

Concentration of Lipschitz functionals of determinantal and other strong Rayleigh measures

Robin Pemantle^{1, 2} and Yuval Peres³

Abstract: Let $\{X_1, \dots, X_n\}$ be a collection of binary valued random variables and let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be a Lipschitz function. Under a negative dependence hypothesis known as the *strong Rayleigh* condition, we show that $f - \mathbb{E}f$ satisfies a concentration inequality. The class of strong Rayleigh measures includes determinantal measures, weighted uniform matroids and exclusion measures; some familiar examples from these classes are generalized negative binomials and spanning tree measures. For instance, any Lipschitz-1 function of the edges of a uniform spanning tree on vertex set V (e.g., the number of leaves) satisfies the Gaussian concentration inequality $\mathbb{P}(f - \mathbb{E}f \geq a) \leq \exp\left(-\frac{a^2}{8|V|}\right)$. We also prove a continuous version for concentration of Lipschitz functionals of a determinantal point process.

Key words and phrases: negative association, stochastic domination, stochastic covering, Azuma-Hoeffding, balance, matroid, gaussian, random zero set, exclusion measure.

Subject classification 60G55.

¹Department of Mathematics, University of Pennsylvania, 209 South 33rd Street, Philadelphia, PA 19104, USA, pemantle@math.upenn.edu

²Research supported in part by NSF grant # DMS-0905937

³Microsoft Research, 1 Microsoft Way, Redmond, WA, 98052, USA, peres@microsoft.com

1 Introduction

Our goal in this paper is to prove concentration inequalities for Lipschitz functions of certain collections of negatively dependent binary-valued random variables. To illustrate our general methods we state our main result in a special case that was motivated by a question of E. Mossel (personal communication).

Theorem 1.1. *Let $G = (V, E)$ be a finite connected graph, let \mathbb{P} be the uniform measure on the spanning trees of G , and for $e \in E$ let X_e be the indicator function of the event that e is in the chosen spanning tree. Let $f : \{0, 1\}^E \rightarrow \mathbb{R}$ be any function with Lipschitz constant 1. Then*

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq \exp\left(-\frac{a^2}{8|V|}\right).$$

For example we might take f to be one half the number of vertices whose degree in the random tree is odd. This result is a consequence of more general results stated in Section 3 and Section 5.

1.1 Classical concentration inequalities

Let $\{X_n : n \geq 1\}$ be independent Bernoulli random variables with respective means $\{p_n\}$. Let $S_n := \sum_{k=1}^n X_k$ denote the partial sums, $\mu_n := \mathbb{E}S_n = \sum_{k=1}^n p_k$ denote the means and $V_n := \sum_{k=1}^n p_k(1-p_k)$ denote the variance of S_n . The simple and well known one-sided tail estimate for S_n is the classical Gaussian bound

$$\mathbb{P}(S_n - \mu_n \geq a) \leq \exp\left(-\frac{2a^2}{n}\right). \quad (1.1)$$

Replacing X_n with $1 - X_n$ gives the two-sided bound

$$\mathbb{P}(|S_n - \mu_n| \geq a) \leq 2 \exp\left(-\frac{2a^2}{n}\right). \quad (1.2)$$

The bound (1.1) may be found, among other places, in [McD89, Corollary 5.2]. The references given there include [Hoe63, (2.3)] as well as [Che52], which proves the result for identically distributed variables.

When p_n and $1 - p_n$ are bounded away from zero, the variance of S_n is of order n and this kind of bound is the best one can expect. However, when $n \gg \mu_n$, one might hope for uniformity in n via bounds in which the exponent depends on μ_n and not on n . For example, if $\max_{j \leq n} p_j$ is small then S_n is well approximated by a Poisson variable with mean μ_n . The upper tail of a Poisson is not as thin as a Gaussian, being $\exp[-\Theta(a \log(a/\mu))]$ rather than $\exp[-\Theta(a^2/\mu)]$. The bound

$$\mathbb{P}(S_n \geq a + \mu) \leq e^a \left(\frac{\mu}{a + \mu}\right)^{a + \mu} \leq \exp\left[-\frac{a^2}{2(a + \mu)}\right] \quad (1.3)$$

is proved in [Hoe63, Theorem 1] and asymptotically matches the Poissonian upper tail.

1.2 Generalizations

Our aim is to generalize (1.1) or its Poissonian version (1.3) in two ways. Instead of S_n we consider arbitrary Lipschitz functions of X_1, \dots, X_n , and instead of independent Bernoullis we consider a more general negatively dependent collection of binary random variables. We will give a number of applications, but before this, we briefly discuss what is known about each of the two generalizations separately.

For the first generalization, let \mathcal{B}_n denote the rank- n Boolean lattice $\{0, 1\}^n$ and let $f : \mathcal{B}_n \rightarrow \mathbb{R}$ be Lipschitz with respect to the Hamming distance. Replacing f by f/c if necessary, we will lose no generality in assuming our Lipschitz functions to have Lipschitz constant 1, and we do so hereafter; thus $|f(x) - f(x')| \leq 1$ whenever x and x' are two strings differing in only one position. When \mathbb{P} is a product measure, a well known generalization of (1.1) [McD89] is

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq e^{-2a^2/n}. \quad (1.4)$$

For the second generalization, we say that a collection of random variables $\{X_j\}$ in $\{0, 1\}$ is **negative cylinder dependent** if

$$\mathbb{P}(X_j = 1 \text{ for all } j \in S) \leq \prod_{j \in S} p_j \quad (1.5)$$

and

$$\mathbb{P}(X_j = 0 \text{ for all } j \in S) \leq \prod_{j \in S} (1 - p_j). \quad (1.6)$$

Negative cylinder dependence implies the inequalities (1.1)–(1.2) (see, e.g., [PS97, Theorem 3.4] with $\lambda = 1$). Lyons [Lyo03, Section 6] lists extensions and applications including one to balls in bins [DR98] and one to determinantal measures [Sos00].

It is not known whether these two generalizations can be combined. The random variables $\{X_j : 1 \leq j \leq n\}$ are said to be **negatively associated** if $\mathbb{E}fg \leq (\mathbb{E}f)(\mathbb{E}g)$ for every pair f, g of increasing functions on $\{0, 1\}^n$ such that $f(X_1, \dots, X_n)$ depends only on the values $\{X_i : i \in S\}$ and $g(X_1, \dots, X_n)$ depends only on the values $\{X_i : i \notin S\}$, for some subset $S \subseteq \{1, \dots, n\}$. By induction, this implies the weaker property of negative cylinder dependence. E. Mossel (personal communication, 2009) asked us whether the following holds.

Conjecture 1.2. *Let X_1, \dots, X_n be negatively associated binary-valued random variables. Let $f : \{0, 1\}^n \rightarrow \mathbb{R}$ be Lipschitz-1 and denote $f_n := f(X_1, \dots, X_n)$. Then (1.4) holds with the bound $\exp(-2a^2/n)$ replaced by $c_0 \exp(-ca^2/n)$ for some positive constants c_0 and c .*

To see why the exponent must be weakened consider the example of Bernoulli random variables X_1, \dots, X_n with n even, $\{X_1, \dots, X_{n/2}\}$ independent with mean $1/2$, and $X_{n/2+j} = 1 - X_j$ for

$1 \leq j \leq n/2$. These are negatively associated and yet the Lipschitz-1 function $f := \sum_{j=1}^{n/2} X_j - \sum_{j=n/2+1}^n X_j$ has tail probabilities on the order of $e^{-a^2/n}$. It is possible that this is the worst example and that the conjecture holds with $c = 1$, but a resolution of the conjecture would be interesting even without the optimal value of c .

Recent investigations of negative dependence properties indicate that negative association may not be sufficiently robust to use as a hypothesis in this context. The problem was posed in [Pem00] to find a more useful negative dependence property; this was answered in [BBL09], who showed that the **strong Rayleigh property** implies negative association and many other desirable consequences, and is stable under probabilistic operations such as conditioning, symmetrizing and reweighting.

