Lecture 23

Prof. Nick Harvey *University of British Columbia*

We continue our theorem from last time on random partitions of metric spaces

1 Review of Previous Lecture

Define the partial Harmonic sum $H(a,b) = \sum_{i=a+1}^{b} 1/i$. Let $B(x,r) = \{\, y \in X \,:\, d(x,y) \leq r \}$ be the ball of radius r around x.

Theorem 1 *Let (X,d) be a metric with $|X| = n$. For every $\Delta > 0$, there is Δ-bounded random partition $\mathcal{P}$ of X with*

$$\Pr[B(x,r) \not\subseteq \mathcal{P}(x)] \;\leq\; \frac{8r}{\Delta} \cdot H\big(\,|B(x,\Delta/4 - r)|,\; |B(x,\Delta/2 + r)|\,\big) \quad \forall x \in X,\ \forall r > 0. \tag{1}$$

The algorithm to construct $\mathcal{P}$ is as follows.

- Pick $\alpha \in (1/4, 1/2]$ uniformly at random.
- Pick a bijection (i.e., ordering) $\pi : \{1,\ldots,n\} \to X$ uniformly at random.
- For $i = 1,\ldots,n$
  - Set $P_i \;=\; B(\pi(i), \alpha\Delta) \setminus \cup_{j=1}^{i-1} P_j$.
- Output the random partition $\mathcal{P} = \{P_1,\ldots,P_n\}$.

We have already proven that this outputs a Δ-bounded partition. So it remains to prove (1).

2 The Proof

Fix any point $x \in X$ and radius $r > 0$. For brevity let $B = B(x,r)$. Let us order all points of X as $\{y_1,\ldots,y_n\}$ where $d(x,y_1) \leq \cdots \leq d(x,y_n)$. The proof involves two important definitions.

- **Sees:** A point y **sees** B if $d(x,y) \leq \alpha\Delta + r$.
- **Cuts:** A point y **cuts** B if $\alpha\Delta - r \leq d(x,y) \leq \alpha\Delta + r$.

Obviously "cuts" implies "sees". To help visualize these definitions, the following claim interprets their meaning in Euclidean space. (In a finite metric, the ball B is not a continuous object, so it doesn't really have a "boundary".)

Claim 2 *Consider the metric (X,d) where $X = \mathbb{R}^n$ and d is the Euclidean metric. Then*

- *y sees B if and only if $B = B(x,r)$ intersects $B(y, \alpha\Delta)$.*
- *y cuts B if and only if $B = B(x,r)$ intersects the boundary of $B(y, \alpha\Delta)$.*

The following claim is in the same spirit, but holds for any metric.

Claim 3 *Let (X,d) be an arbitrary metric. Then*

- *If y does not see B then $B \cap B(y, \alpha\Delta) = \emptyset$.*
- *If y sees B but does not cut B then $B \subseteq B(y, \alpha\Delta)$.*

To illustrate the definitions of "sees" and "cuts", consider the following example. The blue ball around x is B. The points y_1 and y_2 both see B; y_3 does not. The point y_2 cuts B; y_1 and y_3 do not. This example illustrates Claim 3: y_1 sees B but does not cut B, and we have $B \subseteq B(y, \alpha\Delta)$.

