Gravitational allocation on the sphere

Nina Holden[a], **Yuval Peres**[b,1], **and Alex Zhai**[c]

[a]Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA 02139; [b]Microsoft Research, Redmond, WA 98052; and [c]Department of Mathematics, Stanford University, Stanford, CA 94305

This contribution is part of the special series of Inaugural Articles by members of the National Academy of Sciences elected in 2016.

Contributed by Yuval Peres, August 3, 2018 (sent for review December 11, 2017; reviewed by Michel Ledoux and Mikhail Sodin)

Given a collection $\mathcal{L}$ of n points on a sphere $\mathbf{S}^2_n$ of surface area n, a fair allocation is a partition of the sphere into n parts each of area 1, and each is associated with a distinct point of $\mathcal{L}$. We show that, if the n points are chosen uniformly at random and if the partition is defined by a certain "gravitational" potential, then the expected distance between a point on the sphere and the associated point of $\mathcal{L}$ is $O(\sqrt{\log n})$. We use our result to define a matching between two collections of n independent and uniform points on the sphere and prove that the expected distance between a pair of matched points is $O(\sqrt{\log n})$, which is optimal by a result of Ajtai, Komlós, and Tusnády.

bipartite matching | allocation | transportation | gravity

Suppose that we are given n points on the unit sphere $\mathbf{S}^2 \subset \mathbf{R}^3$. We would like to partition the sphere into n equally sized cells, assigning each point to a different cell. How can we make this partition so that each point is close to the points in the cell to which it has been assigned? This natural question, known as the fair allocation problem, has connections to optimal transport and discretization (or "quantization") of continuous measures (1, 2). Allocation is also closely related to the matching problem, in which n red points and n blue points are chosen from the sphere (say, independently at random), and our goal is to pair each red point with a different blue point so as to make the distances between paired points as small as possible. Minimal matching for random points in the plane has generated a substantial literature in its own right (3–5).*

We analyze a particular allocation rule called gravitational allocation and apply it to matchings. Gravitational allocation is based on treating our n points as wells of a potential function. The cell allocated to a given point z is then taken to be the basin of attraction of z with respect to the flow induced by the negative gradient of this potential. When the potential takes a particular form that mimics the gravitational potential of Newtonian mechanics, it is ensured that each cell has area 1 (Fig. 1).

Related Work

The idea of transportation between measures via gradient flows dates back at least to Dacorogna and Moser (6). However, the first analysis that we know of concerning the resulting allocation cells was carried out by Nazarov, Sodin, and Volberg (7), who studied allocation to zeroes of a Gaussian analytic function.

The term gravitational allocation was introduced by Chatterjee, Peled, Peres, and Romik (8, 9), who studied fair allocations to a Poisson process in $\mathbf{R}^d$ with $d \geq 3$. In that setting, they proved exponential tail bounds on the diameter of a typical cell, showing that this diameter is of constant order.

The same does not hold when $d \leq 2$: it was shown in refs. 10 and 11 that, for translation invariant allocation schemes in $\mathbf{R}$ or $\mathbf{R}^2$, the expected allocation distance must be infinite. For this reason, to understand what happens when $d = 2$, it helps to consider a finite setting, such as the sphere. Suppose that we take the scaling where each cell has unit area. Then, it turns out that the typical allocation distance will need to be of at least order $\sqrt{\log n}$, which is also the same asymptotic behavior that is seen in minimal matching (3). In a recent paper, Ambrosio, Stra, and Trevisan (12) proved a more precise estimate of $\frac{\log n}{4\pi}$ for the expectation of the minimum average squared distance between random points and the uniform measure, confirming a prediction of Caracciolo, Lucibello, Parisi, and Sicuro (13).

Other than gravitational allocation, other allocation schemes have been proposed and analyzed, many based on the Gale–Shapley stable matching algorithm (11, 14–16).

Formal Definitions and Main Result

Let $\mathbf{S}^2_n \subset \mathbf{R}^3$ denote the sphere centered at the origin with surface area n, so that we work in the scaling where each cell has unit area. Let λ_n denote the surface area measure on $\mathbf{S}^2_n$, so that $\lambda_n(\mathbf{S}^2_n) = n$.

For any set $\mathcal{L} \subset \mathbf{S}^2_n$ consisting of n points, we say that a measurable function $\psi : \mathbf{S}^2_n \to \mathcal{L} \cup \{\infty\}$ is a fair allocation of λ_n to $\mathcal{L}$ if it satisfies the following:

$$\lambda_n(\psi^{-1}(\infty)) = 0, \qquad \lambda_n(\psi^{-1}(z)) = 1, \quad \forall z \in \mathcal{L}. \qquad [1]$$

For $z \in \mathcal{L}$, we call $\psi^{-1}(z)$ the cell allocated to z.

