Flows, Cuts, and Line Embeddings

Kent Quanrud

October 5, 2020

1 Single commodity flow

Recall the well known (s,t)-max flow problem, which we present as the following **path packing problem**. The input consists of a graph $G = (V, E)$ with positive edge capacities $c : E \to \mathbb{R}_{\geq 0}$, a source vertex s, and a sink vertex t. The **(s,t)-max flow problem** asks for a maximum size fractional packing of (s,t) paths into G. To formalize this problem, let $\mathcal{P}_{s,t}$ denote the collection of all $s \to t$ paths. Then the maximum flow is given by the following packing LP.

$$\text{maximize} \sum_{p \in \mathcal{P}_{s,t}} x_p \text{ over } x : \mathcal{P}_{s,t} \to \mathbb{R}_{\geq 0} \text{ s.t.} \sum_{p \ni e} x_p \leq c(e) \text{ for all } e \in E. \tag{1}$$

Recall also the (s,t)-minimum cut problem. Given the same input as (s,t)-max flow, we interpret the capacities $c : E \to \mathbb{R}_{\geq 0}$ as costs. The **(s,t)-minimum cut** is the minimum cost set of edges whose removal disconnects s from to t. We can think of this is a (discrete) covering problem, where we want to choose "edges covering paths", as follows.

$$\text{minimize} \sum_{e \in C} c(e) \text{ over } C \subset E \text{ s.t. } p \cap C \neq \emptyset \text{ for all } p \in \mathcal{P}_{s,t}. \tag{2}$$

We might also consider the *fractional relaxation* of the (s,t)-minimum cut, where we want the minimum cost

$$\text{minimize} \sum_{e \in E} y_e \text{ over } y : E \to \mathbb{R}_{\geq 0} \text{ s.t. } \sum_{e \in p} y_e \geq 1 \text{ for all } p \in \mathcal{P}_{s,t}. \tag{3}$$

A priori, we know that (3) $\leq$ (2) becausde the former is a relaxation of the latter. We also observed, in our previous discussion on LP-duality, we observed that (1) and (3) are dual LP's, and thus have the same objective value. Thus we have

$$(1) = (3) \leq (2).$$

The following fundamental **max flow min cut theorem** theorem that all three quantities above are equal.

Theorem 1 (Menger (1927) and Ford, Jr. and Fulkerson (1956))**.** *The (s,t)-maximum flow equals the (discrete) (s,t)-minimum cut. Moreover, if the capacities are integral, then there is an optimum maximum flow that is integral.*

Later we will give a simple randomized proof that the max flow equals the min cut.

2 Multicommodity flows and cuts

A natural generalization of (s,t)-flow is to allow for multiple pairs of terminals. In **multicommodity flow problems** problems, the input consists of a graph $G=(V,E)$ with edge capacities $c: E \to \mathbb{R}_{\geq 0}$, and k **terminal pairs** $\{(s_1,t_1),\dots,(s_k,t_k)\}$ and **scalar demands** $d_1,\dots,d_k > 0$. We assume, for the rest of this note, that G is undirected.[1]

Let us define two problems for the multicommopdity flow setting, analogous to the single commodity problems discussed in Section 1.

2.1 Maximum concurrent flow

The maximum concurrent flow problem is to simultaneously route λd_i units of flow from s_i to t_i for every terminal pair (s_i,t_i). "Simultaneous" means that each k different flows must share the capacity of the edges. We can write this problem as an LP, except rather than packing paths, we are packing "bundles" of paths $p=(p_1,\dots,p_k)$, where each p_i is a path from s_i to t_i. To formalize this, for each i, let $\mathcal{P}_i$ be the family of all paths from s_i to t_i, and let $\mathcal{P} = \prod_{i=1}^k \mathcal{P}_i$ be the family of bundles consisting of one path from each family. The maximum concurrent flow problem can be written as

$$\text{maximize} \sum_{p\in\mathcal{P}} x_p \text{ over } x:\mathcal{P}\to\mathbb{R}_{\geq 0} \text{ s.t. } \sum_{p\in\mathcal{P}} x_p \sum_{i:e\in p_i} d_i \leq c(e) \text{ for all } e\in E.$$

Note that for $k=1$, this is exactly the "single commodity" (s,t)-maximum flow problem.

