Random Graphs — Assignment 3

Yuval Filmus

April 6, 2020

Question 1 (Alternative proof of zero-one law)**.** The k-round Ehrenfeucht–Fraïssé game is played on a pair of graphs (G_1, G_2) by two players, Spoiler and Duplicator. We require both graphs to contain at least k vertices.

In the first round, Spoiler chooses a vertex on one of the graphs, and then Duplicator chooses a vertex on the other graph. We let a_1 be the vertex chosen in G_1, and b_1 be the vertex chosen in G_2.

In the second round, Spoiler again chooses a vertex on one of the graphs, and Duplicator chooses a vertex on the other graph. Both players are forced to choose vertices not already chosen (different from a_1, b_1). We let a_2, b_2 be the vertices chosen in G_1, G_2 (respectively).

All subsequent rounds proceed in the same way. After k rounds, we end up with k vertices $a_1, \ldots, a_k$ in G_1 and k vertices $b_1, \ldots, b_k$ in G_2. Duplicator's goal is that the following property is satisfied: for all i, j, there is an edge (a_i, a_j) in G_1 iff there is an edge (b_i, b_j) in G_2. We say that *Duplicator wins* if Duplicator has a strategy which guarantees that her goal is fulfilled.

For every first-order formula ϕ in the language of graphs, there is a constant k such that if Duplicator wins (G_1, G_2) then $G_1 \vdash \phi$ (that is, ϕ is satisfied for G_1) iff $G_2 \vdash \phi$.

(a) Show that for each k there is a function $e(n) = o(1)$ such that if $G_1 \sim G(n_1, 1/2)$ and $G_2 \sim G(n_2, 1/2)$ then Duplicator wins (G_1, G_2) with probability at least $1 - e(\min(n_1, n_2))$.

Answer. The argument is similar to the back-and-forth argument that we saw in class. Suppose without loss of generality that at round ℓ, Spoiler chooses a vertex a_ℓ from G_1. Let $A \subseteq [\ell - 1]$ be the indices of vertices adjacent to a_ℓ, and let $B \subseteq [\ell - 1]$ be the indices of vertices not adjacent to a_ℓ. Duplicator will attempt to find a vertex b_ℓ in G_2 which is adjacent to the vertices $\{b_i : i \in A\}$ and not adjacent to the vertices $\{b_j : j \in B\}$. If Duplicator is always successful, she wins the game.

Let $n = \min(n_1, n_2)$. The probability that a particular vertex in G_2 satisfies the conditions is $2^{-(\ell-1)}$, and so the probability that all remaining vertices do not satisfy the condition is at most $(1 - 2^{1-\ell})^{n-\ell+1} \leq e^{-(n-k)2^{-k}}$. Hence the probability that Duplicator fails to find a proper vertex at any stage is at most

$$e(n) = ke^{-(n-k)2^{-k}}.$$

It is easy to check that $e(n) = o(1)$, completing the proof. □

(b) Show that for each ϕ there is a function $e'(n) = o(1)$ such that for each n, either $\Pr[G(n,1/2) \vdash \phi] \le e'(n)$ or $\Pr[G(n,1/2) \vdash \phi] \ge 1 - e'(n)$.

Answer. Let k be the constant such that if Duplicator wins (G_1, G_2) then $G_1 \vdash \phi$ iff $G_2 \vdash \phi$. Let $p_n = \Pr[G(n,1/2) \vdash \phi]$. If $G_1, G_2 \sim G(n,1/2)$, then according to the previous item, with probability at least $1 - e(n)$ we have $G_1 \vdash \phi$ iff $G_2 \vdash \phi$. Since the probability of the latter event is by definition $p_n^2 + (1-p_n)^2$, we deduce

$$1 - e(n) \le p_n^2 + (1-p_n)^2 = 2p_n^2 - 2p_n + 1.$$

Rearranging, this gives

$$2p_n(1-p_n) \le e(n).$$

If $p_n \ge 1/2$ then $2p_n \ge 1$ and so $1 - p_n \le e(n)$, implying $p_n \ge 1 - e(n)$. Similarly, if $p_n \le 1/2$ then $p_n \le e(n)$. Hence we can take $e'(n) = e(n)$. □

(c) Show that for each ϕ, either $\Pr[G(n,1/2) \vdash \phi] \to 0$ or $\Pr[G(n,1/2) \vdash \phi] \to 1$.