Let us now describe gravitational allocation in particular. First, we define a potential function $U : \mathbf{S}^2_n \to \mathbf{R}$ given by

$$U(x) = \sum_{z \in \mathcal{L}} \log |x - z|, \qquad [2]$$

where $|\cdot|$ denotes Euclidean distance in $\mathbf{R}^3$. For each location $x \in \mathbf{S}^2_n$, let $F(x)$ denote the negative gradient of U with respect

Significance

Given a set $\mathcal{L}$ of n points on the sphere, an allocation is a way to divide the sphere into n cells of equal area, each associated with a point of $\mathcal{L}$. Given two sets of n points $\mathcal{A}$ and $\mathcal{B}$ on the sphere, a matching is a bijective map from $\mathcal{A}$ to $\mathcal{B}$. Allocation and matching rules that minimize the distance between matched points are related to optimal transport and discretization of continuous measures. We define a matching and allocation rule by considering the gravitational field associated with the point configurations and show that they are optimal in expectation up to multiplication by a constant when our points are chosen independently and uniformly at random.

Author contributions: N.H., Y.P., and A.Z. performed research and wrote the paper.

Reviewers: M.L., University of Toulouse; and M.S., Tel Aviv University.

The authors declare no conflict of interest.

See Profile on page 9646.

[1]To whom correspondence should be addressed. Email: peres@microsoft.com.

Published online September 7, 2018.

*We note that many results are stated for points in a square or a 2D torus instead of the sphere. As $n \to \infty$, all of these settings are essentially equivalent. For the sake of consistency, in this article, we will state everything in terms of the sphere.

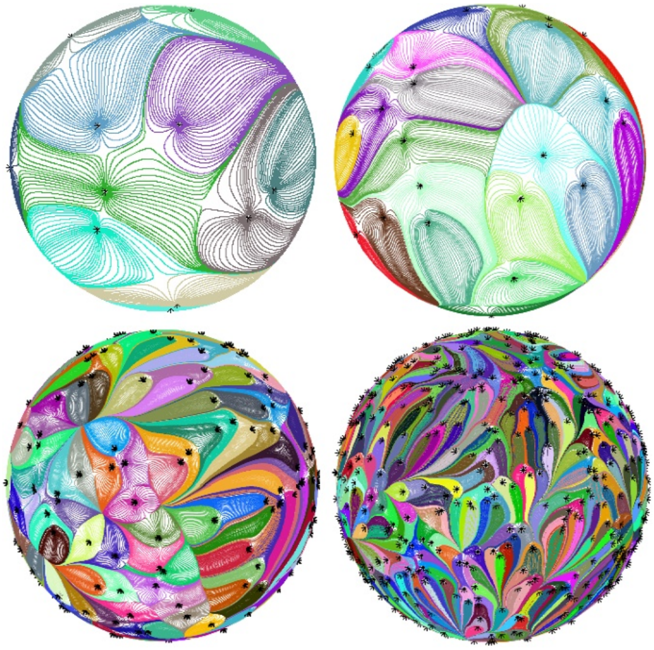

Fig. 1. Gravitational allocation to n uniform and independent points on a sphere with $n = 15, 40, 200$, and 750. The basin of attraction of each point has equal area. The basins become more elongated as n grows, reflecting *Theorem 1*. The MATLAB script used to generate the gravitational allocation figures in this article is based on code written by Manjunath Krishnapur.

to the usual spherical metric (i.e., the one induced from $\mathbf{R}^3$). We can view $F(x)$ as lying in the plane tangent to $\mathbf{S}_n^2$ at x (i.e., the tangent space), so that F is a vector field on $\mathbf{S}_n^2$.

Second, we consider the flow induced by F. For any $x \in \mathbf{S}_n^2$, let $Y_x(t)$ denote the integral curve that solves the differential equation

$$\frac{dY_x}{dt}(t) = F(Y_x(t)), \qquad Y_x(0) = x. \qquad [3]$$

By standard results about ordinary differential equations, the curve $Y_x(t)$ can be defined up until some maximal time τ_x (possibly $\tau_x = \infty$). In fact, τ_x will be finite for all x,[†] because by flowing along F, Y_x will eventually fall into one of the wells of the potential U (i.e., one of the points in $\mathcal{L}$) (Fig. 2).

We thus define the basin of attraction of $z \in \mathcal{L}$ as

$$B(z) = \left\{ x \in \mathbf{S}_n^2 : \lim_{t \uparrow \tau_x} Y_x(t) = z \right\} \qquad [4]$$

(i.e., the set of points that will eventually flow into z). We then define the gravitational allocation function to be

$$\psi(x) = \begin{cases} z & \text{if } x \in B(z) \text{ for } z \in \mathcal{L}, \\ \infty & \text{if } x \notin \bigcup_{z \in \mathcal{L}} B(z). \end{cases} \qquad [5]$$

It turns out that ψ indeed defines a fair allocation of λ_n to $\mathcal{L}$, so that each $B(z)$ has area 1. Before explaining why this is the case, let us first state our main result.

Theorem 1. *Let $n \geq 2$ be a positive integer. Consider any $x \in \mathbf{S}_n^2$, and let $\mathcal{L} \subset \mathbf{S}_n^2$ be a set of n points chosen uniformly and independently at random from $\mathbf{S}_n^2$. Then, there is a constant $C > 0$ such that*

$$\mathbf{E}|\psi(x) - x| \leq C\sqrt{\log n}. \qquad [6]$$

More generally, for any $p > 0$, there is a constant $C_p > 0$ depending only on p such that

$$\mathbf{E}|\psi(x) - x|^p \leq C_p(\log n)^{p/2}. \qquad [7]$$

Why Is Gravitational Allocation a Fair Allocation?

