

Poisson Matching

Alexander E. Holroyd Robin Pemantle Yuval Peres
Oded Schramm

December 11, 2007 (revised February 29, 2008)

Abstract

Suppose that red and blue points occur as independent homogeneous Poisson processes in \mathbb{R}^d . We investigate translation-invariant schemes for perfectly matching the red points to the blue points. For any such scheme in dimensions $d = 1, 2$, the matching distance X from a typical point to its partner must have infinite $d/2$ -th moment, while in dimensions $d \geq 3$ there exist schemes where X has finite exponential moments. The Gale-Shapley stable marriage is one natural matching scheme, obtained by iteratively matching mutually closest pairs. A principal result of this paper is a power law upper bound on the matching distance X for this scheme. A power law lower bound holds also. In particular, stable marriage is close to optimal (in tail behavior) in $d = 1$, but far from optimal in $d \geq 3$. For the problem of matching Poisson points of a single color to each other, in $d = 1$ there exist schemes where X has finite exponential moments, but if we insist that the matching is a deterministic factor of the point process then X must have infinite mean.

1 Introduction

Let \mathcal{R} be a simple point process of finite intensity in \mathbb{R}^d . The **support** of \mathcal{R} is the random set $[\mathcal{R}] := \{x \in \mathbb{R}^d : \mathcal{R}(\{x\}) = 1\}$. Elements of $[\mathcal{R}]$ are called **red points**. A **one-color matching scheme** of \mathcal{R} is a simple point process \mathcal{M} of unordered pairs $\{x, y\} \subset \mathbb{R}^d$, on a shared probability

Key words: Poisson process, point process, matching, stable marriage.

2000 Mathematics Subject Classifications: 60D05, 60G55, 05C70.

Funded by Microsoft Research (AEH, YP & OS), an NSERC discovery grant (AEH), NSF grants DMS-0603821 (RP) and DMS-0605166 (YP).




Figure 1: Matchings of 2000 uniformly random points on a 2-dimensional torus: (i) stable 1-color; (ii) minimum-length 1-color; (iii) stable 2-color; (iv) minimum-length 2-color.

space, such that almost surely $(V, E) = ([\mathcal{R}], [\mathcal{M}])$ is a random graph which is a perfect matching of $[\mathcal{R}]$ (i.e. a simple graph with all degrees 1). Let \mathcal{B} be a second simple point process, and call elements of $[\mathcal{B}]$ **blue points**. A **two-color matching scheme** between \mathcal{R} and \mathcal{B} is a process \mathcal{M} which similarly yields almost surely a perfect bipartite matching between $[\mathcal{R}]$ and $[\mathcal{B}]$ (i.e. a perfect matching of $[\mathcal{R}] \cup [\mathcal{B}]$ where all the edges are from $[\mathcal{R}]$ to $[\mathcal{B}]$). In either case we denote by $\mathcal{M}(x)$ the **partner** of a red or blue point x ; that is the unique point such that $\{x, \mathcal{M}(x)\} \in [\mathcal{M}]$. See Figure 1 for some examples on the finite torus.

We say that a one-color (respectively two-color) matching scheme \mathcal{M} is **translation-invariant** if the law of the joint process $(\mathcal{R}, \mathcal{M})$ (respectively $(\mathcal{R}, \mathcal{B}, \mathcal{M})$) is invariant under translations of \mathbb{R}^d . **Isometry-invariance** is defined analogously. If almost surely $\mathcal{M} = f(\mathcal{R})$ (respectively $\mathcal{M} = f(\mathcal{R}, \mathcal{B})$) for some deterministic function f then we call \mathcal{M} a **factor** matching scheme. We sometimes refer to a matching scheme which is not a factor as **randomized**.

For a translation-invariant one-color or two-color matching scheme, let \mathbb{P} be the probability measure governing $(\mathcal{R}, \mathcal{M})$ or $(\mathcal{R}, \mathcal{B}, \mathcal{M})$, and \mathbb{E} the associated expectation operator. We are interested in the typical distance between matched pairs. Assume without loss of generality that \mathcal{R} has intensity 1 (otherwise rescale). For $r \in [0, \infty]$ it is natural to consider the quantity

$$F(r) := \mathbb{E} \# \{x \in [\mathcal{R}] \cap [0, 1)^d : |x - \mathcal{M}(x)| \leq r\},$$

where $|\cdot|$ denotes the Euclidean norm. It is easy to see that F is a distribution function, therefore we introduce a random variable X with probability measure \mathbb{P}^* and expectation operator \mathbb{E}^* such that

$$\mathbb{P}^*(X \leq r) = F(r). \tag{1}$$

We can think of X as the typical distance from a red point to its partner. In fact, this interpretation can be made rigorous via the technology of Palm processes – see Section 2 below.

We consider the following main questions. For Poisson processes on \mathbb{R}^d , what is the best possible tail behavior (as measured by X) for a translation-invariant (or isometry-invariant) matching scheme, in the one-color and two-color cases? How do the answers depend on dimension? And if we insist on a factor matching scheme? We also address the case of *stable* matchings – see below.

Note the following trivial lower bound on X . In the one-color or two-color case, the partner of a point must be at least as far as the closest other point. In the case when \mathcal{R} (respectively $\mathcal{R} + \mathcal{B}$) is a homogeneous Poisson process, this gives

$$\mathbb{E}^* e^{cX^d} = \infty \tag{2}$$

for some $c = c(d) \in (0, \infty)$.

The following theorems show that the optimal tail behavior of two-color matching schemes depends dramatically on the dimension.

Theorem 1 (2-color upper bounds). *Let \mathcal{R}, \mathcal{B} be independent Poisson processes of intensity 1. There exist isometry-invariant two-color matching schemes satisfying:*

- (i) in $d = 1$: $\mathbb{P}^*(X > r) \leq Cr^{-1/2} \quad \forall r > 0$;
- (ii) in $d = 2$: $\mathbb{P}^*(X > r) \leq Cr^{-1} \quad \forall r > 0$;
- (iii) in $d \geq 3$: $\mathbb{E}^* e^{cX^d} < \infty$.

Here $C = C(d) \in (0, \infty)$ denotes a constant. Furthermore, in (i) the matching scheme is a factor.

Theorem 2 (2-color lower bounds; [14]). *Let \mathcal{R}, \mathcal{B} be independent Poisson processes of intensity 1. In $d = 1$ or $d = 2$, any translation-invariant two-color matching scheme (factor or not) satisfies*

$$\mathbb{E}^* X^{d/2} = \infty.$$

Together with the trivial bound (2), Theorems 1 and 2 settle reasonably accurately the question of optimal tail behavior for randomized two-color matchings. We do not know the optimal tail behavior for translation-invariant *factor* matching schemes of two independent Poisson processes in dimensions $d \geq 2$. Theorem 2 was derived in [14] via results from [17, 12], the proofs of which were quite involved. We will present here a simple direct proof.

Here is a brief heuristic explanation for the above tail behavior, and in particular the sharp difference between dimensions $d \leq 2$ and $d \geq 3$. In a ball of large radius r , the discrepancy between the numbers of red and blue points is typically a random multiple of $r^{d/2}$ (by the central limit theorem). This discrepancy must be accommodated via the boundary of the ball, which has size of order r^{d-1} . When $d \leq 2$, the discrepancy exceeds the boundary

(with substantial probability), so we expect that a fraction $r^{d/2}/r^d = r^{-d/2}$ of points must be matched at distance at least of order r . When $d \geq 3$, the boundary exceeds the discrepancy, so far more efficient matching is possible, with the tails determined by local deviations in the distribution of points.

The following illustrates a case where allowing additional randomization makes a striking difference to tail behavior.

Theorem 3 (1-color, 1 dimension). *Let $d = 1$, let \mathcal{R} be a Poisson process of intensity 1, and consider one-color matching schemes.*

- (i) *Any translation-invariant factor matching satisfies $\mathbb{E}^* X = \infty$.*
- (ii) *There exists an isometry-invariant randomized scheme satisfying $\mathbb{P}^*(X > r) = e^{-r} \quad \forall r > 0$.*

The following shows that the above dichotomy does not extend to higher dimensions.

Theorem 4 (1-color, 2 or more dimensions). *Let \mathcal{R} be a Poisson process of intensity 1. For all $d \geq 2$ there exists a translation-invariant one-color factor matching scheme satisfying*

$$\mathbb{E}^* e^{CX^d} < \infty,$$

for some $C = C(d) \in (0, \infty)$. *The same bound can be attained by a randomized isometry-invariant matching scheme in all $d \geq 2$, and by an isometry-invariant factor matching scheme in $d = 2$.*

We do not know how to construct an isometry-invariant factor matching satisfying the bound in Theorem 4 for $d \geq 3$ (see the remarks in Section 4, however).

Stable matching: Iterated mutually closest matching algorithm.

The following natural “greedy” algorithm gives a matching scheme by trying to optimize locally. When considering a two-color matching, call a pair of points x, y **potential partners** if one is red while the other is blue. In the one-color case, call x and y potential partners if they are distinct points in $[\mathcal{R}]$. We say that potential partners x and y are **mutually closest** if y is the closest potential partner to x and x is the closest potential partner to y . Now, given the point configuration, match all mutually closest pairs to each

other, then remove these points and match all mutually closest pairs in the remaining set of points. Repeat indefinitely. \diamond

It turns out that the above algorithm yields a perfect matching under general conditions (Proposition 9), and in particular this holds almost surely in the case when \mathcal{R} is a Poisson process (and for the two-color case, when \mathcal{B} is an independent Poisson process of the same intensity). Furthermore, it is the unique **stable matching** in the sense of Gale and Shapley [8]. (See Section 2 for the details). Evidently (under the aforementioned conditions) the stable matching gives an isometry-invariant factor matching scheme.

We can accurately describe the tail behavior of stable matchings in the one-color case, but some questions remain in the two-color case.

Theorem 5 (1-color stable matching). *Let \mathcal{R} be a Poisson process of intensity 1. For any $d \geq 1$, the one-color stable matching satisfies*

- (i) $\mathbb{E}^* X^d = \infty$;
- (ii) $\mathbb{P}^*(X > r) \leq Cr^{-d} \quad \forall r > 0$;

for some $C = C(d) \in (0, \infty)$.

Theorem 6 (2-color stable matching). *Let \mathcal{R}, \mathcal{B} be independent Poisson processes of intensity 1. For $d \geq 1$, the two-color stable matching satisfies:*

- (i) $\mathbb{E}^* X^d = \infty$,
(Theorem 2 gives a better bound in $d = 1, 2$);
- (ii) $\mathbb{P}^*(X > r) \leq Cr^{-s} \quad \forall r > 0$,
where $s = s(d) \in (0, 1)$ satisfies $s(1) = 1/2$, and $C = C(d) \in (0, \infty)$.

The power $s(d)$ is given explicitly as the solution of an equation, and for example $s(2) = 0.496 \dots$ and $s(3) = 0.449 \dots$ – see Theorem 19. It is a fascinating unsolved question to determine the correct power law for $d \geq 2$ (see the open problems at the end of the article).

It is interesting that stable matching performs essentially optimally (in terms of tail behavior) among (possibly randomized) matching schemes in the two-color case for $d = 1$, but not for $d \geq 3$, and not in the one-color case.