Our main result implies that Conjecture 1.2 holds with $c = 1/8$ if one assumes the strong Rayleigh property rather than just negative association. The strong Rayleigh property is known to hold for most standard examples in which negative association is known to hold, so this gives up little generality, and moreover the strong Rayleigh property is usually easier to check than is negative association. Indeed for some of the measures described below, the only way we know they are negatively associated is by establishing the strong Rayleigh property. Several classes of measures satisfying the strong Rayleigh property are:

- Determinantal measures and point processes;
- Bernoullis conditioned on the sum;
- Measures obtained by running exclusion dynamics from a deterministic starting state (or more generally, exclusion with birth and death).

An overview of the rest of the paper is as follows. In the next section we introduce the strong Rayleigh property and discuss its consequences. One important consequence for us will be the **stochastic covering property**, which is all we use to derive our basic concentration inequality. In Section 3 we state our results, and these are proved in Section 4. Section 5 contains a number of applications.

2 Strong Rayleigh property, stochastic covering property, and other negative dependence conditions

Let $[n]$ denote $\{1, \dots, n\}$ and let $\mathcal{B}_n := \{0, 1\}^n$ denote the Boolean lattice of rank n , with coordinatewise partial order. The function $N : \mathcal{B}_n \rightarrow \mathbb{Z}^+$ will be used throughout to denote the counting function defined by $N(\omega) := \sum_{j=1}^n \omega_j$. A measure ν on \mathcal{B}_n is said to be **k -homogeneous** if ν is supported on the set of $\{\omega : N(\omega) = k\}$. The probability measure ν on \mathcal{B}_n is said to be **negatively**

associated if $\int fg d\nu \leq (\int f d\nu)(\int g d\nu)$ for every pair of nonnegative monotone functions f and g such that for some set $S \subseteq [n]$, the function f depends only on coordinates $\{\omega_j : j \in S\}$ while the function g depends only on coordinates $\{\omega_j : j \in S^c\}$.

The **strong Rayleigh condition** is said to hold for a measure \mathbb{P} on \mathcal{B}_n if the generating function

$$\sum_{\omega \in \mathcal{B}_n} \mathbb{P}(\omega) \prod_{j=1}^n z_j^{\omega_j}$$

has no roots (z_1, \dots, z_n) all of whose coordinates lie in the (strict) upper half plane. This and many consequences are given in [BBL09], including (implicitly) the **stochastic covering property** (see Proposition 2.2), which was conjectured [Pem00, Conjecture 9] to follow from something a little weaker. Some of the relevant implications are summarized in Figure 1 below.

The definition of the stochastic covering property requires a few preliminary definitions. Recall that a measure ν on a partially ordered set is said to **stochastically dominate** a measure ρ , denoted $\nu \succeq \rho$, if $\nu(A) \geq \rho(A)$ for every upwardly closed set A . An equivalent condition is that there exists a coupling, that is a measure Q on $\mathcal{B}_n \times \mathcal{B}_n$ with respective marginals ν and ρ , supported on the set $\{(x, y) : x \geq y\}$. If \mathbb{P} is a measure on \mathcal{B}_n making the coordinate variables $\{X_i : 1 \leq i \leq n\}$ negatively associated, an immediate consequence of negative association is that the conditional measure $(\mathbb{P} | X_n = 0)$ on \mathcal{B}_{n-1} stochastically dominates the conditional measure $(\mathbb{P} | X_n = 1)$.




Figure 1: relations among negative dependence properties

We say that the probability measure ν on \mathcal{B}_n **stochastically covers** another probability measure ρ if there is a measure on \mathcal{B}_n^2 with first marginal ν and second marginal ρ (in other words, a coupling) supported on the set of pairs (x, y) for which $x = y$ or x covers y in the coordinatewise partial order; here x is said to cover y when $x > y$ but there is no z such that $x > z > y$. We denote the covering relation in \mathcal{B}_n by $x > y$, and one measure covering another by $\nu \triangleright \rho$. Stochastic covering is strictly stronger than stochastic domination, and may be thought of as “stochastic domination, but by at most 1”.

Stochastic covering combines stochastic ordering with closeness in the so-called L^∞ -transportation metric, defined on probability measures on a given metric space as follows: $d_\infty(\mu, \nu)$ is the least ρ such that there is a coupling of μ and ν supported on the set $\{(x, y) : |x - y| \leq \rho\}$. Thus $\mu \triangleright \nu$ implies $d_\infty(\mu, \nu) \leq 1$. This is useful because if $\|f\|_{\text{Lip}}$ denotes the Lipschitz norm on Lipschitz functions,

then

$$\left| \int f d\mu - \int f d\nu \right| \leq \|f\|_{\text{Lip}} d_\infty(\mu, \nu). \quad (2.7)$$

Suppose that $x \geq y$ and we compare the conditional laws $\mathbb{P}_x := (\mathbb{P} | X_j = x_j, j \in S)$ and $\mathbb{P}_y := (\mathbb{P} | X_j = y_j, j \in S)$ on the remaining coordinates, that is as laws on $\{0, 1\}^{S^c}$. If \mathbb{P} and all its conditionalizations are negatively associated, it follows that $\mathbb{P}_x \preceq \mathbb{P}_y$.

Definition 2.1 (stochastic covering property). *We say that a probability measure ν on \mathcal{B}_n has the stochastic covering property if for every $S \subseteq \{1, \dots, n\}$ and for every $x, y \in \{0, 1\}^{S^c}$ with $x \succ y$, the conditional law $(\nu | X_j = x_j, j \in S)$ is covered by the conditional law $(\nu | X_j = y_j, j \in S)$.*

In [BBL09, Theorem 4.2] it was shown that strong Rayleigh property implies the **projected homogeneous Rayleigh** property (PHR), meaning that the measure can be embedded as the first n coordinates of a homogeneous measure ν' on \mathcal{B}_m , for some $m \geq n$, that has the ordinary Rayleigh property; the ordinary Rayleigh property is that the partial derivatives of the generating function $F(z_1, \dots, z_n) := \mathbb{E} \prod_{j=1}^n z_j^{X_j}$ satisfy $F_i F_j \geq F_{ij} F$ at any point with positive real coordinates. We record two further consequences.

Proposition 2.2. *PHR (and hence strong Rayleigh) implies the stochastic covering property.*

PROOF: PHR implies negative association of all conditionalizations (CNA) [BBL09, Theorem 4.10]; the homogeneous extension ν' witnessing the PHR property is also PHR hence also CNA. By negative association, if $x \succ y$ then $(\nu' | X_j = y_j, j \in S) \succeq (\nu' | X_j = x_j, j \in S)$, when viewed as measures on the coordinates in $[m] \setminus S$. Because ν' is homogeneous, the coupling that witnesses this \succeq relation in fact witnesses the relation \triangleright . Restricting to $[n] \setminus S$ we see that $(\nu | X_j = y_j, j \in S) \triangleright (\nu | X_j = x_j, j \in S)$. \square

Proposition 2.3 ([BBL09, Theorem 4.19]). *Let \mathbb{P} be a strong Rayleigh measure on \mathcal{B}_n . Let \mathbb{P}_k denote \mathbb{P} conditioned on $N = k$. Then for every $0 \leq k \leq n-1$ such that $\mathbb{P}(N = k)$ and $\mathbb{P}(N = k+1)$ are both nonzero, we have the covering relation $\mathbb{P}_{k+1} \triangleright \mathbb{P}_k$.* \square

3 Results

The chief consequence of the strong Rayleigh property that we use to prove concentration inequalities is the stochastic covering property. Although all of our examples so far of measures with the SCP are in fact strong Rayleigh, we note that this may not be the case in the future, and with this in mind, we state a result that uses only the SCP.