The reader may find it somewhat surprising that the basins of attraction in gravitational allocation always have equal areas, even if a point in $\mathcal{L}$ is crowded by many other points in $\mathcal{L}$ (Fig. 3). As seen in Fig. 3, the surrounded point will still attract certain faraway points, so that its basin of attraction still has total area 1.

We give two explanations for this phenomenon. Both explanations rely on the fact that our potential U satisfies the Poisson equation

$$\Delta_S U(x) = -2\pi + 2\pi \sum_{z \in \mathcal{L}} \delta_z,$$

where Δ_S denotes the spherical Laplacian (i.e., the Laplace–Beltrami operator on $\mathbf{S}_n^2$).

The first explanation is based on the divergence theorem. Consider any $z \in \mathcal{L}$ and its cell $B(z)$. Since $B(z)$ is a basin of attraction, F must be parallel to $B(z)$ along its boundary. We can then apply the divergence theorem[‡] to obtain

$$0 = -\int_{\partial B(z)} F \cdot \mathbf{n} ds = \int_{B(z)} \operatorname{div} F \, d\lambda_n$$

$$= \int_{B(z)} \Delta_S U d\lambda_n = 2\pi - 2\pi\lambda_n(B(z)).$$

It follows that $\lambda_n(B(z)) = 1$ as desired.

The second explanation is slightly longer, but it also provides a more detailed understanding of the flow under F. Imagine the surface area measure λ_n as representing the density of grains of sand uniformly distributed on the sphere. The sand is flowing along F, so that a grain of sand at x will be moved to location $Y_x(t)$ after time t.

In a small time ϵ, the net change in the density of sand at a point $x \in \mathbf{S}_n^2$ will be approximately

$$-\epsilon \operatorname{div} F(x) = \epsilon \Delta_S U(x) = -2\pi\epsilon + 2\pi\epsilon \sum_{z \in \mathcal{L}} \delta_z(x).$$

Thus, the density is decreasing everywhere at a uniform rate, except at points of $\mathcal{L}$, where sand is accumulating (at the same rate for each point). Integrating this over time, the density of sand at a time t will be given by

$$\lambda_{n,t} := e^{-2\pi t}\lambda_n + (1 - e^{-2\pi t}) \sum_{z \in \mathcal{L}} \delta_z.$$

We find that $\lim_{t \to \infty} \lambda_{n,t} = \sum_{z \in \mathcal{L}} \delta_z$, so that the amount of sand at each point in $\mathcal{L}$ tends to one. Consequently, the area of each basin of attraction must have been one.

Proof Outline of the Main Theorem

The proof of *Theorem 1* is based on estimating the magnitude of the gradient force F. In the previous section, we saw that, after time t, all but a $e^{-2\pi t}$ proportion of the sphere will have reached one of the points in $\mathcal{L}$, and therefore, the average time that it

[†]Except for a set of measure zero.

[‡]Assuming various smoothness properties, which we do not justify here.

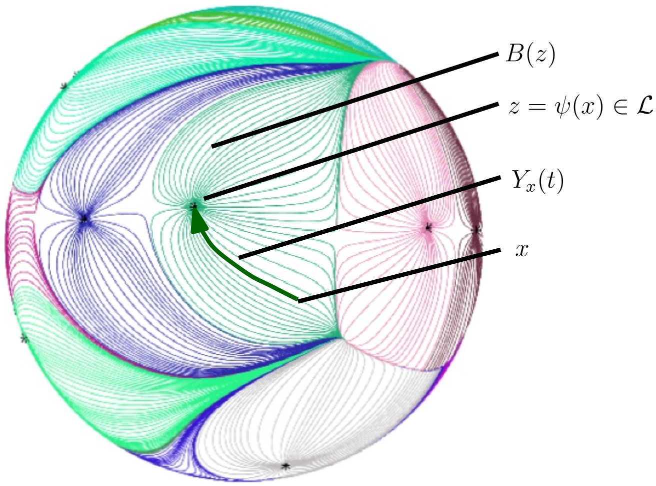

Fig. 2. Illustration of Y_x, $B(z)$, and $\psi(x)$ for $x \in \mathbf{S}^2_n$ and $z \in \mathcal{L}$.

takes for a point to flow into a potential well is $\int_0^\infty e^{-2\pi t}\,dt = 1/2\pi$. We can also estimate the average distance traveled in a similar way:

$$\begin{aligned}\int_{\mathbf{S}^2_n}\int_0^{\tau_x} |F(Y_x(t))|\,dt\,d\lambda_n(x) &= \int_0^\infty \int_{\mathbf{S}^2_n\setminus\mathcal{L}} |F(x)|\,d\lambda_{n,t}(x)\,dt\\ &= \int_0^\infty e^{-2\pi t}\int_{\mathbf{S}^2_n}|F(x)|\,d\lambda_n(x)\,dt\\ &= \frac{1}{2\pi}\int_{\mathbf{S}^2_n}|F(x)|\,d\lambda_n(x).\end{aligned} \quad [8]$$

It remains to estimate the average magnitude of $F(x)$, which is given by the following lemma.