Summary

The following tables summarize the best known results for isometry-invariant matchings of Poisson processes.

1-color matching		Lower bound		Upper bound / best construction
Randomized	All d	$\mathbb{E}^* e^{cX^d} = \infty$	\approx	$\mathbb{E}^* e^{cX^d} < \infty$
Factor	$d = 1$	$\mathbb{E}^* X = \infty$	\approx	$[\mathbb{P}^*(X > r) \leq Cr^{-1}]$
	$d \geq 2$	$\mathbb{E}^* e^{cX^d} = \infty$	\approx	$\mathbb{E}^* e^{cX^d} < \infty$ ((T) for $d \geq 3$)
Stable	All d	$\mathbb{E}^* X^d = \infty$	\approx	$\mathbb{P}^*(X > r) \leq Cr^{-d}$

2-color matching		Lower bound		Upper bound / best construction
Randomized	$d = 1, 2$	$\mathbb{E}^* X^{d/2} = \infty$	\approx	$\mathbb{P}^*(X > r) \leq Cr^{-d/2}$
	$d \geq 3$	$\mathbb{E}^* e^{cX^d} = \infty$	\approx	$\mathbb{E}^* e^{cX^d} < \infty$
Factor	$d = 1$	$[\mathbb{E}^* X^{1/2} = \infty]$	\approx	$[\mathbb{P}^*(X > r) \leq Cr^{-1/2}]$
	$d = 2$	$[\mathbb{E}^* X = \infty]$	\ll	$[\mathbb{P}^*(X > r) \leq Cr^{-0.496\dots}]$
	$d \geq 3$	$\mathbb{E}^* e^{cX^d} = \infty$	\ll	$[\mathbb{P}^*(X > r) \leq Cr^{-s(d)}]$
Stable	$d = 1$	$[\mathbb{E}^* X^{1/2} = \infty]$	\approx	$\mathbb{P}^*(X > r) \leq Cr^{-1/2}$
	$d = 2$	$[\mathbb{E}^* X = \infty]$	\ll	$\mathbb{P}^*(X > r) \leq Cr^{-0.496\dots}$
	$d \geq 3$	$\mathbb{E}^* X^d = \infty$	\ll	$\mathbb{P}^*(X > r) \leq Cr^{-s(d)}$

Notes:	
(T)	Translation-invariant scheme only.
$[\dots]$	Bound follows from line above (lower bounds) or below (upper bounds).
\approx	Indicates reasonably close lower and upper bounds.
\ll	Indicates a substantial gap between the lower and upper bounds.

Extensions to other processes

The results on 2-color matchings in Theorems 1, 2, 6 all extend, with similar proofs, to the following two variant settings.

- (i) Perfect matchings of Heads to Tails for i.i.d. fair coin flips indexed by \mathbb{Z}^d (see [14, 18, 20] for details). (In order to define stable matchings in this context, one must specify a way to break ties between pairs of sites \mathbb{Z}^d which are the same distance apart). Here the random variable X denotes the distance from the origin to its partner.
- (ii) Fair allocations of Lebesgue measure to a Poisson process (see [11] and also [5, 10, 14] for definitions and background). Here X denotes the

distance from the origin to its allocated Poisson point. The resulting upper bounds on the stable allocation in Theorem 6(ii) represent a considerable improvement on the previous best results: X was known to have finite $(1/18)$ th moment in $d = 1$, and no quantitative upper bound was known in $d \geq 2$ [10].

Where appropriate, we make remarks following the proofs regarding the adaptation of our results to these settings.

Some of our results extend easily to point processes other than the Poisson process. In particular, Theorem 5(ii) holds for any translation-invariant simple point process for which the stable matching is well-defined (see Proposition 9). Theorems 5(i) and 6(i) hold provided in addition the point processes are tolerant of local modifications (see Proposition 18).

Another interesting variant concerns matching of random points on large finite boxes; see e.g. [1, 3, 7, 19]. We do not explore this connection in depth, but our proof of Theorem 1(iii) relies on the remarkable results of [19].

2 Preliminaries

In this section we present some useful elementary definitions and results.

Some notation

Let \mathcal{L} denote Lebesgue measure on \mathbb{R}^d , and denote the Euclidean ball $B(x, r) := \{z \in \mathbb{R}^d : |x - z| < r\}$. We denote the unit cube $Q := [0, 1]^d \subset \mathbb{R}^d$, and $Q_u := Q + u$ for $u \in \mathbb{Z}^d$.

Palm processes

Consider a translation-invariant one-color or two-color matching scheme, and let \mathbb{P} be the probability measure governing $(\mathcal{R}, \mathcal{M})$ or $(\mathcal{R}, \mathcal{B}, \mathcal{M})$. We introduce the **Palm process** $(\mathcal{R}^*, \mathcal{M}^*)$ or $(\mathcal{R}^*, \mathcal{B}^*, \mathcal{M}^*)$, with law \mathbb{P}^* and expectation \mathbb{E}^* , in which we condition on the presence of a red point at the origin, while taking \mathcal{M} and (in the two color-case) \mathcal{B} as a stationary background. See e.g. [15, Ch. 11] for details. In the case when \mathcal{R} is a homogeneous Poisson process, it turns out that \mathcal{R}^* has the same distribution as \mathcal{R} with an *added* point at the origin:

$$[\mathcal{R}^*] \stackrel{d}{=} [\mathcal{R}] \cup \{0\}. \tag{3}$$

Also, if \mathcal{R} and \mathcal{B} are independent processes then \mathcal{R}^* and \mathcal{B}^* are independent, and $\mathcal{B}^* \stackrel{d}{=} \mathcal{B}$.

If \mathcal{R} is a translation-invariant measure-valued process of intensity $\lambda \in (0, \infty)$, and (\mathcal{R}^*, Ψ^*) is the Palm version of \mathcal{R} taken together with any jointly translation-invariant random background Ψ (which may be a random function or a random measure on \mathbb{R}^d), then the following properties are standard (see [15, Ch. 11]). Let θ^x denote translation by $x \in \mathbb{R}^d$ (defined to act on measures via $(\theta^x \pi)(S) = \pi(S - x)$ for $S \subseteq \mathbb{R}^d$). For any measurable $S \subseteq \mathbb{R}^d$ and any event A we have

$$\mathbb{E} \int_S \mathbf{1}[\theta^{-x}(\mathcal{R}, \Psi) \in A] d\mathcal{R}(x) = \lambda \cdot \mathcal{L}S \cdot \mathbb{P}^*[(\mathcal{R}^*, \Psi^*) \in A]. \quad (4)$$

(Indeed this may be taken as a *definition* of the Palm process). More generally, for any non-negative measurable f on the appropriate space,

$$\mathbb{E} \int_{\mathbb{R}^d} f(\theta^{-x}(\mathcal{R}, \Psi), x) d\mathcal{R}(x) = \lambda \int_{\mathbb{R}^d} \mathbb{E}^* f((\mathcal{R}^*, \Psi^*), x) dx. \quad (5)$$

Let \mathcal{M} be a translation-invariant (one- or two- color) matching scheme. If we let

$$X := |\mathcal{M}^*(0)| \quad (6)$$

denote the distance from the origin to its partner under the Palm measure, then (4) yields in particular

$$\mathbb{E} \#\{x \in [\mathcal{R}] \cap S : |x - \mathcal{M}(x)| \leq r\} = \lambda \cdot \mathcal{L}S \cdot \mathbb{P}^*(X \leq r). \quad (7)$$

Hence the above definition of X is consistent with the earlier notation in (1).

Note that the tail bound (2) for the Poisson process is now an elementary consequence of (3) and (6).

Partial matching and mass transport

A partial matching of a set U is the edge set m of a simple graph (U, m) in which each vertex has degree at most 1. As before we write $m(x) = y$ if $\{x, y\} \in m$, and in addition we write $m(x) = \infty$ if x is unmatched (i.e. has degree 0). A **one-color** (respectively **two-color**) **partial matching scheme** \mathcal{M} is a point process on pairs which yields almost surely a partial matching of $[\mathcal{R}]$ (respectively between $[\mathcal{R}]$ and $[\mathcal{B}]$).

Proposition 7 (Fairness). *Let \mathcal{R}, \mathcal{B} be simple point processes of finite intensity, and let \mathcal{M} be a translation-invariant two-color partial matching scheme of \mathcal{R} and \mathcal{B} . Then the process of matched red points and the process of matched blue points have equal intensity.*

In particular, Proposition 7 shows that translation-invariant perfect matching schemes are possible only between two point processes of equal intensity. In addition, by applying the result to the matching obtained by deleting all edges longer than r , we see that X is equal in law to the analogous random variable defined in terms of a typical *blue* point.

We prove Proposition 7 via the following lemma which will be useful elsewhere. (See [4] for background).

Lemma 8 (Mass transport principle).

(i) *Suppose $t : \mathbb{Z}^d \times \mathbb{Z}^d \rightarrow [0, \infty]$ satisfies $t(u + w, v + w) = t(u, v)$ for all $u, v, w \in \mathbb{Z}^d$, and write $t(A, B) := \sum_{u \in A, v \in B} t(u, v)$. Then*

$$t(0, \mathbb{Z}^d) = t(\mathbb{Z}^d, 0).$$

(ii) *Suppose T is a random non-negative measure on $\mathbb{R}^d \times \mathbb{R}^d$ such that $T(A, B) := T(A \times B)$ and $T(A + w, B + w)$ are equal in law for all $w \in \mathbb{Z}^d$. Then*

$$\mathbb{E}T(Q, \mathbb{R}^d) = \mathbb{E}T(\mathbb{R}^d, Q).$$

Proof.

(i): $t(0, \mathbb{Z}^d) = \sum_{u \in \mathbb{Z}^d} t(0, u) = \sum_{u \in \mathbb{Z}^d} t(-u, 0) = t(\mathbb{Z}^d, 0).$

(ii): Apply (i) to $t(u, v) := \mathbb{E}T(Q_u, Q_v).$ □

We sometimes call t or T a mass transport, and think of $t(A, B)$ or $T(A, B)$ as the amount of mass sent from A to B .

Proof of Proposition 7. Apply Lemma 8 to the mass transport

$$T(A, B) := \#\{x \in [\mathcal{R}] \cap A : \mathcal{M}(x) \in B\}$$

in which each matched red point sends unit mass to its partner. Then $\mathbb{E}T(Q, \mathbb{R}^d) = \mathbb{E}\#\{x \in [\mathcal{R}] \cap Q : \mathcal{M}(x) \neq \infty\}$, which is the intensity of matched red points, while similarly $\mathbb{E}T(\mathbb{R}^d, Q)$ is the intensity of matched blue points. □

Stable matching

Following Gale and Shapley [8], we say that a partial matching m of a set $U \subset \mathbb{R}^d$ is **stable** if there do *not* exist distinct points $x, y \in U$ satisfying

$$|x - y| < \min \{|x - m(x)|, |y - m(y)|\}, \quad (8)$$

where $|x - m(x)| := \infty$ if x is unmatched. A pair x, y satisfying (8) is called unstable. (The motivation for this definition is that each point prefers to be matched with closer points, so an unstable pair x, y prefer to divorce their current partners and marry each other.) Similarly, a partial bipartite matching between two sets U, V is called **stable** if there do not exist $x \in U$ and $y \in V$ satisfying (8).