Theorem 3.1 (homogeneity and SCP implies Gaussian concentration). *Let \mathbb{P} be a k -homogeneous probability measure on \mathcal{B}_n satisfying the SCP. Let f be a Lipschitz-1 function on \mathcal{B}_n . Then*

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq \exp\left(-\frac{a^2}{8k}\right).$$

Remarks. (i) Replacing f with $-f$ gives immediate two-sided bounds:

$$\mathbb{P}(|f - \mathbb{E}f| > a) \leq 2 \exp\left(\frac{-a^2}{8k}\right). \quad (3.8)$$

(ii) Replacing every X_i by $1 - X_i$ if necessary, we may assume without loss of generality that $k \leq n/2$, whence

$$\mathbb{P}(f - \mathbb{E}f > a) \leq \exp\left(\frac{-a^2}{4n}\right). \quad (3.9)$$

For strong Rayleigh measures that are not necessarily homogeneous, we have the following result.

Theorem 3.2 (Gauss-Poisson bounds for general strong Rayleigh measures). *Let \mathbb{P} be strong Rayleigh with mean $\mu = \mathbb{E}N$. Let $f : \mathcal{B}_n \rightarrow \mathbb{R}$ be Lipschitz-1. Then*

$$\begin{aligned} \mathbb{P}(f - \mathbb{E}f > a) &\leq 3 \exp\left(-\frac{a^2}{16(a + 2\mu)}\right); \\ \mathbb{P}(|f - \mathbb{E}f| > a) &\leq 5 \exp\left(-\frac{a^2}{16(a + 2\mu)}\right). \end{aligned}$$

Remark. Because $a, \mu \leq n$, the denominator in these inequalities may be replaced by $48n$.

Continuous versions

Continuous versions of these results may be stated in terms of point processes, which we now briefly review. Formally, a point process on a space S is a random counting measure on S . In other words, a point process is a map \mathcal{Z} defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ taking values in the space of counting measures on S , a counting measure being one that takes only integer values or $+\infty$. Intuitively, one envisions the sample counting measure $\mathcal{Z}(\omega)$ as a set of points such that the sum of delta functions at these points is the sample counting measure.

If the number k of points in the support of \mathcal{Z} is deterministic, we may dispense with much of the formalism by ordering the points in the support of \mathcal{Z} uniformly at random and identifying the process \mathcal{Z} with the resulting exchangeable probability law on sequences of length k in S . Notationally, if \mathcal{Z} is a k -homogeneous point process on \mathbb{R}^d with law \mathbb{P} , we denote by \mathbb{P}_\uparrow the corresponding exchangeable law on $(\mathbb{R}^d)^k$. For $1 \leq j \leq k$, we use X_j to denote the “ j^{th} random point”, that is, the j^{th} coordinate function on $(\mathbb{R}^d)^k$. The following sampling algorithm for any k -homogeneous point process is almost trivial once one identifies \mathcal{Z} with \mathbb{P}_\uparrow , and yet it is a generalization of an algorithm previously proved only in the case of determinantal point process in [HKPV09, Proposition 4.4.3].

Lemma 3.3 (sampling in k steps). *Let \mathcal{Z} be a k -homogeneous point process on a standard Borel space S and let \mathbb{P}_\uparrow be the corresponding exchangeable measure on S^k . Then for $0 \leq j < k$ there are regular conditional distributions Q_{x_1, \dots, x_j} for the law of X_{j+1} given $X_1 = x_1, \dots, X_j = x_j$ such that the following procedure samples from \mathbb{P}_\uparrow .*

Sample X_1 from Q_\emptyset .

Recursively, conditional on $X_1 = x_1, \dots, X_j = x_j$, sample X_{j+1} from Q_{x_1, \dots, x_j} .

In the case where S is finite, let R denote the random set $\{X_1, \dots, X_n\}$. Then the law Q_{x_1, \dots, x_j} is equal to $1/(k-j)$ times the conditional intensity measure of $R \setminus \{x_1, \dots, x_j\}$ given $x_1, \dots, x_j \in R$.

PROOF: Any standard Borel space admits regular conditional distributions [Dur04, Theorem 4.1.6]. The sampling algorithm essentially restates the definition of regular conditional probabilities for sequential sampling. Because \mathbb{P}_\uparrow is exchangeable, conditioning on $X_{i_1} = x_1, \dots, X_{i_j} = x_j$ gives the same exchangeable measure on the sequence of remaining elements of R for any i_1, \dots, i_j . Thus for any x other than x_1, \dots, x_j , we have

$$\mathbb{P}(x \in R \mid x_1 \in R, \dots, x_j \in R) = (k-j)\mathbb{P}(X_{j+1} = x \mid X_1 = x_1, \dots, X_j = x_j)$$

which is the final conclusion. □

Remark. In the case of a measures on a finite set of size n , the main point of this sampling scheme is to sample in k steps rather than n steps, so as better to control the Azuma martingale. But also, sequential conditioning on $x \in R$ can be easy to compute. For example, conditioning on an edge being in a spanning tree replaces the original graph by a contraction along that edge.

Intuitively, there is a stochastic covering property for point processes defined to hold when conditioning on the presence of a point depresses the process everywhere else but by at most one point. To make this definition precise begin by extending the notion of one measure stochastically covering another to point processes. We say the point process \mathcal{Z} stochastically covers the point process \mathcal{W} if there is a coupling of these two laws on counting measures supported on pairs (μ, ν) such that $\mu = \nu$ or $\mu = \nu + \delta_x$ for some x . Metrizing the space of finite counting measures on S by the total variation distance, we see as before that if $\mathcal{Z} \triangleright \mathcal{W}$ then $d_\infty(\mathcal{L}(\mathcal{Z}), \mathcal{L}(\mathcal{W})) \leq 1$.

Next, given a k -homogeneous point process \mathcal{Z} on a space S , we let $\mathcal{Z}_{x_1, \dots, x_j}$ denote the $(k-j)$ -homogeneous point process whose law is the law of $\{X_{j+1}, \dots, X_k\}$ when sampling according to the procedure in Lemma 3.3 conditional on $X_1 = x_1, \dots, X_j = x_j$. The k -homogeneous point process \mathcal{Z} is said to have the **stochastic covering property** if

$$\mathcal{Z}_{x_1, \dots, x_j} \triangleright \mathcal{Z}_{x_1, \dots, x_{j+1}}$$

for all choices of x_1, \dots, x_{j+1} . Note that the left-hand side is $(k-j)$ -homogeneous while the right-hand side is $(k-j-1)$ -homogeneous.

Theorem 3.4. *Let \mathcal{Z} be a k -homogeneous point process on a standard Borel space S and let f be a Lipschitz-1 function (with respect to the total variation distance) on counting measures with total mass k on \mathbb{R}^d . If \mathcal{Z} has the SCP, then*

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq \exp\left(-\frac{a^2}{8k}\right).$$

For point processes that are not homogeneous, as in the discrete case, we require more than the SCP. Rather than defining a notion of strong Rayleigh here, we will stick to the case of determinantal point processes, this being where all of our examples arise; see Section 6 for definitions.

