Geometric Properties of Poisson Matchings

Alexander E. Holroyd

1 September 2009

Dedicated to Oded Schramm, 10 December 1961 – 1 September 2008

Abstract

Suppose that red and blue points occur as independent Poisson processes of equal intensity in $\mathbb{R}^d$, and that the red points are matched to the blue points via straight edges in a translation-invariant way. We address several closely related properties of such matchings. We prove that there exist matchings that locally minimize total edge length in $d = 1$ and $d \geq 3$, but not in the strip $\mathbb{R} \times [0, 1]$. We prove that there exist matchings in which every bounded set intersects only finitely many edges in $d \geq 2$, but not in $d = 1$ or in the strip. It is unknown whether there exists a matching with no crossings in $d = 2$, but we prove positive answers to various relaxations of this question. Several open problems are presented.

1 Introduction

Let $\mathcal{R}$ and $\mathcal{B}$ be simple point processes in $\mathbb{R}^d$. The **support** of a point process Π is the random set $[\Pi] := \{x : \Pi(\{x\}) = 1\}$; the elements of $[\Pi]$ are called **Π-points**. We call $\mathcal{R}$-points **red points** and $\mathcal{B}$-points **blue points**. A (perfect, two-color) **matching scheme** of $\mathcal{R}$ and $\mathcal{B}$ is a simple point process $\mathcal{M}$ in $(\mathbb{R}^d)^2$ such that almost surely $(V, E) = ([\mathcal{R}] \cup [\mathcal{B}], [\mathcal{M}])$ is a perfect matching of $[\mathcal{R}]$ to $[\mathcal{B}]$ (i.e. a bipartite graph with vertex classes $[\mathcal{R}], [\mathcal{B}]$ and all degrees 1). We call $\mathcal{M}$-points **edges**. Similarly, we say that $\mathcal{M}$ is a **partial** matching

Key words: Poisson process, point process, matching
AMS 2010 Mathematics Subject Classifications: 60D05, 60G55, 05C70
Funded in part by Microsoft and NSERC

scheme if all degrees are at most 1; and $\mathcal{M}$ is a **one-color** matching scheme of $\mathcal{R}$ if $([\mathcal{R}], [\mathcal{M}])$ is a.s. a simple undirected graph will all degrees 1. A matching scheme $\mathcal{M}$ is **translation-invariant** if the law of $(\mathcal{R}, \mathcal{B}, \mathcal{M})$ is invariant under all translations of $\mathbb{R}^d$. We also consider matchings of point processes on the **strip** $\mathbb{R} \times [0, 1)$, in which case translation-invariance refers to all translations in the first coordinate direction.

In $\mathbb{R}^2$ or the strip, we call a matching scheme $\mathcal{M}$ **planar** if a.s. for any distinct matched pairs $(r, b), (r', b') \in [\mathcal{M}]$, the two closed line segments joining r to b and r' to b' do not intersect.

We focus on the case where $\mathcal{R}$ and $\mathcal{B}$ are independent Poisson processes of equal intensity. It is unknown whether there exists a translation-invariant planar matching scheme in $\mathbb{R}^2$ (see the discussion on open problems below), but the following natural variations of this question all lead to positive answers.

Theorem 1 (Planarity)**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1 in $\mathbb{R}^2$. The following all exist.*

(i) *A planar matching scheme of $\mathcal{R}$ to $\mathcal{B}$ that is* not *translation-invariant.*
(ii) *A translation-invariant planar one-color matching scheme of $\mathcal{R}$.*
(iii) *A translation-invariant planar* partial *matching scheme of $\mathcal{R}'$ to $\mathcal{B}$, in which every blue point is matched, where $\mathcal{R}', \mathcal{B}$ are independent Poisson processes of intensities $\lambda, 1$, for any $\lambda > 1$.*
(iv) *A translation-invariant process consisting of a matching of $\mathcal{R}$ to $\mathcal{B}$, together with a polygonal arc in $\mathbb{R}^2$ joining each pair of matched points, such that the arcs do not intersect.*

Furthermore, each of (i)–(iv) exists on the strip.

The following concept has close connections with planarity. We call a matching scheme $\mathcal{M}$ **minimal** if a.s. every finite set of edges is matched in a way that minimizes total edge length, i.e. for any $\{(r_i, b_i)\}_{i=1,\dots,n} \subset [\mathcal{M}]$ (where $r_i \in [\mathcal{R}]$ and $b_i \in [\mathcal{B}]$ for each i), we have

$$\sum_i |r_i - b_i| = \min_\sigma \sum_i |r_i - b_{\sigma(i)}|,$$

where the minimum is over all permutations σ of $1, \dots, n$, and $|\cdot|$ denotes the Euclidean norm. An elementary argument (see the discussion below) shows that any minimal matching scheme in $\mathbb{R}^2$ is planar. However, the concept of minimality is natural in all dimensions, and we have the following surprising facts.

Theorem 2 (Minimality)**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1.*

(i) *There exists a translation-invariant minimal matching of $\mathcal{R}$ to $\mathcal{B}$ in $\mathbb{R}^d$ for $d = 1$ and all $d \geq 3$.*
(ii) *There does not exist a translation-invariant minimal matching of $\mathcal{R}$ to $\mathcal{B}$ in the strip.*

The remaining case of $\mathbb{R}^2$ is open.

Next, we say that an edge $(r, b) \in [\mathcal{M}]$ **crosses** a set $S \subset \mathbb{R}^d$ if the closed line segment from r to b intersects S. We call a matching scheme $\mathcal{M}$ **locally finite** if a.s., every bounded set $S \subset \mathbb{R}^d$ is crossed by only finitely many edges.

Theorem 3 (Local finiteness)**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1.*

(i) *There does not exist a translation-invariant locally finite matching scheme in $\mathbb{R}$, nor in the strip.*
(ii) *There exist translation-invariant locally finite matching schemes in $\mathbb{R}^d$ for all $d \geq 2$.*

Finally, we establish the following conditional result.

Theorem 4 (Minimal implies locally finite)**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1 in $\mathbb{R}^d$, where $d \geq 2$. Any translation-invariant minimal matching scheme must be locally finite.*

Note that Theorems 2(i) and 3(i) show that the assertion of Theorem 4 fails in $d = 1$.

Motivation and open problems

Invariant Poisson matching schemes, and particularly their quantitative properties, were studied extensively in [3]. Other related work appears in [1, 4, 6, 8, 9]. The present work is largely motivated by the following question, which was posed by Yuval Peres in 2002, stated in [3], and remains open.