Lemma 2. *Fix any $x \in \mathbf{S}^2_n$. Then, $\mathbf{E}|F(x)| = O(\sqrt{\log n})$, where the expectation is taken over the randomness of $\mathcal{L}$.*

Taking expectations in Eq. **8** and then integrating *Lemma 2* over all $x \in \mathbf{S}^2_n$ proves *Theorem 1* in the case $p = 1$. Larger values of p can be handled in the same spirit, but it requires more involved estimates for F that we do not reproduce here (ref. 17 has details).

***Proof*:** Let $U_z(x) = \log|x-z|$ and $F_z(x) = \nabla_S U_z(x)$, so that $U(x) = \sum_{z\in\mathcal{L}} U_z(x)$ and $F(x) = \sum_{z\in\mathcal{L}} F_z(x)$. Thus, $F_z(x)$ represents the contribution to $F(x)$ coming from the point $z \in \mathcal{L}$.

To estimate $F(x)$, it is convenient to decompose into the contributions of nearby and faraway points in $\mathcal{L}$. For our purposes, "near" means points within the spherical cap of radius 1 around x, which we denote by $B(x,1)$. Then, we may write

$$F(x) = \overbrace{\sum_{z\in\mathcal{L}\cap B(x,1)} F_z(x)}^{F_{\text{near}}(x)} + \overbrace{\sum_{z\in\mathcal{L}\setminus B(x,1)} F_z(x)}^{F_{\text{far}}(x)}. \quad [9]$$

When $|z - x| = r$, an explicit computation shows that $|F_z(x)|$ is of order $1/r$. It is also not hard to calculate that the expected number of points in $\mathcal{L}$ with distance from x that is between r and $r + dr$ is of order $r\,dr$. By the triangle inequality, we can estimate F_{near} as

$$\begin{aligned}\mathbf{E}|F_{\text{near}}(x)| &\le \mathbf{E}\sum_{z\in\mathcal{L}\cap B(x,1)}|F_z(x)| = \int_{B(x,1)}|F_y(x)|\,dy\\ &= O\left(\int_0^1 \frac{1}{r}\cdot(r\,dr)\right) = O(1).\end{aligned} \quad [10]$$

To estimate the far term, the triangle inequality is too weak, because we expect much cancellation between the $F_z(x)$. In fact, by symmetry, we have $\mathbf{E}[F_{\text{far}}(x)] = 0$. Thus, we instead estimate the second moment

$$\begin{aligned}\mathbf{E}|F_{\text{far}}(x)|^2 &= \mathbf{E}\sum_{z\in\mathcal{L}\setminus B(x,1)}|F_z(x)|^2 = \int_{\mathbf{S}^2_n\setminus B(x,1)}|F_y(x)|^2\,dy\\ &= O\left(\int_1^{\sqrt{n}}\frac{1}{r^2}(r\,dr)\right) = O(\log n).\end{aligned} \quad [11]$$

Combining Eqs. **9–11** yields

$$\mathbf{E}|F(x)| \le \mathbf{E}|F_{\text{near}}(x)| + \sqrt{\mathbf{E}|F_{\text{far}}(x)|^2} = O(\sqrt{\log n}),$$

which is the bound claimed in *Lemma 2*.

A Heuristic Picture

Lemma 2 also provides a good heuristic proof of Eq. **7**. We know by *Lemma 2* that, for a typical point x, we have $F(x) = O(\sqrt{\log n})$, and moreover, our above analysis suggests that the value of $F(x)$ is dominated by contributions from faraway points. Thus, we expect that direction and speed of travel for x under the flow induced by F will remain relatively constant.

However, x will not travel forever in this way; suppose that it passes within $O(1/\sqrt{\log n})$ distance of a point $z \in \mathcal{L}$. Then, the contribution $F_z(x)$ from z to the overall "force" F will be of order $\sqrt{\log n}$, which may overpower the contribution from all other points, causing x to fall into the potential well at z.

Consider a strip of width $1/\sqrt{\log n}$ around the path of x (Fig. 4). If there is a point $z \in \mathcal{L}$ in this strip, then it is likely to "swallow" x (i.e., x will be allocated to z). The probability that any given region contains no points of $\mathcal{L}$ decays exponentially in its area, which suggests the heuristic

$$\begin{aligned}&\mathbf{P}\left(x \text{ travels distance at least } r\sqrt{\log n}\right)\\ &\quad\approx \mathbf{P}\begin{pmatrix}\text{no points of } \mathcal{L} \text{ in a strip of area roughly}\\ r\sqrt{\log n}\cdot(1/\sqrt{\log n}) = r\end{pmatrix}\\ &\quad\approx e^{-r},\end{aligned}$$

giving Eq. **7**, because $|\psi(x) - x|$ is bounded above by the distance traveled by x.