Answer. We use the notation of the preceding item. In that item, we showed that for each n, either $p_n \le e(n)$ or $p_n \ge 1 - e(n)$. Let N be such that for each $n \ge N$, it holds that $e(n) < 1/3$; such an N exists since $e(n) \to 0$.

Suppose that $p_n \le e(n)$ and $p_m \ge 1 - e(m)$ for some $N \le n, m$. If we repeat the argument of the preceding item with $G_1 \sim G(n,1/2)$ and $G_2 \sim G(m,1/2)$, then we get

$$1 - e(\min(n,m)) \le p_n p_m + (1-p_n)(1-p_m) \le e(n) + e(m),$$

contradicting $e(n), e(m) < 1/3$.

It follows that either $p_n \le e(n)$ for all $n \ge N$ or $p_n \ge 1 - e(n)$ for all $n \ge N$. In the former case, $\Pr[G(n,1/2) \vdash \phi] \to 0$, and in the latter case, $\Pr[G(n,1/2) \vdash \phi] \to 1$. □

Question 2 (Failure of zero-one law for colored graphs)**.** A *colored graph* is a graph in which each vertex v has a color $c(v) \in \mathbb{N}$. Given a distribution π on $\mathbb{N}$, let $G_\pi(n,1/2)$ be the colored graph obtained by coloring each vertex in $G(n,1/2)$ according to π independently.

We say that two colored graphs G_1, G_2 are *isomorphic* if there is an isomorphism f of graphs between G_1 and G_2 that respects the coloring, that is, $c(v) = c(f(v))$.

(a) Show that if $G_1, G_2 \sim G_\pi(\aleph_0, 1/2)$ then almost surely, G_1 and G_2 are isomorphic (as colored graphs).

Answer. The argument is very similar to the proof given in class, and so will only be sketched. Let v_i be an enumeration of all vertices in both graphs. We will construct partial isomorphisms $\emptyset = f_0 \subseteq f_1 \subseteq \cdots \subseteq f_i \subseteq \cdots$, with the promise that $v_1, \ldots, v_i$ are in the support of f_i. The function $f = \bigcup_i f_i$ is then an isomorphism between G_1 and G_2.

The base case is $f_0 = \emptyset$. At step i, we have to ensure that v_i is in the support of f_i. If v_i is already in the support of f_{i-1}, then we can take $f_i = f_{i-1}$. Otherwise, suppose that $c(v_i) = c$, so that $\pi(c) > 0$. A given vertex on the other graph can be matched to v_i with probability at least $2^{-(i-1)}\pi(c) > 0$. Since there are infinitely many potential vertices, almost surely one of them can be matched to v_i, thus forming f_i.

The failure probability at each step is zero. Since there are only countably many steps, the total failure probability is also zero. In other words, the construction succeeds almost surely. □

(b) Let π have a Poisson distribution with expectation 1: $\Pr[\pi = k] = e^{-1}/k!$. Show that

$$\Pr[G_\pi(k!, 1/2) \text{ contains a color appearing exactly once}] \to e^{-1-e^{-1}}.$$

Answer. First, note that

$$\Pr[\pi > k] = e^{-1} \sum_{\ell=k+1}^{\infty} \frac{1}{\ell!} = \frac{e^{-1}}{(k+1)!} \sum_{\ell=0}^{\infty} \frac{1}{(k+1)\cdots(k+1+\ell)} \leq \frac{e^{-1}}{(k+1)!} \sum_{\ell=0}^{\infty} \frac{1}{\ell!} = \frac{1}{(k+1)!}.$$

It follows that the probability that any vertex has a color other than $0, \ldots, k$ is at most $k!/(k+1)! = 1/(k+1) = o(1)$.