Question 1. *For $\mathcal{R}$ and $\mathcal{B}$ independent Poisson processes of intensity 1 in $\mathbb{R}^2$, does there exist a translation-invariant planar matching scheme?*

It is far from clear what answer to guess to the above question, with several of the results presented here perhaps suggesting opposite answers. We propose the following natural variants.

Question 2. *For $\mathcal{R}$ and $\mathcal{B}$ independent Poisson processes of intensity 1 in the strip, does there exist a translation-invariant planar matching scheme?*

Question 3. *For $\mathcal{R}$ and $\mathcal{B}$ independent Poisson processes of intensity 1 in $\mathbb{R}^2$, does there exist a translation-invariant minimal matching scheme?*

Note on the other hand that Theorem 2(ii) gives a negative answer to the analogous question of minimal matchings in the strip. The following simple consequence of the triangle inequality implies immediately that a minimal matching in $\mathbb{R}^2$ is planar, so a positive answer to Question 3 would imply a positive answer to Question 1. (A positive answer to Question 2 would also imply a positive answer to Question 1, by a simple argument – see Section 3). We say that a set of points $K \subset \mathbb{R}^d$ is **parallel-free** if there do not exist $x, y, u, v \in K$ and $a \in \mathbb{R}$ with $\{x, y\} \neq \{u, v\}$ and $x - y = a(u - v) \neq 0$.

Observation 5 (Finite minimum matchings are planar)**.** *Let $R, B \in \mathbb{R}^2$ be disjoint finite sets of equal cardinality, and suppose $R \cup B$ is parallel-free. Then in any perfect matching of R to B that minimizes the total length, the line segments joining matched pairs do not intersect.*

In the light of this observation, the following possible approach to constructing a translation-invariant planar matching seems natural. Take n red and n blue points uniformly at random in a square of area n, randomly translated so that the origin is uniformly distributed in the square. Consider the matching of minimum total length, and take suitable a limit in distribution as $n \to \infty$. For such an approach to be successful, the limit must be a genuine matching – it is possible that instead the partner of a point goes to infinity. If it exists, the limiting matching would be minimal. This motivates Question 3. We will employ a somewhat similar limiting argument in the proof of Theorem 2(i) in $d \geq 3$.

Question 3 remains open if the invariance requirement is dropped.

Question 4. *For $\mathcal{R}$ and $\mathcal{B}$ independent Poisson processes of intensity 1 in $\mathbb{R}^2$, does there exist a minimal matching scheme?*

The notion of locally finite matching (see Theorem 3) becomes particularly interesting in $\mathbb{R}^2$, owing to the following result proved in [3, Proof of Theorem 2, case $d = 2$].

Theorem 6 (Infinite mean crossings; [3])**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1 in $\mathbb{R}^2$. In any translation-invariant matching scheme, for any fixed bounded set $S \subset \mathbb{R}^2$ that contains some line segment, the number of edges that cross S has infinite expectation.*

Notwithstanding Theorem 3(ii), one might speculate that if a *minimal* matching scheme exists in $\mathbb{R}^2$, the infinite expectation in Proposition 6 should in some sense be "spread around evenly", so that the matching is not locally finite. Combined with Theorem 4, this perhaps suggests a negative answer to Question 3.

Returning to the issue of planarity, we propose the following question.

Question 5. *Do there exist jointly ergodic point processes $\mathcal{R}$ and $\mathcal{B}$ in $\mathbb{R}^2$, both of intensity 1, for which there is (provably) no planar translation-invariant matching scheme?*

Finally, note that our definition of a matching scheme requires only that $\mathcal{R}, \mathcal{B}, \mathcal{M}$ are all defined on some joint probability space, so the matching may involve additional randomization besides that of the red and blue processes. If instead $\mathcal{M} = f(\mathcal{R}, \mathcal{B})$ for some deterministic function f, the matching scheme is called a **factor**. See e.g. [3] for more on this distinction. Most of the matching schemes we construct will not be factors. Another interesting line of enquiry (which we do not pursue here) is to determine whether there exist factor matching schemes satisfying the various conditions under consideration.

Some notation

We write $\mathcal{L}$ for Lebesgue measure on $\mathbb{R}^d$, and $|\cdot|$ for the Euclidean norm. The ball is denoted $B(r) := \{x \in \mathbb{R}^d : |x| < r\}$. If $\mathcal{M}$ is a matching scheme, we write $\mathcal{M}(x)$ for the **partner** of a red or blue point x, i.e. the unique point such that $(x, \mathcal{M}(x)) \in [\mathcal{M}]$. Similarly in a deterministic matching m we write $m(x)$ for the partner of x.

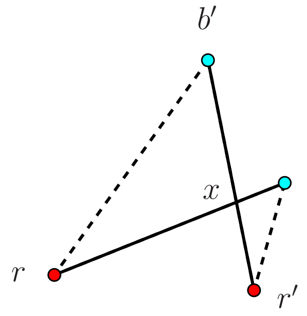

Figure 1: Uncrossing a pair of edges: the triangle inequality gives $rb' + r'b < rx + xb' + r'x + xb = rb + r'b'$.

2 Finite matchings

In this section we verify some elementary facts.

Proof of Observation 5. Suppose on the contrary that, in some length-minimizing matching, two edges intersect. By the parallel-free assumption, they must intersect non-trivially, i.e. in a single point that is not one of their endpoints. But now an application of the triangle inequality shows that the other possible matching of these four points has strictly smaller total length; see Figure 1. ☐

Lemma 7. *In a homogeneous Poisson process* Π *in* $\mathbb{R}^d$ *with* $d \geq 2$*,* $[\Pi]$ *is a.s. parallel-free.*

Proof. It is enough to check this for the Poisson process restricted to a ball, and in this case we may also condition on the number of points in the ball. So it suffices to check that if $x_1, \ldots, x_4$ are independent uniform points in a fixed ball, then a.s. the vectors $x_1 - x_2$, $x_1 - x_3$ and $x_3 - x_4$ are pairwise non-parallel, which is elementary. ☐

As remarked earlier we can deduce the following.