2.2 Sparsest cut

In the **sparset cut** problem, the goal is to partition the vertices into $(S,\bar{S})$ (where $\bar{S} = V\setminus S$) to minimize the ratio

$$\frac{\sum_{e\in\partial(S)} c(e)}{\sum_{i:s_i\in S,t_i\in\bar{S}} d_i + \sum_{i:t_i\in S,s_i\in\bar{S}} d_i}.$$

Here $0/0$ is interpreted as $+\infty$, the worst possible score. Again observe that for.

2.3 A multicommodity "max flow min cut"?

We might take $\{s,t\}$ max-flow min-cut theorem as an inspiration for a similar relation between concurrent flow and sparsest cut. Let us take an LP based approach, starting from the LP for maximum concurrent flow. The *dual* of that packing LP is the following covering LP.

$$\text{minimize} \sum_e c(e)y_e \text{ over } y:E\to\mathbb{R}_{\geq 0} \text{ s.t. } \sum_{s,t\in V} d(s,t) \sum_{e\in p(s,t)} y_e \geq 1 \text{ for all } p\in\mathcal{P}.$$

To interpret this LP, observe that for every $p\in\mathcal{P}$ and $s,t\in V$, $\sum_{e\in p(s,t)} y_e$ is the length of the path $p(s,t)$ w/r/t the edge lengths y_e. For a fixed vector y, let $\delta_y : V\times V\to\mathbb{R}_{\geq 0}$ give the shortest path distances w/r/t y. Then the LP can be written as

$$\text{minimize} \sum_e c(e)y_e \text{ over } y:E\to\mathbb{R}_{\geq 0} \text{ s.t. } \sum_{s,t\in V} d(s,t)\delta_y(s,t) \geq 1.$$

[1]The problems here can also be formulated for directed graphs, and are also important. However, this note will be about techniques for undirected graphs that do not extend to directed graphs.

Let us consider the special case of *undirected graphs*. In the LP above, δ_y gives a metric, and in any optimal solution, we have $\delta_y(s,t) = y_e$ for any edge $e = \{s,t\}$. This allows us to rewrite the LP (again!) as follows. For a metric $\delta : \binom{V}{2}$ and an edge $e = \{s,t\}$, let us denote $\delta(e) = \delta(s,t)$. Consider the following equivalent LP (for unidrected graphs).

$$\text{minimize } \sum_{e} c(e)\delta(e) \text{ over all metrics } \delta : E \to \mathbb{R}_{\geq 0} \text{ s.t. } \sum_{s,t \in V} d(s,t)\delta(s,t) \geq 1.$$

By scaling, this is equivalent to the following LP.

$$\text{minimize } \frac{\sum_e c(e)\delta(e)}{\sum_{s,t\in V} d(s,t)\delta(s,t)} \text{ over all metrics } \delta : \binom{V}{2} \to \mathbb{R}_{\geq 0}. \tag{4}$$

Finally, we recognize this last LP as a relaxation to the sparsest cut problem. Indeed, any cut $(S, \bar{S})$ gives rise to a (boring) metric δ_S defined by

$$\delta_S(s,t) = \begin{cases} 0 & \text{if } s,t \in S \\ 1 & \text{if } s \in S \text{ and } t \in \bar{S} \\ 1 & \text{if } s \in \bar{S} \text{ and } t \in S \\ 0 & \text{if } s,t \in \bar{S}. \end{cases}$$

Let us the value of the LP (4) the "minimum fractional sparsest cut". We now have the following.

$$\text{(max. concurrent flow)} \overset{\text{(a)}}{=} \text{(min. fractional sparsest cut)} \overset{\text{(b)}}{\leq} \text{(min. sparsest cut)}.$$

Here (a) applies LP-duality (**??**). (b) is because minimum fractional sparsest cut is a relaxation of minimum sparsest cut.

Recall that with $\{s,t\}$ max flow, the maximum flow equals the minimum cut. With concurrent flow, we are consider the multiple-source-sink generalization of $\{s,t\}$-max flow. In this sense, we have "max flow $\leq$ min cut". One might ask if also "min cut $\leq$ max flow", or more generally, for the smallest factor $\alpha \geq 1$ such that

$$\text{(min. sparsest cut)} \leq \alpha\text{(max. concurrent flow)}.$$

This is the goal of our discussion today.

