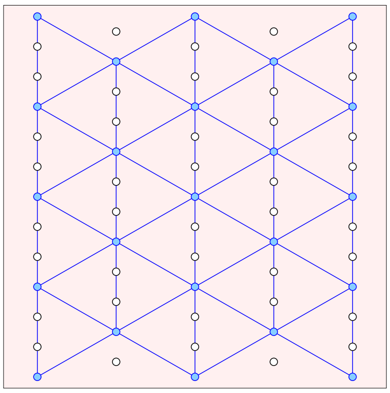

Figure 2: The portion of the horizontally stretched hexagonal lattice, Λ, and the subset of *blue* points, B, inside a square centered at the origin. The edges show the union of all possible minimum spanning trees of the *blue* points.

$[-r, r]^2$. Hence, Λ_r consists of $p_r = 2\lfloor 2r/9 \rfloor + 1$ vertical columns, which alternate between $q_r = 2\lfloor r/\sqrt{3} \rfloor + 1$ and $q_r - 1$ or $q_r + 1$ points. Observe that p_r and q_r are both odd, and that $n_r = q_r p_r \pm (p_r - 1)/2$ is the cardinality of Λ_r. The number of points of B_r in the columns alternates between $b_r = 2\lfloor r/(3\sqrt{3})\rfloor + 1$ and $b_r - 1$ or $b_r + 1$, so $m_r = b_r p_r \pm (p_r - 1)$ is the cardinality of B_r. It is easy to see that $n_r - 2p_r \leq 3m_r \leq n_r + 2p_r$.

By choice of the stretch factor, B is a hexagonal lattice with distance $3\sqrt{3}$ between closest points. Hence, $|\mathrm{MST}(B_r)| = 3\sqrt{3}(m_r - 1)$. Compare this with a minimum spanning tree of Λ_r, which first connects the points in each column and second connects neighboring columns with one edge for each pair. Hence,

$$|\mathrm{MST}(\Lambda_r)| = \sqrt{3}(n_r - p_r) + \sqrt{21}(p_r - 1), \tag{7}$$

because every point, except the last in each column, has a neighbor at distance $\sqrt{3}$ below, and any two neighboring columns have points at distance $\sqrt{21}$ from each other. Similarly, any minimum spanning tree of $\Lambda_r \setminus B_r$ first connects the points in each column and second connects neighboring columns with one edge for each pair. Its length is therefore the same as that of $\mathrm{MST}(\Lambda_r)$. Using $3m_r = n_r + o(n_r)$, this implies

$$\frac{|\mathrm{MST}(B_r)| + |\mathrm{MST}(\Lambda_r \setminus B_r)|}{|\mathrm{MST}(\Lambda_r)|} = \frac{3\sqrt{3}(m_r - 1) + \sqrt{3}(n_r - p_r) + \sqrt{21}(p_r - 1)}{\sqrt{3}(n_r - p_r) + \sqrt{21}(p_r - 1)} \tag{8}$$

$$= \frac{2\sqrt{3}n_r + o(n_r)}{\sqrt{3}n_r + o(n_r)} \xrightarrow{r\to\infty} 2.0. \tag{9}$$

For any $\varepsilon > 0$, we can choose r sufficiently large such that the MST-ratio exceeds $2.0 - \varepsilon$, which implies the claimed lower bound. ◀

3.2 Upper Bound for Sup-Max

This subsection proves the upper that matched the lower bound on the supremum maximum MST-ratio established in the preceding subsection. Given any lattice and any partition of this lattice into two sets, we show that for any $\varepsilon > 0$, the MST-ratio cannot exceed $2 + \varepsilon$. We begin with a bound for the length of the minimum spanning tree of any finite set in a square.

▶ **Lemma 3.3.** *The length of the minimum spanning tree of any n or fewer points in $[0, n]^2$ is at most $2n\sqrt{n}$.*

Proof. Assuming the number of points is n, the minimum spanning tree has $n-1$ edges, and we write $\ell_1, \ell_2, \ldots, \ell_{n-1}$ for their lengths The sum of the squares of these lengths is at most $4n^2$, as proved in [9]. By the Cauchy–Schwarz inequality, the sum of the ℓ_i is maximized when all terms are the same, namely $\ell_i^2 = 4n^2/(n-1)$ for all i. This implies

$$\sum_{i=1}^{n-1} \ell_i \leq (n-1)\sqrt{4n^2/(n-1)} = 2n\sqrt{n-1}, \tag{10}$$

from which the claimed bound follows. ◀

Lemma 3.3 will provide a crucial step in the proof of the upper bound for the supremum maximum MST-ratio, which we present next.

▷ Claim 3.4. $C_0 \leq 2.0$.

Proof. We show that the MST-ratio of any lattice $\Lambda \subseteq \mathbb{R}^2$ and any subset $B \subseteq \Lambda$ is at most the claimed upper bound. Let $\mathbf{u}$ and $\mathbf{v}$ be the shortest two vectors that span Λ, breaking ties arbitrarily if necessary. Suppose their lengths satisfy $1 = \|\mathbf{u}\| \leq \|\mathbf{v}\| = \nu$. To simplify language, we call the points on a line parallel to $\mathbf{u}$ a *row* of Λ. For every positive integer, n, let $\Lambda_n \subseteq \Lambda$ contain all points $\alpha\mathbf{u} + \beta\mathbf{v}$, with $0 \leq \alpha, \beta \leq n$. The minimum spanning tree of Λ_n first connects the points in each row and then the neighboring rows, so

$$|\mathrm{MST}(\Lambda_n)| = (n+1)n + n\nu. \tag{11}$$

Set $B_n = B \cap \Lambda_n$. We construct a spanning tree, $T(B_n)$, by first connecting the points within the rows. This allows for the possibility that some rows do not contain any points of B_n. In each of the other rows, we choose an arbitrary but fixed point of B_n, write $B'_n \subseteq B_n$ for the chosen points, construct $\mathrm{MST}(B'_n)$, and add its edges to $T(B_n)$. Since $T(B_n)$ spans B_n but is not necessarily the shortest such tree, so $|\mathrm{MST}(B_n)| \leq |T(B_n)|$. To bound the latter, recall that there are $n+1$ rows, each of length at most n. Furthermore, B'_n consists of at most $n+1$ points that fit inside a square of side length $n(\nu+1)$, in which ν is independent of n. Lemma 3.3 implies $|\mathrm{MST}(B'_n)| \leq 2(\nu+1)\sqrt{\nu+1} \cdot n\sqrt{n}$. Hence,

$$|\mathrm{MST}(B_n)| \leq (n+1)n + 2(\nu+1)\sqrt{\nu+1} \cdot n\sqrt{n}. \tag{12}$$

By symmetry, we have the same upper bound for the length of $\mathrm{MST}(\Lambda_n \setminus B_n)$. Comparing this with the minimum spanning tree of Λ_n, we get

$$\frac{|\mathrm{MST}(B_n)| + |\mathrm{MST}(\Lambda_n \setminus B_n)|}{|\mathrm{MST}(\Lambda_n)|} \leq \frac{2n^2 + 2n + 4(\nu+1)^{3/2} \cdot n\sqrt{n}}{n^2 + n + \nu n} \xrightarrow{n\to\infty} 2.0. \tag{13}$$

For every $\varepsilon > 0$, we can choose n large enough so that the MST-ratio is less than $2.0 + \varepsilon$. This works for every lattice and partition, which implies the claimed upper bound. ◀

3.3 Lower Bound for Inf-Max

This subsection establishes the lower bound for the infimum maximum MST-ratio. We do this by establishing a partition into one and three quarters that can be defined for any lattice and has MST-ratio at least the infimum of the maximum MST-ratio claimed in Theorem 3.1.

▷ Claim 3.5. $c_0 \geq 1.25$.

Proof. Let $\mathbf{u}$ and $\mathbf{v}$ be two vectors spanning Λ, and let B be the sublattice spanned by $2\mathbf{u}$ and $2\mathbf{v}$. Assuming the minimum distance between two points in Λ is 1, most edges of $\mathrm{MST}(\Lambda)$ have length 1, while most edges of $\mathrm{MST}(B)$ have length 2. Since B contains only a quarter of the points, this implies $|\mathrm{MST}(B)| = \frac{1}{2}|\mathrm{MST}(\Lambda)|$. The complement of the sublattice contains three quarters of the points, and the edges in its MST have length at least 1, which implies $|\mathrm{MST}(\Lambda \setminus B)| \geq \frac{3}{4}|\mathrm{MST}(\Lambda)|$. Hence, the MST-ratio satisfies $\mu(\Lambda, B) \geq \frac{1}{2} + \frac{3}{4} = 1.25$. ◀

3.4 Upper Bound for Inf-Max

The upper bound for the infimum of the maximum MST-ratio will be proved in Section 4. This proof is carefully constructed from a network of inequalities that require attention to detail. This subsection makes an argument why it is not unreasonable to believe that significant short-cuts may be difficult to find.