We call a set $U \subset \mathbb{R}^d$ **non-equidistant** if there do not exist $w, x, y, z \in U$ with $\{w, x\} \neq \{y, z\}$ and $|w - x| = |y - z| > 0$. A **descending chain** is an infinite sequence $x_0, x_1, \dots \in U$ for which the distances $|x_i - x_{i+1}|$ form a strictly decreasing sequence.

Proposition 9 (Unique stable matching). *Let \mathcal{R} be a translation-invariant homogeneous point processes of finite intensity. (Respectively, let \mathcal{R}, \mathcal{B} be point processes of equal finite intensity, jointly ergodic under translations). Suppose that almost surely $[\mathcal{R}]$ (respectively $[\mathcal{R}] \cup [\mathcal{B}]$) is non-equidistant, and has no descending chains. Then almost surely there is a unique stable partial matching of $[\mathcal{R}]$ (respectively between $[\mathcal{R}]$ and $[\mathcal{B}]$). Furthermore, it is almost surely a perfect matching, and it is produced by the iterated mutually closest matching algorithm described in the introduction.*

Under the conditions in Proposition 9, the stable matching has the following additional interpretation. Grow a ball centered at each red point (respectively, each red point and each blue point) simultaneously, so that at time t all the balls have radius t . Whenever two balls touch (respectively, whenever an \mathcal{R} -ball and a \mathcal{B} -ball touch), match their centers to each other, and remove the two balls.

The conditions on the point processes in Proposition 9 hold in particular for homogeneous Poisson processes, as proved in [9]. Clearly, under the conditions of the proposition, the unique stable matching gives an isometry-invariant factor matching scheme. We postpone the proof of Proposition 9 to Section 5.

3 Two-color matching

In this section we give proofs of Theorems 1 and 2.

Proof of Theorem 1(i). Use the stable matching (see Proposition 9 and Theorem 6(ii)). \square

Proof of Theorem 1(ii). We shall give a construction which works in all dimensions, and gives a matching scheme with tails $\mathbb{P}^*(X > r) \leq Cr^{-d/2}$.

First note that it is sufficient to give a translation-invariant matching satisfying the required bound, for then we may obtain a (randomized) isometry-invariant version by applying a uniformly random isometry preserving the origin (i.e., chosen according to the Haar measure on the compact group of such isometries) to $(\mathcal{R}, \mathcal{B}, \mathcal{M})$. Indeed, it suffices to give a matching scheme which is invariant under translations by elements of \mathbb{Z}^d and which satisfies

$$\mathbb{E}\#\{x \in [\mathcal{R}] \cap Q : |x - \mathcal{M}(x)| > r\} \leq Cr^{-d/2}, \quad (9)$$

for then we may achieve a translation-invariant version satisfying the same bound by similarly applying a uniformly random translation in the unit cube Q .

We start by defining a sequence of successively coarser random partitions of \mathbb{R}^d into boxes in a \mathbb{Z}^d -invariant way. Let τ_0, τ_1, \dots be i.i.d. uniformly random elements of the discrete cube $\{0, 1\}^d$, independent of the point processes \mathcal{R}, \mathcal{B} . For each $k = 0, 1, \dots$, define a k -**box** to be any subset of \mathbb{R}^d of the form

$$[0, 2^k)^d + 2^k z + \sum_{i=0}^{k-1} 2^i \tau_i,$$

where $z \in \mathbb{Z}^d$.

Now, given the point processes \mathcal{R}, \mathcal{B} and the partitioning into boxes, define a matching as follows. Within each 0-box, match as many red/blue pairs as possible in some arbitrary pre-determined way. (For definiteness, choose from among the bipartite partial matchings of maximum cardinality the one which minimizes the total edge length.) Remove those points which have been matched. Now match as many red/blue pairs of the remaining points as possible within each 1-box, remove these matched points, and repeat for 2-boxes and so on. The union of all these partial matchings clearly gives a \mathbb{Z}^d -invariant *partial* matching scheme \mathcal{M} between \mathcal{R} and \mathcal{B} .

We shall prove that the partial matching \mathcal{M} satisfies (9). From this it follows by taking $r \rightarrow \infty$ that almost surely every red point is matched, and hence by applying Lemma 7 (to the \mathbb{R}^d -invariant version) that every blue point is matched also.

Fix k , and call a red point **k -bad** if it has not been matched within its k -box by stage k of the matching algorithm. Suppose each k -bad red point distributes mass 1 uniformly to its k -box, i.e. let

$$T(A, B) := \sum_{x \in A \cap [\mathcal{R}] : x \text{ is } k\text{-bad}} 2^{-dk} \mathcal{L}\{y \in B : x \text{ and } y \text{ lie in the same } k\text{-box}\}.$$

Then we have

$$\begin{aligned} \mathbb{E}T(Q, \mathbb{R}^d) &= \mathbb{E}\#\{k\text{-bad red points in } Q\} \\ &\geq \mathbb{E}\#\{x \in [\mathcal{R}] \cap Q : |x - \mathcal{M}(x)| > 2^k \sqrt{d}\} \end{aligned} \quad (10)$$

(since a k -box has diameter $2^k \sqrt{d}$).

On the other hand, writing W for the random k -box containing Q ,

$$\begin{aligned} \mathbb{E}T(\mathbb{R}^d, Q) &= 2^{-dk} \mathbb{E}\#\{k\text{-bad red points in } W\} \\ &= 2^{-dk} \mathbb{E}(\mathcal{R}(W) - \mathcal{B}(W))^+ = 2^{-dk} \mathbb{E}S^+, \end{aligned} \quad (11)$$

where $S := \mathcal{R}[0, 2^k)^d - \mathcal{B}[0, 2^k)^d$, since the location of W is independent of \mathcal{R}, \mathcal{B} .

The central limit theorem gives $\mathbb{E}S^+ / \sqrt{2^{dk+1}} \rightarrow \mathbb{E}\chi^+$ as $k \rightarrow \infty$, where χ is a standard Gaussian. Combining this with (10),(11) and applying Lemma 8 we deduce (9) for some $C = C(d) \in (0, \infty)$ and all $r = 2^k \sqrt{d}$ with $k = 0, 1, 2, \dots$. Hence by taking $2^k \sqrt{d} \leq r < 2^{k+1} \sqrt{d}$ the same holds (with a modified constant) for all $r > 0$. \square

Proof of Theorem 1(iii). We will deduce the result by a limiting argument from a result in [19] on matchings of finite sets of points; a similar argument was used in [14].

The following is proved in [19, equation (1.8)]. Let $d \geq 3$ and let \mathcal{R}_n and \mathcal{B}_n each consist of n point masses whose locations are all independent and uniformly distributed on $[0, 1]^d$. Then for each n there exists a two-color matching scheme \mathcal{F}_n between \mathcal{R}_n and \mathcal{B}_n such that

$$\mathbb{P}(G_n) \geq 1 - n^{-2}, \text{ where } G_n := \left\{ n^{-1} \sum_{x \in [\mathcal{R}_n]} \exp(Cn|x - \mathcal{F}_n(x)|^d) \leq 2 \right\}. \quad (12)$$

Here the constant C depends on d but not n .

Now let $\widetilde{\mathcal{F}}_n$ be \mathcal{F}_n conditioned on the event G_n (this corrects a minor error in [14]). We construct a translation-invariant matching scheme \mathcal{M}_n by scaling $\widetilde{\mathcal{F}}_n$ to cover a cube of volume n , and tiling \mathbb{R}^d with identical copies of this matching, with the origin chosen uniformly at random. More formally, regarding a two-color matching scheme \mathcal{M} as a simple point process (i.e. a random point measure) of *ordered* pairs $(x, y) \in \mathbb{R}^d \times \mathbb{R}^d$ in which the presence of a point $(r, b) \in [\mathcal{M}]$ indicates a matched pair $r \in [\mathcal{R}]$ and $b \in [\mathcal{B}]$, we define

$$\mathcal{M}_n(A \times B) := \sum_{z \in \mathbb{Z}^d} \widetilde{\mathcal{F}}_n \left(n^{1/d}(A + U + z) \times n^{1/d}(B + U + z) \right),$$

where U is uniformly distributed on $[0, 1]^d$ and independent of $\widetilde{\mathcal{F}}_n$. Then (12) implies that for any Borel $A \subseteq \mathbb{R}^d$ we have

$$\mathbb{E} \iint \exp(C|x - y|^d) \mathbf{1}_{x \in A} \mathcal{M}_n(dx \times dy) \leq 2\mathcal{L}A. \quad (13)$$

(To check this, we first use invariance to deduce that the left side must be a linear multiple of $\mathcal{L}A$, and then take $A = [0, n^{1/d}]^d$ to find the constant).

By (13), the random sequence (\mathcal{M}_n) is tight in the vague topology of measures on $\mathbb{R}^d \times \mathbb{R}^d$ (see [15, Lemma 16.15]). Therefore let \mathcal{M} be any subsequential limit in distribution, and note that it has the following properties. It is a two-color matching scheme between the marginal point processes $\mathcal{M}(\cdot, \mathbb{R}^d)$ and $\mathcal{M}(\mathbb{R}^d, \cdot)$. These processes are independent Poisson point processes of intensity 1 (this would clearly be true for any limit constructed in the same way from the unconditioned matchings \mathcal{F}_n , therefore it holds for \mathcal{M} because $\mathbb{P}(G_n) \rightarrow 1$). The process \mathcal{M} inherits the translation-invariance of \mathcal{M}_n . Finally, it satisfies (13) (with \mathcal{M} replacing \mathcal{M}_n), which implies $\mathbb{E}^* e^{CX^d} \leq 2$ as required. \square

Remarks. The analogous results to Theorem 1(iii) for matchings of coin flips on \mathbb{Z}^d and for fair allocations of \mathbb{R}^d may be proved by following a similar limiting argument in the appropriate space (see [14] for another variant). The required matchings and allocations on finite cubes exist by the results of [19].

The result (12) is a special case of a much more general result in [19], proved by deep (and indirect) methods. A remark is made in [19] that the

bound (12) can also be proved for an explicit matching obtained from the construction in [1]. Also, by results in [5], in $d \geq 3$ it is possible to obtain an explicit “*transport*” between two independent Poisson processes (i.e. a translation-invariant random measure T on $\mathbb{R}^d \times \mathbb{R}^d$ with marginals \mathcal{R} and \mathcal{B}), satisfying a tail bound that is exponential in a power of distance. (Specifically, construct the “gravitational allocations” – see [5] – for \mathcal{R} and \mathcal{B} independently, and let a red point r send to a blue point b a mass equal to the volume of the intersection of the cell of r and the cell of b .)

We do not know the optimal tail behavior of translation-invariant *factor* matching schemes of two independent Poisson processes in dimensions $d \geq 2$. We suspect that there exist factor matchings satisfying the same bounds as in the randomized case (see Theorem 1). For the analogous questions concerning coin flips on \mathbb{Z}^d , substantial progress has been made by T. Soo [18] and A. Timar [20]. \diamond

Theorem 2 is Corollary 9 of [14], which in turn was deduced from results on “extra head schemes” in [12, 17]. Here we give a simple direct argument.