Theorem 3.5. *Let \mathcal{Z} be a determinantal point process with $\mathbb{E}N = \mu < \infty$. Let f be a Lipschitz-1 function on finite counting measures. Then*

$$\begin{aligned} \mathbb{P}(f - \mathbb{E}f \geq a) &\leq 3 \exp\left(-\frac{a^2}{16(a + 2\mu)}\right); \\ \mathbb{P}(|f - \mathbb{E}f| \geq a) &\leq 5 \exp\left(-\frac{a^2}{16(a + 2\mu)}\right). \end{aligned}$$

4 Proofs

4.1 The classical proofs

To prove bounds such as (1.1), one obtains an upper bound for $\mathbb{E}e^{\lambda S_n}$, and then applies Markov's inequality, choosing λ optimally. Underlying the bounds on $\mathbb{E}e^{\lambda S_n}$ are corresponding bounds for compensated increments. Let Δ denote a variable with mean zero. Three classical exponential bounds are as follows.

$$|\Delta| \leq 1 \quad \Rightarrow \quad \mathbb{E}e^{\lambda\Delta} \leq e^{\lambda^2/2} \tag{4.10}$$

$$\Delta \in [r, s] \quad \Rightarrow \quad \mathbb{E}e^{\lambda\Delta} \leq e^{\lambda^2(s-r)^2/8} \tag{4.11}$$

$$\Delta \in [r, s] \quad \Rightarrow \quad \mathbb{E}e^{\lambda\Delta} \leq \exp[(e^\lambda - 1 - \lambda) |rs|], \tag{4.12}$$

These are used together with the following two special cases of Markov's inequality.

$$\mathbb{E}e^{\lambda X} \leq e^{c\lambda^2/2} \quad \Rightarrow \quad \mathbb{P}(X \geq a) \leq e^{-a^2/(2c)} \tag{4.13}$$

$$\mathbb{E}e^{\lambda X} \leq e^{b(e^\lambda - \lambda - 1)} \quad \Rightarrow \quad \mathbb{P}(X > a) \leq e^a \left(\frac{b}{a+b}\right)^{a+b} \leq \exp\left[-\frac{a^2}{2(a+b)}\right] \tag{4.14}$$

These inequalities have appeared many times in the literature. Inequalities (4.10) and (4.13)

constitute the classical Azuma-Hoeffding inequality and imply

$$\mathbb{E}e^{\lambda(S_n - \mu_n)} \leq e^{\lambda^2 n/2}; \tag{4.15}$$

$$\mathbb{P}(S_n - \mu_n \geq a) \leq e^{-a^2/(2n)}. \tag{4.16}$$

This is valid for any martingale with differences bounded by 1; an exposition can be found in [AS08, Theorem 7.2.1]. The improvement to (4.11) is present already in [Hoe63], though the exposition in [McD89] is clearer (see Lemma 5.8 therein). When the increments of $S_n - \mu_n$ are compensated Bernoullis, one may take $b - a = 1$ rather than 2, resulting in an improvement by a factor of 4 in the exponent,

$$\mathbb{E}e^{\lambda(S_n - \mu_n)} \leq e^{\lambda^2 n/8}. \tag{4.17}$$

which together with (4.13) yields (1.1). Finally, (4.12) and induction yield

$$\mathbb{E}e^{\lambda(S_n - \mu_n)} \leq e^{(e^\lambda - \lambda - 1) V_n} \leq e^{(e^\lambda - \lambda - 1) \mu_n} \tag{4.18}$$

where $V_n := \sum_{k=1}^n p_k(1 - p_k)$ is the variance of S_n ; together with (4.14) this implies (1.3); these results appear, for instance, in [Fre75, (1.3)–(1.6)].

To prove the generalization to Lipschitz functions, let

$$M_k := \mathbb{E}(f(X_1, \dots, X_n) \mid X_1, \dots, X_k) - \mathbb{E}f(X_1, \dots, X_n).$$

It is immediate that $\{M_k\}$ is a martingale and that conditional on X_1, \dots, X_{k-1} , the two possible values of M_k differ by at most 1. Hence, conditional on X_1, \dots, X_{k-1} , the increment $\Delta_k := M_k - M_{k-1}$ is constrained to an interval of length at most 1. Applying (4.11) then yields (1.4).

The extension of inequalities (1.1)–(1.2) to negatively cylinder dependent random variables is established by examining the power series for $e^{\lambda S_n}$. This may be expanded into positive sums of expectations of products of powers of the variables $\{X_j : 1 \leq j \leq n\}$. Negative cylinder dependence implies that these are bounded from above by the corresponding products of expectations. Therefore, (4.17) and (4.18) hold when the assumption of independence is replaced by negative cylinder dependence, whence the probability inequalities (1.1) and (1.3) hold as well. This and more is shown in [PS97, Theorem 3.4], specializing their more general negative cylinder property to $\lambda = 1$. We remark that only the first inequality (1.5) in the definition of negative cylinder dependence is used to obtain bounds on $\mathbb{E}e^{\lambda S_n}$ for $\lambda > 0$, which suffices for the upper tail bounds. Lower tail bounds require these inequalities for $\lambda < 0$, for which the second inequality (1.6) is required.

4.2 Proof of Theorems 3.1 and 3.4

Theorem 3.1 is a special case of Theorem 3.4. This is because any probability measure μ on \mathcal{B}_n may be viewed as the law of a point process on the n element set $[n]$, where the random counting measure

$\mathcal{Z}(\omega)$ is defined by $\mathcal{Z}(\omega)(S) = \sum_{j \in S} \omega_j$. Informally, the points of the process are the coordinates of the ones in the sample ω . With this interpretation, the SCP on \mathcal{B}_n is inherited from the SCP for the point process \mathcal{Z} , whence Theorem 3.4 with $S = \mathbb{R}$ (or any other standard Borel space containing $[n]$) implies Theorem 3.1. It remains to prove Theorem 3.4.

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space on which is constructed the generalized sampling scheme described in Lemma 3.3. Let $\mathcal{F}_j := \sigma(X_1, \dots, X_j)$ and let

$$M_j := \mathbb{E}(f | \mathcal{F}_j) - \mathbb{E}f. \quad (4.19)$$

denote the martingale of sequential revelation. Applying the method of bounded differences is now mostly a matter of bookkeeping. At a sample point where $X_i = x_i, 1 \leq i \leq k$, the quantity M_j may be written as the integral of f against the law of the point process

$$\mathcal{Z}_{x_1, \dots, x_j} + \sum_{i=1}^j \delta_{x_i}.$$

By the SCP, we have $\mathcal{Z}_{x_1, \dots, x_j} \triangleright \mathcal{Z}_{x_1, \dots, x_{j+1}}$ whence

$$d_\infty \left(\mathcal{Z}_{x_1, \dots, x_j} + \sum_{i=1}^j \delta_{x_i}, \mathcal{Z}_{x_1, \dots, x_{j+1}} + \sum_{i=1}^{j+1} \delta_{x_i} \right) \leq 2.$$

By the Lipschitz assumption on f , it follows that $|M_{j+1} - M_j| \leq 2$. We now apply the basic Azuma-Hoeffding inequality (4.16) to $\{M_j/2\}_{1 \leq j \leq k}$ yielding

$$\mathbb{P}(f - \mathbb{E}f \geq a) = \mathbb{P}\left(\frac{M_k}{2} > \frac{a}{2}\right) \leq \exp\left(-\frac{a^2}{8k}\right).$$

□

4.3 Proof of Theorem 3.2

In this section we assume \mathbb{P} is the law of either a strong Rayleigh measure on \mathcal{B}_n with finite mean $\mathbb{E}N = \mu$. We also let f denote an arbitrary but fixed Lipschitz-1 function on configurations and define a function ϕ on \mathbb{Z}^+ by

$$\phi(k) := \mathbb{E}(f | N = k).$$

Lemma 4.1. *The variable N is distributed as the (possibly infinite) sum of independent Bernoullis.*

PROOF: In the definition of the strong Rayleigh property, setting the variables z_1, \dots, z_n equal produces a univariate polynomial with no roots in the upper half plane. As pointed out at the beginning of Section 3 of [BBL09], such a polynomial with real coefficients must have all its roots real. Since the coefficients are nonnegative, this implies it is the generating function for a convolution of Bernoullis. □

Lemma 4.2. *The variable N satisfies*

$$\mathbb{E}e^{\lambda(N-\mu)} \leq \exp [\mu(e^\lambda - 1 - \lambda)]$$

and consequently for any $a > 0$,

$$\mathbb{P}\left(N \geq \mu + \frac{a}{2}\right) \leq \exp\left(-\frac{(a/2)^2}{2(\mu + a/2)}\right).$$