The lattice that is most resistant to large MST-ratios is the hexagonal lattice, Λ, of which four different subsets, $B \subseteq \Lambda$, are illustrated as packings of hexagonal neighborhoods in Figure 3. Starting at the upper middle, then left, then right, and finally the lower middle, the density of the packing decreases monotonically as the minimum distance between points of B increases from $\sqrt{3}$ to 2, to $\sqrt{7}$, and finally to 3. Corresponding, B contains one third, one quarter, one seventh, and one ninth of the points. Perhaps surprisingly, the MST-ratio

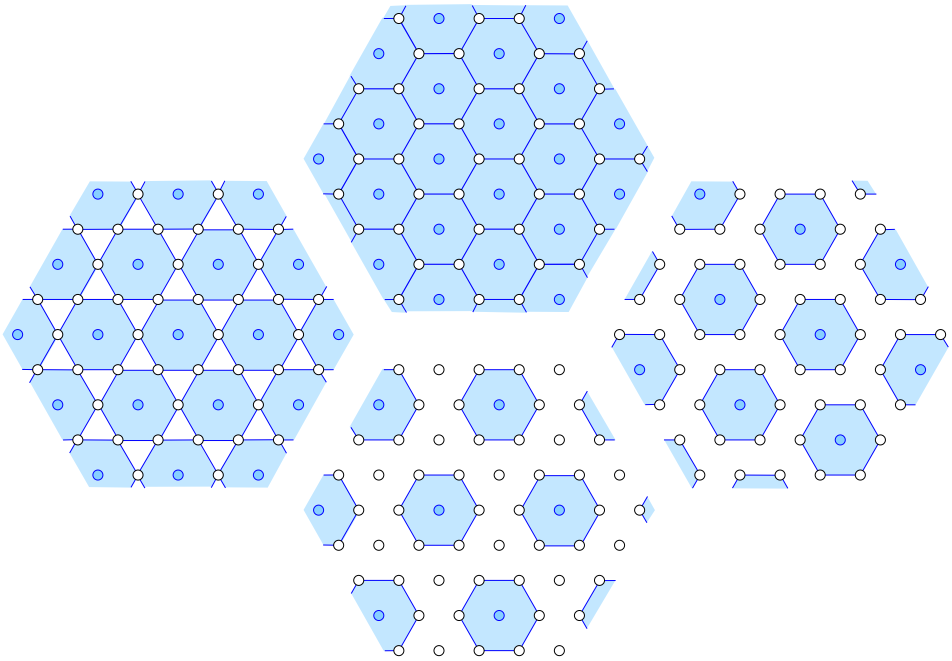

Figure 3: Four partitions of the hexagonal lattice into two sets, in which we draw each *(blue)* point of the smaller set with its hexagonal neighborhood. The proportions of *blue* versus *white* points are 1 : 2 in the *upper middle*, 1 : 3 on the *left*, 1 : 6 on the *right*, and 1 : 8 in the *lower middle*. The corresponding MST-ratios are approximately 1.245, 1.25, 1.236, and 1.222, in this sequence.

does not vary monotonically and attains the largest value for the subset B that contains one quarter of the points. The purpose of Section 4 is to prove that no other subset of Λ achieves a larger MST-ratio; that is: 1.25 is the maximum MST-ratio of the hexagonal lattice.

▷ Claim 3.6. $c_0 \leq 1.25$.

Because the value matches the lower bound stated in Claim 3.5, this implies that 1.25 is indeed the infimum maximal MST-ratio over all 2-dimensional lattices. Prior to studying the hexagonal lattice, the authors of this paper proved that the maximal MST-ratio of the integer lattice is $\sqrt{2}$—which happens to match the ratio found for random sets [7]—and the maximizing subset are the points whose coordinates add up to even integers. The proof is similar to the one for the hexagonal lattice presented in Section 4, and almost as long. If instead we consider the points whose coordinates add up to odd integers, we get the same MST-ratio, so the integer lattice has at least two global maxima. Similarly, the hexagonal lattice has at least four global maxima, and moving from one to the other means walking a path along which the MST-ratio is sometimes barely below 1.25. To support the hypothesis of a rugged but shallow landscape between the global minima, we conducted computational experiments, which identified many local maxima that prevent local improvement strategies from reaching any global maximum. We feel that these findings justify the exhaustive case analysis in Section 4, and the many delicate inequalities in that section give evidence for how close the paths get to the maximum MST-ratio.

4 Hexagonal Lattice on Torus

In this section, we prove Claim 3.6 for the hexagonal lattice on the torus. We begin by constructing this lattice from a portion of the hexagonal lattice in the plane and proving that the minimum spanning trees in the two topologies are not very different in length. In the remaining subsections, we give a precise statement of the theorem that implies Claim 3.6 and prove the theorem with a packing argument in six steps.

4.1 Plane versus Torus

We consider the hexagonal lattice on the torus rather than in $\mathbb{R}^2$ in order to eliminate boundary effects, which appear when we study a finite portion of the hexagonal lattice. Let $\mathbf{u}$ and $\mathbf{v}$ be two unit vectors with a 60° degree angle between them, and write $\Lambda \subseteq \mathbb{R}^2$ for the hexagonal lattice they span. For every positive $n \in \mathbb{Z}$, let $\Lambda_n \subseteq \Lambda$ contain the n^2 points $a = \alpha\mathbf{u} + \beta\mathbf{v}$ with $0 \leq \alpha, \beta \leq n-1$. We write Λ_n' for the same n^2 points but with the topology of the torus, which we get by identifying a with $a + in\mathbf{u} + jn\mathbf{v}$ for all $i, j \in \mathbb{Z}$, and defining the distance as the minimum Euclidean distance between any two representatives. Equivalently, consider the rhombus of points $\varphi\mathbf{u} + \psi\mathbf{v}$ for real coefficients $-\frac{1}{2} \leq \varphi, \psi \leq n - \frac{1}{2}$,

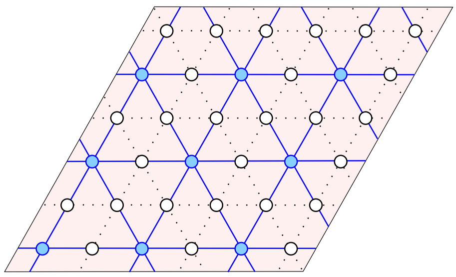

Figure 4: The hexagonal lattice of 36 points on the torus, obtained by gluing opposite sides of the rhombus. The sublattice with twice the distance between neighboring points in shown in *blue*.

and glue this rhombus along opposite sides as illustrated for $n = 6$ in Figure 4. Call the boundary of this rhombus the *seam*. Its length is $4n$ in the plane but only $2n$ on the torus since the sides are glues in pairs. Note also that every point of Λ has distance at least $\sqrt{3}/4$ from the nearest point in the seam.

▶ **Lemma 4.1.** *Let $\Lambda \subseteq \mathbb{R}^2$ be the hexagonal lattice, $\Lambda_n \subseteq \Lambda$ the subset of n^2 points, and Λ_n' the same n^2 points but on the torus, as described above. For any subset $B_n \subseteq \Lambda_n$ and the corresponding subset $B_n' \subseteq \Lambda_n'$ on the torus, the lengths of the minimum spanning trees satisfy* $|\mathrm{MST}(B_n')| \leq |\mathrm{MST}(B_n)| \leq |\mathrm{MST}(B_n')| + 32\sqrt{2} \cdot n\sqrt{n}$.

Proof. Fix two minimum spanning trees, T of B_n in $\mathbb{R}^2$ and T' of B_n' on the torus. Since the distances on the torus are smaller than or equal to those in $\mathbb{R}^2$, we have $|T'| \leq |T|$, which is the first claimed inequality. Let E' be the edges of T' that have the same length in both topologies, and let E'' be the other edges of T', which are shorter on the torus than in $\mathbb{R}^2$. To draw an edge of E'' in the plane so its length matches the length on the torus, we need to connect representatives of the endpoints that lie in different rhombi. Assuming one endpoint is in Λ_n, this edge crosses the seam. In contrast, every edge in E' can be drawn between two points of Λ_n, so without crossing the seam. We will prove shortly that the distance between two crossings measured along the seam is at least $\frac{1}{2}$. Since the length of the seam is

$2n$, this implies that E'' contains at most $4n$ edges. Let $V'' \subseteq \Lambda_n$ be the set of at most $8n$ endpoints of the edges in E'', and let T'' be a minimum spanning tree of V'', with distances measured in $\mathbb{R}^2$. Since Λ_n easily fits inside a square with sides of length $8n$, Lemma 3.3 implies $|T''| \leq 32\sqrt{2} \cdot n\sqrt{n}$. The edges in E' together with the edges of T'' form a connected graph with vertices Λ_n. Hence,

$$|T| \leq |T'| + |T''| \leq |T'| + 32\sqrt{2} \cdot n\sqrt{n}, \tag{14}$$

which is the second claimed inequality. It remains to show that the distance between two crossings along the seam is at least $\frac{1}{2}$. Let ab and xy be two edges in E'', and recall that the greedy construction of the minimum spanning tree prohibits x and y to lie inside the smallest circle that passes through a and b, and vice versa. If the edges share an endpoint, then the angle between them is at least $60°$. Since the common endpoint is at distance at least $\sqrt{3}/4$ from the seam, the implies the claimed lower bound on the distance between the two crossings. So assume a, b, x, y are distinct, and let $c \in ab$ and $z \in xy$ be the points that minimize the distance between the edges, and observe that $\|c - z\|$ is a lower bound for the distance between the crossings. At least one of c and z must be an endpoint, so suppose $z = x$. But since x lies outside the smallest circle of a and b, and outside the unit circles centered at a and b, the distance of x to any point of ab is at least 1. ◀

The inequalities in Lemma 3.3 generalize to all 2-dimensional lattices. Letting $\mathbf{u}$ and $\mathbf{v}$ be two shortest vectors that span a lattice, and assuming $1 = \|\mathbf{u}\| \leq \|\mathbf{v}\| = \nu$, we get $2(4 + 4\nu)^{3/2} \cdot n\sqrt{n}$ as an upper bound for the difference in length, in which we compare a rhombus of $n \times n$ points in $\mathbb{R}^2$ and on the torus, as before.