Proof of Theorem 2 (case $d=1$). Using (5) and Fubini’s theorem, for $t > 0$ we have

$$\begin{aligned} & \mathbb{E}\#\{x \in [\mathcal{R}] \cap [0, 2t] : \mathcal{M}(x) \notin [0, 2t]\} \\ & \leq \mathbb{E}\#\{x \in [\mathcal{R}] \cap [0, 2t] : |\mathcal{M}(x) - x| > x \wedge (2t - x)\} \\ & = \int_0^{2t} \mathbb{P}^*[X > x \wedge (2t - x)] dx \\ & = 2\mathbb{E}^*(X \wedge t). \end{aligned}$$

Now the central limit theorem gives

$$\mathbb{E}\#\{x \in [\mathcal{R}] \cap [0, 2t] : \mathcal{M}(x) \notin [0, 2t]\} \geq \mathbb{E}[(\mathcal{R}[0, 2t] - \mathcal{B}[0, 2t])^+] \sim Ct^{1/2}$$

as $t \rightarrow \infty$ for some $C \in (0, \infty)$. On the other hand, if $\mathbb{E}^*X^{1/2} < \infty$ then the dominated convergence theorem gives $t^{-1/2}\mathbb{E}^*(X \wedge t) \rightarrow 0$ as $t \rightarrow \infty$, so we obtain a contradiction. \square

For any $x, y \in \mathbb{R}^d$, we define the **line-segment** $\langle x, y \rangle := \{\lambda x + (1 - \lambda)y : \lambda \in [0, 1]\} \subset \mathbb{R}^d$. The following lemma will be used to derive a contradiction in the proof of Theorem 2 in the case $d = 2$.

Lemma 10 (Edge intersections). *For any translation-invariant 1-color or 2-color matching scheme \mathcal{M} (of any translation-invariant point process(es)) which satisfies $\mathbb{E}^*X < \infty$, the number of matching edges $\{x, y\} \in [\mathcal{M}]$ such that the line segment $\langle x, y \rangle$ intersects the unit cube Q has finite expectation.*

Proof. Consider the mass transport

$$t(u, v) := \mathbb{E}\#\{x \in [\mathcal{R}] \cap Q_u : \langle x, \mathcal{M}(x) \rangle \text{ intersects } Q_v\}.$$

Since an edge of length ℓ intersects at most $d(1 + \ell)$ cubes of the form $z + Q$, where $z \in \mathbb{Z}^d$, we have

$$t(0, \mathbb{Z}^d) \leq d + d\mathbb{E} \sum_{x \in Q \cap [\mathcal{R}]} |x - \mathcal{M}(x)| = d(1 + \mathbb{E}^*X) < \infty,$$

hence by Lemma 8,

$$\mathbb{E}\#\{\text{matching edges intersecting } Q\} = t(\mathbb{Z}^d, 0) < \infty. \quad \square$$

Proof of Theorem 2 (case $d=2$). Without loss of generality we may assume that the matching scheme \mathcal{M} is ergodic with respect to the full group of translations of \mathbb{R}^2 ; if not we apply the claimed result to the ergodic components. Therefore suppose for a contradiction that \mathcal{M} is an ergodic matching scheme satisfying $\mathbb{E}^*X < \infty$.

For an ordered pair of distinct sites $x, y \in \mathbb{R}^2$, we define the random variable $K(x, y)$ to be the number of matching edges which intersect the directed line segment $\langle x, y \rangle$ with the red point on the left and the blue on the right. More precisely, $K(x, y)$ is the number of pairs $\{r, b\} \in [\mathcal{M}]$ with $r \in [\mathcal{R}]$ and $b \in [\mathcal{B}]$ such that $\langle r, b \rangle$ intersects $\langle x, y \rangle$, and $\det \begin{pmatrix} b-r \\ y-x \end{pmatrix} > 0$. Also define $F(x, y) := K(x, y) - K(y, x)$, the “net red-blue flow” across the line segment.

Lemma 10 and the assumption of finite mean imply that $\mathbb{E}K(x, y) < \infty$ and hence $\mathbb{E}F(x, y) < \infty$ for any fixed x, y . Note also the antisymmetry $F(x, y) = -F(y, x)$ and the additivity property that if $y \in \langle x, z \rangle$ (and these sites are deterministic), then a.s. $F(x, y) + F(y, z) = F(x, z)$. Fix vectors $x, u \in \mathbb{R}^2$ with $u \neq 0$. Using the ergodic theorem we deduce

$$\frac{F(x, x + nu)}{n} \xrightarrow{\text{a.s. and } L^1} \Phi(x, u) \quad \text{as } n \rightarrow \infty, \quad (14)$$

for some random variable $\Phi(x, u)$ with finite mean. We will show next that a.s. $\Phi(x, u)$ is constant in x , and deduce that it is deterministic.

For linearly independent vectors $a, b \in \mathbb{R}^2$, assume them ordered so that $\det \begin{pmatrix} a \\ b \end{pmatrix} > 0$, and let $S = S(a, b)$ be the interior of the parallelogram with vertices $0, a, b, a + b$. The difference $\mathcal{R}(S) - \mathcal{B}(S)$ equals the number of edges crossing the boundary of S with the red point inside and the blue point outside minus the number crossing in the opposite orientation. Hence, a.s.,

$$\mathcal{R}(S) - \mathcal{B}(S) = F(0, a) - F(b, a + b) + F(a, a + b) - F(0, b). \quad (15)$$

(Note in particular that matching edges which intersect two sides of S do not contribute to the right side of (15), since the positive and negative contributions cancel.)

Applying this to $S(a, b) = S(nu, x)$ with x and u as above (and assuming $\det \begin{pmatrix} u \\ x \end{pmatrix} > 0$), we obtain

$$\frac{\mathcal{R}(S) - \mathcal{B}(S)}{n} = \frac{F(0, nu)}{n} - \frac{F(x, x + nu)}{n} + \frac{F(nu, x + nu)}{n} - \frac{F(0, x)}{n}.$$

As $n \rightarrow \infty$, the left side converges a.s. to 0 by the strong law of large numbers, while the last term converges a.s. to 0 because $F(0, x)$ is a.s. finite. Also, since for each n we have $F(nu, x + nu) \stackrel{d}{=} F(0, x)$, and the latter has finite mean, the third term on the right converges a.s. to 0. Thus, using (14),

$$0 = \Phi(0, u) - \Phi(x, u) \quad \text{a.s.}$$

Similarly, the same conclusion holds for x satisfying $\det \begin{pmatrix} u \\ x \end{pmatrix} \leq 0$. Thus $\Phi(0, u)$ is a translation-invariant function of \mathcal{M} , so the ergodicity assumption implies that it is an a.s. constant, say $\phi(u)$. Furthermore, since $F(x, x + nu)$ has the same law as $F(0, nu)$, it now follows from (14) that

$$\frac{F(x_n, x_n + nu)}{n} \xrightarrow{L^1} \phi(u) \quad \text{as } n \rightarrow \infty, \quad (16)$$

for any $u \in \mathbb{R}^2$ and any deterministic sequence $x_n \in \mathbb{R}^2$.

Now write $h = (1, 0)$ and $v = (0, 1)$, and apply (15) to the square $S(nh, nv) = (0, n)^2$. Dividing by n and using (16) gives

$$\frac{\mathcal{R}(S) - \mathcal{B}(S)}{n} \xrightarrow{L^1} \phi(h) - \phi(h) + \phi(v) - \phi(v) = 0$$

as $n \rightarrow \infty$; that is $\mathbb{E}|\mathcal{R}(S) - \mathcal{B}(S)| = o(n)$. But the central limit theorem gives $\mathbb{E}|\mathcal{R}(S) - \mathcal{B}(S)| = \Omega(n)$, a contradiction. \square

Remark. We are grateful to Scott Sheffield for correcting an error in an earlier version of the above proof. \diamond

4 One-color matching

In this section we prove Theorems 3 and 4.

Consider the case $d = 1$. An **adjacent** matching of a discrete set $U \subset \mathbb{R}$ is one in which for every edge $\{x, y\}$, the interval (x, y) contains no points of U . Clearly for any given infinite discrete U there are exactly two adjacent matchings (one in which the origin lies in (x, y) for some edge $\{x, y\}$, and one in which it does not).

Proof of Theorem 3(ii). Conditional on \mathcal{R} , choose one of the two adjacent matchings each with probability $1/2$. It is elementary to check that the resulting matching scheme is isometry-invariant, and that $\mathcal{M}^*(0)$ is a symmetric two-sided exponential random variable. \square

The following lemma states that, unsurprisingly, we cannot achieve an adjacent matching without randomization. In what follows we say that an edge $\{x, y\}$ of \mathcal{M} **crosses** a site $z \in \mathbb{R}$ if $z \in (x, y)$.

Lemma 11. *Let \mathcal{R} be a homogeneous Poisson process on \mathbb{R} . There does not exist a one-color factor matching scheme where the matching is almost surely adjacent.*

Proof. Suppose on the contrary that \mathcal{M} is such a matching scheme. Write \mathcal{F}_S for the σ -algebra generated by the restriction of \mathcal{R} to $S \subseteq \mathbb{R}$. Since the event

$$A := \{0 \text{ is crossed by some edge}\}$$

lies in $\mathcal{F}_{\mathbb{R}}$, for every $\epsilon > 0$ there exists $r = r(\epsilon) < \infty$ and an event $A_\epsilon \in \mathcal{F}_{[-r, r]}$ such that $\mathbb{P}(A \Delta A_\epsilon) < \epsilon$. Moreover, by translation-invariance we can find $B_\epsilon \in \mathcal{F}_{[-2r, 0]}$ such that $\mathbb{P}(\{-r \text{ is crossed}\} \Delta B_\epsilon) < \epsilon$. For an adjacent matching, observing the \mathcal{R} -points in an interval together with whether some deterministic point in the interval is crossed determines the matching on the interval a.s. Hence there exists $L_\epsilon \in \mathcal{F}_{[-2r, 0]} \subset \mathcal{F}_{(-\infty, 0]}$ with $\mathbb{P}(A \Delta L_\epsilon) < \epsilon$. Since this is true for every $\epsilon > 0$ we deduce that $A \in \overline{\mathcal{F}_{(-\infty, 0]}}$, where the bar denotes completion under \mathbb{P} . Similarly we have $A \in \overline{\mathcal{F}_{[0, \infty)}}$, so A is independent of itself and has probability 0 or 1. But now it is easy to see that neither of the two resulting matching schemes is translation-invariant. \square

Remark. Lemma 11 may be strengthened to the (equally unsurprising) fact that the *only* a.s. adjacent matching scheme is the one in the proof of Theorem 3(ii). Indeed, for such a scheme consider the a.e.-defined function of point configurations $f(\pi) := \mathbb{P}(O \text{ is crossed} \mid \mathcal{R} = \pi)$ – we must show that it equals $1/2$ a.e. By translation-invariance, for a translation θ^{-t} we have a.e. $f(\theta^{-t}\pi) \in \{f(\pi), 1 - f(\pi)\}$, according to the parity of $\pi[0, t]$. Hence $|f(\pi) - 1/2|$ is a translation-invariant function of π , so by ergodicity of \mathcal{R} it is a.e. constant. If the constant is $c \neq 1/2$ then the translation-equivariant function $h_\pi(t) := f(\theta^{-t}\pi)$ assigns (a.e. with respect to π and t) values c and $1 - c$ to alternate intervals between points of the configuration π , and we can clearly use this to construct an adjacent factor matching, contradicting Lemma 11. \diamond

Proof of Theorem 3(i). Suppose that \mathcal{R} is a Poisson point process of intensity 1 and \mathcal{M} is a translation-invariant factor matching scheme. We claim that

$$\mathbb{P}(0 \text{ is crossed by infinitely many edges}) = 1. \quad (17)$$

Suppose (17) is false. On the event that 0 is crossed by finitely many edges, a.s. the same is true for any other $r \in \mathbb{R}$, because the difference between the number of edges crossing r and the number of edges crossing 0 is at most the number of red points between 0 and r . By ergodicity, it then follows that a.s. every $r \in \mathbb{R}$ is crossed only finitely many times. We can now define a new matching scheme, by matching two adjacent red points x and y if and only if the sites r between them are crossed an odd number of times in the matching \mathcal{M} . The new matching is an adjacent factor matching, which contradicts Lemma 11 and proves (17).