Let $\ell \in \{0, \ldots, k-1\}$. The probability that at most one vertex is colored ℓ is at most

$$\left(1 - \frac{e^{-1}}{\ell!}\right)^{k!} + k! \frac{e^{-1}}{\ell!} \left(1 - \frac{e^{-1}}{\ell!}\right)^{k!-1} \leq e^{-e^{-1}k!/\ell!} + (1 - e^{-1})^{-1} e^{-1} \frac{k!}{\ell!} e^{-e^{-1}k!/\ell!} = O(ke^{-k/e}).$$

Since $k^2 e^{-k/e} = o(1)$, we conclude that with probability $1 - o(1)$, all colors in $\{0, \ldots, k-1\}$ appear at least twice.

Finally, the expected number of vertices colored k is e^{-1}, and so, as shown in class, the distribution of the number of vertices colored k tends to a Poisson distribution with expectation e^{-1}. Hence the probability that exactly one vertex is colored k tends to $e^{-1-e^{-1}}$. □

(c) The *first-order language of colored graphs* is defined similarly to the first-order language of graphs, together with the additional basic predicate $c(x) = c(y)$. Show that the zero-one law doesn't hold for the first-order language of colored graphs with respect to the sequence G_π, where π is the distribution from the preceding item.

Answer. Consider the following sentence:

$$\phi = \exists x \forall y (c(x) = c(y)) \Leftrightarrow (x = y).$$

The sentence expresses the fact that there exists a color which appears exactly once. As the previous item shows, $\Pr[G_\pi(k!, 1/2)]$ tends to a limit different from $0, 1$. □

Question 3 (Quasirandom permutations[1]). The symmetric group S_n consists of all permutations of $[n] := \{1, \ldots, n\}$. We think of permutations as sequences of length n.

For a permutation $\pi \in S_n$ and a "pattern" $\tau \in S_k$, the density $t(\pi, \tau)$ is the probability that if we sample k distinct indices $i_1, \ldots, i_k \in [n]$ then the relative order of $\pi(i_1), \ldots, \pi(i_k)$ is the same as τ. For example,

$$t(13245, 123) = \frac{7}{10}, \quad t(13245, 213) = \frac{2}{10}, \quad t(13245, 132) = \frac{1}{10},$$

since $134, 135, 124, 125, 145, 345, 245$ have relative order 123; $324, 325$ have relative order 213; and 132 has relative order 132.

For each n, let $\pi_n \in S_n$. The sequence $\vec{\pi}$ is *k-quasirandom* if for all $\tau \in S_k$,

$$t(\pi_n, \tau) \to \frac{1}{k!}.$$

The sequence $\vec{\pi}$ is *quasirandom* if it is k-quasirandom for each k. As an example, if π_n is chosen uniformly random permutation for each n, then $\vec{\pi}$ is quasirandom almost surely.

(a) Show that if $\vec{\pi}$ is $(k+1)$-quasirandom then it is k-quasirandom.

Answer. We can sample a k-tuple of indices by first sampling a $(k+1)$-tuple of indices and then removing a random index. This implies that if $\tau \in S_k$ then

$$t(\pi, \tau) = \sum_{\sigma \in S_{k+1}} t(\pi, \sigma) t(\sigma, \tau).$$

If we choose σ at random from S_{k+1} then $\mathbb{E}_\sigma[t(\sigma, \tau))] = 1/k!$. Hence

$$t(\pi_n, \tau) = \sum_{\sigma \in S_{k+1}} t(\pi_n, \sigma) t(\sigma, \tau) \to \sum_{\sigma \in S_{k+1}} \frac{1}{(k+1)!} t(\sigma, \tau) = \frac{1}{k!}.$$ □

(b) Give an example of a 2-quasirandom sequence which is not 3-quasirandom.[2]

[1] After Král' and Pikhurko, Quasirandom permutations are characterized by 4-point densities, GAFA vol. 23, pp. 570–579, 2013.