4.2 Statement of Theorem

We fix n to an even integer and write $\Delta = \Lambda'_n$ for the hexagonal lattice on the torus. Since n is even, $\Delta_1 = \{2x \mid x \in \Delta\}$ is a hexagonal sublattice of Δ, and we set $\Delta_3 = \Delta \setminus \Delta_1$; see Figure 4. The lengths of the three minimum spanning trees are easy to determine because they use only the shortest available edges, which have length 1 for Δ and Δ_3, and length 2 for Δ_1. The MST-ratio is therefore

$$\mu(\Delta, \Delta_1) = \frac{|\mathrm{MST}(\Delta_1)| + |\mathrm{MST}(\Delta_3)|}{|\mathrm{MST}(\Delta)|} = \frac{2\left(n^2/4 - 1\right) + \left(3n^2/4 - 1\right)}{n^2 - 1} \xrightarrow{n \to \infty} 1.25. \tag{15}$$

Call an edge *short* if its length is 1. All other edges have length larger than the desired average, which is $\frac{5}{4} = 1.25$, so we call them *long*. While $\mathrm{MST}(\Delta_3)$ has only short edges, and $\mathrm{MST}(\Delta_1)$ uses only the shortest edges connecting its points, we claim that their combined length is as large as it can be.

▶ **Theorem 4.2.** *Let Δ be a hexagonal lattice on the torus. Then the maximum MST-ratio of Δ converges to $\frac{5}{4} = 1.25$ from below.*

The proof consists of six steps, which are presented in the same number of subsections: 4.3 introduces the hexagonal distance, compares its MST with the Euclidean MST, and uses the former to formulate the proof strategy; 4.4 introduces the main tool, which are hexagonal-neighborhoods of the lattice points; 4.5 constructs a hierarchy of such neighborhoods aimed at counting the short edges; 4.6 introduces so-called satellites, which provide additional short edges needed in the proof; 4.7 forms loop-free subgraphs of short edges and bounds their sizes; and 4.8 does the final accounting while paying special attention to the cases in which all long edges have length between $\sqrt{3}$ and 3. Throughout this proof, we use the fact that the minimum spanning tree can be computed by greedily adding the shortest available edge that does not form a cycle to the tree [1, 11].

4.3 Hexagonal Distance and Proof Strategy

It is convenient to write the points in Δ with three integer coordinates. To explain this, let

$$\mathbf{x} = \tfrac{1}{\sqrt{3}}(0,1), \quad \mathbf{y} = \tfrac{1}{\sqrt{3}}\left(-\tfrac{\sqrt{3}}{2}, -\tfrac{1}{2}\right), \quad \mathbf{z} = \tfrac{1}{\sqrt{3}}\left(\tfrac{\sqrt{3}}{2}, -\tfrac{1}{2}\right) \tag{16}$$

be three vectors, each of length $\sqrt{3}/3$, that mutually enclose an angle of 120°. These are the projections of the unit coordinate vectors of $\mathbb{R}^3$ onto the plane normal to the diagonal direction, scaled such that the three points are mutually one unit of distance apart. The plane consists of all points $u = a\mathbf{x} + b\mathbf{y} + c\mathbf{z}$ for which $a + b + c = 0$, and such a point belongs to the hexagonal lattice iff $a, b, c \in \mathbb{Z}$; see Figure 5. Given a second point, $v = \alpha\mathbf{x} + \beta\mathbf{y} + \gamma\mathbf{z}$, we write $i = a - \alpha$, $j = b - \beta$, $k = c - \gamma$ to compute the squared Euclidean distance between u and v. Since $\mathbf{x}^2 = \mathbf{y}^2 = \mathbf{z}^2 = \frac{1}{3}$ and $\mathbf{xy} = \mathbf{yz} = \mathbf{zx} = -\frac{1}{6}$, we get

$$\|u - v\|^2 = \|i\mathbf{x} + j\mathbf{y} + k\mathbf{z}\|^2 = \tfrac{1}{3}(i^2 + j^2 + k^2) - \tfrac{1}{3}(ij + ik + jk) = i^2 + ij + j^2, \tag{17}$$

in which we get the final expression using $k = -(i + j)$. For points of the hexagonal lattice, i and j are integers, and so is the squared Euclidean distance between them. It follows that the minimum distance between two points in Δ is 1.

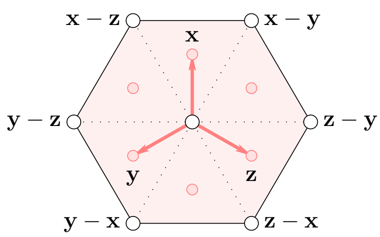

Figure 5: The unit disk under the hexagonal distance in the plane. The edges that connect the origin to the corners at $\pm(\mathbf{x} - \mathbf{y})$, $\pm(\mathbf{y} - \mathbf{z})$, $\pm(\mathbf{z} - \mathbf{x})$ decompose the hexagon into six equilateral triangles, whose barycenters are $\pm\mathbf{x}$, $\pm\mathbf{y}$, $\pm\mathbf{z}$.

We adapt the notion of distance to construct neighborhoods in the hexagonal lattice. By definition, the *hexagonal distance* between points $u = a\mathbf{x} + b\mathbf{y} + c\mathbf{z}$ and $v = \alpha\mathbf{x} + \beta\mathbf{y} + \gamma\mathbf{z}$ is

$$\|u - v\|_{hex} = \max\{|a - \alpha|, |b - \beta|, |c - \gamma|\} = \max\{|i|, |j|, |i + j|\}. \tag{18}$$

The *unit disk* under this distance consists of all points with hexagonal distance at most 1 from the origin: $\mathbb{H} = \{u \in \mathbb{R}^2 \mid \|u - 0\|_{hex} \leq 1\}$. It is the regular hexagon with unit length sides that is the convex hull of the points $\pm(\mathbf{x} - \mathbf{y})$, $\pm(\mathbf{y} - \mathbf{z})$, $\pm(\mathbf{z} - \mathbf{x})$; see Figure 5. For $B \subseteq \Delta$, we write $\mathrm{MST_{hex}}(B)$ for the spanning tree that minimizes the hexagonal length. We construct it by adding the edges in sequence of non-decreasing hexagonal length, breaking ties with Euclidean length, and breaking the remaining ties arbitrarily. Since $\mathrm{MST_{hex}}(B)$ is a spanning tree but not necessarily the one that minimizes Euclidean length, we have

$$|\mathrm{MST}(B)| \leq |\mathrm{MST_{hex}}(B)|, \tag{19}$$

in which we measure the Euclidean length on both sides. To prove Theorem 4.2, we show that for every $B \subseteq \Delta$, the average (Euclidean) length of the long edges in $\mathrm{MST_{hex}}(B)$ and the short edges in $\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)$ is at most $\frac{5}{4}$. Interchanging B and $\Delta \setminus B$, we get the same relation by symmetry. Using (19), this implies

$$|\mathrm{MST}(B)| + |\mathrm{MST}(\Delta \setminus B)| \leq |\mathrm{MST}_{\mathrm{hex}}(B)| + |\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)| \leq \tfrac{5}{4}(n^2 - 2). \tag{20}$$

Compare this with (15), which establishes $|\mathrm{MST}(\Delta_1)| + |\mathrm{MST}(\Delta_3)| = \frac{5}{4}n^2 - 3$ for the partition $\Delta = \Delta_1 \sqcup \Delta_3$. The right-hand side differs from the upper bound in (20) by only a small additive constant. We thus conclude that the maximum MST-ratio of Δ converges to $\frac{5}{4}$ from below, as claimed by Theorem 4.2.

4.4 Hierarchy of Habitats

Let T_ℓ be the subset of edges in $\mathrm{MST}_{\mathrm{hex}}(B)$ whose hexagonal lengths are at most ℓ, together with the endpoints of these edges. For example, T_0 has zero edges, T_1 consist of all short edges, and $T_\ell = \mathrm{MST}_{\mathrm{hex}}(B)$ for sufficiently large ℓ. All edges connecting points in different components of T_ℓ have hexagonal length $\ell + 1$ or larger. We thus write $k\mathbb{H}$ for the scaled copy of the unit disk and call

$$D_k(B) = \bigcup\nolimits_{u \in B} (k\mathbb{H} + u) \tag{21}$$

the *k-th thickening* of B, in which $k\mathbb{H} + u$ is the translate of $k\mathbb{H}$ whose center is u. As illustrated in Figure 6, the k-th thickenings of points u and v *overlap*, *touch*, are *disjoint* if the hexagonal distance between u and v is less than, equal to, larger than $2k$, respectively.

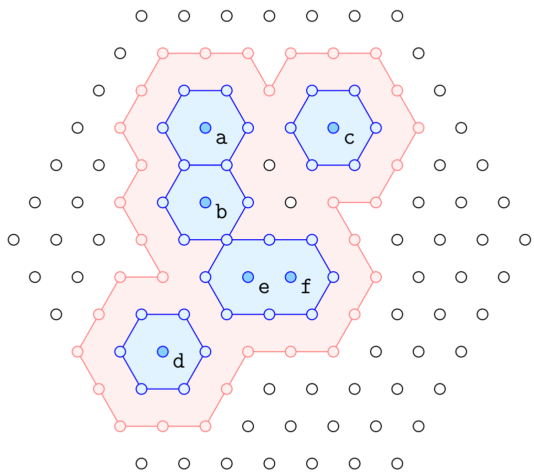

Figure 6: The *blue* 1-st thickening and the *pink* 2-nd thickening of $B = \{\mathtt{a}, \mathtt{b}, \mathtt{c}, \mathtt{d}, \mathtt{e}, \mathtt{f}\}$ in the hexagonal lattice. $\mathbb{H} + \mathtt{a}$ and $\mathbb{H} + \mathtt{b}$ share an edge and therefore form two rooms in a common house, while $\mathbb{H} + \mathtt{e}$ and $\mathbb{H} + \mathtt{f}$ overlap and thus form a one-room house in $D_1(B)$. These two houses form a block, and together with $\mathbb{H} + \mathtt{d}$, they form a compound of two blocks. $\mathbb{H} + \mathtt{c}$ is a room, a house, a block, and a compound by itself. The two compounds lie in the interior of a room in $D_2(B)$.