Corollary 8 (Minimal implies planar)**.** *For* $\mathcal{R}$ *and* $\mathcal{B}$ *independent Poisson processes of intensity 1 in* $\mathbb{R}^2$*, any minimal matching scheme (if such exists) is planar.*

Proof. This is immediate from Observation 5 and Lemma 7; in fact we need the minimality property only for sets of two edges. ☐

In order to use Observation 5, we will often wish to consider the perfect matching of minimum total length between two finite sets of points $R, B \subset \mathbb{R}^d$. If $d \geq 2$ and the points are any subset of the points of a Poisson process, one may show that such a minimum matching is a.s. unique. Formally, this fact is not needed, because we can choose among minimum-length matchings according to some fixed rule, such as the earliest in lexicographic order with respect to the coordinates of the points.

3 Matchings in strips and lines

In this section we prove Theorem 1. Each part will be proved in the strip, and the case of $\mathbb{R}^2$ then follows as an easy consequence (see below). We also prove Theorem 2(i) in the case $d = 1$.

Proof of Theorem 1, $\mathbb{R}^2$ case. Divide $\mathbb{R}^2$ into the disjoint strips

$$\mathbb{R} \times [i, i+1), \quad i \in \mathbb{Z}.$$

Within each strip, take an independent copy of the matching scheme (together with the associated red and blue points) on the strip from the appropriate part (i)–(iv) of the theorem (see the proofs below). This yields a matching scheme $\mathcal{M}$ in $\mathbb{R}^2$ which inherits the appropriate properties of the matching on the strip and is invariant under translations in $\mathbb{R} \times \mathbb{Z}$ (for (ii)–(iv)). To achieve full translation-invariance in $\mathbb{R}^2$ (for (ii)–(iv)), let U be a uniform random variable in $[0, 1)$, independent of $(\mathcal{R}, \mathcal{B}, \mathcal{M})$, and translate the entire process $(\mathcal{R}, \mathcal{B}, \mathcal{M})$ by the vector $(0, U)$. $\square$

Remark. By the above argument, a positive answer to Question 2 would imply a positive answer to Question 1. $\Diamond$

We will make frequent use of the following object. Given $\mathcal{R}, \mathcal{B}$ in the strip, define a function $F : \mathbb{R} \to \mathbb{Z}$ by

$$\begin{aligned} F(0) &= 0; \\ F(y) - F(x) &= (\mathcal{R} - \mathcal{B})\big((x, y] \times [0, 1)\big), \quad x < y. \end{aligned} \tag{1}$$

Thus F is a right-continuous continuous-time simple symmetric random walk on the integers, with its up-steps and down-steps corresponding to red and blue points respectively. See Figure 2.

Proof of Theorem 1(i), strip case. With F as in (1), define

$$Z := \{z \in \mathbb{R} : F(z) = 0 \text{ and } F(z-) \neq 0\};$$

i.e. the set of left endpoints of the intervals where F is zero. Almost surely, Z is a discrete set, and because F is a recurrent random walk, Z is unbounded in both the positive and negative directions. If $z_1 < z_2$ are two consecutive elements of Z, then the rectangle $(z_1, z_2] \times [0, 1)$ contains equal numbers of red and blue points. Therefore, within each such rectangle, take the matching of minimum total edge length, and appeal to Observation 5 and Lemma 7. $\square$

Remark. The above construction cannot be adapted to give a translation-invariant matching on the strip, because the random walk F is null-recurrent, therefore Z has zero density. $\diamondsuit$

Proof of Theorem 1(ii), strip case. Let $(r_i)_{i\in\mathbb{Z}} = ((x_i, y_i))_{i\in\mathbb{Z}}$ be the points of $[\mathcal{R}]$, ordered so that their first coordinates are in increasing order (i.e. $x_i < x_{i+1}\ \forall i$), and so that r_0 is the first point to the right of the origin (so $x_{-1} < 0 < x_0$). Conditional on $\mathcal{R}$, choose one of the two matchings

$$\{\ldots, (r_{-2}, r_{-1}), (r_0, r_1), (r_2, r_3), \ldots\}; \quad \{\ldots, (r_{-1}, r_0), (r_1, r_2), (r_3, r_4), \ldots\};$$

each with probability 1/2. $\square$

Proof of Theorem 1(iii), strip case. Let F be as in (1), but with $\mathcal{R}'$ in place of $\mathcal{R}$. Now F is a biased random walk with positive drift. Hence the set of **cut-times** given by

$$C := \left\{x \in \mathbb{R} : \sup_{t<x} F(t) = F(x-) < F(x) = \inf_{t\geq x} F(t)\right\}$$

forms a translation-invariant ergodic point process of positive intensity in $\mathbb{R}$. If $c_1 < c_2$ are two consecutive cut-times, then the rectangle $(c_1, c_2] \times [0, 1)$ contains strictly more red than blue points. Therefore, within each such rectangle, take the matching which minimizes the total edge length from among all possible partial red-blue matchings of maximum cardinality (i.e. those in which all blue points are matched). By Observation 5 and Lemma 7, the resulting matching scheme has the required properties. $\square$

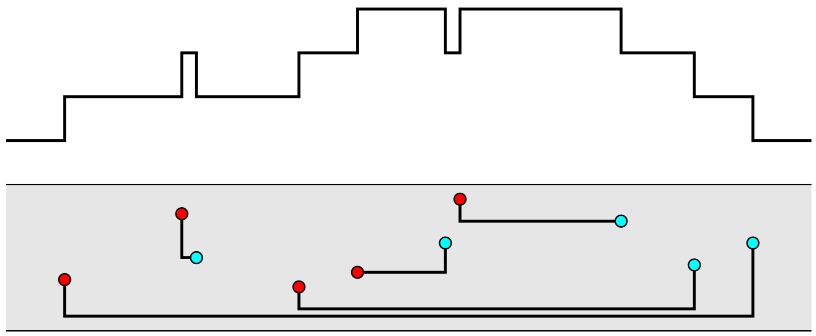

Figure 2: The random walk F associated with red and blue points in the strip, and the matching with disjoint arcs from the proof of Theorem 1(iv).

Proof of Theorem 1(iv), strip case. We first construct the matching scheme $\mathcal{M}$. Given $\mathcal{R}, \mathcal{B}$ on the strip, let F be as in (1). Suppose $r = (r_1, r_2) \in [\mathcal{R}]$ is a red point, so that $F(r_1) = F(r_1-) + 1$. We match r to the blue point $b = (b_1, b_2)$, where

$$b_1 := \inf\{t > r_1 : F(t) = F(r_1-)\}.$$

Thus, b marks the end of F's upward excursion starting at r, or equivalently b is the first point to the right of r such that the red and blue points in the intervening rectangle equalize. It is easy to check that $\mathcal{M}$ is indeed a translation-invariant matching scheme. Furthermore, if $(r, b), (r', b')$ are two distinct edges of the matching then the intervals $[r_1, b_1]$ and $[r_1', b_1']$ are either disjoint or nested one inside the other.