To conclude, using (5) and Fubini's theorem we have

$$\begin{aligned} \mathbb{E}\#\{\text{edges crossing } 0\} &\leq \frac{1}{2}\mathbb{E}\#\{x \in [\mathcal{R}] : |\mathcal{M}(x) - x| > |x|\} \\ &= \frac{1}{2} \int_{-\infty}^{\infty} \mathbb{P}^*(X > |r|) dr \\ &= \mathbb{E}^* X, \end{aligned}$$

so (17) implies $\mathbb{E}^* X = \infty$. \square

Matching from a forest. The following construction will be used in the proof of Theorem 4. Let U be a countable infinite set, and suppose we are

given a locally finite forest with vertex set U and one end per tree (that is, a simple acyclic graph with finite degrees where from each vertex there is exactly one singly-infinite self-avoiding path). The **parent** of vertex x is the vertex y such that edge (x, y) lies on the infinite path from x . If y is the parent of x then x is a **child** of y . Suppose we are also given, for each vertex, a total ordering of its children.

Under the above conditions, the following construction gives a perfect matching of U . (Similar constructions were used in [6] and [13]). A **leaf** is a vertex with no children, and a **twig** is a vertex which is not a leaf but whose only children are leaves. Suppose x is a twig, and let x_1, x_2, \dots, x_k be its ordered children. Match the pairs $\{x_1, x_2\}, \{x_3, x_4\}, \dots$. If k is odd, also match x_k to x . Do this for all twigs. Now remove those vertices which have been matched, together with their incident edges, and repeat the construction on the remaining graph. Repeat indefinitely.

The above construction clearly gives a partial matching of U . To see that it is a perfect matching, observe that for any given vertex, at least one of its finitely many descendants is matched and removed at every stage, so it is eventually matched itself. (Vertex x is called a **descendant** of y if y lies on the infinite path from x .)

Note the additional property that any matched pair were at graph-theoretic distance at most 2 in the original forest. This will enable us to derive tail bounds for the matching from tail bounds on the forest. \diamond

Let \mathcal{R} be a Poisson process of intensity 1. The **minimal spanning forest** \mathcal{F} of $[\mathcal{R}]$ is the graph obtained from the complete graph on $[\mathcal{R}]$ by deleting every edge which is the longest (in Euclidean distance) in some cycle. In dimension $d = 2$ it is known [2, p. 94] that \mathcal{F} is almost surely a locally finite one-ended tree. Now, ordering the children of each vertex x by Euclidean distance from x (say), and applying the above construction, we get an isometry-invariant one-color factor matching scheme. This will enable us to prove Theorem 4 for $d = 2$. To prove the required tail bound we need the lemmas below.

Let $S := \{y \in \mathbb{R}^d : |y| = 1\}$ be the unit sphere. A **cap** is a proper subset of S of the form $H = S \cap B(y, r)$, where $y \in S$. A **cone** of **width** w is a set of the form

$$V = V_H := \{\alpha h : h \in H \text{ and } \alpha \in (0, \infty)\},$$

where H is some cap of diameter w .

Lemma 12 (Cones). *If $V = V_H$ is a cone of width 1, then for any $x, y \in V$,*

$$|x| \geq |y| \text{ implies } |x - y| \leq |x|.$$

Proof. Let $\beta := |y|/|x|$, so that $|\beta x| = |y|$, and hence $\beta x, y \in |y|H$. Then

$$|x - y| \leq |x - \beta x| + |\beta x - y| \leq (1 - \beta)|x| + \text{diam}(|y|H) = |x| - |y| + |y|. \quad \square$$

Lemma 13. *In the minimal spanning forest \mathcal{F} of a Poisson process \mathcal{R} of intensity 1 on \mathbb{R}^d ,*

$$\mathbb{P}^*(0 \text{ has an incident edge of length } > r) \leq C' e^{-Cr^d} \quad \forall r > 0,$$

for $C, C' \in (0, \infty)$ depending only on d .

Proof. For distinct $x, y \in \mathbb{R}^d$, define a set $S_{x,y} := (V + x) \cap B(x, |x - y|)$, where V is some cone of width 1 such that $y \in V + x$. For red points x and y , if there exists another red point $z \in S_{x,y}$ then \mathcal{F} cannot have an edge from x to y . This is because $|x - z| < |x - y|$, while $|y - z| \leq |y - x|$ by Lemma 12, so $\{x, y\}$ is the longest edge in the cycle (x, y, z) . Hence the probability in question is at most

$$\mathbb{P}^* \left[\mathcal{R}^*(V \cap B(0, r)) = 0 \text{ for some cone } V \text{ of width } 1 \right].$$

Now let H_1, \dots, H_k be a finite subcover of the cover of the unit sphere by caps of diameter $1/3$, and let $V_i := V_{H_i}$ be the associated cones of width $1/3$. Any cone of width 1 must completely contain at least one of the V_i , so the above probability is at most

$$\begin{aligned} & \mathbb{P}^* \left[\mathcal{R}^*(V_i \cap B(0, r)) = 0 \text{ for some } i \in \{1, \dots, k\} \right] \\ & \leq k \mathbb{P}^* \left[\mathcal{R}(V_1 \cap B(0, r)) = 0 \right] \leq k e^{-Cr^d}. \end{aligned}$$

□

Lemma 14. *If non-negative random variables Y, Z satisfy $\mathbb{E} e^{cY^d} < \infty$ and $\mathbb{E} e^{cZ^d} < \infty$, then $\mathbb{E} e^{c'(Y+Z)^d} < \infty$ where $c' := 2^{-d}c$.*

Proof. $\mathbb{E} e^{c'(Y+Z)^d} \leq \mathbb{E} e^{c(Y \vee Z)^d} = \mathbb{E}[e^{cY^d} \vee e^{cZ^d}] \leq \mathbb{E} e^{cY^d} + \mathbb{E} e^{cZ^d}$. □

Proof of Theorem 4 (case $d = 2$). Construct a matching \mathcal{M} from the minimum spanning forest \mathcal{F} as described above. To prove the tail bound note that any matched pair are either siblings or a parent and child. Hence if we write $D(x)$ for the distance from x to its parent in \mathcal{F} then $|x - \mathcal{M}(x)| \leq D(x) + D(\mathcal{M}(x))$. Also, Lemma 13 implies that $\mathbb{E}^* e^{CD(0)^2} < \infty$ for a certain constant $C > 0$. Now defining

$$T(A, B) := \sum_{x \in [\mathcal{R}] \cap A: \mathcal{M}(x) \in B} e^{CD(x)^2},$$

the mass transport principle (Lemma 8) gives $\mathbb{E}^* e^{CD(\mathcal{M}^*(0))^2} = \mathbb{E}^* e^{CD(0)^2}$, and by Lemma 14 we obtain $\mathbb{E}^* e^{C'X^2} \leq \mathbb{E}^* e^{C'[D(0)+D(\mathcal{M}^*(0))]^2} < \infty$. \square

Remark. In dimensions $d \geq 3$ the minimal spanning forest is believed to have one end per tree, but this is not proved. Therefore we use a different forest (see below), which is translation-invariant but not isometry-invariant. Given sufficient effort, it seems probable that a suitable one-ended isometry-invariant forest could be constructed, giving an isometry-invariant factor matching satisfying a similar bound. \diamond

Proof of Theorem 4 (case $d \geq 3$). Let \mathcal{R} be a Poisson process of intensity 1. As in the proof of Theorem 1(ii), it suffices to give a translation-invariant factor matching scheme, for then we may obtain an isometry-invariant randomized version by applying a random isometry preserving the origin.

The following construction is inspired by [6]. For $z = (z_1, z_2, \dots, z_d) \in \mathbb{R}^d$ we write $\underline{z} := (z_2, \dots, z_d) \in \mathbb{R}^{d-1}$. Define the cone $V := \{z \in \mathbb{R}^d : z_1 > |\underline{z}|\}$. Now for each red point x , let $S(x)$ be the a.s. unique red point in $x + V$ for which the first coordinate $S(x)_1$ is least, and put a directed edge from x to $S(x)$. Let \mathcal{G} denote the resulting graph.

Since a.s. the out-degree of each vertex is 1 and there are no oriented cycles, the graph \mathcal{G} is clearly a forest. The mass-transport principle (Lemma 8) shows that, under the Palm measure, the expected in-degree of the origin is also 1, hence \mathcal{G} is locally finite. Furthermore, it is immediate that

$$\mathbb{E}^* e^{-C|S(0)|^d} < \infty \tag{18}$$

for some $C = C(d) > 0$. We claim that \mathcal{G} has one end per tree. Once this is proved, we can use it to construct a matching as described above (ordering

children by distance), and the required tail bound may then be deduced from (18) in the same way as in the above proof for the case $d = 2$.

Turning to the proof that \mathcal{G} has one end per tree, it suffices to prove that \mathcal{G} a.s. has no backward infinite path (that is, no sequence of vertices and directed edges $\cdots \rightarrow x_2 \rightarrow x_1 \rightarrow x_0$), for clearly from each vertex there is exactly one forward infinite path. Call a red point **bad** if it lies on some backward infinite path. It is proved in [2] that no translation-invariant random forest in \mathbb{R}^d can have a component with more than 2 ends. Assuming that bad points exist we shall obtain a contradiction to this result.

Consider the ‘hyperplane’ of cubes $L := \{Q_u : u \in \mathbb{Z}^d, u_1 = -1\}$. If bad points exist then by invariance and ergodicity there exist two fixed cubes $Q_u, Q_v \in L$ at distance at least $d + 1$ from each other such that the event

$$A := \{Q_u \text{ and } Q_v \text{ each contain a bad point}\}$$

has positive probability. Note that whenever \mathcal{G} has a directed path from x to y then $y \in x + V$, or equivalently $x \in y + (-V)$. Hence for any red points $x \in Q_u$ and $y \in Q_v$, there cannot be a directed path from x to y or from y to x . Recalling also that all out-degrees equal 1 we deduce that on A there exist two *disjoint* backward infinite paths to red points $x \in Q_u$ and $y \in Q_v$. Also note that the event A is measurable with respect to the restriction of \mathcal{R} to the half-space $\mathbb{H}_- := \mathbb{R}_- \times \mathbb{R}^{d-1}$.