The boundary of $k\mathbb{H}$ passes through $6k$ points of the hexagonal lattice, which we call the *vertices* of $k\mathbb{H}$. Furthermore, we call the $6k$ (short) edges that connect these points in cyclic order the *edges* of $k\mathbb{H}$. Let $B_k \subseteq B$ be the vertex set of a component of T_{2k-1}, and observe that for all $u, v \in B_k$ there is a sequence of points $u = x_1, x_2, \ldots, x_m = v$ in B_k such that $k\mathbb{H} + x_i$ and $k\mathbb{H} + x_{i+1}$ overlap for all $1 \leq i \leq m - 1$. We define the *frontier* of the component, denoted $\partial D_k(B_k)$, as the lattice points and the connecting (short) edges in the boundary of $D_k(B_k)$. Furthermore, $\partial D_k(B)$ is the union of frontiers of the components of T_{2k-1}. These notions are illustrated in Figure 6, which shows $\partial D_1(B)$ and $\partial D_2(B)$ for six marked points. Note that the edge shared by $\mathbb{H} + \mathtt{a}$ and $\mathbb{H} + \mathtt{b}$ is part of $\partial D_1(B)$.

4.5 Subdivided Foreground and Background

Consider the 1-st thickening of B, which for the time being we call the *foreground*. Letting $B_1 \subseteq B_2$ be the vertex sets of two nested components of T_1 and T_2, we call $D_1(B_1)$ a *room* and $D_1(B_2)$ a *block* of the foreground. We say two rooms are *adjacent* if they share at least one edge. In Figure 6, there are five rooms, two of which are adjacent, and three blocks, one of which contains three rooms.

To make a finer distinction, observe that for any edge, its Euclidean length is smaller than or equal to the hexagonal length. The two notions agree on edges with slope 0, 2, and -2. Consider T_2 and T_3 after removing all edges whose Euclidean length equals 2 and 3, respectively, and let B_2' and B_3' be the vertex sets of the components that satisfy $B_1 \subseteq B_2' \subseteq B_2 \subseteq B_3'$. Observe that any two rooms in $D_1(B_2')$ have a sequence of pairwise adjacent rooms connecting them. We therefore call $D_1(B_2')$ a *house*. For comparison, any two rooms in $D_1(B_2)$ have a sequence of room connecting then such that any two consecutive rooms share at least a vertex but not necessarily a full edge. Similarly, for any two blocks in $D_1(B_3')$, there is a sequence of blocks connecting them such that the channel separating any two consecutive blocks at its narrowest place is only $\sqrt{3}/2$ wide. We therefore call $D_1(B_3')$ a *compound*; see Figure 6 for examples. For comparison, the channel that separates two compounds is at its narrowest place at least one unit of distance wide. A few observations:

- **(i)** all vertices of $\partial D_1(B)$ are points in $\Delta \setminus B$;
- **(ii)** all edges of $\partial D_1(B)$ are short;
- **(iii)** the frontier of a room consists of at least six (short) edges.

We call the complement of the foreground the *background*, and the components of the background its *backyards*. We say a backyard is *adjacent* to a house if the two share a non-empty portion of their boundary. There are configurations in which the number of backyards is twice the number of houses; see Figure 3 on the left, where each backyard is adjacent to three houses, and each house is adjacent to six backyards. In general, we distinguish between backyards adjacent to at most two and at least three houses, denoting their numbers α_1 and β_1, respectively. We prove an upper bound for β_1 in terms of the number of houses and blocks.

▶ **Lemma 4.3.** *Given h_1 houses arranged in b_1 blocks, the number of backyards adjacent to three or more houses satisfies $\beta_1 \leq 2h_1 - 2b_1 + 2$.*

Proof. We construct a graph $G = G(B)$ on the torus by placing a node inside each house, and whenever two houses meet at a boundary vertex, we connect the corresponding nodes with a curved arc that passes through the shared vertex. This can be done such that no two of the arcs cross and each face of G contains one backyard. A face bounded by a single arc (*loop*) or two arcs (*multi-arcs*) contains a backyard adjacent to at most two houses and thus does not count toward β_1. We remove this face by deleting the loop or one of the two multi-arcs. The resulting graph has h_1 nodes, b_1 components, and β_1 faces. Write a_1 for the number of arcs. If the graph is connected and all faces are bounded by three arcs, we have $h_1 - a_1 + \beta_1 = 0$ because the Euler characteristic of the torus is 0. Whenever we remove an arc from this graph, we either merge two faces or split a component, but it is also possible that the removal of the arc has neither of those two side-effects. Hence, we have

$h_1 - a_1 + \beta_1 \geq b_1 - 1$ in the general case. Since $2a_1 \geq 3\beta_1$, this implies $\beta_1 \leq 2h_1 - 2b_1 + 2$, as claimed. ◀

4.6 Satellites

By definition, compounds cannot be packed as tightly as blocks; see Figure 3 with lattice points between the compounds in the lower middle but no such points between the blocks on the right. Recall that each component of $D_1(B)$ is contained in a room of $D_2(B)$. For each such room, we single out the largest compound it contains—breaking ties arbitrarily—and call this the *big compound* of the room. All others are *small compounds* of the room. We identify satellites for each compound and distinguish between small and big compounds because of differences in the construction. The targeted lattice points are at distance $\sqrt{3}/2$ outside $D_1(B)$ and either on the boundary or in the interior of $D_2(B)$.

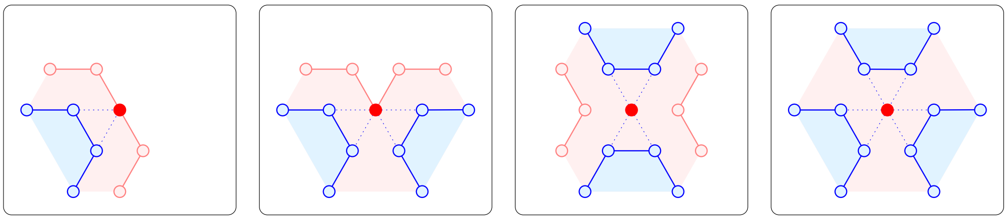

Figure 7: From *left* to *right*: a single, a double, another double, and a triple satellite in *red*. In the *left* two cases, the satellite belongs to the frontier of a room of the 2-nd thickening of B, while in the *right* two cases, the satellite lies in the interior of such a room.

For each small compound we find three satellites as follows: sandwich the compound between three lines with slopes $0, 2, -2$, choose a (short) edge as the basis of an equilateral triangle outside the compound on each line, and pick the vertex of this triangle opposite to the basis as a *satellite*. Observe that the Euclidean distance between any two satellites of the same compound is at least 3. In contrast, we pick six lattice points as the satellites of the big compound by sandwiching it between six lines, two each of slope $0, 2, -2$, choosing one basis on each line, and picking the vertex of the equilateral triangle opposite to the basis as a satellite. The Euclidean distance between any two such satellites is at least $\sqrt{3}$.

As illustrated in Figure 7, a lattice point can be a satellite of one, two, or three compounds in the same room. Accordingly, we call the point a *single*, *double*, or *triple satellite* of the room, respectively. A single satellite is necessarily a vertex on the frontier of the room, a triple satellite is necessarily in the interior of the room, and a double satellite can be one or the other. For a room, R, we write $s(R)$ and $d(R)$ for the number of single and double satellites on its frontier, and $e(R)$ and $t(R)$ for the number of double and triple satellites in its interior. Summing over all rooms in $D_2(B)$, we set $s_1 = \sum s(R)$, $d_1 = \sum d(R)$, $e_1 = \sum e(R)$, $t_1 = \sum t(R)$, and refer to s_1, d_1, e_1, t_1 as the *satellite sums* of $D_2(B)$. Since $s(R) + 2d(R) + 2e(R) + 3t(R)$ is three times the number of small compounds in R plus six for the big compound, the satellite sums satisfy a linear relation, which we state together with a property of short edges connecting satellites in the interior:

(iv) if $c_1 > 1$, then the satellite sums of $D_2(B)$ satisfy $s_1 + 2d_1 + 2e_1 + 3t_1 = 3c_1 + 3r_2$;

(v) any unit length edge connecting blocks of $D_1(B)$ inside a room of $D_2(B)$ with each other or to satellites in the interior of this room is contained in the interior of this room.

By construction, there are $s(R) + d(R)$ satellites that are vertices of R. We prove a stronger lower bound on the number of vertices, which also strengthens Claim (iii).

▶ **Lemma 4.4.** *Assume $r_2 \geq 2$ and let R be a room of $D_2(B)$. Then the frontier of R has at least $6 + \frac{2}{3}s(R) + \frac{4}{3}d(R)$ vertices.*

Proof. Let p, s, d be the number of non-satellite lattice points, single satellites, double satellites, and write $\mathrm{per}(R)$ for the *perimeter*, which is the length *of* or the number of (short) edges *in* the frontier of R. To begin note that a satellite in the frontier of R is in the boundary of at most one backyard. This is because the external angle is 180° at a single satellite and 60° at a double satellite. The internal angle at any vertex of another room is at least 120°, so there is not enough space for two backyards around a satellite; see the left two panels in Figure 7. This implies that we may assume that the frontier of R is a simple polygon, or a collection of such. Indeed, if the polygon touches itself at a vertex, this must be a non-satellite, which we can duplicate, and if the polygon touches itself along a sequence of edges, we can remove these edges and their shared vertices. This operation neither changes the number of single and double satellites, nor does it increase the perimeter. A room that contains only one compound can have perimeter as small as 12, but a room with at least two compounds has significantly larger perimeter, certainly larger than 15. For $\mathrm{per}(R) \leq 15$, we thus get only one compound and, by construction, only 6 single and no double satellites. This implies the claimed inequality. We therefore assume (22), aim at proving (23), and note that (24) follows as the convex combination of (22) and (23) with coefficients $\frac{1}{3}$ and $\frac{2}{3}$:

$$\mathrm{per}(R) \geq 16; \tag{22}$$

$$\mathrm{per}(R) \geq 1 + s + 2d; \tag{23}$$

$$\mathrm{per}(R) \geq \tfrac{1}{3}16 + \tfrac{2}{3}(1 + d + 2d) = 6 + \tfrac{2}{3}s + \tfrac{4}{3}d. \tag{24}$$

It remains to prove (23). Call the endpoints of an edge in the frontier of R *neighbors*. Two neighbors cannot both be double satellites, else they would belong to a common compound, which contradicts that the distance between them is at least $\sqrt{3}$. Furthermore, if a double satellite neighbors a single satellite, then this is only possible if they are vertices of an equilateral triangle bounding a backyard, as in Figure 8 on the left. For lack of space around this triangle, its third vertex is a non-satellite. The contribution of these three vertices to the right-hand side of (23) is $2 + 1 + 0 = 3$. Hence, we can remove the three edges from the left-hand side and the three vertices from the right-hand side of (23) without affecting the validity of the inequality. As illustrated in Figure 8 on the left, two such triangles may touch at a non-satellite vertex, but this does not matter and we can remove the edges and vertices of both triangles from (23).