Now we construct the polygonal arcs. If r and b are two matched points, we will join them by a polygonal arc with vertices

$$r = (r_1, r_2), (r_1, H), (b_1, H), (b_1, b_2) = b;$$

we need only choose H smaller (say) than the heights of all the intervening arcs. This is achieved by taking for example $H := L/D$, where

$$L = \min\Big\{y : (x, y) \in ([\mathcal{R}] \cup [\mathcal{B}]) \cap ([r_1, b_1] \times [0, 1))\Big\}$$

is the height of the lowest point between r and b (including r and b), and

$$D = \max_{t \in [r_1, b_1]} F(t) - F(r_1-)$$

is the maximum nesting depth between r and b. $\square$

Remark. The matching $\mathcal{M}$ constructed in the above proof has an interpretation as the unique Gale-Shapley stable matching with preferences based on one-sided horizonal distance, and also as a version of a matching introduced by Meshalkin in the construction of finitary isomorphisms. See [2, 4, 7] for details. The same matching is used our next proof. $\diamondsuit$

Proof of Theorem 2(i), case $d = 1$. Given $\mathcal{R}$ and $\mathcal{B}$ Possion processes on $\mathbb{R}$, define F as in (1) except replacing the rectangle $(x, y] \times [0, 1)$ with the interval $(x, y]$, so F is again a simple symmetric random walk. Construct a matching scheme exactly as in the proof of Theorem 1(iv) by matching the red point r to the blue point

$$b := \inf\{t > r : F(t) = F(r-)\}.$$

We claim that this matching scheme $\mathcal{M}$ is minimal. To prove this, let $(r_1, b_1), \ldots, (r_n, b_n) \in [\mathcal{M}]$ be any finite set of edges of the matching – we must prove that no other matching of these red and blue points has smaller total length. Since any two edges of $\mathcal{M}$ span disjoint or nested intervals, there exists a bounded interval $(u, v) \subset \mathbb{R}^d$ containing $\{r_i, b_i\}_{i=1,\ldots,n}$, and such that every point in (u, v) has its partner in (u, v). (To prove this, let I be the smallest interval containing the original points, then take (u, v) just large enough to contain all the points in I and their partners). Therefore, it suffices to prove the above minimality statement under the assumption that $\{r_i, b_i\}_{i=1,\ldots,n}$ are the *only* red and blue points in (u, v).

For any perfect matching m of $\{r_i\}_{i=1,\ldots,n}$ to $\{b_i\}_{i=1,\ldots,n}$, and any $t \in (u, v)$, define

$$h_m(t) := \#\big\{i : \ r_i \le t \le m(r_i) \text{ or } m(r_i) \le t \le r_i\big\};$$

(i.e. the number of edges that cross t). Note that the total edge-length of the matching m may be expressed thus:

$$\sum_i |r_i - m(r_i)| = \int_u^v h_m(t)\, dt. \tag{2}$$

We claim that, writing $M = \{(r_i, b_i)\}_{i=1,\ldots,n}$ for the restriction of $\mathcal{M}$ to these points, we have $h_M(t) \le h_m(t)$ for all $t \in (u, v) \setminus \{r_i, b_i\}_{i=1,\ldots,n}$. Once this is proved, the required minimality follows from (2). To prove this inequality, first note that for all such t,

$$h_M(t) = F(t) - F(u).$$

Indeed, this holds for any matching in which every edge has the red point to the left of the blue point: as t increases from u to v, the quantity $h_M(t)$ increases by 1 at each red point, and decreases by 1 at each blue point, so it equals the right side. On the other hand, for any m and all such t,

$$h_m(t) \geq F(t) - F(u),$$

because the right side equals the excess of red points minus blue points in $(u, t]$, so at least this many edges must cross t. $\square$

4 Impossibility results

In this section we prove Theorem 3(i), Theorem 2(ii), and Theorem 4.

Proof of Theorem 3(i). The proofs for the strip and the line are nearly identical. We first consider the strip.

Let $\mathcal{M}$ be a translation-invariant matching scheme on the strip. Assume without loss of generality that the matching is ergodic under the group of translations of $\mathbb{R}$ (otherwise consider its ergodic components). We will prove that a.s. infinitely many edges cross the line segment $\{0\} \times [0, 1)$. Suppose on the contrary that for some $k < \infty$, exactly k edges cross $\{0\} \times [0, 1)$ with positive probability. Define the random set

$$S := \big\{x \in \mathbb{R} : \{x\} \times [0, 1) \text{ does not intersect } [\mathcal{R}] \cup [\mathcal{B}], \\ \text{and is crossed by exactly } k \text{ edges}\big\}.$$

Then the above assumptions imply that S is a translation-invariant ergodic random set of positive intensity, say λ, and thus

$$\lim_{n\to\infty} \frac{\mathcal{L}(S \cap [0, n))}{n} = \lambda \quad \text{a.s.} \tag{3}$$

If $s < t$ are any two elements of S, then there are at most $2k$ edges from the rectangle $[s, t) \times [0, 1)$ to its complement, therefore the difference between the numbers of red and blue points in this rectangle is at most $2k$. Hence, with F defined as in (1), $|F(s) - F(t)| \leq 2k$. This implies that there exists some random integer H such that a.s.,

$$F(s) \in [H, H + 2k] \text{ for all } s \in S.$$

However, since F is a null-recurrent random walk, we have for every integer h,

$$\lim_{n\to\infty} \frac{\mathcal{L}\{s \in \mathbb{R} : F(s) \in [h, h+2k]\}}{n} = 0 \quad \text{a.s.},$$

giving a contradiction to (3), and completing the proof.