[2] There are also examples of 3-quasirandom sequences which are not 4-quasirandom, but they are more complicated. One example is described in the paper of Král' and Pikhurko mentioned above, and another one in Cooper and Petrarca, Symmetric and asymptotically symmetric permutations.

Answer. Let π_n be the following permutation:

$$\lfloor \tfrac{n}{2} \rfloor + 1, \lfloor \tfrac{n}{2} \rfloor + 1, n, 1, 2, \ldots, \lfloor \tfrac{n}{2} \rfloor.$$

This permutation is composed of two halves, of sizes $\lceil \frac{n}{2} \rceil, \lfloor \frac{n}{2} \rfloor$. If we choose two indices $i < j$ at random, they fall on the same half with probability $1/2 \pm O(1/n)$ and on different halves with probability $1/2 \pm O(1/n)$. In the former case, $\pi(i) < \pi(j)$, and in the latter case, $\pi(i) > \pi(j)$. Therefore the sequence is 2-quasirandom.

In contrast, $t(\pi_n, 132) = t(\pi_n, 213) = t(\pi_n, 321) = 0$, and so the sequence is not 3-quasirandom. □

A *permuton* is a probability distribution μ over $[0,1]^2$ such that if $(x,y) \sim \mu$ then the marginal distributions of x and y are uniform over $[0,1]$. Given a permuton μ, for each n we can draw a random permutation $\pi \sim P(n,\mu)$ as follows. Let $(x_1, y_1), \ldots, (x_n, y_n)$ be n independent samples of μ. We arrange the x_i in order, and let π consist of the relative order of the y_i. (Since the marginal distributions are uniform over $[0,1]$, almost surely all x_i and all y_i are distinct.) For example, if μ is the uniform distribution over $[0,1]^2$ then $P(n,\mu)$ is a uniformly random permutation in S_n.

For $\tau \in S_k$, let $t(\mu,\tau)$ be the probability that if we take k samples (x_i, y_i) from μ and arrange the x_i in order, then the relative order of the y_i is τ.

(c) Show that $\mathbb{E}_{\pi \sim P(n,\mu)}[t(\pi,\tau)] = t(\mu,\tau)$. (In fact, more is true: if $\pi_n \sim P(n,\mu)$ for each n independently, then almost surely $t(\pi,\tau) \to t(\mu,\tau)$.)

Answer. We can choose π together with a random k-subset of π by drawing n samples (x_i, y_i) from μ, choosing k of them, arranging the x_i in order, and letting π be the relative order of the y_i. The expected value of $t(\pi,\tau)$ is then the probability that the relative order of the chosen samples is τ. But this process is the same as the one used to define $t(\mu,\tau)$. □

A permuton μ is *k-quasirandom* if for each $\tau \in S_k$, $t(\mu,\tau) = 1/k!$. A permuton μ is quasirandom if it is k-quasirandom for all k. As in the case of individual distributions, it is not hard to show that a $(k+1)$-quasirandom permuton is also k-quasirandom. Furthermore, if μ is a (k-)quasirandom permuton and for each n we sample $\pi_n \sim P(n,\mu)$, then almost surely $\vec{\pi}$ is (k-)quasirandom (where we think of π_n as a constant random variable).

In the rest of this exercise, we show that if μ is a 4-quasirandom permuton, then it is in fact quasirandom (the constant 4 is optimal). This implies that if $\vec{\boldsymbol{\pi}}$ is a 4-quasirandom sequence of random permutations, then it is in fact quasirandom.