PROOF: By Lemma 4.1 N is distributed as the sum of independent Bernoullis, which implies the first inequality; this implies the second inequality by (4.14). \square

Lemma 4.3. *The function ϕ is Lipschitz-1.*

PROOF: By Proposition 2.3 in the case of strong Rayleigh measures on \mathcal{B}_n we know that $\mathbb{P}_{k+1} \triangleright \mathbb{P}_k$. By definition of the stochastic covering relation, $(\phi(k+1), \phi(k))$ may be written as $\mathbb{E}(f(\eta), f(\xi))$ where $d(\eta, \xi) = 1$ almost surely. The conclusion then follows from the fact that f is Lipschitz-1. \square

Lemma 4.4. *The random variable $\phi(N)$ satisfies the concentration inequality*

$$\mathbb{E}e^{\lambda(\phi(N) - \mathbb{E}\phi(N))} \leq e^{\mu(e^\lambda - 1 - \lambda)}.$$

Consequently, the upper tails of $\phi(N)$ obey the bound

$$\mathbb{P}(\phi(N) - \mathbb{E}\phi(N) > t) \leq e^{-\frac{t^2}{2(t+\mu)}}.$$

PROOF: Pursuant to Lemma 4.1, let $\{Y_j\}$ be a finite or countably infinite collection of independent Bernoulli variables whose sum has the same law as N ; we may therefore prove the statements with N replaced by $\sum_j Y_j$. Write $\phi(\sum_j Y_j) - \mathbb{E}\phi(\sum_j Y_j)$ as the final term of a martingale $\{M_\ell\}$ where $M_\ell := \mathbb{E}(\phi(\sum_j Y_j) | \mathcal{F}_\ell) - \mathbb{E}\phi$ and $\mathcal{F}_\ell := \sigma(Y_1, \dots, Y_\ell)$. If the number of Bernoullis is infinite, the final term is a limit almost surely and in L^2 . The martingale $\{M_\ell\}$ is a binary martingale, meaning that conditional on \mathcal{F}_ℓ , the distribution of $M_{\ell+1}$ is concentrated on two values. In other words,

$$(M_{\ell+1} | \mathcal{F}_\ell) = p\delta_r + (1-p)\delta_s$$

where p is the mean of the Bernoulli variable $Y_{\ell+1}$. More importantly, $r = \int f d\mu$ and $s = \int f d\nu$ where μ is the conditional law of $\sum_j Y_j$ given the values of Y_1, \dots, Y_ℓ (which are measurable with respect to \mathcal{F}_ℓ) and given $Y_{\ell+1} = 1$, and ν is the conditional law of $\sum_j Y_j$ given the values of Y_1, \dots, Y_ℓ and given $Y_{\ell+1} = 0$. Clearly μ and ν are probability measures on \mathbb{Z}^+ satisfying $d_\infty(\mu, \nu) \leq 1$, whence because ϕ is Lipschitz-1, we see that $|r - s| \leq 1$. From (4.12) we then obtain

$$\mathbb{E}\left(e^{\lambda(M_{\ell+1} - M_\ell)} | \mathcal{F}_\ell\right) \leq \exp(p(1-p)(e^\lambda - 1 - \lambda)).$$

The lemma follows by induction. \square

PROOF OF THEOREMS 3.2 The event $\{f - \mathbb{E}f > a\}$ is contained in the union of three events:

$$\left\{N > \mu + \frac{a}{2}\right\} \cup \left\{\phi(N) - \mathbb{E}f > \frac{a}{2}\right\} \cup \left\{f - \phi(N) > \frac{a}{2}, N \leq \mu + \frac{a}{2}\right\}.$$

Thus $\mathbb{P}(f - \mathbb{E}f > a)$ is bounded above by the sum of the corresponding probabilities. Each of the first two pieces is bounded above by $\exp[-a^2/(4(a + 2\mu))]$: the first follows from Lemma 4.2 and the second uses Lemma 4.4, noting that $\mathbb{E}\phi = \mathbb{E} \mathbb{E}(f | N) = \mathbb{E}f$. The last piece is bounded above by $\exp[-a^2/(16(a + 2\mu))]$. To see this, observe that the measures \mathbb{P}_k are all strong Rayleigh (this is [BBL09, Corollary 4.18]). For any $k \leq \mu + a/2$, we can apply Theorem 3.1 to the homogeneous measure \mathbb{P}_k , obtaining

$$\mathbb{P}\left(f - \phi(N) > \frac{a}{2} \mid N = k\right) \leq \exp\left(-\frac{(a/2)^2}{8k}\right) \leq \exp\left(-\frac{a^2}{16(a + 2\mu)}\right).$$

Reassembling these gives the upper bound

$$\mathbb{P}\left(f - \phi(N) > \frac{a}{2}, N \leq \mu + \frac{a}{2}\right) \leq \exp\left(-\frac{a^2}{16(a + 2\mu)}\right).$$

This last piece has the worst bound; using it for all three pieces gives the first inequality of the theorem; we remark that the better upper bound of $2 \exp[-a^2/(4(a + 2\mu))] + \exp[-a^2/(16(a + 2\mu))]$ is in fact valid.

For the two-sided bound, we need to consider two more events in addition to the three already considered, namely the events $\{\phi - \mathbb{E}\phi < -a/2\}$ and $\{f - \phi(N) < -a/2, N \leq \mu + a/2\}$. The arguments for these two extra events are exactly analogous to two of the three arguments we have already seen, leading to a bound of $\exp[-a^2/(16(a + 2\mu))]$ for each of the two new summands and establishing the two-sided bounds. \square

5 Applications

In this section we discuss some classes of measures known to satisfy the hypotheses of our concentration results. The following Venn diagram gives a sense of how these classes intersect each other.

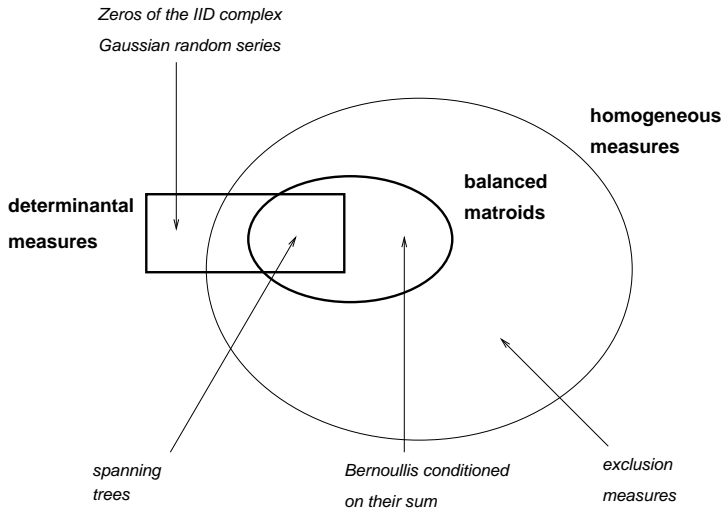


Figure 2: some classes of strong Rayleigh measures

5.1 Matroids

A collection \mathcal{C} of subsets of a finite set E , all of a given cardinality, k , is said to be the set of bases of a matroid if it satisfies the base exchange axiom (see, e.g., [Tut71]): if A and B are distinct members of \mathcal{C} and $a \in A \setminus B$, then there exists $b \in B \setminus A$ such that $A \cup \{b\} \setminus \{a\} \in \mathcal{C}$. Given a matroid, it is natural to consider the uniform measure on \mathcal{C} . More generally, the **weighted random base** is chosen from the probability measure

$$\nu_w(B) := C \prod_{e \in B} w(e),$$

where $\{w(e) : e \in E\}$ is a collection of nonnegative real numbers (weights) and C is a normalizing constant. Identifying E with the set $\{1, \dots, |E|\}$, the measure ν_w and the random variables $X_e := 1_{e \in B}$ can be thought of as living on $\mathcal{B}_{|E|}$.