Now we may construct an event B of positive probability, measurable with respect to the restriction of \mathcal{R} to the half-space $\mathbb{H}_+ := \mathbb{R}_+ \times \mathbb{R}^{d-1}$, such that on B , the forward infinite paths from any red points $x \in Q_u$ and $y \in Q_v$ must coalesce. Specifically, this will hold if a sufficiently large region of \mathbb{H}_+ is empty except for one red point which lies in $\bigcap_{z \in [(Q_u \cup Q_v) + V] \cap \mathbb{H}_-} (z + V)$ (see Figure 2). Now A and B are independent so $\mathbb{P}(A \cap B) > 0$, but on the latter event, \mathcal{G} has a component with at least 3 ends (formed by the backward paths from the bad points in Q_u and Q_v together with their joint forward path). This contradicts the result from [2] noted above. \square

5 Stable Matching

In this section we prove Theorems 5 and 6 as well as Proposition 9. We start with some relatively straightforward cases.

Proof of Theorem 5(ii). Call a red point x t -**bad** if $|\mathcal{M}(x) - x| > t$, and note that no two t -bad points may be within distance t of each other, for they




Figure 2: An event B which forces coalescence of paths starting in two given cubes: the light gray region should be empty while the dark gray region should contain exactly one red point.

would form an unstable pair. Hence, using (4),

$$1 \geq \mathbb{E} \# \{t\text{-bad red points in } B(0, \frac{t}{2})\} = \mathcal{L}B(0, \frac{t}{2}) \cdot \mathbb{P}^*(X > t). \quad \square$$

Proof of Theorem 6(ii) (case $d = 1$). Fix $t > 0$. We say that a red point x is **t -bad** if $|\mathcal{M}(x) - x| > t$. Let W be the set of t -bad red points in the interval $[0, t]$, and denote the random variables

$$a = \min W; \quad b = \max W.$$

We claim that, provided $W \neq \emptyset$, every blue point in $[a, b]$ is matched to a red point in $[a, b]$. To prove this, suppose on the contrary that $x \in [a, b]$ is a blue point matched outside $[a, b]$. Without loss of generality suppose $\mathcal{M}(x) < a$. Then since $|\mathcal{M}(x) - x| > |a - x|$ and $|\mathcal{M}(a) - a| > t > |a - x|$, the pair a, x would be unstable, a contradiction.

Since elements of W are t -bad, they cannot be matched within $[a, b] \subseteq$

$[0, t]$, so the above claim implies that $\mathcal{R}[a, b] - \mathcal{B}[a, b] \geq \#W$. Hence

$$\begin{aligned} \#W &\leq \max_{[\alpha, \beta] \subseteq [0, t]} \left(\mathcal{R}[\alpha, \beta] - \mathcal{B}[\alpha, \beta] \right) \\ &\leq \max_{z \in [0, t]} F(z) - \min_{z \in [0, t]} F(z), \end{aligned}$$

where $F(z) := \mathcal{R}[0, z] - \mathcal{B}[0, z]$. But F is a continuous-time simple symmetric random walk, so (by e.g. [15, Proposition 13.13 and Theorem 14.6]) taking the expectation of the above inequality and using (4), we obtain

$$t \mathbb{P}^*(X > t) = \mathbb{E} \#W \leq C\sqrt{t}$$

for a fixed constant $C \in (0, \infty)$. □

Proof of Theorem 6(ii) (case $d \geq 2$). See Theorem 19 at the end of this section. □

To prove the lower bounds Theorems 5(i) and 6(i) we need the following simple properties of stable matchings and Poisson processes. The proofs are given after the statement of Lemma 18.

Lemma 15 (Stable partial matchings). *Let U (respectively, U, V) be (disjoint) subset(s) of \mathbb{R}^d , and suppose that U (respectively $U \cup V$) is discrete, non-equidistant, and has no descending chains. Then there is a unique stable partial matching of U (respectively between U and V), and it is produced by the iterated mutually closest matching algorithm described in the introduction. In the one-color case, there is at most one unmatched point, while in the two-color case, all unmatched points are of the same color.*

It should be noted that stable marriage problems do not in general have unique solutions; see [8]. The key to uniqueness in our setting is that preferences are based on distance, and are therefore symmetric. Note also that some condition on U (or $U \cup V$) is needed in order to guarantee existence and uniqueness of the stable matching. For example, in the one-color case with $d = 1$, the set $U = \{0\} \cup \{3^{-n} : n \in \mathbb{Z}^+\}$ has no stable partial matching, while $U = \{\sum_{k=1}^n k^{-1} : n \in \mathbb{Z}^+\}$ has more than one.

Lemma 16 (Modifications for 1-color stable matching). *Let $U \subset \mathbb{R}^d$ be a discrete, non-equidistant set with no descending chains, and let m be its unique stable partial matching.*

- (i) If $\{x, y\} \in m$ is a matched pair then $m \setminus \{\{x, y\}\}$ is the unique stable partial matching of $U \setminus \{x, y\}$.
- (ii) If $z \in \mathbb{R}^d \setminus U$ is such that $U \cup \{z\}$ is non-equidistant, and $|m(x) - x| < |z - x|$ for all $x \in U$, then m is the unique stable partial matching of $U \cup \{z\}$ (in particular, z is unmatched).

Lemma 17 (Monotonicity for 2-color stable matching). *Let $U, V, \{v'\} \subset \mathbb{R}^d$ be disjoint sets, and suppose that $U \cup V \cup \{v'\}$ is discrete, non-equidistant, and has no descending chains. Let m be the stable bipartite partial matching between U and V , and let m' be the stable bipartite partial matching between U and $V \cup \{v'\}$. Then*

$$|u - m'(u)| \leq |u - m(u)| \quad \forall u \in U,$$

(where as usual $|x - m(x)| := \infty$ if x is unmatched).

Lemma 17 states that adding an extra blue point makes the matching no worse for red points. Such results are well-known for finite stable marriage problems – see e.g. [8, 16].

Lemma 18 (Modifications for Poisson process). *Let \mathcal{R} be a homogeneous Poisson process.*

- (i) *Let U be a uniform random point in a set S with $\mathcal{L}S \in (0, \infty)$, independent of \mathcal{R} . The law of the point process $\mathcal{R} + \delta_U$ obtained by adding a point at U is absolutely continuous with respect to the law of \mathcal{R} .*
- (ii) *Let \mathcal{F} be a simple point process whose support $[\mathcal{F}]$ is a.s. a random finite subset of $[\mathcal{R}]$. The law of the process $\mathcal{R} - \mathcal{F}$ obtained by removing the points of \mathcal{F} is absolutely continuous with respect to the law of \mathcal{R} .*

Proof of Lemma 15. Consider the iterated mutually closest matching algorithm. Non-equidistance ensures that it is well-defined. We first claim that every pair matched by the algorithm are matched to each other in every stable partial matching. This is proved by induction on the stage of the algorithm: supposing the claim holds for all pairs matched so far, any remaining mutually closest pair cannot be matched to points removed earlier (by the inductive hypothesis) and they cannot be matched further away than each other (by stability).

Now consider the set N of points left unmatched by the algorithm. We must show that these points are unmatched in every stable partial matching. This is clear if $\#N \leq 1$. Therefore, suppose that $\#N \geq 2$ and consider first the one-color case. Let $x_0 \in N$, and let x_1 be the closest point to x_0 in $N \setminus \{x_0\}$, (which exists, by discreteness). Inductively, let x_{n+1} be the point in $N \setminus \{x_n\}$ closest to x_n . Clearly, $|x_n - x_{n+1}| \geq |x_{n+1} - x_{n+2}|$ for all $n \in \mathbb{N}$. Since there is no descending chain, it follows that there is some first $m \in \mathbb{N}$ such that $x_{m+1} = x_j$ with $j \in \{0, 1, \dots, m-1\}$. Because of non-equidistance, we must have $j = m-1$. But then x_{m-1} and x_m are mutually closest in N , which implies that they would have been matched by the algorithm right after all other points had been removed from $B(x_{m-1}, |x_{m-1} - x_m|) \cup B(x_m, |x_{m-1} - x_m|)$ (which would happen after a finite number of stages by discreteness). This contradicts $x_{m-1}, x_m \in N$ and shows that $\#N \leq 1$.

In the two-color case, let N_U be the unmatched points in U and let N_V be the unmatched points in V . If $N_U = \emptyset$, then clearly all points in N_V are unmatched in every stable matching. The case $N_V = \emptyset$ is similar. If both N_U and N_V are nonempty, we choose $x_0 \in N_U$ and inductively let x_{2n+1} be the point in N_V closest to x_{2n} and let x_{2n+2} be the point in N_U closest to x_{2n+1} . The argument is then completed as in the one-color case. \square

Proof of Proposition 9. Apply Lemma 15 and consider the random process \mathcal{N} of unmatched points in the unique stable partial matching - we must show that $[\mathcal{N}]$ is almost surely empty.

In the one-color case, the lemma implies that $\#[\mathcal{N}] \leq 1$. But if \mathcal{N} has exactly one point with positive probability then (after conditioning on this event) its location would be a translation-invariant \mathbb{R}^d -valued random variable, which is impossible.

In the two-color case, the lemma implies that $[\mathcal{N}]$ must be empty, or consist entirely of red points or entirely of blue points. By ergodicity, one of these three possibilities must have probability 1. But Lemma 7 implies that the processes of unmatched red points and unmatched blue points have equal intensity, so the latter two possibilities are ruled out. \square

Proof of Lemma 16. By Lemma 15, we need only check that the claimed matching is stable. In (i) this is immediate, since any unstable pair would have been unstable in the original matching. Similarly in (ii), the given condition ensures that z does not form an unstable pair with any $x \in U$. \square

Proof of Lemma 17. Suppose on the contrary that for some $u \in U$ we have $|u - m'(u)| > |u - m(u)|$. In particular $m(u) \neq \infty$, so write $v := m(u) \in V$. Stability of (u, v) in m' implies $|v - m'(v)| \leq |v - u|$, but $m'(v) \neq u$ so non-equidistance implies that the previous inequality is strict; we write $u_1 := m'(v) \in U$. Similarly, by stability of (u_1, v) in m we have $|u_1 - m(u_1)| < |u_1 - v|$; write $v_1 := m(u_1) \in V$. Iterating this argument gives a descending chain $u, v, u_1, v_1, u_2, v_2, \dots$ \square

Proof of Lemma 18. (i): It is elementary to check that the Radon-Nikodym derivative of the laws is $p(N-1)/p(N)$, where $N := \mathcal{R}(S)$, and $p(k)$ is the probability that $N = k$.