We can therefore assume that both neighbors of a double satellite are non-satellites. Hence, between any two double satellites there is at least one non-satellite, which implies $p \geq d$. But $p = d$ only if $p = d = 0$ or there is strict alternation between double satellites and non-satellites. It is not possible that all vertices in the frontier are single satellites, because this contradicts that the distance between any two of them is at least $\sqrt{3}$. Strict alternation is possible, but only for the polygon of 12 edges shown in Figure 8 on the right. By assumption, $D_2(B)$ has at least two rooms, so not all backyards of R can be bounded by such 12-gons. But this implies $p \geq d + 1$, so $\mathrm{per}(R) = p + s + d \geq 1 + s + 2d$, as claimed. ◀

To generalize the above concepts to $k \geq 1$, we let $B_{2k-1} \subseteq B_{2k}$ be the vertex sets of two nested components of T_{2k-1} and T_{2k}, and call $D_k(B_{2k-1})$ a *room* and $D_k(B_{2k})$ a *block* of $D_k(B)$. The rooms that share edges join to form *houses*, and the blocks separated by channels that are only $\sqrt{3}/2$ wide join to form *compounds*. Write r_k, h_k, b_k, c_k for the number of rooms, houses, blocks, compounds of $D_k(B)$, α_k, β_k for the number of backyards adjacent to at most 2, at least 3 houses, and s_k, d_k, e_k, t_k for the satellite sums of $D_{k+1}(B)$. We can now extend Claims (i) to (v) and Lemmas 4.3 and 4.4 merely by substituting $D_k(B)$ for $D_1(B)$, β_k for β_1, c_k for c_1, etc. In particular, the extension of Claim (iv) to

$$s_k + 2d_k + 2e_k + 3t_k = 3c_k + 3r_{k+1} \tag{25}$$

assuming $c_k > 1$ will be needed shortly. We note that (25) and the extension of Lemma 4.4 can be strengthened, but it is not necessary for the purpose of proving Theorem 4.2.

4.7 Loop-free Subgraphs

Let V_k be the vertices of $D_k(B)$ together with all double and triple satellites that lie in the interior of rooms in $D_{k+1}(B)$, and note that $V_j \cap V_k = \emptyset$ whenever $j \neq k$. Let V_k' be V_k together with the remaining satellites of $D_k(B)$, and note that $V_j \cap V_k' = \emptyset$ if $j < k$, but V_k' and V_{k+1} may share points. To account for this difference, let ℓ be the smallest integer such that $r_{\ell+1} = 1$, and define

$$V = \begin{cases} V_1 & \text{if } \ell = 0; \\ V_1 \sqcup \ldots \sqcup V_{\ell-1} \sqcup V_\ell & \text{if } \ell \geq 1 \text{ and } c_\ell = 1; \\ V_1 \sqcup \ldots \sqcup V_{\ell-1} \sqcup V_\ell' & \text{if } \ell \geq 1 \text{ and } c_\ell > 1. \end{cases} \tag{26}$$

By construction, all points in V belong to $\Delta \setminus B$, and all unit length edges connecting these points are candidates for $\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)$. We therefore let U be a loop-free graph whose vertices are the points in V and whose edges all have unit length. Since U has no loops, there is an $\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)$ that contains U as a subgraph. We are therefore motivated to study the number of edges in U. Using a slight abuse of notation, we denote this number $\#U$. For every k, let U_k and U_k' be the subgraphs of U induced by V_k and V_k', respectively. We first count the edges in U_1 and U_1'.

▶ **Lemma 4.5.** *Let $r_1 \geq h_1 \geq b_1 \geq c_1$ be the number of rooms, houses, blocks, and compounds of $D_1(B)$, and s_1, d_1, e_1, t_1 the satellite sums of $D_2(B)$. Then*

$$\#U_1 \geq 2r_1 + h_1 + 3b_1 + (e_1 + t_1) - r_2 - 4; \tag{27}$$

$$\#U_1' \geq 2r_1 + h_1 + 3b_1 + (s_1 + d_1 + e_1 + t_1) - 5, \tag{28}$$

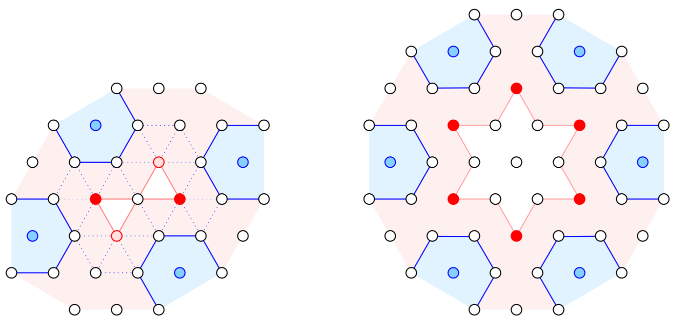

Figure 8: *Left:* two touching triangular backyards. Their shared vertex is a non-satellite, the two *red* vertices are double satellites, and the two *pink* vertices are single satellites. *Right:* unique polygon with strictly alternating double satellites and non-satellites. On *both sides*, all (partially drawn) *blue* compounds are different and belong to the same (partially drawn) *pink* room.

in which we assume $c_1 > r_2 = 1$ for the second inequality.

Proof. We argue in three steps: first counting edges in $\partial D_1(B)$, second counting edges connecting blocks, and third counting edges connecting the satellites. In each case, we count only unit length edges, and we make sure that the edges we count do not form loops.

For the first step, it is convenient to count *half-edges*, which are the two sides of an edge. These two sides either face two rooms, or one faces a room and the other faces the background. For a house, H, we make its $r(H)$ rooms accessible from the outside by removing $r(H) - 1$ edges shared by adjacent rooms plus 1 edge shared with the background. By (iii), each room was originally faced by at least 6 half-edges, so we still have at least $4r(H) + 1$ of them left. Doing this for each house, we make all r_1 rooms accessible from the background, and we have at least $4r_1 + h_1$ half-edges left facing these rooms.

Observe that the convex hull of a house contains at least six of the (short) edges that bound the house. One may have been removed, so we still have at least 5 half-edges facing the background. Keeping in mind that the cycles that bound backyards still need to be opened, we now have at least $4r_1 + h_1 + 5h_1$ half-edges and therefore at least $2r_1 + 3h_1$ edges. If a backyard is adjacent to at most two houses, then it has two consecutive (short) edges that enclose an angle less than π and that are both shared with the same house. Hence, the complementary angle on the side of the house is larger than π, which implies that these two edges cannot belong to the convex hull of the house. We remove one of them and use the half-edge facing the backyard of the other to compensate for the removed half-edge facing the room. Since both edges have not yet been accounted for, we still have at least $2r_1 + 3h_1$ edges. If a backyard is adjacent to three or more houses, we also remove one edge, but this time count one less. Recalling that β_1 is the number of such backyards, we still have at least $2r_1 + 3h_1 - \beta_1 \geq 2r_1 + h_1 + 2b_1 - 2$ edges, in which we get the right-hand side from Lemma 4.3.

For the second step, we connect the $b(R)$ blocks inside a common room of $D_2(B)$ with $b(R) - 1$ short edges. A total of b_1 blocks are hierarchically organized in r_2 rooms, so we add $b_1 - r_2$ short edges to those counted in the first step. Similarly, we add $e_1 + t_1$ short edges that connect the double and triple satellites in the interiors of the rooms to the vertices in the frontier of $D_1(B)$. Finally, we remove two edges to open the meridian and longitudinal cycles of the graph, if they exist. The final count is therefore at least $2r_1 + h_1 + 3b_1 + (e_1 + t_1) - r_2 - 4$, which is the claimed lower bound for $\#U_1$.

For the third step, we assume $c_1 > r_2 = 1$. Since there is only one room, there are no shared satellites between different rooms, and we can connect them to the frontier of $D_1(B)$ with $s_1 + d_1$ short edges without creating any loop. This implies that the number of edges in U'_1 is at least $2r_1 + h_1 + 3b_1 + (s_1 + d_1 + e_1 + t_1) - 5$, as claimed. ◀

The bounds in Lemma 4.5 generalize to $k > 1$, but there are differences. Most important is the existence of a loop-free graph for thickness $k - 1$. In particular, we have satellites that affect the structure and size of U_k and U'_k.

▶ **Lemma 4.6.** *Let $r_k \geq h_k \geq b_k \geq c_k$ be the number of rooms, houses, blocks, compounds of $D_k(B)$, and s_k, d_k, e_k, t_k the satellite sums of $D_{k+1}(B)$. Then for $k \geq 2$, we have*

$$\#U_k \geq (3r_k + \tfrac{1}{3}s_{k-1} + \tfrac{2}{3}d_{k-1}) + 4h_k + 3b_k + (e_k + t_k) - r_{k+1} - 4; \tag{29}$$

$$\#U'_k \geq (3r_k + \tfrac{1}{3}s_{k-1} + \tfrac{2}{3}d_{k-1}) + 4h_k + 3b_k + (s_k + d_k + e_k + t_k) - 5, \tag{30}$$

in which we assume $c_k > r_{k+1} = 1$ for the second inequality.