From Allocations to Matchings

We now turn to the connection between fair allocations and optimal matchings. Suppose that $\mathcal{A} = \{a_1, \ldots, a_n\}$ and $\mathcal{B} = \{b_1, \ldots, b_n\}$ are two sets of n points in $\mathbf{S}^2_n$. A matching from $\mathcal{A}$ to $\mathcal{B}$ is a bijective function $\varphi : \mathcal{A} \to \mathcal{B}$. Recall that the matching

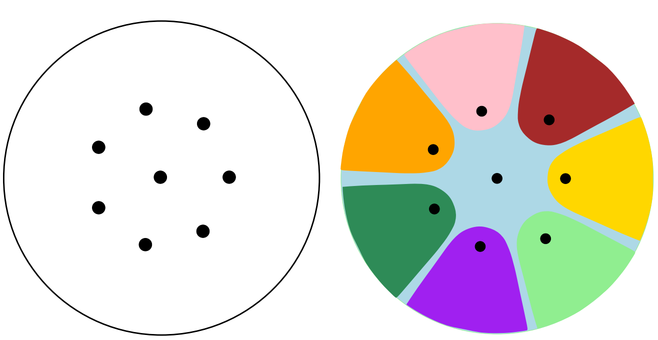

Fig. 3. The center point is surrounded by seven other nearby points (*Left*). Nevertheless, it turns out that its basin of attraction (light blue; *Right*) can slip past its neighbors in certain places.

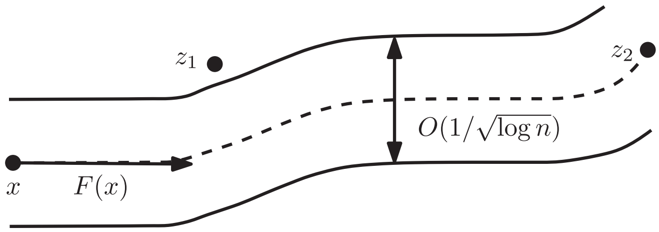

Fig. 4. The speed F is mainly determined by points far away and is approximately constant in large regions, except very near points of $\mathcal{L}$. A typical point, therefore, travels in an approximately straight line until it gets within distance $O(1/\sqrt{\log n})$ of some point in $\mathcal{L}$.

problem is to find the matching that minimizes the total distance between matched points.

When the points of $\mathcal{A}$ and $\mathcal{B}$ are drawn uniformly at random, the asymptotic behavior of the minimal matching distance was identified by Ajtai, Komlós, and Tusnády (3), who proved the following theorem.

Theorem 3 (Ajtai–Komlós–Tusnády). *Suppose that $\mathcal{A}$ and $\mathcal{B}$ each consist of n points drawn uniformly and independently at random from $[0, \sqrt{n}]^2$. Let*

$$d_{\text{match}}(\mathcal{A}, \mathcal{B}) = \min_{\substack{\varphi: \mathcal{A} \to \mathcal{B} \\ \text{bijective}}} \frac{1}{n} \sum_{a \in \mathcal{A}} |\varphi(a) - a|.$$

Then, there are constants $C_1, C_2 > 0$ for which

$$\lim_{n \to \infty} \mathbf{P}\left(C_1 \sqrt{\log n} \leq d_{\text{match}}(\mathcal{A}, \mathcal{B}) \leq C_2 \sqrt{\log n}\right) = 1. \qquad \textbf{[12]}$$

It turns out that the average displacement of a fair allocation gives an upper bound on the matching distance, as the next proposition shows.

Proposition 4. *Let $\mathcal{A}, \mathcal{B} \subset \mathbf{S}_n^2$ be two sets of n points, and let $\psi_{\mathcal{A}}$ and $\psi_{\mathcal{B}}$ be fair allocations of λ_n to $\mathcal{A}$ and $\mathcal{B}$, respectively. Then, there exists a matching $\varphi: \mathcal{A} \to \mathcal{B}$ such that*

$$\sum_{a \in \mathcal{A}} |a - \varphi(a)| \leq \int_{\mathbf{S}_n^2} |x - \psi_{\mathcal{A}}(x)| \, d\lambda_n(x) + \int_{\mathbf{S}_n^2} |x - \psi_{\mathcal{B}}(x)| \, d\lambda_n(x). \qquad \textbf{[13]}$$

Remark 5: Consider the case where $\mathcal{A}$ and $\mathcal{B}$ are drawn uniformly at random, and suppose that we use gravitational allocation for $\psi_{\mathcal{A}}$ and $\psi_{\mathcal{B}}$ in *Proposition 4*. Then, the $p = 1$ case of *Theorem 1* implies that the right-hand side of Eq. **13** has expectation of order $n\sqrt{\log n}$. Comparing with *Theorem 3*, this implies that the asymptotic rate of $\sqrt{\log n}$ in *Theorem 1* is the best possible up to a constant factor. By Eq. **8**, we also get that $\mathbf{E}|F(x)|$ is at least of order $\sqrt{\log n}$ for any fixed $x \in \mathbf{S}_n^2$.