(ii): Let A be some measurable set such that $\mathbb{P}(\mathcal{R} - \mathcal{F} \in A) > 0$. We need to show that $\mathbb{P}(\mathcal{R} \in A) > 0$. Since a.s. $[\mathcal{F}]$ is finite and $[\mathcal{R}]$ is discrete, there is a.s. a finite random set of balls with rational centers and radii such that $[\mathcal{F}]$ is the intersection of $[\mathcal{R}]$ with the union of these balls. Therefore, there is a deterministic finite union of open balls W such that $\delta := \mathbb{P}(\mathcal{R} - \mathcal{F} \in A, [\mathcal{R}] \cap W = [\mathcal{F}]) > 0$. Let \mathcal{R}_1 denote the restriction of \mathcal{R} to the complement of W , and let A_1 be the event that $\mathbb{P}(\mathcal{R} - \mathcal{F} \in A, [\mathcal{R}] \cap W = [\mathcal{F}] \mid \mathcal{R}_1) > \delta/2$. Note that $\mathbb{P}(A_1) \geq \delta/2$. On the event A_1 , with positive probability we have $\mathcal{R} - \mathcal{F} = \mathcal{R}_1$ and therefore $\mathcal{R}_1 \in A$. But A_1 is $\sigma(\mathcal{R}_1)$ -measurable, so we must have $\mathcal{R}_1 \in A$ a.s. on A_1 . Since \mathcal{R}_1 and $\mathcal{R}(W)$ are independent, we deduce

$$\mathbb{P}(\mathcal{R} \in A) \geq \mathbb{P}(A_1, \mathcal{R}(W) = 0) = \mathbb{P}(A_1) \mathbb{P}(\mathcal{R}(W) = 0) > 0. \quad \square$$

We now turn to the proofs of the lower bounds.

Proof of Theorem 5(i). Let \mathcal{M} be the one-color stable matching, and consider the random set

$$H = H(\mathcal{R}) := \{x \in [\mathcal{R}] : |x - \mathcal{M}(x)| > |x| - 1\}. \quad (19)$$

This is the set of red points that would prefer some red point in the unit ball $B(0, 1)$ (if one were present in the correct location) over their current partners. We claim that

$$\mathbb{P}(\#H = \infty) = 1. \quad (20)$$

Once this is proved, we obtain the required result as follows, using (5) and

Fubini's theorem:

$$\begin{aligned} \infty = \mathbb{E}\#H &= \int_{\mathbb{R}^d} \mathbb{P}^*(X > |x| - 1) dx \\ &= \int_0^\infty \mathbb{P}^*(X + 1 > t) c t^{d-1} dt \\ &= \frac{c}{d} \mathbb{E}^*[(X + 1)^d], \end{aligned}$$

hence $\mathbb{E}^*X^d = \infty$.

Returning to the claim (20), suppose on the contrary that H is finite with positive probability. For each point configuration \mathcal{R} , construct a modified configuration $\widehat{\mathcal{R}}$ as follows:

- (i) if $\#H < \infty$, remove all the points in $H \cup \{\mathcal{M}(x) : x \in H\}$;
- (ii) add a uniformly random point in $B(0, 1)$, independently of \mathcal{R} .

Using Lemma 18, the law of the random configuration $\widehat{\mathcal{R}}$ is absolutely continuous with respect to that of \mathcal{R} . Now by Lemma 16, whenever $\#H(\mathcal{R}) < \infty$, the stable partial matching of $[\widehat{\mathcal{R}}]$ has an unmatched point (the one added in (ii)), hence this happens with positive probability. Absolute continuity therefore implies that with positive probability the stable partial matching of $[\mathcal{R}]$ has an unmatched point, contradicting Proposition 9. \square

Proof of Theorem 6(i). Define the random set H exactly as in (19) (now it is the set of red points which would prefer a blue point in $B(0, 1)$). We will prove that $\mathbb{P}(\#H = \infty) = 1$, whereupon the result follows as in the proof of Theorem 5(i).

Fix any $k < \infty$; we will prove that $\mathbb{P}(\#H \geq k) = 1$. Let \mathcal{B}' be obtained from \mathcal{B} by adding k independent uniformly random points in $B(0, 1)$. By Lemma 18(i), the law of $(\mathcal{R}, \mathcal{B}')$ is absolutely continuous with respect to that of $(\mathcal{R}, \mathcal{B})$. Hence, by Proposition 9, almost surely all the k added points are matched in the stable matching between $[\mathcal{R}]$ and $[\mathcal{B}']$. By Lemma 17, it follows that the partners of the added points were matched as far away or further in the stable matching with $[\mathcal{B}]$, so these partners lie in H , and thus $\#H \geq k$ as required. \square

Remark. As stated in the introduction, Theorem 6 holds also for the stable matching of heads (red) to tails (blue) on \mathbb{Z}^d (given some tie-breaking rule).

In order to adapt the proof of (i) to that setting, we claim that the set H of red sites $v \in \mathbb{Z}^d$ which would prefer the origin to their current partner (if the origin was blue) must be infinite. Indeed, if H is finite, then a contradiction is obtained by considering the configuration in which the sites in H are recolored blue and the origin is colored blue.

Finally we prove the upper bound for the two-color stable matching in $d \geq 2$.

Theorem 19. *In the two-color stable matching of two independent Poisson processes of intensity 1 in \mathbb{R}^d , $d \geq 2$, we have*

$$\mathbb{P}^*(X > r) \leq C r^{-s}, \quad (21)$$

where $C = C(d) \in (0, 1)$, and $s = s(d)$ is the unique solution in $(0, 1)$ of the equation

$$2 \omega_d \int_1^2 (t-1)^{d-1} t^{-s} dt = \frac{\omega_{d-1}}{d-1} \int_0^2 (1 - (t/2)^2)^{\frac{d-1}{2}} t^{-s} dt, \quad (22)$$

and ω_d denotes the $(d-1)$ -dimensional volume of the unit sphere in \mathbb{R}^d .

For $d = 2$, (22) simplifies to

$$\sqrt{\pi} (2^s - 2s) \Gamma((2-s)/2) = \Gamma((3-s)/2),$$

and for general d the integrals can be evaluated in terms of hypergeometric functions. The numerical values (rounded to the nearest 10^{-3}) of s at $d = 2, 3, 10$ and 100 are 0.496, 0.449, 0.339 and 0.224, respectively. It is not hard to see that $s_d \log d$ stays bounded away from 0 and ∞ as $d \rightarrow \infty$.

Proof of Theorem 19. Set $\alpha(r) := \mathbb{P}^*(X > r)$. Fix some $R > 0$ and consider the ball $B = B(0, R)$ of radius R about 0. Set

$$Y_{\mathcal{R}} := \{x \in [\mathcal{R}] \cap B : |x - \mathcal{M}(x)| > R + |x|\},$$

and similarly

$$Y_{\mathcal{B}} := \{x \in [\mathcal{B}] \cap B : |x - \mathcal{M}(x)| > R + |x|\}.$$

First, observe that if $x \in Y_{\mathcal{R}}$, then $\mathcal{M}(x) \notin B$. Next, note that if $x \in Y_{\mathcal{R}}$, then x prefers any blue point in B to its partner. Since the corresponding statement also holds for $Y_{\mathcal{B}}$, we have

$$\text{if } Y_{\mathcal{R}} \neq \emptyset \text{ then } Y_{\mathcal{B}} = \emptyset. \quad (23)$$

(This is the principal observation on which the proof rests.) Let Q_B denote the set of blue points in B that are matched outside of B , and similarly for $Q_{\mathcal{R}}$.

Let $Z := \mathcal{R}(B) - \mathcal{B}(B)$, and note that $\#Q_{\mathcal{R}} - \#Q_B = Z$. Therefore, $\#Y_{\mathcal{R}} \leq \#Q_B + Z$, and (23) gives

$$\#Y_{\mathcal{R}} + \#Y_B \leq \#Q_B + Z^+.$$

Our bound will follow by taking the expectation of both sides of this inequality. Since $\mathbb{E}\#Y_{\mathcal{R}} = \mathbb{E}\#Y_B$ and $\mathbb{E}(Z^+) \leq C R^{d/2}$ for some fixed constant $C = C_d$ (which may depend only on d), we get

$$2 \mathbb{E}\#Y_{\mathcal{R}} \leq C R^{d/2} + \mathbb{E}\#Q_B. \quad (24)$$

By (5), it is easy to express the left hand side in terms of α , namely,

$$\begin{aligned} \mathbb{E}\#Y_{\mathcal{R}} &= \int_B \alpha(R + |x|) dx \\ &= \omega_d \int_0^R \alpha(R + r) r^{d-1} dr = \omega_d \int_R^{2R} \alpha(r) (r - R)^{d-1} dr. \end{aligned} \quad (25)$$

The proof will proceed by expressing $\mathbb{E}\#Q_B$ in terms of α and using (24). Before embarking on the full argument we note the following simplified version which already gives a power law upper bound on α . If a blue point in B is matched outside B then the length of its edge is at least its distance to the boundary of B , hence (5) gives

$$\begin{aligned} \mathbb{E}\#Q_B &\leq \int_B \alpha(R - |x|) dx \\ &= \omega_d \int_0^R \alpha(R - r) r^{d-1} dr = \omega_d \int_0^R \alpha(r) (R - r)^{d-1} dr. \end{aligned} \quad (26)$$

Substituting (25) and (26) into (24) and using the fact that α is decreasing yields a bound for $\alpha(2R)$ in terms of $\alpha(r)$ for $r \in [0, R]$, and it is straightforward to deduce (by induction on k) that $\alpha(2^k) \leq C'(2^k)^{-s'}$ for some $C' = C'(d) \in (0, \infty)$ and $s' = s'(d) \in (0, 1)$.

In order to get a better power we will instead use an exact expression for $\mathbb{E}\#Q_B$, and analyze the resulting inequality more carefully. Denote the unit sphere $S^{d-1} := \{z \in \mathbb{R}^d : |z| = 1\}$. The intensity of the process of

pairs $(x, u) \in \mathbb{R}^d \times S^{d-1}$ such that $x \in [\mathcal{B}]$, $|x - \mathcal{M}(x)| > r$ and $(\mathcal{M}(x) - x)/|\mathcal{M}(x) - x| = u$ is precisely $\alpha(r)/\omega_d$. (That is, the expected number of such pairs in any set $A \subset \mathbb{R}^d \times S^{d-1}$ is $\alpha(r)/\omega_d$ times the volume of A .) For $x \in B$ and $u \in S^{d-1}$, let $q(x, u) := \inf\{t \geq 0 : x + t u \notin B\}$, and fix some $u_0 \in S^{d-1}$. Then

$$\mathbb{E}\#Q_B = \frac{1}{\omega_d} \int_{S^{d-1}} \int_B \alpha(q(x, u)) dx du = \int_B \alpha(q(x, u_0)) dx, \quad (27)$$

where the last equality is a consequence of rotational symmetry. Let L denote the orthogonal projection onto the subspace of \mathbb{R}^d orthogonal to u_0 ; that is, $Lz = z - (z \cdot u_0)u_0$. For $x \in B$ define $f(x) = Lx + q(x, u_0)u_0$. Since $Lf(x) = Lx$ and $f(x + tu_0) = f(x) - tu_0$, it follows by differentiation that f is measure preserving. This allows us to use the substitution $z = f(x)$ and write

$$\mathbb{E}\#Q_B = \int_{f(B)} \alpha(z \cdot u_0) dz = \int_0^{2R} \mu_{d-1}\{z \in f(B) : z \cdot u_0 = r\} \alpha(r) dr,$$

where μ_{d-1} is $(d-1)$ -dimensional Lebesgue measure and the last equality follows by Fubini. Now,

$$\mu_{d-1}\{z \in f(B) : z \cdot u_0 = r\} = \mu_{d-1}\{Lx : x \in B, q(x, u_0) = r\}.$$