The proof in the case of $\mathbb{R}$ is identical except that we consider edges crossing the site x rather than the line segment $\{x\} \times [0, 1)$. □

Proof of Theorem 4. Let $d \geq 2$ and let $\mathcal{M}$ be a matching in $\mathbb{R}^d$ that is not locally finite. On the event that some bounded set is crossed by infinitely many edges, we will derive a contradiction to minimality. Suppose that $S \subset \mathbb{R}^d$ is a bounded set that is crossed by infinitely many edges. Suppose also that $[\mathcal{R}] \cup [\mathcal{B}]$ is parallel-free, and locally finite as a subset of $\mathbb{R}^d$ – both of these properties hold a.s. for a Poisson process (Lemma 7). The remainder of the proof will be a deterministic geometric argument given these assumptions.

Let $L := \{\{au : a \in \mathbb{R}\} : u \in \mathbb{R}^d \setminus \{0\}\}$ be the projective space of all lines passing through the origin; L is a compact metric space under the angle metric. The **direction** of an edge (r, b) is the line $\{a(r-b) : a \in \mathbb{R}\} \in L$. Since L is compact, the set of all edges that cross S contains a sequence whose directions converge, say to some $\ell \in L$. By the local finiteness of $[\mathcal{R}] \cup [\mathcal{B}]$, the set of edges that intersect S and have both endpoints outside any given bounded set contains a sequence whose directions converge to ℓ.

Now fix some edge $(r, b) \in [\mathcal{M}]$ whose direction is not equal to ℓ (this is possible, otherwise all edges would be parallel), and let θ be the acute angle between (r, b) and ℓ. By the above observations, for any $t > 0$ there exists an edge $(r', b') \in [\mathcal{M}]$ that crosses S, makes angle less than $\theta/2$ with ℓ, and has both endpoints r', b' outside $B(t)$. We claim that if t is sufficiently large, any such edge satisfies

$$|r - b'| + |r' - b| < |r - b| + |r' - b'|,$$

contradicting minimality.

Figure 3 illustrates the proof of this claim: take any (doubly infinite) line Λ intersecting S and making angle less than $\theta/2$ with ℓ. If r' and b' are points on Λ going to infinity in opposite directions, such that both are at distance t from O, where $t \to \infty$, then the difference $|r-b'|+|r'-b|-|r'-b'|$ converges to the orthogonal projection of (r, b) onto (r', b'). The latter equals $\pm|r-b| \cos \alpha$, where α is the angle between (r, b) and Λ (and the sign depends on the order of r' and b'). Furthermore, this convergence is uniform in the choice of the

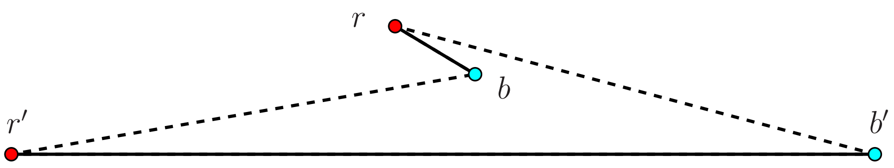

Figure 3: Contradicting minimality: as r' and b' go to infinity along a fixed line, the difference between the total dashed length and $|r'-b'|$ converges to the projection of (r,b) on (r',b').

line Λ, because the distance between Λ and (r,b) is bounded for lines that intersect S. Since $\alpha > \theta/2$ we have $\pm|r-b|\cos\alpha \le |r-b|\cos(\theta/2) < |r-b|$, and the above claim follows. □

Proof of Theorem 2(ii). The proof of Theorem 4 above applies unchanged on the strip (the direction ℓ must of course be horizonal), and shows that any minimal matching scheme must be locally finite. On the other hand, Theorem 3(i) states that a locally finite translation-invariant matching scheme cannot exist on the strip. □

5 Locally finite matching

In this section we prove Theorem 3(ii).

Lemma 9. *Let X, X', Y be independent Poisson random variables with respective means λ, λ, μ, where $\mu \le \lambda$. Then*

$$\mathbb{P}\big(X - X' \ge Y\big) \le \exp -\frac{\mu^2}{6\lambda}.$$

Proof. Write $Z := Y + X'$ and $\nu := \lambda + \mu$, so that Z is Poisson with mean ν and independent of X, and $X - X' \ge Y$ is equivalent to $X \ge Z$. Now we apply a Chernoff bound: taking $s = \sqrt{\nu/\lambda}$ we have

$$\begin{aligned}\mathbb{P}(X \ge Z) &= \mathbb{P}(s^{X-Z} \ge 1) \le \mathbb{E}s^{X-Z} \\ &= \exp\big[\lambda(s-1) + \nu(s^{-1}-1)\big] = \exp\big[2\sqrt{\lambda\nu} - \lambda - \nu\big].\end{aligned}$$

Writing $\delta = \mu/\lambda$, the last expression equals

$$\begin{aligned} &\exp 2\lambda\big[\sqrt{1+\delta} - 1 - \delta/2\big] \\ &\le \exp 2\lambda\Big[-\frac{\delta^2}{12}\Big] = \exp -\frac{\mu^2}{6\lambda}, \end{aligned}$$

(where we used the fact that $\sqrt{1+\delta} \le 1 + \delta/2 - \delta^2/12$ for $\delta \in [0,1]$). $\square$

Proof of Theorem 3(ii). We first argue that it is suffices to construct a locally finite matching scheme in the case $d = 2$. Starting from such a scheme, we may obtain a matching scheme in the slab $\mathbb{R}^2 \times [0,1)^{d-2}$ (where $d \ge 3$) by assigning each red or blue point $x \in \mathbb{R}^2$ a location (x, U_x), where the (U_x) are independent and uniformly random in $[0,1)^{d-2}$. Now take independent copies of this matching in each of the slabs $\mathbb{R}^2 \times (z + [0,1)^{d-2})$, for $z \in \mathbb{Z}^{d-2}$; the resulting matching in $\mathbb{R}^d$ clearly inherits the locally finite property, and is invariant under all translations in $\mathbb{R}^2 \times \mathbb{Z}^{d-2}$. To obtain a fully translation-invariant version, translate by a uniformly random element of $\{0\}^2 \times [0,1)^{d-2}$.

Similarly, it now suffices to find a matching scheme in $\mathbb{R}^2$ that is invariant under translations of $\mathbb{Z}^2$; then we obtain a fully translation-invariant version by applying a translation by a uniformly random element of $[0,1)^2$.