For general matroids, $\mathbb{E}X_e X_f$ may be greater than $(\mathbb{E}X_e)(\mathbb{E}X_f)$. Some speculation has been given to the most natural class of matroids for which negative correlation or negative association must hold. Feder and Mihail [FM92] define a balanced matroid to be a matroid all of whose minors

satisfy pairwise negative correlation. Their proof of the following fact was the basis for the original proof of negative association for determinantal processes [Lyo03, Theorem 6.5].

Proposition 5.1 ([FM92, Theorem 3.2]). *The law ν_w of a random base of a balanced matroid, multiplicatively weighted by the weighting function w , has the SCP.* \square

Because measures supported on the bases of a matroid are homogeneous, there is nothing gained by improving the SCP to the strong Rayleigh property, and we have the following immediate corollary.

Corollary 5.2. *Let f be a Lipschitz-1 function with respect to Hamming distance on the bases of a balanced matroid of rank k on n elements. Then*

$$\mathbb{P}(f - \mathbb{E}f > a) \leq \exp\left(-\frac{a^2}{8k}\right).$$

\square

Example 5.3 (spanning trees). *One of the most important examples of a matroid is the set of spanning trees of a finite, connected, undirected graph. To spell this out, a spanning tree for a finite graph $G = (V, E)$ is a subset $E' \subseteq E$ such that (V, E') is a connected and acyclic. The set of spanning trees is a matroid on E . The weighted random spanning tree was shown to be a balanced matroid by [BP93, Theorem 1]. In fact they showed it is determinantal (see also [Lyo03, Example 1.1] and [HKPV09, Example 4.3.2]), though at the time consequences of being determinantal, such as the Strong Rayleigh property, had not been developed. Spanning trees are the only well known class of matroid whose uniform (or weighted) measure is determinantal.*

Let $f_0 : \{0, 1\}^E \rightarrow \mathbb{Z}$ count the number of vertices of odd degree in the graph defined by any subset of the edges. Deleting or adding an edge changes f_0 by at most 2. Let f be the random variable resulting from applying $(1/2)f_0$ to a the weighted random spanning tree on a graph G . Thus f is a Lipschitz 1 function that counts half the number of vertices that have odd degree in the random tree. Random variables that count local properties such as this are of natural graph theoretic interest. Parity counting variables similar to f play a role, for example in the randomized TSP approximation algorithm of [GSS11]. The number of edges in any spanning tree is $|V| - 1$. An application of Corollary 5.2 immediately gives the concentration inequality in Theorem 1.1; note that $|V|$, rather than $|E|$, appears in the denominator of the exponent.

Example 5.4 (conditioned Bernoullis, weighted matroids). *Let $\lambda_1, \dots, \lambda_n > 0$ be real numbers and let μ be the measure on the subsets of $[n]$ of cardinality k given by*

$$\mu(x_1, \dots, x_n) = \frac{\prod_{j=1}^n \lambda_j^{x_j}}{\sum_{N(\mathbf{y})=k} \prod_{j=1}^n \lambda_j^{y_j}}.$$

We may think of μ in two ways. The first is as a special case of the weighted random base, specialized to $M(n, k)$, the matroid whose bases are all the subsets of $[n]$ that have cardinality k . The second

is that it is the law of independent Bernoulli variables X_j with $\mathbb{E}X_j = \lambda_j/(1 + \lambda_j)$, conditioned on $\sum_{j=1}^n X_j = k$. We see from Proposition 5.1 that μ is strong Rayleigh. Alternatively, we may deduce this from the strong Rayleigh property for product measures along with closure under conditioning on the sum [BBL09, Corollary 4.18]. Restricting to $[m]$, for $m < n$, gives a joint distribution on \mathcal{B}_m which may be thought of as a multivariate generalization of the hypergeometric distribution. Because the strong Rayleigh property is inherited, this restriction is strong Rayleigh as well. The resulting concentration properties of these measures have been exploited in [GSS11] in connection with TSP approximation. We remark that more general conditioning, such as conditioning $N := \sum_j X_j$ to lie in an interval of more than two points, does not preserve the strong Rayleigh property.

5.2 Exclusion measures

The symmetric group \mathcal{S}_n acts on \mathcal{B}_n by permuting the coordinates. Suppose a nonnegative rate $r(\tau)$ is given for each transposition $\tau \in \mathcal{S}_n$. Define a random evolution on \mathcal{B}_n by letting each pair of coordinates (i, j) transpose independently at rate $r(\tau_{ij})$. In other words, we have a continuous time chain on \mathcal{B}_n which jumps from x to $\tau(x)$ at rate $r(\tau)$ for each transposition τ . This process is known as the symmetric exclusion process.

Borcea, Branden and Liggett [BBL09, Proposition 5.1] prove that the strong Rayleigh property is preserved under this evolution. In particular, because the point mass at a single state is always strong Rayleigh, it follows that the time t distribution of a symmetric exclusion process started from a deterministic state is strong Rayleigh. The stochastic covering property follows, as do PHR and negative association. Interestingly, before the publication of [BBL09], all that was known about this model was negative cylinder dependence [Lig77, Lemma 2.3.4].

Recently, it was shown by [Wag11] that one can add birth and death to the exclusion dynamics and still preserve the strong Rayleigh property. More specifically, let $\{\alpha_i, \beta_i : 1 \leq i \leq n\}$ be positive real numbers and let ω_i change to one at rate α_i and to zero at rate β_i , along with the exclusion dynamics. Then the evolution preserves the strong Rayleigh property and in particular, if the starting state is deterministic, all time t marginals are strong Rayleigh.

Corollary 5.5. *Let \mathbb{P} be the law on \mathcal{B}_n resulting from running an exclusion process for a fixed time, starting from a deterministic state with k sites occupied. Then*

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq e^{-a^2/(8k)}.$$

Example 5.6. *Let $n > 0$ be even and populate a $n \times n$ square of the integer lattice in \mathbb{Z}^2 (with torus boundary conditions) by filling all sites in the left half and leaving empty all sites in the right half. Run the symmetric exclusion process for time t with rate 1 on each edge. Let $f_t(\omega)$ denote the number of edges at time t with exactly one endpoint occupied. The mean of f_t starts at n at time 0 and approaches its limiting value of $n^2 - O(1)$ as $t \rightarrow \infty$. Once $t = \Theta(n^2)$, the variance of f_t*

becomes $\Theta(n^2)$ and the concentration inequality

$$\mathbb{P}(f - \mathbb{E}f \geq a) \leq e^{-a^2/(4n^2)},$$

which holds for all t , becomes a meaningful Gaussian tail bound (here $k = n^2/2$).

5.3 Determinantal measures on a finite Boolean lattice

We say that a probability measure \mathbb{P} on \mathcal{B}_n is **determinantal** (in the general sense) if there is an $n \times n$ real or complex matrix K such that that for every $S \subseteq \{1, \dots, n\}$,

$$\mathbb{E} \prod_{j \in S} X_j = \det K_S \tag{5.1}$$

where K_S is the submatrix of K obtained by choosing only those rows and columns whose index is in S . In this definition, the phrase “general sense” refers to the lack of further assumptions on K . An important subclass is the **Hermitian** determinantal measures, for which the matrix K is Hermitian. In this paper we will be interested only in the Hermitian case and will use the term **determinantal** hereafter to refer only to the case where K is Hermitian. Determinantal measures are known to be negatively associated [Lyo03, Theorem 6.5]. In fact they are strong Rayleigh [BBL09, proof of Theorem 3.4] and therefore satisfy the stochastic covering property.

Example 5.7 (uniform or weighted spanning tree). *As previously remarked, the uniform or weighted random spanning tree is a determinantal measure.*

In the next section we will extend the notion of a determinantal measure to the continuous setting. The extension to a countably infinite set of variables is more straightforward: the kernel K is now indexed by a countably infinite set, but (5.1) may be interpreted as holding for all finite sets S . The following example of a determinantal process on \mathbb{Z} appeared first in [Joh03].