Proof. We argue again in three steps: first counting edges in $\partial D_k(B)$, second counting edges connecting blocks, and third counting edges connecting to the satellites. Each of these three steps is moderately more involved than the corresponding step in the proof of Lemma 4.5, and we emphasize the differences.

The first step starts the construction with Lemma 4.4, which implies that the rooms in $D_k(B)$ are faced by a total of at least $6r_k + \frac{2}{3}s_{k-1} + \frac{4}{3}d_{k-1}$ half-edges. After making all rooms accessible to the background, we still have at least $(4r_k + \frac{2}{3}s_{k-1} + \frac{4}{3}d_{k-1}) + h_k$ half-edges. Adding the at least 11 half-edges per house facing the background, we have at least $(4r_k + \frac{2}{3}s_{k-1} + \frac{4}{3}d_{k-1}) + 12h_k$ half-edges and thus at least $(2r_k + \frac{1}{3}s_{k-1} + \frac{2}{3}d_{k-1}) + 6h_k$ edges. Let α_k and β_k be the number of backyards adjacent to at most two and at least three houses, respectively. By extension of Lemma 4.3, we have $\beta_k \leq 2h_k - 2b_k + 2$. We remove an edge per backyard, which for the first type does not affect the current edge count, while the backyards of the second type reduce the count to $(2r_k + \frac{1}{3}s_{k-1} + \frac{2}{3}d_{k-1}) + 4h_k + 2b_k - 2$.

For the second step, we connect the blocks of $D_k(B)$ inside a common room of $D_{k+1}(B)$ with $b_k - r_{k+1}$ edges. Furthermore, we add r_k edges to connect the blocks of $D_{k-1}(B)$ inside a common room of $D_k(B)$—which inductively are already connected to each other—to the frontier of this room, and we add at least $e_k + t_k$ edges connecting to the triple satellites of compounds inside the rooms of $D_{k+1}(B)$. After removing two additional edges to break the meridian and longitudinal loops, if they exist, we arrive at a lower bound of at least $(3r_k + \frac{1}{3}s_{k-1} + \frac{2}{3}d_{k-1}) + 4h_k + 3b_k + (e_k + t_k) - r_{k+1} - 4$ edges in U_k.

For the third step, we assume $c_k > r_{k+1} = 1$, in which case we can add at least $s_k + d_k$ edges connecting to the single and double satellites. This implies $\#U'_k \geq (3r_k + \frac{1}{3}s_{k-1} + \frac{2}{3}d_{k-1}) + 4h_k + 3b_k + (s_k + d_k + e_k + t_k) - 5$. ◀

4.8 Book-keeping

The goal is to show that the average (Euclidean) length of the long edges in $\mathrm{MST}_{\mathrm{hex}}(B)$ and the short edges in $\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)$ is at most $\frac{5}{4}$. We thus assign a *credit* of $\alpha = \frac{1}{4}$ to every short edge and set the *cost* of a long edge to be its Euclidean length minus $\frac{5}{4}$. For convenience, we set the value of α to 1 Euro and convert the costs into Euros; see Table 1.

hex	2	2	3	3	4	4	4	5	5	5
L_2	$\sqrt{3}$	$\sqrt{4}$	$\sqrt{7}$	$\sqrt{9}$	$\sqrt{12}$	$\sqrt{13}$	$\sqrt{16}$	$\sqrt{19}$	$\sqrt{21}$	$\sqrt{25}$
cost	**1.92**	**3.00**	**5.58**	**7.00**	8.85	**9.42**	**11.00**	12.43	**13.33**	**15.00**

Table 1: The Euclidean lengths of the edges with hexagonal lengths 2 to 5, and their costs in Euros, each truncated beyond the first two digits after the decimal point.

For the accounting, we need the costs of the last two edges for each hexagonal length. Letting w_k, x_k and y_k, z_k be the costs of the two longest edges with hexagonal length $2k$ and $2k+1$, respectively, we have

$$w_k = \tfrac{1}{\alpha}\left[\sqrt{4k^2 - 2k + 1} - \tfrac{5}{4}\right], \; x_k = \tfrac{1}{\alpha}\left[2k - \tfrac{5}{4}\right], \tag{31}$$

$$y_k = \tfrac{1}{\alpha}\left[\sqrt{4k^2 + 2k + 1} - \tfrac{5}{4}\right], \; z_k = \tfrac{1}{\alpha}\left[(2k+1) - \tfrac{5}{4}\right]; \tag{32}$$

see Table 1, which shows the values of $w_1, x_1, y_1, z_1, w_2, x_2, y_2, z_2$ in boldface. Listing the edges in sequence, we need bounds for the cost differences between consecutive edges:

$$2 \le w_k - z_{k-1} \le 2.928\ldots; \quad 1.071\ldots \le x_k - w_k \le 2; \tag{33}$$

$$2 \le \quad y_k - x_k \quad \le 2.583\ldots; \quad 1.414\ldots \le z_k - y_k \le 2, \tag{34}$$

which are not difficult to prove using elementary computations. We use accounting with credits and costs to prove that the average (Euclidean) edge length of the two minimum spanning trees is less than $\frac{5}{4}$:

▶ **Lemma 4.7.** *Let Δ be the hexagonal lattice with $12n^2$ points and unit minimum distance on the torus, and $B \subseteq \Delta$. Then $|\mathrm{MST}(B)| + |\mathrm{MST}(\Delta \setminus B)| \le 15n^2 - \frac{5}{2}$.*

Proof. By (19), it suffices to prove the inequality for $\mathrm{MST}_{\mathrm{hex}}(B)$ and $\mathrm{MST}_{\mathrm{hex}}(\Delta \setminus B)$. For $k \ge 1$, we compare the edges of hexagonal length $2k$ and $2k+1$ in $\mathrm{MST}_{\mathrm{hex}}(B)$ with the (short) edges in U_k or possibly in U'_k. Since $T_{2k+1} \setminus T_{2k-1}$ is the set of these long edges, we can do this in one step by comparing $T_{2\ell+1}$ with U, for sufficiently large ℓ and U as defined right after the definition of V in (26). Recall that r_k is the number of components of T_{2k-1} or, equivalently, the number of rooms of $D_k(B)$. These rooms are organized hierarchically into h_k houses, b_k blocks, and c_k compounds. Hence, $r_1 \ge h_1 \ge b_1 \ge c_1 \ge r_2$, etc. This implies that there are

- $r_1 - h_1$ edges of hexagonal length 2 and Euclidean length less than 2 that connect the rooms pairwise inside the h_1 houses;
- $h_1 - b_1$ edges of hexagonal and Euclidean length 2 that connect the houses pairwise inside the b_1 blocks;
- $b_1 - c_1$ edges of hexagonal length 3 and Euclidean length less than 3 that connect the blocks pairwise inside the c_1 compounds;
- $c_1 - r_2$ edges of hexagonal and Euclidean length 3 that connect the compounds pairwise inside the r_2 rooms of $D_2(B)$, etc.

The costs for these edges are w_1, x_1, y_1, z_1, respectively. Setting $z_0 = 0$, and generalizing to $k \ge 1$, we observe that the total cost satisfies

$$\mathrm{cost} \le \sum\nolimits_{k\ge 1} [w_k(r_k - h_k) + x_k(h_k - b_k) + y_k(b_k - c_k) + z_k(c_k - r_{k+1})] \tag{35}$$

$$= \sum\nolimits_{k\ge 1} [(w_k - z_{k-1})r_k + (x_k - w_k)h_k + (y_k - x_k)b_k + (z_k - y_k)c_k] \tag{36}$$

$$\le [2r_1 + h_1 + 3b_1 + c_1 - 7] + \sum\nolimits_{k\ge 2} [3r_k + h_k + 3b_k + c_k - 8]. \tag{37}$$

To see how (37) derives from (36), we first make the sums finite by letting ℓ be the smallest integer such that $r_{\ell+1} = 1$. Then the last non-zero term in (35) is $z_\ell(c_\ell - r_{\ell+1})$ and, correspondingly, the last term in (36) is $z_\ell r_{\ell+1} = z_\ell$, which by (32) is equal to $8\ell - 1$. But this is the same as the sum of constants in (37). Furthermore, we note that if $r_k = h_k = b_k = c_k = 1$, for every k, then (36) vanishes because (35) vanishes, and (37) vanishes because for any k the corresponding sum of four terms minus the constant vanishes. Hence, the difference between (37) and (36) vanishes. To prove the inequality, we reintroduce the variables, which satisfy $r_1 \ge h_1 \ge \ldots \ge c_\ell$, and look at their coefficients. The first is $2 - w_1 + z_0$, which is positive because $w_1 < 2$ and $z_0 = 0$. Indeed, using the inequalities in (33) and (34), we observe that the coefficients alternate between positive and negative. For example, $3 - w_k + z_{k-1}$ is positive because $w_k - z_{k-1} < 3$, and $1 - x_k + w_k$ is negative because $x_k - w_k > 1$. This implies that the difference is non-negative, so (37) follows.