(d) Let $F_\mu(X,Y) = \Pr_{(x,y)\sim\mu}[x \le X, y \le Y]$ be the CDF of μ. Show that

$$\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] = \Pr_{(x_i,y_i)\sim\mu}[x_1, x_2 \le x_3;\ y_1, y_2 \le y_3].$$

Answer. For fixed X, Y,

$$F_\mu(X,Y)^2 = \Pr_{(x,y)\sim\mu}[x \le X, y \le Y]^2 =$$
$$\Pr_{(x_1,y_1)\sim\mu}[x_1 \le X, y_1 \le Y] \Pr_{(x_2,y_2)\sim\mu}[x_2 \le X, y_2 \le Y] = \Pr_{\substack{(x_1,y_1)\sim\mu\\(x_2,y_2)\sim\mu}}[x_1, x_2 \le X, y_1, y_2 \le Y].$$

The required formula follows by taking expectation over $(X, Y) \sim \mu$. □

(e) Deduce that $\mathbb{E}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] = 1/9$, using only the fact that μ is 3-quasirandom.

Answer. We can write

$$\Pr_{(x_i,y_i)\sim\mu}[x_1, x_2 \le x_3, y_1, y_2 \le y_3] =$$
$$\frac{1}{6}\Pr_{(x_i,y_i)\sim\mu}[y_1, y_2 \le y_3 \mid x_1 \le x_2 \le x_3] + \frac{1}{6}\Pr_{(x_i,y_i)\sim\mu}[y_2, y_1 \le y_3 \mid x_2 \le x_1 \le x_3],$$

since the probability that x_1, x_2, x_3 have any particular order is $1/6$. Since μ is 3-quasirandom, each of the probabilities above is $2/6 = 1/3$, for a total of $1/6 \cdot 1/3 + 1/6 \cdot 1/3 = 1/9$. □

(f) Show that

$$\mathbb{E}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY] = \Pr_{(x_i,y_i)\sim\mu}[x_1, x_2 \le x_4;\ y_1, y_3 \le y_4].$$

Hint: if $(x, y) \sim \mu$ then since the marginals x and y are uniform over $[0, 1]$, then $\Pr[x \le X] = X$ and $\Pr[y \le Y] = Y$.

Answer. For fixed X, Y,

$$F_\mu(X,Y)XY = \Pr_{(x_1,y_1)\sim\mu}[x_1 \le X, y_1 \le X] \Pr_{(x_2,y_2)\sim\mu}[x_2 \le X] \Pr_{(x_3,y_3)\sim\mu}[y_3 \le Y] =$$
$$\Pr_{(x_i,y_i)\sim\mu}[x_1, x_2 \le X, y_1, y_3 \le Y].$$

The required formula follows by taking expectation over $(X, Y) \sim \mu$. □

(g) Deduce that $\mathbb{E}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY] = 1/9$, using the fact that μ is 4-quasirandom.

Answer. We use the formula from the preceding item. The probability that $x_1, x_2 \le x_4$ is $1/3$. Fixing any order of the x_i, the relative order of the y_i is uniform, since μ is 4-quasirandom. Hence the probability that $y_1, y_3 \le y_4$ is $1/3$. In total, the probability is $1/3 \cdot 1/3 = 1/9$. □

(h) Let λ be the permuton corresponding to two independent samples of the uniform distribution over $[0,1]$. Show that λ is quasirandom and $F_\lambda(X,Y) = XY$.

Answer. Let $(x_1, y_1), \ldots, (x_n, y_n)$ be independent samples from λ. The relative order of the y_i after ordering the x_i has the same distribution as the relative order of the y_i conditioned on $x_1 < x_2 < \cdots < x_n$. Since the y_i are independent from the x_i, this relative order is uniformly random. This implies that λ is quasirandom.

If z is sampled from the uniform distribution over $[0,1]$ then $\Pr[z \le Z] = Z$, and this implies the formula $F_\lambda(X,Y) = XY$. □

(i) Show that

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2] = \Pr_{(X_i,Y_i)\sim\mu}[x_1, x_2 \le x_3;\ y_1, y_2 \le y_4].$$

Hint: use two samples of μ to generate one sample of λ.