Example 5.8 (positions of non-colliding RW’s). *Let $\{Y^{(k)} : 1 \leq k \leq n\}$ be n independent time homogeneous nearest neighbor random walks on \mathbb{Z} . Begin the walks at locations y_1, \dots, y_n and suppose the event that the walks are all at their starting positions at time $2n$ and have not intersected has positive probability. Conditional on this event, the positions at time n form a determinantal measure. That is, the indicator functions $\{X_j\}$ have a determinantal law, where $X_j = 1$ if some $Y^{(k)}$ is at position j at time n , and zero otherwise.*

Remark. The positions of non-colliding random walks are given by a determinant under more general conditions (see [KM59]). The present situation is arranged so as to make the kernel Hermitian.

6 Determinantal point processes

We consider here only simple point processes and often assume $\mathbb{E}N < \infty$ as well. If $\rho_k : (\mathbb{R}^d)^k \rightarrow \mathbb{R}^+$ are measurable functions, then the simple point process \mathcal{Z} is said to have joint intensities $\{\rho_k\}$ if for any k and any family D_1, \dots, D_k of disjoint Borel subsets of \mathbb{R}^d ,

$$\mathbb{E} \left[\prod_{j=1}^k \mathcal{Z}(D_j) \right] = \int_{\prod_j D_j} \rho_k(x_1, \dots, x_k) dx_1 \cdots dx_k.$$

In particular,

$$\mathbb{E}N = \int_{\mathbb{R}^d} \rho_1(x) dx$$

so under the assumption $\mathbb{E}N < \infty$, we see that $\rho_1(x) dx$ is a finite measure on \mathbb{R}^d . If ρ_1 is not finite, we will assume it is σ -finite. In any case, ρ_1 is called the **first intensity measure**; see [HKPV09, Sections 1.2 and 4.2] for further discussion of joint intensities and determinantal measures.

Definition 6.1 (determinantal point process). *A point process \mathcal{Z} is said to be determinantal if it has joint intensities $\{\rho_k\}$ and there is a measurable kernel $K : (\mathbb{R}^d)^2 \rightarrow \mathbb{C}$ such that*

$$\rho_k(x_1, \dots, x_k) = \det (K(x_i, x_j))_{1 \leq i, j \leq k}. \quad (6.2)$$

If $K(y, x) = \overline{K(x, y)}$ for every x, y , then the process is said to be Hermitian. When discussing determinantal processes below, we will always assume they are Hermitian.

Stochastic covering carries over to the continuous case. To state the relevant results we invoke the notion of the **Palm process**. This is a version of the process conditioned on the (measure zero) event of a point at a specified location, x . It may be obtained by conditioning on there being a point within distance ϵ of a given location x , then taking a weak limit. A more complete treatment may be found in [Kal86]. The following proposition is proved in [Gol10].

Proposition 6.2 ([Gol10]). *Suppose \mathcal{Z} is a determinantal point process with continuous kernel K and finite trace. Fix x and let \mathcal{Z}_x denote the Palm process that conditions on a point at x . Let \mathcal{Z}'_x denote the result of removing the point at x from \mathcal{Z}_x . Then*

- (i) *Whenever $K - L$ is positive semi-definite, the process with kernel K stochastically dominates the process with kernel L (this is [Gol10, Theorem 3]).*
- (ii) *\mathcal{Z}'_x is determinantal with kernel L such that $K - L$ is positive semi-definite.*
- (iii) *Consequently, $\mathcal{Z} \succeq \mathcal{Z}'_x$.*

The continuous analogue of Proposition 2.3 is

Proposition 6.3. *Let \mathcal{Z} be a determinantal point process with finite mean $\mathbb{E}N = \mu < \infty$. Then for any k for which $\mathbb{P}(N = k + 1)$ and $\mathbb{P}(N = k)$ are both nonzero, the conditional distributions of \mathcal{Z} given N satisfy*

$$(\mathcal{Z} \mid N = k + 1) \triangleright (\mathcal{Z} \mid N = k).$$

PROOF: The following facts may be found in [HKPV06, Theorem 7]. A determinantal point process \mathcal{Z} with mean $\mu < \infty$ has a kernel K whose spectrum is countable, contained in $[0, 1]$, and sums to μ . Furthermore, \mathcal{Z} may be represented as a mixture of homogeneous determinantal processes as follows. Let $\{\lambda_i : i \geq 1\}$ enumerate the eigenvalues with multiplicities and let $\{\phi_i\}$ be a corresponding eigenbasis. For each i , flip an independent coin with success probability λ_i . Let I denote the set of i for which the coin-flip was successful. Let K_I be the (random) projection operator onto the subspace spanned by the eigenvectors ϕ_i for which the coin-flip was successful. Then K_I is almost surely a projection of finite dimension $|I|$ and is the kernel of a $|I|$ -homogeneous determinantal point process. Choosing K_I at random and then sampling from the corresponding process recovers the law of \mathcal{Z} .

Several consequences are apparent. First, conditioning on $N = k$ is the same as conditioning on exactly k successes among the Bernoulli trials. Secondly, the conditional law of I given $|I| = k + 1$ stochastically dominates the conditional law of I given $|I| = k$. When the number of Bernoullis is finite, this follows from the strong Rayleigh property for independent Bernoullis; an easy limit argument extends the conclusion to the infinite case. This fact about stochastic domination is equivalent to saying that the conditional law of the random subspace K_I given $|I| = k + 1$ stochastically dominates the conditional law of the random subspace K_I given $|I| = k$, in the sense that the two laws can be coupled as (K, K') so that $K' \subseteq K$. When $K' \subseteq K$, the operator $\pi_K - \pi_{K'}$ is positive semi-definite. By (ii) of Proposition 6.2, we conclude that $(\mathcal{Z} \mid N = k + 1) \succeq (\mathcal{Z} \mid N = k)$ which is equivalent to stochastic covering in this case. \square

PROOF OF THEOREMS 3.5:

With f as in the statement of the theorem, and I the collection of indices described in the previous proposition, define $\psi(I)$ to be the expectation of f applied to a configuration chosen from the determinantal process with kernel K_I . Recall the notation $N = |I|$.

The event $\{f - \mathbb{E}f > a\}$ is contained in the union of three events:

$$\left\{N > \mu + \frac{a}{2}\right\} \cup \left\{\psi(I) - \mathbb{E}f > \frac{a}{2}\right\} \cup \left\{f - \psi(I) > \frac{a}{2}, N \leq \mu + \frac{a}{2}\right\}.$$

Thus $\mathbb{P}(f - \mathbb{E}f > a)$ is bounded above by the sum of the corresponding probabilities. Each of the first two pieces is bounded above by $\exp[-a^2/(4(a + 2\mu))]$: the first follows from Lemma 4.2 and the second follows from the proof of Lemma 4.4 because ψ is Lipschitz in the Bernoulli variables

$Y_i := \mathbf{1}_{i \in I}$. The last piece is bounded above by $\exp[-a^2/(16(a+2\mu))]$. To see this, apply Theorem 3.4 to the homogenous determinantal processes \mathbb{P}_I with kernels K_I , obtaining, when $|I| = k$, that

$$\mathbb{P}_I \left(f - \psi(I) > \frac{a}{2} \right) \leq \exp \left(-\frac{(a/2)^2}{8k} \right) \leq \exp \left(-\frac{a^2}{16(a+2\mu)} \right).$$

Reassembling these gives the upper bound

$$\mathbb{P} \left(f - \psi(I) > \frac{a}{2}, N \leq \mu + \frac{a}{2} \right) \leq \exp \left(-\frac{a^2}{16(a+2\mu)} \right).$$

The rest of the argument is identical to the conclusion of the proof of Theorem 3.2. \square

Example 6.4 (Ginibre’s translation invariant process). *Ginibre [Gin65] considers the distribution of eigenvalues of an $k \times k$ matrix with independent complex Gaussian entries. In the limit as $k \rightarrow \infty$, the density becomes constant over the whole plane. The limiting process \mathcal{Z} turns out to be a (Hermitian) determinantal point process with kernel*

$$K(z_1, z_2) := \frac{1}{\pi} e^{z_1 \bar{z}_2} \exp \left(-\frac{|z_1|^2 + |z_2|^2}{2} \right);$$

see, e.g. [Sos00, (2.16)]. The process \mathcal{Z} is ergodic and invariant under all rigid transformations of the plane. It was suggested [LCH90] to use this process as the set of centers for a random Voronoi tessellation because the mutual repulsion of the points makes the resulting tessellation more realistic than the standard Poisson-Voronoi tessellation for many purposes. Some rigorous results along these lines were obtained in [Gol10].