We start by defining a random sequence of successively coarser partitions of $\mathbb{R}^2$ into rectangles. Let $a_n = n!$. For each $n = 1, 2 \ldots$, an ***n*-block** will be an a_n-by-a_{n-1} (respectively a_{n-1}-by-a_n) rectangle if n is even (respectively odd). Each n-block A will be a disjoint union of a_n/a_{n-2} $(n-1)$-blocks, called the **children** of A. The left-most (respectively bottom-most) child is called the **heir** of A. See Figure 4. We choose the blocks in a $\mathbb{Z}^2$-invariant way as follows. The 1-blocks are all the squares $z + [0,1)^2$ for $z \in \mathbb{Z}^2$. Let r_n be a uniformly random integer in $[0, a_n/a_{n-2})$, where the r_n are independent of each other and of $\mathcal{R}, \mathcal{B}$. Let $t_n = r_n a_{n-2} + r_{n-2} a_{n-4} + \ldots$, where the last term in this sum is $r_2 a_0$ (respectively $r_3 a_1$), and define an n-block to be any rectangle of the form $[x a_n + t_n, (x+1) a_n + t_n) \times [y a_{n-1} + t_{n-1}, (y+1) a_{n-1} + t_{n-1})$ where $(x, y) \in \mathbb{Z}^2$ if n is even (respectively, the same with the coordinates reversed if n is odd).

We now construct a matching via a sequence of stages $n = 1, 2, \ldots$. At the end of stage n we will have a partial matching with each of its edges confined within some n-block.

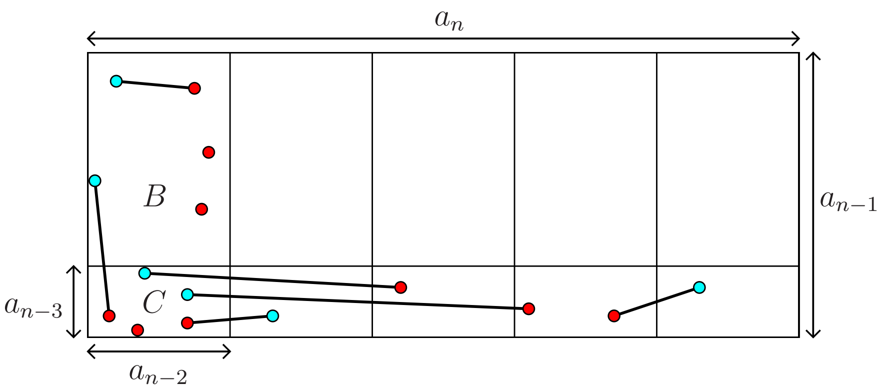

Figure 4: An n-block A, its children including its heir B, and their heirs including B's heir C. Also shown are the new edges added in stage n assuming A is neither bad nor dodgy.

Stage 1. Within each 1-block, match as many red-blue pairs as possible (for definiteness, choose from among the partial matchings of maximal cardinality of the points in the block the one with minimum total length).

Stage n ($n \geq 2$). For each n-block A, let B be the heir of A, and let C be the heir of B (or let $C = B$ if $n = 2$). Now:

(i) unmatch all points in B;

(ii) if possible, match all currently non-matched points in $A \setminus B$ to non-matched points in $(A \setminus B) \cup C$;

(iii) match as many of the remaining non-matched points in A as possible.

(In (ii) and (iii), for definiteness take the matching of minimum length among those with the required property. The unmatching step (i) is a matter of convenience only - an alternative would be to match only in blocks that lie in no higher-order heir; see the finiteness claims below.) We call the block A **bad** if step (ii) does not succeed. We call a block **dodgy** if at least one of its children is bad. $\lozenge$

Note that at the end of stage n, in any n-block A, the number of non-matched points equals the excess $|\mathcal{R}(A) - \mathcal{B}(A)|$. Moreover, if $n \geq 2$ and A is not bad, then all the non-matched points lie in the heir of A. Furthermore (and this is the key observation), if $n \geq 3$ and A is neither bad nor dodgy then all the new edges added at stage n are confined to heirs (the heir B of A and the heirs of the children of A). See Figure 4.

Let $Q := [0,1)^2$ be the unit square; Q is contained in exactly one n-block for each n. We will prove below that a.s. Q lies in only finitely many heirs, bad blocks and dodgy blocks. Once this is proved, the same also holds for every integer unit square, and we deduce the following. Each point becomes unmatched (in step (i)) only finitely many times, so we can define a limiting partial matching $\mathcal{M}$. This is in fact a perfect matching, since the only non-matched points in a non-bad block lie in its heir. Furthermore, from the key observation above we deduce that Q intersects only finitely many edges a.s., so the same holds for any bounded set as required.

Finally we turn to the proofs of the finiteness claims above. Since the partitions into blocks are independent of $\mathcal{R}, \mathcal{B}$ we have

$$\mathbb{P}(\text{the } n\text{-block containing } Q \text{ is an heir}) = \frac{a_n a_{n-1}}{a_{n+1} a_n} = \frac{1}{n(n+1)}.$$

Since $\sum_n \frac{1}{n(n+1)} < \infty$, the Borel-Cantelli lemma implies that Q lies in only finitely many heirs a.s.

Also, the block A is bad only if the net excess of one color of points over the other in $A \setminus B$ cannot be accommodated in C. Thus, using Lemma 9,

$$\begin{aligned}
&\mathbb{P}(Q \text{ lies in a bad } n\text{-block}) \\
&\le \mathbb{P}\big[(\mathcal{R}-\mathcal{B})(A\setminus B) > \mathcal{B}(C)\big] + \mathbb{P}\big[(\mathcal{B}-\mathcal{R})(A\setminus B) > \mathcal{R}(C)\big] \\
&\le 2\exp - \frac{(a_{n-2}a_{n-3})^2}{6(a_n a_{n-1} - a_{n-1}a_{n-2})} \le 2 \exp - \frac{n!^2}{6n^9} \le c_1 e^{-c_2 n},
\end{aligned}$$

for some constants $c_i \in (0,\infty)$. Therefore

$$\mathbb{P}(Q \text{ lies in a dodgy } n\text{-block}) \le \frac{a_n}{a_{n-2}}\, c_1 e^{-c_2(n-1)} \le c_3 e^{-c_4 n}.$$

Hence by the Borel-Cantelli lemma again, a.s. Q lies in only finitely many bad blocks and dodgy blocks. □

6 The minimum matching

In this section we prove Theorem 2(i) in the case $d \ge 3$. The approach was suggested by Yuval Peres.