Note that the set $\{Lx : x \in B, q(x, u_0) = r\}$ is precisely the set of sites $z \in L\mathbb{R}^d$ such that $z - (r/2)u_0 \in B$, which is $\{z \in L\mathbb{R}^d : |z| < \sqrt{R^2 - (r/2)^2}\}$. The $(d-1)$ -volume of this set is just $(R^2 - (r/2)^2)^{(d-1)/2}$ times the volume of the $(d-1)$ -dimensional unit ball. Since the volume of the $(d-1)$ -dimensional unit ball is $\omega_{d-1}/(d-1)$, we get

$$\mathbb{E}\#Q_B = \frac{\omega_{d-1}}{d-1} \int_0^{2R} (R^2 - (r/2)^2)^{(d-1)/2} \alpha(r) dr. \quad (28)$$

Now, taking into account (24), (25) and (28), we obtain

$$\int_0^\infty g(r/R) \alpha(r) dr \leq C R^{1-d/2},$$

where

$$g(t) := 2\omega_d \mathbf{1}_{[1,2]}(t) (t-1)^{d-1} - \frac{\omega_{d-1}}{d-1} \mathbf{1}_{[0,2]}(t) (1 - (t/2)^2)^{\frac{d-1}{2}}.$$

This bound will be useful when R is large. For R small, we use the trivial estimate

$$\int_0^\infty g(r/R) \alpha(r) dr \leq \int_0^\infty g(r/R) dr \leq \int_0^{2R} \|g\|_\infty dr = 2R \|g\|_\infty.$$

Combining these two estimates, we get

$$\int_0^\infty g(r/R) \alpha(r) dr \leq \min\{C R^{1-d/2}, 2R \|g\|_\infty\}. \quad (29)$$

We will get our desired bound on α by taking an appropriate average of (29) with respect to R .

Note that the set of s satisfying $\int_0^\infty g(t) t^{-s} dt = 0$ is precisely the set of s satisfying (22). Observe that $g(t)$ is supported on $[0, 2]$ and is continuous and monotone increasing there. Moreover, $g(0) < 0 < g(2)$. Therefore, there is some $s = s_d \in (-\infty, 1)$ such that $\int_0^\infty g(t) t^{-s} dt = 0$. We claim that s is unique. Indeed, let $s' < s$ and let t_0 be the unique solution of $g(t) = 0$ in $(0, 2)$. Then $t_0^{s'-s} t^{-s'} < t^{-s}$ precisely when $t < t_0$. Therefore $t_0^{s'-s} \int_0^\infty g(t) t^{-s'} dt > \int_0^\infty g(t) t^{-s} dt = 0$, proving uniqueness. Next, we claim that $s > 0$. Observe that if we replace α by 1 we get the volume of B (that is, $R^d \omega_d/d$) in (25) and (27). The algebraic manipulations within and following these equalities are valid for any measurable bounded function in place of α . Therefore $\int_0^\infty g(t) dt = \mu_d(B(0, 1)) = \omega_d/d > 0$, which implies $s > 0$.

Since $\int_0^\infty g(t) t^{-s} dt = 0$, a change of variables $t = r/\rho$ gives

$$\int_0^\infty g(r/\rho) \rho^{s-2} d\rho = 0. \quad (30)$$

Set

$$G_R(r) := \int_0^R g(r/\rho) \rho^{s-2} d\rho.$$

We claim that $G_R(r) \geq 0$ for all $r > 0$, and that

$$C_0 := \inf\{G_R(r) R^{1-s} : R > 0, r \in [R/2, R]\} > 0.$$

As before, let t_0 be the unique solution of $g(t) = 0$ in $(0, 2)$. If $r/R \geq t_0$, then $g(r/\rho) \geq 0$ for $\rho < R$, and hence $G_R(r) \geq 0$. On the other hand, if $r/R < t_0$, then $g(r/\rho) \leq 0$ for all $\rho > R$ and (30) gives $G_R(r) = -\int_R^\infty g(r/\rho) \rho^{s-2} d\rho \geq 0$. Since $g(t) > 0$ for $t \in (t_0, 2)$ and $g(t) < 0$ for $t \in (0, t_0)$, the above reasoning

actually gives $G_R(r) > 0$ for $r \in (0, 2R)$. Since G_R is continuous, this implies $\inf_{r \in [1/2, 1]} G_1(r) > 0$. A change of variables gives $G_R(r) = R^{s-1} G_1(r/R)$, which now proves $C_0 > 0$.

From the monotonicity of α , the definition of C_0 and from $G_R(r) \geq 0$, we now get

$$C_0 R^s \alpha(R)/2 \leq C_0 R^{s-1} \int_{R/2}^R \alpha(r) dr \leq \int_0^\infty G_R(r) \alpha(r) dr. \quad (31)$$

Note that

$$\int_0^R \int_0^\infty \rho^{s-2} |g(r/\rho)| \alpha(r) dr d\rho \leq \int_0^R \int_0^\infty \rho^{s-2} \|g\|_\infty 1_{r \leq 2\rho} dr d\rho < \infty.$$

Therefore, Fubini and (29) give

$$\begin{aligned} \int_0^\infty G_R(r) \alpha(r) dr &= \int_0^R \left(\rho^{s-2} \int_0^\infty g(r/\rho) \alpha(r) dr \right) d\rho \\ &\leq \int_0^R \rho^{s-2} \min\{C \rho^{1-d/2}, 2\rho \|g\|_\infty\} d\rho \\ &\leq \int_1^R C \rho^{s-1-d/2} d\rho + 2 \|g\|_\infty \int_0^1 \rho^{s-1} d\rho. \end{aligned} \quad (32)$$

Since the right hand side is bounded in R (this is where we use $d > 1$), this and (31) imply (21). \square

Remarks. In order to adapt the proof of Theorem 19 to the stable allocation of Lebesgue to Poisson, we replace $\#Y_{\mathcal{R}}$ with the volume of sites $z \in B$ whose Poisson point is at distance greater than $R + |z|$, and replace $\#Y_{\mathcal{B}}$ with the sum over Poisson points $x \in B$ of the volume of x 's territory that is at distance greater than $R + |x|$. The mass transport principle easily shows that these two quantities have the same expectation. A similar remark applies to $\#Q_{\mathcal{R}}$ and $\#Q_{\mathcal{B}}$. This allows us to obtain the analog of (24).

To adapt the proof to the setting of a stable matching of heads to tails in \mathbb{Z}^d , we apply a uniform random translation in $[0, 1]^d$, and then apply a random isometry preserving the origin. Then the law of the matching is invariant under isometries of \mathbb{R}^d , and the above proof applies.

Open Problems

- (i) For the two-color stable matching of two independent Poisson processes, what is the correct power law for the tail behavior of X in dimensions $d \geq 2$? We conjecture that $\mathbb{E}^* X^\alpha < \infty$ if and only if $\alpha < d/2$.
- (ii) Does there exist a translation-invariant matching of two independent Poisson processes in \mathbb{R}^2 such that the line segments connecting matched pairs do not cross?

References

- [1] M. Ajtai, J. Komlós, and G. Tusnády. On optimal matchings. *Combinatorica*, 4(4):259–264, 1984.
- [2] K. S. Alexander. Percolation and minimal spanning forests in infinite graphs. *Ann. Probab.*, 23(1):87–104, 1995.
- [3] D. Avis, B. Davis, and J. M. Steele. Probabilistic analysis of a greedy heuristic for euclidean matching. *Probability in the Engineering and Information Sciences*, 2:143–156, 1988.
- [4] I. Benjamini, R. Lyons, Y. Peres, and O. Schramm. Group-invariant percolation on graphs. *Geom. Funct. Anal.*, 9(1):29–66, 1999.
- [5] S. Chatterjee, R. Peled, Y. Peres, and D. Romik. Gravitational allocation to Poisson points. *Ann. Math.*, arXiv:math.PR/0611886. To appear.
- [6] P. A. Ferrari, C. Landim, and H. Thorisson. Poisson trees, succession lines and coalescing random walks. *Ann. Inst. H. Poincaré Probab. Statist.*, 40(2):141–152, 2004.
- [7] A. Frieze, C. McDiarmid, and B. Reed. Greedy matching on the line. *SIAM J. Comput.*, 19(4):666–672, 1990.
- [8] D. Gale and L. Shapley. College admissions and stability of marriage. *Amer. Math. Monthly*, 69(1):9–15, 1962.

- [9] O. Häggström and R. Meester. Nearest neighbor and hard sphere models in continuum percolation. *Random Structures Algorithms*, 9(3):295–315, 1996.
- [10] C. Hoffman, A. E. Holroyd, and Y. Peres. Tail bounds for the stable marriage of Poisson and Lebesgue. *Canadian Journal of Mathematics*, arXiv:math.PR/0507324. To appear.
- [11] C. Hoffman, A. E. Holroyd, and Y. Peres. A stable marriage of Poisson and Lebesgue. *The Annals of Probability*, 34(4):1241–1272, 2006, arXiv:math.PR/0505668.
- [12] A. E. Holroyd and T. M. Liggett. How to find an extra head: optimal random shifts of Bernoulli and Poisson random fields. *Ann. Probab.*, 29(4):1405–1425, 2001.
- [13] A. E. Holroyd and Y. Peres. Trees and matchings from point processes. *Electron. Comm. Probab.*, 8:17–27 (electronic), 2003.
- [14] A. E. Holroyd and Y. Peres. Extra heads and invariant allocations. *Ann. Probab.*, 33(1):31–52, 2005.
- [15] O. Kallenberg. *Foundations of modern probability*. Probability and its Applications (New York). Springer-Verlag, New York, second edition, 2002.
- [16] D. E. Knuth. *Stable marriage and its relation to other combinatorial problems*, volume 10 of *CRM Proceedings & Lecture Notes*. American Mathematical Society, Providence, RI, 1997.
- [17] T. M. Liggett. Tagged particle distributions or how to choose a head at random. In *In and out of equilibrium (Mambucaba, 2000)*, volume 51 of *Progr. Probab.*, pages 133–162. Birkhäuser Boston, Boston, MA, 2002.
- [18] T. Soo. Translation-invariant matchings of coin-flips on Z^d . arXiv:math/0610334. Preprint.
- [19] M. Talagrand. The transportation cost from the uniform measure to the empirical measure in dimension ≥ 3 . *Ann. Probab.*, 22(2):919–959, 1994.

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

Alexander E. Holroyd: `holroyd(at)math.ubc.ca`
University of British Columbia, 121-1984 Mathematics Rd,
Vancouver BC V6T 1Z2, Canada

Robin Pemantle: `pemantle(at)math.upenn.edu`
David Rittenhouse Laboratories, 209 S 33rd St,
Philadelphia PA 19104, USA

Yuval Peres: `peres(at)microsoft.com`
Oded Schramm:
Microsoft Research, One Microsoft Way,
Redmond WA 98052, USA