The mean number of points in any region D is $1/\pi$ times the area $|D|$, so the restriction of \mathcal{Z}_D to such a region of finite area is a determinantal process with finite mean number of points. Fix a finite region, D , and let f count the number of “lonely” points in D , these being such that no other point of \mathcal{Z} in D is within distance 1. We claim that f is Lipschitz with constant equal to 6. Clearly if a point z is added to the configuration η then f can increase by at most 1. It is well known that the maximum number of points in a unit disk that can be at mutual distance of at least 1 from one another is 6, which implies that the addition of z can result in the loss of at most 6 lonely points. Applying Theorem 3.5 to $f/6$ yields the concentration inequality

$$\mathbb{P}(|f - \mathbb{E}f| \geq a) \leq 5 \exp \left(-\frac{a^2}{96(a+12|D|/\pi)} \right).$$

Example 6.5 (Zeros of random polynomials). *Let $\{X_n\}$ be IID standard complex Gaussian random variables and define the random power series*

$$h(z) := \sum_{n=0}^{\infty} X_n z^n.$$

It is easy to see that h is almost surely analytic on the open unit disk and the number of zeros on any disk of radius $\rho < 1$ has finite mean. The remarkable properties of the point process \mathcal{Z} on the

unit disk that is the zero set of h are detailed in [PV05]. It is a determinantal process whose kernel is the Bergman kernel $\pi^{-1}(1 - z\bar{w})^{-2}$. It is invariant under Möbius transformations of the unit disk and has intensity measure $\pi^{-1}/(1 - |z|^2)^2$. Endowing the unit disk with the hyperbolic metric, the Möbius transformations become isometries, whence \mathcal{Z} is hyperbolic isometry invariant.

Fix $\rho < 1$ and $r > 0$ and let f count the number of zeros of the restriction \mathcal{Z}_ρ of \mathcal{Z} to the disk of radius $1 - \rho$ that are “hyperbolically lonely”, meaning that no other point of \mathcal{Z}_ρ is within a hyperbolic distance r . Let c_r denote the maximum number of points at mutual hyperbolic distance r that may be placed in a disk of hyperbolic radius r . Arguing as in Example 6.4 we see that f is Lipschitz with constant c_r . The mean number of points in \mathcal{Z}_ρ is $\rho^2/(1 - \rho^2)$ which for simplicity we can bound from above by $1/(1 - \rho^2)$. An application of Theorem 3.5 to f/c_r now yields

$$\mathbb{P}(|f - \mathbb{E}f| \geq a) \leq 5 \exp\left(-\frac{a^2}{16c_r a + 32c_r^2(1 - \rho^2)^{-1}}\right).$$

Acknowledgments

We thank Subhroshekhar Ghosh and an anonymous referee for detailed comments leading to the correction and improvement of the exposition and results.

References

- [AS08] N. Alon and J. Spencer. *The probabilistic method*. John Wiley & Sons, Hoboken, New Jersey, third edition, 2008.
- [BBL09] J. Borcea, P. Brändén, and T. Liggett. Negative dependence and the geometry of polynomials. *J. AMS*, 22:521–567, 2009.
- [BP93] R. Burton and R. Pemantle. Local characteristics, entropy and limit theorems for uniform spanning trees and domino tilings via transfer-impedances. *Ann. Probab.*, 21:1329–1371, 1993.
- [Che52] H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Statist.*, 23:493–507, 1952.
- [DR98] D. Dubhashi and D. Ranjan. Balls and bins: a study in negative dependence. *Rand. Struct. Alg.*, 13:99–124, 1998.
- [Dur04] R. Durrett. *Probability: Theory and Examples*. Duxbury Press, Belmont, CA, third edition, 2004.

- [FM92] T. Feder and M. Mihail. Balanced matroids. *Annual ACM Symposium on Theory of Computing*, pages 26–38, 1992.
- [Fre75] D. Freedman. On tail probabilities for martingales. *Ann. Probab.*, 3:100–118, 1975.
- [Gin65] J. Ginibre. Statistical ensembles of complex, quaternion, and real matrices. *J. Math. Phys.*, 6:440–449, 1965.
- [Gol10] A. Goldman. The palm measure and the Voronoi tessellation for the Ginibre process. *Ann. Appl. Prob.*, 20:90–128, 2010.
- [GSS11] S. Gharan, A. Saberi, and M. Singh. A randomized rounding approach to the Traveling Salesman Problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE Computer Science: Los Alamitos, CA, 2011.
- [HKPV06] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virag. Determinantal processes and independence. *Probability Surveys*, 3:206–229, 2006.
- [HKPV09] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virag. *Zeros of Gaussian analytic functions and determinantal point processes*, volume 51 of *University Lecture Series*. American Mathematical Society, Providence, RI, 2009.
- [Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Stat. Assoc.*, 58:13–30, 1963.
- [Joh03] K. Johansson. Discrete polynuclear growth and determinantal processes. *Comm. Math. Phys.*, 242:277–329, 2003.
- [Kal86] O. Kallenberg. *Random measures*. Academic Press, London, fourth edition, 1986.
- [KM59] S. Karlin and J. McGregor. Confidence probabilities. *Pacific J. Math.*, 9:1141–1164, 1959.
- [LCH90] G. Le Caer and J. S. Ho. The Voronoi tessellation generated from eigenvalues of complex random matrices. *J. Phys. A Math. Gen.*, 23:3279–3295, 1990.
- [Lig77] T. Liggett. The stochastic evolution of an infinite system of interacting particles. In *Ecole de Probabilites de Saint Flour VI, 1976*, volume 598 of *Lecture notes in Mathematics*, pages 187–248. Springer-Verlag, New York, 1977.
- [Lyo03] R. Lyons. Determinantal probability measures. *IHES*, 98:167–212, 2003.
- [McD89] C. McDiarmid. On the method of bounded differences. In J. Siemons, editor, *Surveys in Combinatorics*, volume 41 of *London Mathematical Society Lecture Note Series*, pages 148–188. Cambridge University Press, Cambridge, 1989.
- [Pem00] R. Pemantle. Toward a theory of negative dependence. *J. Math. Phys.*, 41:1371–1390, 2000.

- [PS97] A. Panconesi and A. Srinivasan. Randomized distributed edge coloring via an extension of the Chernoff–Hoeffding bounds. *SIAM J. Comput.*, 26:350–368, 1997.
- [PV05] Y. Peres and B. Virag. Zeros of the i.i.d. Gaussian power series: a conformally invariant determinantal process. *Acta Math.*, 194:1–35, 2005.
- [Sos00] A. Soshnikov. Determinantal random point fields. *Uspekhi Mat. Nauk*, 55:107–160, 2000.
- [Tut71] W. Tutte. *Introduction to the Theory of Matroids*, volume 37 of *Modern Analytic and Computational Methods in Science and Mathematics*. American Elsevier, New York, 1971.
- [Wag11] D. G. Wagner. Multivariate stable polynomials: theory and application. *Bull. Amer. Math. Soc.*, 48:53–84, 2011.