3 A line embedding proof of the max flow min cut theorem

In this section, we prove the max flow min cut theorem.[2]. Let $y \in \mathbb{R}^E_{\geq 0}$ be an optimum solution to the minimum fractional cut LP.

We embed the vertices V on the line by assigning values $\alpha : V \to [0, +\infty)$ as follows. For each vertex v, let α_v be the length of the shortest $s \rightsquigarrow v$ path w/r/t the edge lengths $y \in \mathbb{R}^n_{\geq 0}$. We have $\alpha_s = 0$. We also have

$$\alpha_t = \min_{p \in \mathcal{P}_{s,t}} \sum_{e \in p} y_e \geq 1$$

because of the LP. Consider the following *random cut*. We pick a value $\theta \in (0,1)$ uniformly at random. Let $S = \{v : x_v \leq \theta\}$, and let $\bar{S} = V \setminus S$. Since $\alpha_s = 0$ and $\alpha_t \geq 1$, $(S, \bar{S})$ is always an (s,t)-cut.

[2] The proof is colorfully illustrated in a video by the author available at `https://youtu.be/J4yUdABv1tE`

Let us bounded the cost of the directed cut from S to $\bar{S}$, *in expectation.* We have

$$\mathrm{E}\left[\sum_{e\in\partial^+(S)} c(e)\right] = \sum_{e\in E} c(e)\mathrm{P}\big[e\in\partial^+(S)\big] \le \sum_{e\in E} c(e)y_e = \text{(fractional min cut)}.$$

Here (c) is by linearity of expectation. (d) is by the following argument. For an edge $e=\{u,v\}$, $e\in\partial^+(S)$ iff $\alpha_u \le \theta \le \alpha_v$, which happens with probability $\le \alpha_v - \alpha_u$. But concatenating the shortest path from s to u with the edge e,

$$s \overset{\alpha_u}{\rightsquigarrow} u \overset{y_e}{\to} v,$$

gives a walk of length $\alpha_u + y_e$, so $\alpha_v \le \alpha_u + y_e$.

Consider now the inequality obtained above,

$$\mathrm{E}\left[\sum_{e\in\partial^+(S)} c(e)\right] \le \text{(fractional min cut)}.$$

We have generated a *randomized (discrete) cut* that is *on average* no worse then the minimum fractional minimum cut. *By the probabilistic method, there exists a value θ where the (s,t)-cut has value at most this average.* If not, then the average would have to be higher. This establishes the existence of a minimum cut with value equal to the minimum fractional minimum cut, hence the max flow min cut theorem. To extract the cut, one can simply scan θ over the interval $(0,1)$ and check all $n-1$ possible cuts. (In fact any $\theta\in(0,1)$ will work, see Exercise **??**.)

A second part of the maximum flow minimum cut theorem asserts that when the capacities are integral, so is the maximum flow. The line embedding proof above does not capture this. However, knowledge of the fact that there is always a discrete cut with the same size as the maximum flow is enough to prove integrality (see Exercise **??**).

4 A line embedding proof of a multicommodity max-flow min-cut theorem

Recall the (undirected) sparsest cut problem introduced earlier. The input contains an undirected graph $G=(V,E)$ with positive edge capacities $c: E\to\mathbb{R}_{\ge 0}$, k terminal pairs $(s_1,t_1),\dots,(s_k,t_k)\in\binom{V}{2}$, and k scalar demands $d_1,\dots,d_k>0$. The goal is to partition the vertices into two sets $(S,\bar{S})$ as to minimize the ratio of the capacities of edges cut to the demands of terminals separated:

$$\text{minimize } \frac{\sum_{e\in\partial(S)} c(e)}{\sum_{i:|\{s_i,t_i\}\cap S|=1} d_i} \text{ over } S\subset V,$$

where $0/0=+\infty$. By observing that each set S induces a simple $\{0,1\}$-cut indicating whether or not two vertices are separated by S, a linear relaxation of the above problem

$$\text{minimize } \frac{\sum_{\{u,v\}\in E} c(u,v)y(u,v)}{\sum_{i=1}^k d_i y(s,t)} \text{ over all metrics}$$

We call this the sparsest metric. By our discussion before, leveraging LP duality, we concluded that

$$\text{(max concurrent flow)} = \text{(sparsest metric)} \le \text{(sparsest cut)}$$

The rest of our discussion is dedicated to proving the following converse.

$$\text{(sparsest cut)} \leq O(\log(k))(\text{max concurrent flow}).$$

This establishes the following **approximate max-flow min-cut theorem for multicommodity flows**.