Answer. If $(x_3, y_3), (x_4, y_4) \sim \mu$ then $(x_3, y_4) \sim \lambda$. Hence

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2] = \mathop{\mathbb{E}}_{\substack{(x_3,y_3)\sim\mu\\(x_4,y_4)\sim\mu}}[F(x_3,y_4)^2].$$

For fixed x_3, y_4,

$$F_\mu(x_3,y_4)^2 = \mathop{\mathbb{E}}_{\substack{(x_1,y_1)\sim\mu\\(x_2,y_2)\sim\mu}}[x_1, x_2 \le x_3;\ y_1, y_2 \le y_4].$$

The required formula follows by taking expectation over $(x_3, y_3), (x_4, y_4)$. □

(j) Deduce that $\mathbb{E}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2] = 1/9$, using the fact that μ is 4-quasirandom.

Answer. The probability that $x_1, x_2 \le x_3$ is $1/3$. Since μ is 4-quasirandom, conditioned on any fixed ordering of the x_i, the probability that $y_1, y_2 \le y_4$ is $1/3$. In total, the probability is $1/3 \cdot 1/3 = 1/9$. □

(k) Show that

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] = \Pr_{\substack{(x,y)\sim\mu\\(z_i,w_i)\sim\lambda}}[x, z_1 \le z_2;\ y, w_1 \le w_2] = \frac{1}{4}\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[(1-X^2)(1-Y^2)].$$

Answer. Let $(x,y) \sim \mu$ and $(z_1, w_1), (z_2, w_2) \sim \lambda$. Then $F_\mu(z_2, w_2) = \Pr[x \le z_2, y \le w_2]$ and $z_2 w_2 = \Pr[z_1 \le z_2, w_1 \le w_2]$. This explains the first equality.

For the second equality, note that for fixed z,

$$\Pr[z, z_1 \le z_2] = \Pr[z_2 \ge z]\mathop{\mathbb{E}}_{z_2}[\Pr[z_1 \le z_2 \mid z_2 \ge z]] = (1-z)\mathop{\mathbb{E}}_{z_2}[z_2 \mid z_2 \ge z] = \frac{1-z^2}{2},$$

since $\mathbb{E}[z_2 \mid z_2 \ge z] = \frac{1+z}{2}$. Similarly, $\Pr[w, w_1 \le w_2] = \frac{1-w^2}{2}$. Taking expectation over $(x,y) \sim \mu$, we obtain the second equality. □

(l) Show that

$$A := \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)XY]^2 = \frac{1}{81}.$$

Answer. This follows directly from (g). □

(m) Show that

$$B := \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[F_\mu(X,Y)^2] \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2Y^2] = \frac{4}{9} \mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] - \frac{1}{27}.$$

Hint: if $(X,Y) \sim \mu$ then $\mathbb{E}[X^2] = \mathbb{E}[Y^2] = 1/3$ since X, Y are individually uniform over $[0,1]$.

Answer. The first factor is $1/9$ by (e). For the second factor, use (k) to obtain

$$\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] = \frac{1}{4} - \frac{1}{4} \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2] - \frac{1}{4} \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[Y^2] + \frac{1}{4} \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2Y^2] = \frac{1}{12} + \frac{1}{4} \mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2Y^2],$$

using the hint. Rearranging,

$$\mathop{\mathbb{E}}_{(X,Y)\sim\mu}[X^2Y^2] = 4 \mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)ZW] - \frac{1}{3}.$$

Multiplying the two factors, we obtain the stated formula. □

(n) Show that

$$C := \frac{4}{9}\sqrt{\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[F_\mu(Z,W)^2]}\sqrt{\mathop{\mathbb{E}}_{(Z,W)\sim\lambda}[Z^2W^2]} - \frac{1}{27} = \frac{1}{81}.$$

Answer. The first expectation equals $1/9$ due to (j). The same holds for the second expectation, since $ZW = F_\lambda(Z,W)$ and λ is quasirandom. Altogether, we get

$$C = \frac{4}{9} \cdot \frac{1}{9} - \frac{1}{27} = \frac{4-3}{81} = \frac{1}{81}.$$ □

(o) Explain why always $A \le B \le C$.

Answer. Both follow from the Cauchy–Schwarz inequality. □

(p) Since $A = C = 1/81$, both inequalities are tight. Show that this implies that $F_\mu(XY) = XY$ and so $\mu = \lambda$ is quasirandom.