The triangle inequality for the linear Wasserstein distance justifies why we can pass from an allocation to a matching, but we choose to describe the connection explicitly. Let $A_i = \psi_{\mathcal{A}}^{-1}(a_i)$ denote the cell allocated to a_i, and similarly, let $B_i = \psi_{\mathcal{B}}^{-1}(b_i)$. Consider the $n \times n$ matrix $M = (M_{ij})_{i,j=1}^n$ given by

$$M_{ij} = \lambda_n(A_i \cap B_j).$$

We see that M is a doubly stochastic matrix:

$$\sum_{j=1}^n M_{ij} = \sum_{j=1}^n \lambda_n(A_i \cap B_j) = \lambda_n(A_i) = 1,$$

$$\sum_{i=1}^n M_{ij} = \sum_{i=1}^n \lambda_n(A_i \cap B_j) = \lambda_n(B_j) = 1.$$

By the Birkhoff–von Neumann theorem (ref. 18, theorem 5.5), any doubly stochastic matrix is a convex combination of permutation matrices. For a permutation σ, we write P^σ to denote the corresponding permutation matrix, so that $P^\sigma_{ij} = 1$ if $j = \sigma(i)$ and $P^\sigma_{ij} = 0$ otherwise. Then, we may write

$$M = \sum_{k=1}^N c_k P^{\sigma_k}, \qquad \textbf{[14]}$$

where c_k are nonnegative numbers summing to one and σ_k are permutations.

Let X be chosen uniformly at random from $\mathbf{S}_n^2$. Observe that $n\mathbf{P}[X \in A_i \cap B_j] = M_{ij}$ and that $|\psi_{\mathcal{A}}(X) - \psi_{\mathcal{B}}(X)| = |a_i - b_j|$ on the event $X \in A_i \cap B_j$. By Eq. **14** and this observation,

$$\begin{aligned} \min_\sigma \sum_{i=1}^n \sum_{j=1}^n P^\sigma_{ij} |a_i - b_j| &\leq \sum_{i=1}^n \sum_{j=1}^n M_{ij} |a_i - b_j| \\ &= n\mathbf{E}|\psi_{\mathcal{A}}(X) - \psi_{\mathcal{B}}(X)|. \qquad \textbf{[15]} \end{aligned}$$

By the triangle inequality, the right side of Eq. **15** is bounded above by the right side of Eq. **13**, which implies *Proposition 4*.

Online Matching

One can also consider an "online" version of the matching problem, in which we initially see only the points in $\mathcal{B}$, and we are given the points in $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$ one by one. As soon as a_i is revealed to us, we must immediately match it to a point $\varphi(a_i)$ in $\mathcal{B}$ (that has not already been matched). In particular, we make this decision without knowing the locations of the remaining points in $\mathcal{A}$.

There is a natural online matching algorithm using gravitational allocation. When a point a_k is revealed, let $\mathcal{B}'$ be the set of points in $\mathcal{B}$ that have not yet been matched. We then consider the gravitational allocation $\psi_{\mathcal{B}'}$ to $\mathcal{B}'$ and match a_k to $\psi_{\mathcal{B}'}(a_k)$.

The analysis of this procedure is particularly simple if the points of $\mathcal{A}$ and $\mathcal{B}$ are sampled uniformly and independently at random. Consider what happens when we pair the first point a_1. According to *Theorem 1*, the expected distance between a_1 and its pair is bounded by

$$\mathbf{E}|a_1 - \varphi(a_1)| = \mathbf{E}|a_1 - \psi_{\mathcal{B}}(a_1)| \leq C\sqrt{\log n}.$$

Since $\psi_{\mathcal{B}}$ gives a fair allocation and the first point a_1 is drawn uniformly at random, each of the points in $\mathcal{B}$ is an equally likely match for a_1 under our scheme. It follows that the remaining points $\mathcal{B} \setminus \{\varphi(a_1)\}$ will still be distributed uniformly and independently at random. Thus, we have reduced the problem to matching two sets of $n-1$ independent random points on $\mathbf{S}_n^2$ after incurring a cost of $C\sqrt{\log n}$ for matching the first pair (Fig. 5).

We may iterate this analysis for each point in $\mathcal{A}$. When we receive a_k, there will be $m := n - k + 1$ remaining unpaired points in $\mathcal{B}$ (still uniformly distributed), so that a typical distance

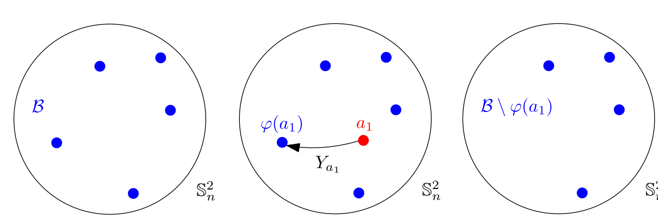

Fig. 5. Illustration of the online matching algorithm. The set $\mathcal{B} \setminus \varphi(a_1)$ consists of $n-1$ uniform and independent points on the sphere $\mathbf{S}_n^2$ of area n.

in gravitational allocation will be $O\left(\sqrt{n/m \cdot \log m}\right)$, where the factor $\sqrt{n/m}$ comes from rescaling $\mathbf{S}_m^2$ to $\mathbf{S}_n^2$. Thus,

$$\sum_{k=1}^{n} \mathbf{E}|a_k - \varphi(a_k)| \leq \sum_{m=2}^{n} O\left(\sqrt{n/m \cdot \log m}\right)$$

$$\leq O(\sqrt{n \log n}) \sum_{m=2}^{n} \frac{1}{\sqrt{m}} = O(n\sqrt{\log n}),$$

which shows that, even in the online setting, one has similar asymptotics as in *Theorem 3*.