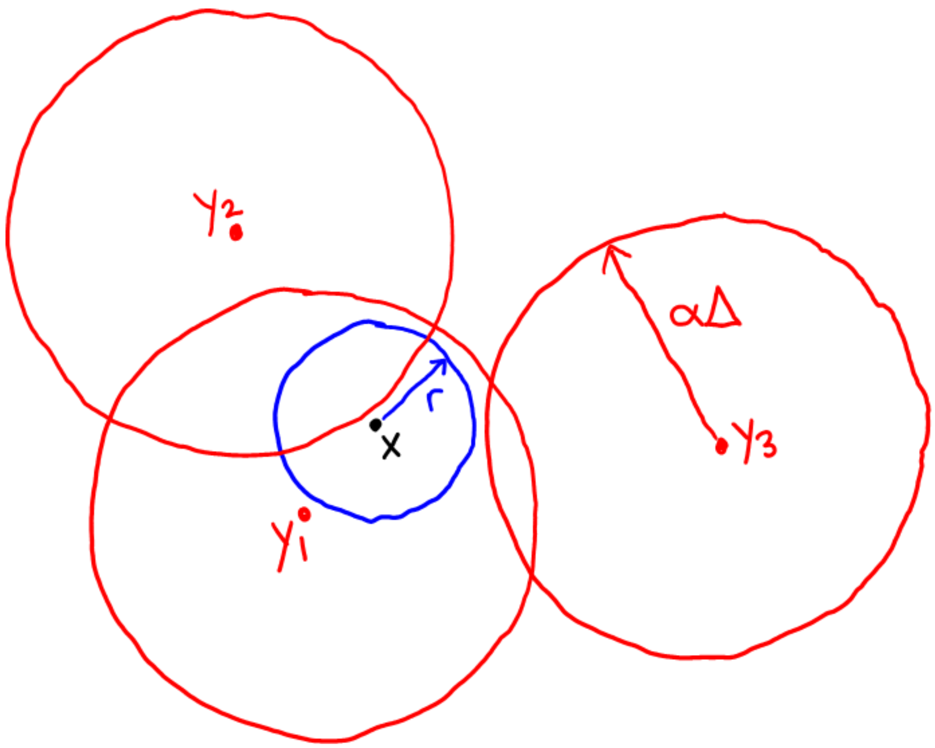

The most important point for us to consider is the *first* point under the ordering π that sees B. We call this point $y_{\pi(k)}$.

The first $k-1$ iterations of the algorithm did not assign any point in B to any P_i. To see this, note that $y_{\pi(1)}, \ldots, y_{\pi(k-1)}$ do not see B, by choice of k. So Claim 3 implies that $B \cap B(y_{\pi(i)}, \alpha\Delta) = \emptyset \ \forall i < k$. Consequently

$$B \cap P_i = \emptyset \quad \forall i < k. \tag{2}$$

The point $y_{\pi(k)}$ sees B by definition, but it may or may not cut B. If it does not cut B then Claim 3 shows that $B \subseteq B(y_{\pi(k)}, \alpha\Delta)$. Thus

$$B \cap P_k \;=\; \Big(\underbrace{B \cap B(y_{\pi(k)}, \alpha\Delta)}_{=B} \Big) \;\setminus\; \bigcup_{i=1}^{k-1} \underbrace{B \cap P_i}_{=\emptyset} \;=\; B,$$

i.e., $B \subseteq P_k$. Since $\mathcal{P}(x) = P_k$, we have shown that

$$y \text{ does not cut } B \quad \Longrightarrow \quad B \subseteq \mathcal{P}(x).$$

Taking the contrapositive of this statement, we obtain

$$\Pr[B \not\subseteq \mathcal{P}(x)] \;\leq\; \Pr[y_{\pi(k)} \text{ cuts } B] \;=\; \sum_{i=1}^{n} \Pr[y_{\pi(k)} = y_i \;\wedge\; y_i \text{ cuts } B].$$

Let us now simplify that sum by eliminating terms that are equal to 0.

Claim 4 *If $y \notin B(x, \Delta/2 + r)$ then y does not see B.*

Claim 5 *If $y \in B(x, \Delta/4 - r)$ then y sees B but does not cut B.*

So define $a = |B(x, \Delta/4 - r)|$ and $b = |B(x, \Delta/2 + r)|$. Then we have shown that

$$\Pr[B \not\subseteq \mathcal{P}(x)] \;\leq\; \sum_{i=a+1}^{b} \Pr[y_{\pi(k)} = y_i \;\wedge\; y_i \text{ cuts } B].$$

The remainder of the proof is quite interesting. The main point is that these two events are "nearly independent", since α and π are independent, "y_i cuts B" depends only on α, and "$y_{\pi(k)} = y_i$" depends primarily on π. Formally, we write

$$\Pr[B \not\subseteq \mathcal{P}(x)] \;\leq\; \sum_{i=a+1}^{b} \Pr[y_i \text{ cuts } B] \cdot \Pr[\, y_{\pi(k)} = y_i \mid y_i \text{ cuts } B]$$

and separately upper bound these two probabilities.