Answer. For the Cauchy–Schwarz inequality $\mathbb{E}[fg]^2 \le \mathbb{E}[f^2]\,\mathbb{E}[g^2]$ to be tight, we need $f \propto g$. In our case, this means that we need

$$F_\mu(X,Y) \propto XY.$$

Clearly $F_\mu(1,1) = 1$, and so the constant of proportionality is 1, that is, $F_\mu(X,Y) = XY = F_\lambda(X,Y)$. Since this holds for all X, Y, we conclude that $\mu = \lambda$ (up to measure zero), and so μ is quasirandom. □

Question 4 (Dijkstra's algorithm on the uniform weight distribution)**.** In this question we will generalize the analysis of Dijkstra's algorithm from exponential weights to uniform weights. We will use $U([0,1])$ to denote the uniform distribution over $[0,1]$, and $\mathrm{Exp}(1)$ to denote the unit mean exponential distribution, given by $\Pr[\mathrm{Exp}(1) \ge t] = e^{-t}$.

(a) Let $X \sim U([0,1])$ and $Y = \log \frac{1}{1-X}$. Show that $Y \sim \mathrm{Exp}(1)$.

Answer. Clearly $Y \ge 0$. Since $x \mapsto \log \frac{1}{1-x}$ is monotone increasing with inverse $y \mapsto 1 - e^{-y}$,

$$\Pr[Y \ge t] = \Pr\left[\log \frac{1}{1-X} \ge t\right] = \Pr[X \ge 1 - e^{-t}] = e^{-t}.$$ □

(b) Consider the coupling (X,Y) from the preceding item. Show that

$$Y(1 - Y/2) \le X \le Y.$$

Answer. Since $X = 1 - e^{-Y}$, this follows from $Y - Y^2/2 \le X \le Y$. □

(c) Suppose that w_1, w_2 are two sets of edge weights that satisfy $w_1(e) \le C w_2(e)$ for all edges e. Let $d_1(x,y), d_2(x,y)$ be the shortest distance from x to y according to the two sets of edge weights. Show that $d_1(x,y) \le C d_2(x,y)$.

Answer. Consider any shortest path from x to y according to w_2. The cost of this path according to w_1 is at most $C d_2(x,y)$. □

(d) For a distribution $\mathcal{D}$ supported on $\mathbb{R}_+$, let $T_{\mathcal{D}}$ be the expected distance from vertex 1 to the farthest vertex, when weights are chosen according to $\mathcal{D}$ independently. Show that

$$T_{U([0,1])} \le T_{\mathrm{Exp}(1)}.$$

Answer. Using the coupling described above, choose two sets of weights w_1, w_2 such that the weights in w_1 have distribution $U([0,1])$, the weights in w_2 have distribution $\mathrm{Exp}(1)$, and $w_1 \le w_2$ pointwise. Then $d_1(x,y) \le d_2(x,y)$ for all x, y. This implies that the distance from vertex 1 to the farthest vertex under w_1 is at most the same under w_2. The item follows by taking expectations. □

Proving a bound in the other direction is more tricky. For each edge e, we construct three different weights $w_1(e), w_2(e), w_3(e)$ as follows. We choose $w_1(e) \sim U([0,1])$ and let $w_2(e) = \log \frac{1}{1-w_1(e)}$, so that $w_2(e) \sim \mathrm{Exp}(1)$. For an $\epsilon \in (0, 1/2)$ to be determined, we choose $w_3(e) = \max(w_1(e), (1-\epsilon)w_2(e))$.

(e) We showed in class that $T_{\mathrm{Exp}(1)} \sim \frac{2\log n}{n}$. Let $\delta = \omega\big(\frac{\log n}{n}\big)$. Show that with high probability, the shortest path tree for vertex 1 under edge weights w_1 only contains edges of weight at most δ.