We remark that our online matching algorithm can be implemented efficiently using the well-known "fast multipole method" introduced by Rokhlin (19) and Greengard and Rokhlin (20). This entails precomputing estimates of the gravitational potential from clusters of points in $\mathcal{B}$, and these computations can be reused as new points of $\mathcal{A}$ are introduced.

Gravitational Allocation for Other Point Processes

So far, we have focused on the setting where our n points on $\mathbf{S}_n^2$ are taken independently at random. However, one may also analyze other random point processes where the points are not independent, which allows them to be distributed more evenly over the sphere.

One example is given by the roots of a certain Gaussian random polynomial. Specifically, we look at the polynomial

$$p(z) = \sum_{k=0}^{n} \zeta_k \frac{\sqrt{n(n-1)\cdots(n-k+1)}}{\sqrt{k!}} z^k,$$

where $\zeta_1, \ldots, \zeta_n$ are independent standard complex Gaussians. The roots $\lambda_1, \ldots, \lambda_n$ of p are then n random points in the complex plane, which we can bring to the sphere via stereographic projection. More explicitly, let $x_0 = (0, 0, 1)$. The function

$$P : z \mapsto \sqrt{\frac{n}{4\pi}} \left(x_0 + \frac{2(z - x_0)}{|z - x_0|^2} \right)$$

maps the horizontal plane in $\mathbf{R}^3$ to $\mathbf{S}_n^2$. Then, viewing the λ_k as lying in the horizontal plane,

$$\mathcal{L} = \{P(\lambda_k)\}_{k=1}^{n}$$

is a rotationally equivariant random set of n points on $\mathbf{S}_n^2$. [The rotational equivariance comes from the particular choice of coefficients for p (ref. 21, chapter 2.3).]

Heuristically, the points of $\mathcal{L}$ are distributed more evenly than independent uniformly random points, because roots of random polynomials tend to "repel" each other (Fig. 6). This can be quantified as follows. Let $\psi : \mathbf{S}_n^2 \to \mathcal{L}$ be the gravitational allocation. Then, we claim that

$$\frac{1}{n} \mathbf{E} \int_{\mathbf{S}_n^2} |x - \psi(x)| \, d\lambda_n(x) = O(1). \qquad \textbf{[16]}$$

To prove this, by Eq. **8** and rotational symmetry, it suffices to show that $\mathbf{E}|F(x)| = O(1)$ for any point $x \in \mathbf{S}_n^2$. It is convenient to pick $x = (0, 0, -\sqrt{n/4\pi})$. Then, in the notation of the *Proof of Lemma 2*, we may calculate that

$$F_{\lambda_k}(x) = \sqrt{\frac{\pi}{n}} \cdot \bar{\lambda}_k^{-1},$$

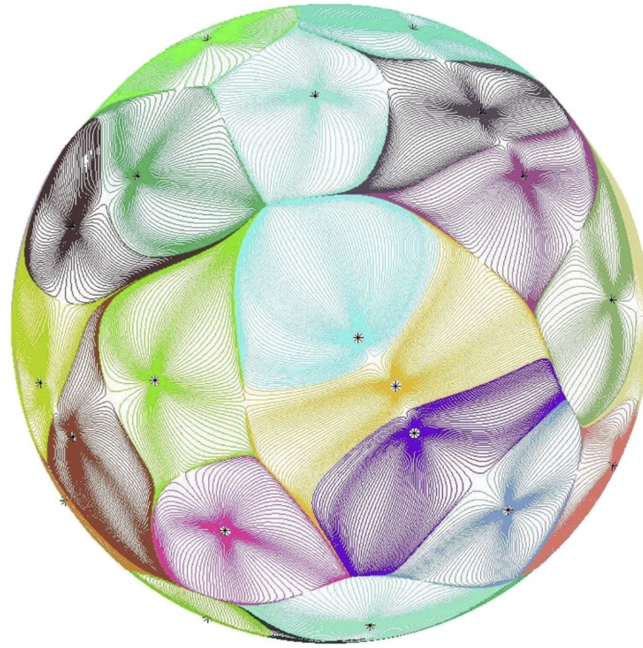

Fig. 6. A simulation of gravitational allocation to the zeroes of a Gaussian random polynomial. The cells are evenly proportioned, in contrast with the more elongated shapes seen in Fig. 1.

where we interpret the complex number on the right-hand side as a 2D vector. Thus, we have

$$F(x) = \sqrt{\frac{\pi}{n}} \sum_{k=1}^{n} \bar{\lambda}_k^{-1} = \sqrt{\frac{\pi}{n}} \cdot \frac{\bar{\zeta}_1 \cdot \sqrt{n}}{\bar{\zeta}_0 \cdot 1} = \sqrt{\pi} \cdot \frac{\bar{\zeta}_1}{\bar{\zeta}_0},$$

which gives a simple expression for F in terms of two independent complex Gaussians. Taking expectations of the magnitude, we obtain

$$\mathbf{E}|F(x)| = \sqrt{\pi}\mathbf{E}\frac{|\bar{\zeta}_1|}{|\bar{\zeta}_0|} = \frac{\pi\sqrt{\pi}}{2},$$

which establishes Eq. **16**.