The difficult cases are the edges of hexagonal lengths 2 and 3. We therefore consider the special cases in which all edges in $\mathrm{MST_{hex}}(B)$ have Euclidean length at most $\sqrt{3}, \sqrt{4}, \sqrt{7}, \sqrt{9}$, so $h_1 = 1, b_1 = 1, c_1 = 1, r_2 = 1$, respectively; see Figure 3. From (37), we get

$$\text{cost} \le \begin{cases} 2r_1 - 2 & \text{if } r_1 > h_1 = 1; \\ 2r_1 + h_1 - 3 & \text{if } h_1 > b_1 = 1; \\ 2r_1 + h_1 + 3b_1 - 6 & \text{if } b_1 > c_1 = 1; \\ 2r_1 + h_1 + 3b_1 + c_1 - 7 & \text{if } c_1 > r_2 = 1. \end{cases} \tag{38}$$

The cost needs to be paid from the credit contributed by the (short) edges in U, which in these four cases is either U_1 or U_1'. Recall that after the conversion, each short edge contributes one Euro of credit, so Lemma 4.5 provides lower bounds:

$$\text{credit} \ge \begin{cases} 2r_1 - 1 & \text{if } r_1 > h_1 = 1; \\ 2r_1 + h_1 - 2 & \text{if } h_1 > b_1 = 1; \\ 2r_1 + h_1 + 3b_1 - 5 & \text{if } b_1 > c_1 = 1; \\ 2r_1 + h_1 + 3b_1 + (s_1 + d_1 + e_1 + t_1) - 5 & \text{if } c_1 > r_2 = 1. \end{cases} \tag{39}$$

Comparing (39) with (38), we get cost $\le$ credit trivially in the first three cases. Using Claim (iv), we get use $s_1 + d_1 + e_1 + t_1 \ge c_1 \ge (s_1 + 2d_1 + 2e_1 + 3t_1) - r_2 = c_1$, which supports the same in the fourth case. To compare the cost with the credit in the remaining cases, we use Lemmas 4.5 and 4.6 to compute a lower bound for the latter, assuming that $\ell > 1$ is the smallest integer for which $r_{\ell+1} = 1$:

$$\begin{aligned} \text{credit} &\ge \#U_1 + \sum\nolimits_{k=2}^{\ell-1} \#U_k + \#U_\ell' && (40) \\ &\ge \left[2r_1 + h_1 + 3b_1 + (\tfrac{1}{3}s_1 + \tfrac{2}{3}d_1 + e_1 + t_1) - r_2 - 4\right] \\ &\quad + \sum\nolimits_{k=2}^{\ell-1} \left[3r_k + 4h_k + 3b_k + (\tfrac{1}{3}s_k + \tfrac{2}{3}d_k + e_k + t_k) - r_{k+1} - 4\right] \\ &\quad + \left[3r_\ell + 4h_\ell + 3b_\ell + (s_\ell + d_\ell + e_\ell + t_\ell) - 5\right], && (41) \end{aligned}$$

in which we group the terms with index $k-1$ that appear in the bounds for $\#U_k$ and $\#U_k'$ with the terms that have the same index. Using the extension of Claim (iv) to $k \ge 1$ stated in (25), we get $\frac{1}{3}s_k + \frac{2}{3}d_k + e_k + t_k \ge \frac{1}{3}(s_k + 2d_k + 2e_k + 3t_k) = c_k + r_{k+1}$, so the lower bound in (41) exceeds the upper bound in (37). Hence, cost $\le$ credit. In other words, the average Euclidean length of the edges in $\mathrm{MST_{hex}}(B)$ and $\mathrm{MST_{hex}}(\Delta \setminus B)$ is at most $\frac{5}{4}$. It follows that their total Euclidean length is at most $\frac{5}{4}(n^2 - 2)$, which by (19) implies the same for $\mathrm{MST}(B)$ and $\mathrm{MST}(\Delta \setminus B)$. ◀

By Lemma 4.7, the average Euclidean length of the edges in $\mathrm{MST}(B)$ and $\mathrm{MST}(\Delta \setminus B)$ is less than $\frac{5}{4}$. Together with (15), this implies Theorem 4.2.

5 Discussion

This paper proves bounds on the supremum and infimum maximum MST-ratio for finite sets in the plane as well as for lattices in the plane. There are many directions of generalization, and their connection to the topological analysis of colored point sets started in [6] provides a potential path to relevance outside of mathematics.

- What about sets in the plane that are less restrictive than lattices but still disallow arbitrarily dense clusters of points, such as periodic sets or Delone sets? A first result in this direction is the lower bound of $1 + 1/(11(2c+1)^2)$ for the maximum MST-ratio of a set of n points with spread at most $c\sqrt{n}$.

- What about partitions of $A \subseteq \mathbb{R}^2$ into three or more sets? For example, is it true that the maximum MST-ratio of the hexagonal lattice partitioned into three subsets is $\sqrt{3}$, as realized by the unique partition into three congruent hexagonal grids? Is $\sqrt{3}$ the infimum, over all lattices in $\mathbb{R}^2$, of the maximum, over all partitions into three subsets?
- What about three and higher dimensions? Consider for example the FCC lattice in $\mathbb{R}^3$ (all integer points whose sums of coordinates are even), and partition it into 2FCC and the rest. The MST-ratio of this example is $\frac{9}{8} = 1.125$. Is it true that this is the maximum MST-ratio of the FCC lattice? Is 1.125 the infimum, over all lattices in $\mathbb{R}^3$, of the maximum, over all partitions into two subsets?

Beyond these extensions in discrete geometry, it would be interesting to study the MST-ratio stochastically, to determine the computational complexity of the maximum MST-ratio, and to frame notions of mingling as measured by homology classes of dimension 1 and higher in elementary geometric terms.

References

1 O. Borůvka. O jistém problému minimálním. *Práce Mor. Přírodověd Spol. v Brně (Acta Societ. Scient. Natur Moravicae)* **3** (1926), 37–58.

2 G. Carlsson. Topology and data. *Bull. Amer. Math. Soc.* **46** (2009), 255–308.

3 F.R.K. Chung and R.L. Graham. A new bound for Euclidean Steiner minimal trees. In *Discrete Geometry and Convexity*, Annals N.Y. Acad. Sci. **440**, New York, 1985, 328–346.

4 D. Cohen-Steiner, H. Edelsbrunner and J. Harer. Stability of persistence diagrams. *Discrete Comput. Geom.* **37** (2007), 103–120.

5 D. Cohen-Steiner, H. Edelsbrunner, J. Harer and D. Morozov. Persistent homology for kernels, images, and cokernels. In "Proc. 20th Ann. ACM-SIAM Sympos. Discrete Alg., 2009", 1011–1020.

6 S. Cultrera di Montesano, O. Draganov, H. Edelsbrunner and M. Saghafian. Chromatic alpha complexes. `arXiv:2212.03128 [math.AT]`, 2023.

7 A. Dumitrescu, J. Pach and G. Tóth. Two trees are better than one. `arXiv:2312.09916 [math.CO]`, 2023.

8 H. Edelsbrunner and J.L. Harer. *Computational Topology. An Introduction.* Amer. Math. Soc., Providence, Rhode Island, 2010.

9 E.N. Gilbert and H.O. Pollack. Steiner minimal trees. *SIAM J. Appl. Math.* **16** (1968), 1–29.

10 V. Jarník and M. Kössler. O minimálních grafech, obsahujících n daných bodů. *Časopis pro Pěstováni Matematiky a Fysiky* **63** (1934), 223–235.

11 J.B. Kruskal. On the shortest spanning tree of a graph and the traveling salesman problem. *Prof. Amer. Math. Soc.* **7** (1956), 48–50.

12 B. Zhilinskii. *Introduction to Louis Michel's Lattice Geometry through Group Action.* EDP Sciences, ENRS Editions, Paris, France, 2015.

A Connection to Chromatic Persistence

As mentioned in the introduction, the study of the MST-ratio is motivated by a recent topological data analysis method for measuring the "mingling" of points in a colored configuration; see Figure 9, which shows six persistence diagrams measuring various aspects of the mingling in a bi-colored configuration. This appendix addresses the meaning of some of these diagrams and explains the connection to the MST-ratio, while referring to [6] for a detailed account of the method. In particular, we short-cut the description by ignoring the discrete structures that are necessary for the algorithm. We first sketch the general background from [8] and [5], and then explain the specific setting that motivates the MST-ratio.

Let $A \subseteq \mathbb{R}^2$ be a finite set of points, $\chi \colon A \to \{0,1\}$ a bi-coloring, and write $B = \chi^{-1}(0)$ and $C = A \setminus B = \chi^{-1}(1)$. Let $\mathbf{a} \colon \mathbb{R}^2 \to \mathbb{R}$ be the function that maps every $x \in \mathbb{R}^2$ to the minimum Euclidean distance between x and the points in A, and let $\mathbf{b} \colon \mathbb{R}^2 \to \mathbb{R}$ and $\mathbf{c} \colon \mathbb{R}^2 \to \mathbb{R}$ be the similarly defined functions for B and C. Furthermore, write $A_r = \mathbf{a}^{-1}[0, r]$,

$B_r = \mathbf{b}^{-1}[0, r]$, and $C_r = \mathbf{c}^{-1}[0, r]$ for the sublevel sets at distance threshold $r \geq 0$. Each is a union of disks with radius r centered at the points of A, B, and C, respectively. The inclusions $B_r \subseteq A_r$ and $C_r \subseteq A_r$ induce homomorphisms in p-th homology, $b_r \colon H_p(B_r) \to H_p(A_r)$ and $c_r \colon H_p(C_r) \to H_p(A_r)$, for each dimension $p \in \mathbb{Z}$ and every threshold $r \geq 0$. Assuming field coefficients in the construction of the homology groups, the latter are vector spaces and the homomorphisms are linear maps.