For a translation-invariant matching scheme $\mathcal{M}$ of processes $\mathcal{R}$ and $\mathcal{B}$ both of intensity 1, define the average edge length

$$\eta(\mathcal{M}) := \frac{1}{\mathcal{L}S}\,\mathbb{E}\int_S |x-\mathcal{M}(x)|\,d\mathcal{R}(x),$$

where $S \subset \mathbb{R}^d$ is any set with $\mathcal{L}S \in (0,\infty)$. The translation-invariance implies that $\eta(\mathcal{M})$ is independent of the choice of S. (The quantity $\eta(\mathcal{M})$ also equals $\mathbb{E}|\mathcal{M}^*(0)|$, where $\mathcal{M}^*$ is the Palm process obtained by conditioning on the presence of a red point at the origin – see e.g. [3] or [5, Ch. 11]).

The key ingredient is the following fact, proved in [3, Theorem 1].

Theorem 10 (Mean edge length; [3])**.** *Let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1 in $\mathbb{R}^d$. There exists a translation-invariant matching scheme $\mathcal{M}$ satisfying $\eta(\mathcal{M}) < \infty$ if and only if $d \geq 3$.*

Corollary 11 (Minimum matching)**.** *Let $d \geq 3$ and let $\mathcal{R}$ and $\mathcal{B}$ be independent Poisson processes of intensity 1 in $\mathbb{R}^d$. There exists a translation-invariant matching scheme $\widehat{\mathcal{M}}$ such that*

$$\eta(\widehat{\mathcal{M}}) = \min_{\mathcal{M}} \eta(\mathcal{M}),$$

where the minimum is over all possible translation-invariant matching schemes of $\mathcal{R}$ to $\mathcal{B}$ (on arbitrary probability spaces).

Corollary 11 follows from Theorem 10 by an abstract argument, which we postpone to the end of the section.

Proof of Theorem 2(i), case $d \geq 3$. We claim that the matching scheme $\widehat{\mathcal{M}}$ from Corollary 11 is minimal. Suppose it is not. Call a finite set of edges **reducible** if the incident red and blue points can be rematched to give a strictly lower total length; so with positive probability there exists a reducible set of edges. Therefore for some fixed t, with positive probability there is some reducible set lying entirely in the cube $[-t/2, t/2)^d$.

Construct a modified matching scheme $\mathcal{M}'$ from $\widehat{\mathcal{M}}$ as follows. Within each of the disjoint cubes $([0,t)^d + tz)_{z\in\mathbb{Z}^d}$, unmatch all the edges that lie entirely within the cube, and replace them with the matching of minimum length for this finite set of red and blue points. The resulting matching scheme satisfies the strict inequality

$$\frac{1}{t^d}\,\mathbb{E}\int_{[0,t)^d} |x-\mathcal{M}'(x)|\,d\mathcal{R}(x) < \eta(\widehat{\mathcal{M}}) \tag{4}$$

(since the edges between the cube and its complement are unaffected by the modification, while the total length of those within it is never increased and sometimes decreased). This matching scheme $\mathcal{M}'$ is not translation-invariant, but we obtain a translation-invariant version $\mathcal{M}''$ by translating it by an independent uniform element of $[0,t)^d$. The left side of (4) is unchanged if we replace $\mathcal{M}'$ with $\mathcal{M}''$, so we obtain $\eta(\mathcal{M}'') < \eta(\widehat{\mathcal{M}})$, contradicting Corollary 11. $\square$

Remarks. The same argument may be applied for instance to prove the existence of a minimal one-color matching scheme for a Poisson process (using [3, Theorem 4]). Of course, the approach cannot work for two-color matching in $\mathbb{R}^2$, since there is no matching scheme with $\eta(\mathcal{M}) < \infty$ (Theorem 10). $\Diamond$

Proof of Corollary 11. Recall that a matching scheme is a simple point process $\mathcal{M}$ in $(\mathbb{R}^d)^2$, where the presence of an ordered pair $(r, b) \in [\mathcal{M}]$ signifies matched points $r \in [\mathcal{R}]$ and $b \in [\mathcal{B}]$. Note that we can recover the red and blue processes from $\mathcal{M}$ as $\mathcal{R}_{\mathcal{M}}(\cdot) = \mathcal{R}(\cdot) := \mathcal{M}(\cdot \times \mathbb{R}^d)$, and similarly for $\mathcal{B}_{\mathcal{M}} = \mathcal{B}$. In particular $\eta(\mathcal{M})$ is a function only of the law of $\mathcal{M}$.

Let $I := \inf_{\mathcal{M}} \eta(\mathcal{M})$, where the infimum is over all translation-invariant matching schemes of two independent Poisson processes of intensity 1, and let $\mathcal{M}_1, \mathcal{M}_2, \ldots$ be a sequence of such schemes such that

$$\eta(\mathcal{M}_n) \searrow I \quad \text{as } n \to \infty.$$

We claim that the sequence $(\mathcal{M}_n)$ is relatively compact in distribution with respect to the vague topology on simple point measures in $(\mathbb{R}^d)^2$. This follows from [5, Lemma 16.15], since any bounded set $A \subset (\mathbb{R}^d)^2$ is a subset of some $S \times \mathbb{R}^d$, where S is Borel and bounded, and $\mathcal{M}_n(S \times \mathbb{R}^d) \stackrel{d}{=} \mathrm{Poi}(\mathcal{L}S)$ for each n, thus $(\mathcal{M}_n(S \times \mathbb{R}^d))$ is a tight sequence. Therefore, by passing to a subsequence, we may assume that for some simple point process $\widehat{\mathcal{M}}$ on $(\mathbb{R}^d)^2$,

$$\mathcal{M}_n \xrightarrow{d} \widehat{\mathcal{M}} \quad \text{as } n \to \infty$$

in the aforementioned topology. Since $I < \infty$, we may also assume that

$$\eta(\mathcal{M}_n) \le C \quad \text{for all } n \tag{5}$$

for some $C < \infty$. Clearly $\widehat{\mathcal{M}}$ is a matching scheme between the point processes $\mathcal{R} = \mathcal{R}_{\widehat{\mathcal{M}}}$ and $\mathcal{B} = \mathcal{B}_{\widehat{\mathcal{M}}}$. We will prove that it has all the required properties. The details will be largely routine, with the crucial step being the use of the uniform bound (5) to preclude points being 'matched to infinity' in the limit.