Open Problems

We conclude by describing two other matching algorithms for which we do not know a precise analysis.

First, one may consider a dynamic electrostatic version of gravitational allocation. Suppose that the points in $\mathcal{A}$ ($\mathcal{B}$) are positive (negative) and that points of different (similar) kinds attract (repulse) each other. After some time, it seems that each point in $\mathcal{A}$ will collide with a point in $\mathcal{B}$, forming a matching. What will be the average distance between the original positions of matched pairs?

Second, in the online matching problem, instead of matching each new point a_k to a point in $\mathcal{B}$ according to gravitational allocation, suppose that we simply match a_k to the closest point in $\mathcal{B}$ that has not been matched already. Alternatively, we can reveal $\mathcal{A}$ and $\mathcal{B}$ simultaneously and iteratively match closest pairs of points. In other words, we choose $i, j \in \{1, \ldots, n\}$ such that $|a_i - b_j|$ is minimized, we define $\varphi(a_i) = b_j$, and we repeat with the sets $\mathcal{A} \setminus \{a_i\}$ and $\mathcal{B} \setminus \{b_j\}$. What will be the average matching distance in these settings? In the latter setting, ref. 16, theorem 6 suggests an upper bound for the matching distance of $\int_0^{\sqrt{n}} r^{-0.496\ldots} \, dr = \Theta(n^{0.252\ldots})$. Can this be improved?

ACKNOWLEDGMENTS. We thank Manjunath Krishnapur for useful discussions as well as for sharing his code for producing simulations. Most of this work was carried out while N.H. and A.Z. were visiting Microsoft Research; we thank Microsoft Research for the hospitality.

1. Huesmann M, Sturm KT (2013) Optimal transport from Lebesgue to Poisson. *Ann Probab* 41:2426–2478.
2. Dereich S, Scheutzow M, Schottstedt R (2013) Constructive quantization: Approximation by empirical measures. *Ann Inst Henri Poincaré Probab Stat* 49:1183–1203.
3. Ajtai M, Komlós J, Tusnády G (1984) On optimal matchings. *Combinatorica* 4:259–264.
4. Leighton T, Shor P (1989) Tight bounds for minimax grid matching with applications to the average case analysis of algorithms. *Combinatorica* 9:161–187.
5. Talagrand M (1994) Matching theorems and empirical discrepancy computations using majorizing measures. *J Am Math Soc* 7:455–537.
6. Dacorogna B, Moser J (1990) On a partial differential equation involving the Jacobian determinant. *Ann de l'Institut Henri Poincaré Analyse Non Linéaire* 7:1–26.
7. Nazarov F, Sodin M, Volberg A (2007) Transportation to random zeroes by the gradient flow. *Geom Funct Anal* 17:887–935.
8. Chatterjee S, Peled R, Peres Y, Romik D (2010) Gravitational allocation to Poisson points. *Ann Math* 172:617–671.
9. Chatterjee S, Peled R, Peres Y, Romik D (2010) Phase transitions in gravitational allocation. *Geom Funct Anal* 20:870–917.
10. Holroyd AE, Liggett TM (2001) How to find an extra head: Optimal random shifts of Bernoulli and Poisson random fields. *Ann Probab* 29:1405–1425.
11. Hoffman C, Holroyd AE, Peres Y (2006) A stable marriage of Poisson and Lebesgue. *Ann Probab* 34:1241–1272.
12. Ambrosio L, Stra F, Trevisan D (2018) A PDE approach to a 2-dimensional matching problem. *Probab Theory Relat Fields*. Available at https://link.springer.com/article/10.1007%2Fs00440-018-0837-x#citeas. Accessed September 4, 2018.
13. Caracciolo S, Lucibello C, Parisi G, Sicuro G (2014) Scaling hypothesis for the Euclidean bipartite matching problem. *Phys Rev E* 90:012118.
14. Holroyd AE, Peres Y (2005) Extra heads and invariant allocations. *Ann Probab* 33:31–52.
15. Hoffman C, Holroyd AE, Peres Y (2009) Tail bounds for the stable marriage of Poisson and Lebesgue. *Canad J Math* 61:1279–1299.
16. Holroyd AE, Pemantle R, Peres Y, Schramm O (2009) Poisson matching. *Ann Inst Henri Poincaré Probab Stat* 45:266–287.
17. Holden N, Peres Y, Zhai A (2017) Gravitational allocation for uniform points on the sphere. arXiv:170408238. Preprint, posted April 26, 2017.
18. van Lint JH, Wilson RM (2001) *A Course in Combinatorics* (Cambridge Univ Press, Cambridge, UK), 2nd Ed.
19. Rokhlin V (1985) Rapid solution of integral equations of classical potential theory. *J Comput Phys* 60:187–207.
20. Greengard L, Rokhlin V (1987) A fast algorithm for particle simulations. *J Comput Phys* 73:325–348.
21. Hough JB, Krishnapur M, Peres Y, Virág B (2009) *Zeros of Gaussian Analytic Functions and Determinantal Point Processes*, University Lecture Series (American Mathematical Society, Providence, RI), Vol 51.