We also have $A_r \subseteq A_s$ whenever $r \leq s$, so there are also linear maps from $H_p(A_r)$ to $H_p(A_s)$. By now it is tradition in the field to consider the *filtration* of the A_r, for r from 0 to ∞, and the corresponding sequence of homology groups together with the linear maps between them. Reading this sequence from left to right, we see homology classes being born and dying. There is a unique way to pair the births with the deaths that regards the identity of the classes, and the *persistence diagram* summarizes this information by drawing a point $(r, s) \in \mathbb{R}^2$ for every homology class that is born at A_r and dies entering A_s; see e.g. [8, Chapter VII]. Every death is paired with a birth, but it is possible that a birth remains unpaired—when the homology class is of the domain—in which case the corresponding point is at infinity. We write $\mathrm{Dgm}_p(\mathbf{a})$ for the persistence diagram defined by the sublevel sets of $\mathbf{a}$, noting that it is a multi-set of points vertically above the diagonal.

Besides $\mathrm{Dgm}_p(\mathbf{a})$, we consider $\mathrm{Dgm}_p(\mathbf{b})$ and $\mathrm{Dgm}_p(\mathbf{c})$, which are the persistence diagrams of the sublevel sets of $\mathbf{b}$ and $\mathbf{c}$, respectively, and work with the disjoint union, $B_r \sqcup C_r$. Conveniently, the p-th persistence diagram of $\mathbf{b} \sqcup \mathbf{c} \colon \mathbb{R}^2 \sqcup \mathbb{R}^2 \to \mathbb{R}$ is the disjoint union of $\mathrm{Dgm}_p(\mathbf{b})$ and $\mathrm{Dgm}_p(\mathbf{c})$, for all p. Write $b_r \oplus c_r \colon H_p(B_r) \oplus H_p(C_r) \to H_p(A_r)$ for the corresponding map in homology. As proved in [5], the sequence of images of the $b_r \oplus c_r$ admit linear maps between them and thus define another persistence diagram, denoted $\mathrm{Dgm}_p(\mathrm{im}\ \mathbf{b} \sqcup \mathbf{c} \to \mathbf{a})$. Similarly, the kernels of the $b_r \oplus c_r$ define a persistence diagram, denoted $\mathrm{Dgm}_p(\ker \mathbf{b} \sqcup \mathbf{c} \to \mathbf{a})$. To simplify the notation, we write $\kappa_r = b_r \oplus c_r$ and use mnemonic notation to indicate whether a persistence diagram belongs to the domain, image, or kernel of the map:

$$\mathrm{Dom}_p(\kappa) = \mathrm{Dgm}_p(\mathbf{b} \sqcup \mathbf{c}), \tag{42}$$

$$\mathrm{Im}_p(\kappa) = \mathrm{Dgm}_p(\mathrm{im}\ \mathbf{b} \sqcup \mathbf{c} \to \mathbf{a}), \tag{43}$$

$$\mathrm{Ker}_p(\kappa) = \mathrm{Dgm}_p(\ker \mathbf{b} \sqcup \mathbf{c} \to \mathbf{a}). \tag{44}$$

The 1-*norm* of a persistence diagram, D, is the sum of the absolute differences between birth- and death-coordinates over all points in D, denoted $\|D\|_1$. To cope with points at infinity, we use a cut-off—e.g. the maximum finite homological critical value, denoted ω_0—so that the contribution of a point at infinity to the 1-norm is finite.

The kernel, domain, and image form a short exact sequence that splits, which implies $\|\mathrm{Im}_p(\kappa)\|_1 + \|\mathrm{Ker}_p(\kappa)\|_1 = \|\mathrm{Dom}_p(\kappa)\|_1$; see [6, Theorem 5.3]. For dimension $p = 0$, all three 1-norms can be rewritten in terms of minimum spanning trees. Indeed, $\|\mathrm{Dgm}_0(\mathbf{b})\|_1 = \frac{1}{2}|\mathrm{MST}(B)| + \omega_0$ because every edge in the minimum spanning tree of B marks the death of a connected component in the sublevel set, and ω_0 is contributed by the one component that never dies. Similarly, $\|\mathrm{Dgm}_0(\mathbf{c})\|_1 = \frac{1}{2}|\mathrm{MST}(C)| + \omega_0$, which implies (45):

$$\|\mathrm{Dom}_0(\kappa)\|_1 = \|\mathrm{Dgm}_0(\mathbf{b})\|_1 + \|\mathrm{Dgm}_0(\mathbf{c})\|_1 = \tfrac{1}{2}|\mathrm{MST}(B)| + \tfrac{1}{2}|\mathrm{MST}(C)| + 2\omega_0; \tag{45}$$

$$\|\mathrm{Im}_0(\kappa)\|_1 = \tfrac{1}{2}|\mathrm{MST}(A)| + \omega_0. \tag{46}$$

Since persistence diagrams are stable, as originally proved in [4], these relations imply that minimum spanning trees are similarly stable. (46) deserves a proof. There are two ways a connected component of B_r can die in the image: by merging with a component of C_r

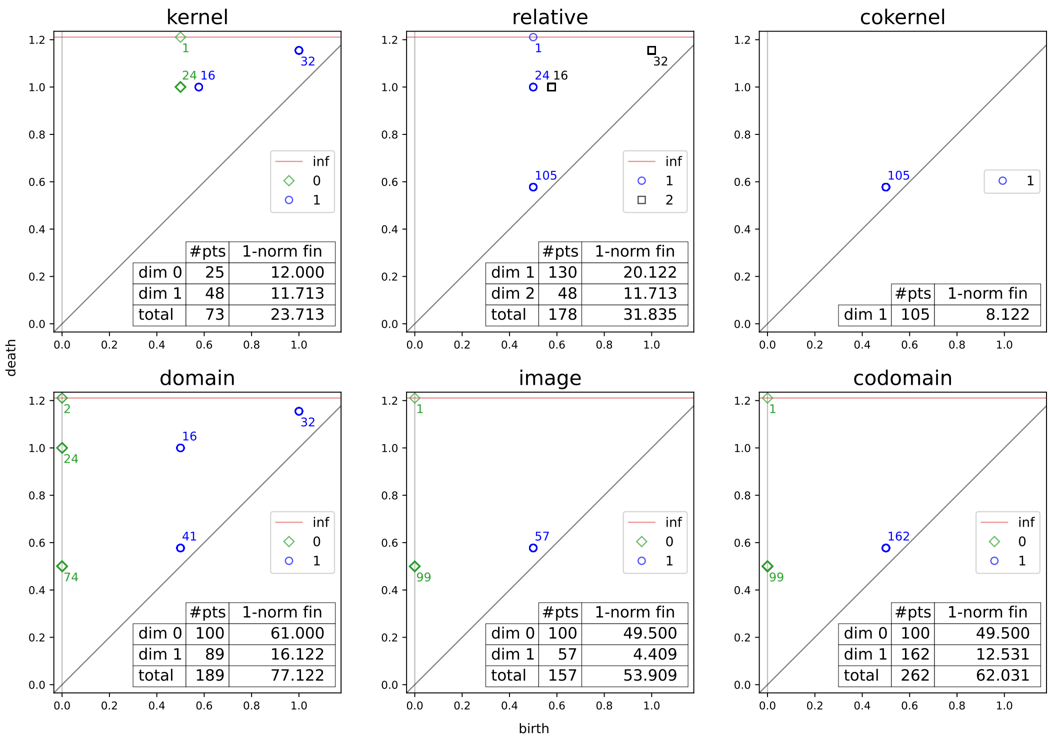

Figure 9: The six-pack for the 10×10 portion of the hexagonal lattice with coloring as in Figure 4. Important for the current discussion are the *diamond-shaped* points in the domain, image, and kernel diagrams. To get the MST-ratio, the 1-norms of the diagrams are computed while ignoring the points at infinity, giving giving 61.0 and 49.5 for the domain and the image diagrams, respectively. Compare the ratio of $1.232\ldots$ with the upper bound of 1.25 proved in Theorem 4.2.

or with another component of B_r. In the first case, the death corresponds to an edge of $\mathrm{MST}(A)$ that connects a point in B with a point in C, and in the second case, it corresponds to an edge of $\mathrm{MST}(A)$ that connects two points in B. There is also the symmetric case in which the edge connects two points in C. This establishes a bijection between the deaths in $\mathrm{Im}_0(\kappa)$ and the edges of $\mathrm{MST}(A)$. There is one component that never dies, which accounts for the extra cut-off term and implies (46).

The 1-norm of the kernel diagram is the difference between the 1-norms of the domain diagram and the image diagram: $\|\mathrm{Ker}_0(\kappa)\|_1 = \|\mathrm{Dom}_0(\kappa)\|_1 - \|\mathrm{Im}_0(\kappa)\|_1$. It thus makes sense to call $\|\mathrm{Im}_0(\kappa)\|_1 / \|\mathrm{Dom}_0(\kappa)\|_1$ and $\|\mathrm{Ker}_0(\kappa)\|_1 / \|\mathrm{Dom}_0(\kappa)\|_1$ the *image share* and *kernel share*, respectively. Observe that both are real numbers between 0.0 and 1.0 and that they add up to 1.0. The intuition is that the kernel share is a measure of the amount of "0-dimensional mingling" of B and C. In other words, the smaller the image share, the more the two colors mingle. We therefore get

$$\mu(A, B) = \frac{|\mathrm{MST}(B)| + |\mathrm{MST}(C)|}{|\mathrm{MST}(A)|} = \frac{\|\mathrm{Dom}_0(\kappa)\|_1 - 2\omega_0}{\|\mathrm{Im}_0(\kappa)\|_1 - \omega_0}, \tag{47}$$

for the MST-ratio, which besides the cut-off terms is the reciprocal of the image share. Hence, the larger the MST-ratio the more the two colors mingle. In this interpretation, Theorem 3.1 says that among all lattices in $\mathbb{R}^2$, the hexagonal lattice is most restrictive to mingling as it does not permit MST-ratios larger than the inf-max, which for 2-dimensional lattices is 1.25.

Acknowledgments

Part of the research described in this paper was conducted during the sabbatical of the third author at the University of Kyoto in Japan. We thank the Institute of Advanced Study in Kyoto for the hospitality afforded to the authors of this paper.