The above convergence implies that for any continuous, compactly supported $f : (\mathbb{R}^d)^2 \to [0, \infty)$ we have $\int f d\mathcal{M}_n \xrightarrow{d} \int f d\widehat{\mathcal{M}}$ [5, Lemma 16.16(i)]. We first check that $\widehat{\mathcal{M}}$ inherits the translation-invariance of $\mathcal{M}_n$ – this holds because for any such f and its image f' under the diagonal action of some translation of $\mathbb{R}^d$ we have $\int f d\widehat{\mathcal{M}} = \int f' d\widehat{\mathcal{M}}$. Next note that if $D \subset \mathbb{R}^d$ is any closed, bounded, $\mathcal{L}$-null set then $\widehat{\mathcal{M}}(D \times \mathbb{R}^d) = 0$ a.s., because we may choose $\mathbf{1}_{D\times\mathbb{R}^d} \le f \le \mathbf{1}_{S\times\mathbb{R}^d}$ with $\mathcal{L}S$ arbitrarily small, thus $\widehat{\mathcal{M}}(D \times \mathbb{R}^d)$ is stochastically dominated by a Poi($\mathcal{L}S$) random variable. Therefore by [5, Lemma 16.16(iii)] we have $\mathcal{M}_n(S_1 \times S_2) \xrightarrow{d} \widehat{\mathcal{M}}(S_1 \times S_2)$ for any bounded Borel $S_1, S_2 \subset \mathbb{R}^d$ with $\mathcal{L}$-null boundaries.

Next we show that $\mathcal{R} = \mathcal{R}_{\widehat{\mathcal{M}}}$ is a Poisson process of intensity 1. It is enough to show that $\widehat{\mathcal{M}}(S \times \mathbb{R}^d) \overset{d}{=} \mathrm{Poi}(\mathcal{L}S)$ for any bounded Borel S with null boundary, but the problem is that $S \times \mathbb{R}^d$ is not bounded. Suppose $S \subset B(t)$ and take $T > t$. We will approximate using $S \times B(T)$. We have for any T,

$$\mathcal{M}_n(S \times B(T)) \xrightarrow{d} \widehat{\mathcal{M}}(S \times B(T)) \quad \text{as } n \to \infty,$$

and also

$$\widehat{\mathcal{M}}(S \times B(T)) \xrightarrow{a.s.} \widehat{\mathcal{M}}(S \times \mathbb{R}^d) \quad \text{as } T \to \infty.$$

We bound the approximation errors using Markov's inequality and (5):

$$\begin{aligned}
\mathbb{P}\Big[\mathcal{M}_n(S \times B(T)) \ne \mathcal{M}_n(S \times \mathbb{R}^d)\Big] &= \mathbb{P}\Big[\mathcal{M}_n(S \times B(T)^c) > 0\Big] \\
&\le \mathbb{P}\Big[\int_S |x - \mathcal{M}_n(x)|\, d\mathcal{R}_{\mathcal{M}_n}(x) \ge T - t\Big] \\
&\le \frac{\eta(\mathcal{M}_n)\,\mathcal{L}S}{T-t} \le \frac{C\,\mathcal{L}S}{T-t},
\end{aligned}$$

so this probability converges to 0 as $T \to \infty$, uniformly in n. (This uniformity is the key point of the proof). By [5, Theorem 4.28] it follows that $\mathcal{M}_n(S \times \mathbb{R}^d) \xrightarrow{d} \widehat{\mathcal{M}}(S \times \mathbb{R}^d)$, so the latter has distribution Poi($\mathcal{L}S$) as required.

The same argument shows also that $\mathcal{B} = \mathcal{B}_{\widehat{\mathcal{M}}}$ is a Poisson process of intensity 1 (here we use the fact that $\eta(\mathcal{M}_n)$ is equal to the analogous quantity with the roles of red and blue reversed – see e.g. [3, Proposition 7]). We can prove that $\mathcal{R}$ and $\mathcal{B}$ are independent by applying similar reasoning to the joint law of $\mathcal{R}(S_1)$ and $\mathcal{B}(S_2)$ for bounded S_1, S_2.

We have established that $\widehat{\mathcal{M}}$ is a translation-invariant matching scheme of two independent intensity-1 Poisson processes, and it follows that $\eta(\widehat{\mathcal{M}}) \geq I$. On the other hand for any $\mathcal{M}$ we have $\eta(\mathcal{M}) = \sup_{k\to\infty} \eta_k(\mathcal{M})$, where

$$\eta_k(\mathcal{M}) := \mathbb{E} \int_{[0,1)^d} k \wedge |x - \mathcal{M}(x)| \, d\mathcal{R}_{\mathcal{M}}(x),$$

and also $\eta_k(\mathcal{M}_n) \to \eta_k(\widehat{\mathcal{M}})$ for each k, so $\eta(\widehat{\mathcal{M}}) \leq I$. Thus $\eta(\widehat{\mathcal{M}}) = I$ as required. $\square$

Acknowledgements

I thank Yuval Peres, Oded Schramm and Terry Soo for many valuable conversations.

References

[1] M. Deijfen. Stationary random graphs with prescribed iid degrees on a spatial Poisson process. *Electron. Commun. Probab.*, 14:81–89, 2009.

[2] D. Gale and L. Shapley. College admissions and stability of marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.

[3] A. E. Holroyd, R. Pemantle, Y. Peres, and O. Schramm. Poisson matching. *Ann. Inst. Henri Poincaré Probab. Stat.*, 45(1):266–287, 2009.

[4] A. E. Holroyd and Y. Peres. Trees and matchings from point processes. *Electron. Comm. Probab.*, 8:17–27 (electronic), 2003.

[5] O. Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.

[6] M. Krikun. Connected allocation to Poisson points in R^2. *Electron. Comm. Probab.*, 12:140–145 (electronic), 2007.

[7] L. D. Meshalkin. A case of isomorphism of Bernoulli schemes. *Soviet Math. Dokl.*, 3, 1962.

[8] T. Soo. Translation-invariant matchings of coin-flips on Z^d. *Adv. in Appl. Probab.*, arXiv:math/0610334. To appear.

[9] A. Timar. Invariant matchings of exponential tail on coin flips in Z^d. Preprint.

Microsoft Research, 1 Microsoft Way, Redmond, WA 98052, USA
and
Department of Mathematics, University of British Columbia,
121-1984 Mathematics Road, Vancouver, BC V6T 1Z2, Canada
`holroyd(at)math.ubc.ca` `www.math.ubc.ca/~holroyd`