Answer. With high probability, the distance from vertex 1 to the farthest vertex under w_2 is at most δ. Hence the same holds under w_1. In particular, the shortest path tree for vertex 1 only contains edges whose weight is at most δ. □

(f) Show that if $\epsilon = \omega\big(\frac{\log n}{n}\big)$ then with high probability, the shortest path trees for vertex 1 under w_1 and under w_3 are identical.

Answer. According to the preceding item, with high probability the shortest path tree for vertex 1 under w_1 contains edges of weight at most ϵ.

If the shortest path trees are not identical, then since $w_1(e) \leq w_3(e)$, the shortest path tree under w_1 must contain an edge e such that $w_1(e) < (1-\epsilon)w_2(e)$. In view of item (b), this implies that $w_2(e) < 2\epsilon$, and so $w_1(e) = 1 - e^{-w_2(e)} > 1 - e^{-2\epsilon} > \epsilon$ (since $2\epsilon < 1$). We conclude that the shortest path trees are identical with high probability. □

(g) Let τ_1, τ_2, τ_3 be the distances from vertex 1 to the farthest vertex under w_1, w_2, w_3, respectively. We showed in class that $\mathbb{V}[\tau_2] = O(1/n^2)$. Using Cauchy–Schwarz, show that if $\delta = \omega\big(\frac{\log n}{n}\big)$ then

$$\mathbb{E}[\tau_2 1_{\tau_2 > \delta}] = o(\mathbb{E}[\tau_2]).$$

(Here $1_{\tau_2 > \delta}$ is the indicator variable for the event $\tau_2 > \delta$.)

Answer. Since $\mathbb{V}[\tau_2] = O(1/n^2)$, it follows that $\mathbb{E}[\tau_2^2] \leq (1+o(1))\,\mathbb{E}[\tau_2]^2$. Hence

$$\mathbb{E}[\tau_2 1_{\tau_2 > \delta}] \leq \sqrt{\mathbb{E}[\tau_2^2]} \cdot \sqrt{\Pr[\tau_2 > \delta]} \leq (1+o(1))\,\mathbb{E}[\tau_2] \cdot o(1) = o(\mathbb{E}[\tau_2]). \qquad \square$$

(h) Show that if $\epsilon = \omega\big(\frac{\log n}{n}\big)$ then $T_{U([0,1])} \geq (1-\epsilon-o(1))T_{\mathrm{Exp}(1)}$, and conclude that $T_{U([0,1])} \sim \frac{2\log n}{n}$.

Answer. We know that $\Pr[\tau_2 \leq \epsilon] = 1 - o(1)$, and given that event, $\tau_1 = \tau_3 \geq (1-\epsilon)\tau_2$ (see the proof of item (f)). Therefore

$$T_{U([0,1])} = \mathbb{E}[\tau_1] \geq \Pr[\tau_2 \leq \epsilon]\,\mathbb{E}[\tau_1 \mid \tau_2 \leq \epsilon] \geq (1-\epsilon)\Pr[\tau_2 \leq \epsilon]\,\mathbb{E}[\tau_2 \mid \tau_2 \leq \epsilon].$$

Now

$$\Pr[\tau_2 \leq \epsilon]\,\mathbb{E}[\tau_2 \mid \tau_2 \leq \epsilon] = \mathbb{E}[\tau_2 \cdot 1_{\tau_2 \leq \epsilon}] = \mathbb{E}[\tau_2] - \mathbb{E}[\tau_2 \cdot 1_{\tau_2 > \epsilon}] = (1-o(1))\,\mathbb{E}[\tau_2],$$

using the preceding item. In total,

$$T_{U([0,1])} \geq (1 - \epsilon - o(1))T_{\mathrm{Exp}(1)}.$$

Choosing $\epsilon = o(1)$, for example $\epsilon = 1/\sqrt{n}$, we deduce that $T_{U([0,1])} \geq (1 - o(1))T_{\mathrm{Exp}(1)}$. Since also $T_{U([0,1])} \leq T_{\mathrm{Exp}(1)}$, we conclude that $T_{U([0,1])} \sim T_{\mathrm{Exp}(1)}$. $\square$