The first probability is easy to bound:

$$\Pr[y_i \text{ cuts } B] \;=\; \Pr[\, \alpha\Delta \in [d(x,y) - r, d(x,y) + r]\,] \;\leq\; \frac{2r}{\Delta/4},$$

because $2r$ is the length of the interval $[d(x,y) - r, d(x,y) + r]$ and $\Delta/4$ is the length of the interval from which $\alpha\Delta$ is randomly chosen.

Next we bound the second probability. Recall that $y_{\pi(k)}$ is defined to be the first element in the ordering π that sees B. Since y_i cuts B, we know that $d(x, y_i) \leq \alpha/2 + r$. Every y_j coming earlier in the ordering has $d(x, y_j) \leq d(x, y_i) \leq \alpha/2 + r$, so y_j also sees B. This shows that there are at least i elements that see B. So the probability that y_i is the *first* element in the random ordering to see B is at most $1/i$.

Combining these bounds on the two probabilities we get

$$\Pr[B \not\subseteq \mathcal{P}(x)] \;\leq\; \sum_{i=a+1}^{b} \frac{8r}{\Delta} \cdot \frac{1}{i} \;=\; \frac{8r}{\Delta} \cdot H(a,b),$$

as required.

3 Optimality of these partitions

Theorem 1 from the previous lecture shows that there is a universal constant $L = O(1)$ such that every metric has a $\log(n)/10$-bounded, L-Lipschitz random partition. We now show that this is optimal.

Theorem 6 *There exist graphs G whose shortest path metric (X, d) has the property that any $\log(n)/10$-bounded, L-Lipschitz random partition must have $L = \Omega(1)$.*

The graphs we need are expander graphs. In Lecture 20 we defined *bipartite* expanders. Today we need *non-bipartite* expanders. We say that $G = (V, E)$ is a non-bipartite expander if, for some constants $c > 0$ and $d \geq 3$:

- G is d-regular, and
- $|\delta(S)| \geq c|S|$ for all $|S| \leq |V|/2$.

It is known that expanders exist for all $n = |V|$, $d = 3$ and $c \geq 1/1000$. (The constant c can of course be improved.)

PROOF: Suppose (X, d) has a $\log(n)/10$-bounded, L-Lipschitz random partition. Then there exists a particular partition P that is $\log(n)/10$-bounded and cuts at most an L-fraction of the edges. Every part P_i in the partition has diameter at most $\log(n)/10$. Since the graph is 3-regular, the number of vertices in P_i is at most $3^{\log(n)/10} < n/2$. So every part P_i has size less than $n/2$. By the expansion condition, the number of edges cut is at least

$$\frac{1}{2} \sum_i c \cdot |P_i| \;=\; cn/2 \;=\; \Omega(|E|).$$

So $L = \Omega(1)$. □

4 Appendix: Proofs of Claims

PROOF:(of Claim 3) Suppose y does not see B. Then $d(x, y) > \alpha\Delta + r$. Every point $z \in B$ has $d(x, z) \leq r$, so $d(y, z) \geq d(y, x) - d(x, z) > \alpha\Delta + r - r$, implying that $z \notin B(y, \alpha\Delta)$.

Suppose y sees B but does not cut B. Then $d(x, y) < \alpha\Delta - r$. Every point $z \in B$ has $d(x, z) \leq r$. So $d(y, z) \leq d(y, x) + d(x, z) < \alpha\Delta - r + r$, implying that $z \in B(y, \alpha\Delta)$. □

PROOF:(of Claim 4) The hypothesis of the claim is that $d(x, y) > \Delta/2 + r$, which is at least $\alpha\Delta + r$. So $d(x, y) \geq \alpha\Delta + r$, implying that y does not see B. □

PROOF:(of Claim 5) The hypothesis of the claim is that $d(x, y) \leq \Delta/4 - r$, which is strictly less than $\alpha\Delta - r$. So $d(x, y) < \alpha\Delta - r$, which implies that y sees B but does not cut B. □