Theorem 2 (Leighton and Rao (1999))**.** *For any multicommodity flow instance with k terminal pairs,*

$$(\textit{max concurrent flow}) \leq O(\log k)(\textit{sparsest cut}).$$

This theorem has a number of applications ranging from approximation algorithms to fast graph algorithms.

4.1 $k = 1$

We will prove the multicommodity max-flow min-cut theorem in a similar fashion to the proof of the single-commodity max flow min cut theorem in Section **??**. There, we embedded the vertices on a line, based on the distance form the source s with respect to the edge weights produced by the LP. With this in mind, let us consider the sparsest cut problem for the special case where $k = 1$. If we inline $k = 1$ into sparsest cut, we get

$$\text{minimize } \frac{\sum_{e \in \partial(S)} c(e)}{1} \text{ over } S \subset V \text{ s.t. } |\{s_1, t_1\} \cap S| = 1.$$

(As expected, this is the minimum $\{s_1, t_1\}$-cut.) The sparsest metric problem becomes

$$\text{minimize } \frac{\sum_{\{u,v\} \in E} c(u,v) y(u,v)}{y(s_1, t_1)} \text{ over all metrics } y : \binom{V}{2} \to \mathbb{R}_{\geq 0}.$$

Now, let y be the sparsest metric, computed by an LP. Rescaling y, we can assume that $y(s,t) = 1$. Then the sparsity of y equals $\sum_{u,v \in E} c(u,v) y(u,v)$.

We can repeat the same randomized rounding from Section **??**. Pick $\theta \in (0, 1)$ uniformly at random, and let $S = \{v : y(s, v) \leq 1\}$. Then S is an $\{s_1, t_1\}$-cut, and

$$\mathrm{E}\left[\sum_{e \in \partial(S)} c(e)\right] \leq \sum_{\{u,v\} \in E} y(u,v) c(u,v).$$

By the probabilistic method, there is a cut $(S, \bar{S})$ with sparsity equal to the sparsest metric.

4.2 Line embeddings and cuts

Thus $k = 1$ works out nicely, which is to be expected since it is jsut the single commmodity max flow min cut. What was so special about $k = 1$? Having a single commodity (s_1, t_1) made it easy to embed the vertices on a line, here by distance from s_1.

Suppose (by some coincidence) that the sparsest metric y happened to be obtained from a line embedding. That is, there is a mapping $\alpha : V \to \mathbb{R}$ such that for all $u, v \in V$,

$$y(u,v) = |\alpha_u - \alpha_v|.$$

Rescaling and reshifting, we may assume that $\min_u \alpha_u = 0$, and $\max_v \alpha_v = 1$.

Consider the following familiar random cut S. Pick $\theta \in (0,1)$ uniformly at random, and let $S = \{u : \alpha_u \leq \theta\}$. Observe that for each edge $e = \{u,v\}$, we have

$$\mathrm{P}[e \in \partial(S)] = |\alpha_u - \alpha_v| = y(u,v).$$

Thus we can rewrite sparsity of y as

$$\frac{\sum_{\{u,v\}\in E} c(u,v)y(u,v)}{\sum_{(s_i,t_i)} d_i y(s_i,t_i)} = \frac{\mathrm{E}\left[\sum_{e\in\partial(S)} c(e)\right]}{\mathrm{E}\left[\sum_{i:|\{s_i,t_i\}\cap S|=1} d_i\right]}.$$

Observe that S can only be one of $n-1$ different sets $S_1,\ldots,S_{n-1}$ where $\emptyset \subsetneq S_1 \subset S_2 \subset S_3 \cdots \subset S_{n-1} \subsetneq V$. For each i, let $p_i = \mathrm{P}[S = S_i]$. Then

$$\frac{\mathrm{E}[e \in \partial(S)c(e)]}{\mathrm{E}\left[\sum_{i:|\{s_i,t_i\}\cap S|=1} d_i\right]} = \frac{\sum_{i=1}^{n-1} p_i \sum_{e\in\partial(S)} c(e)}{\sum_{i=1}^{n-1} p_i \sum_{i:|\{s_i,t_i\cap S\}=1|} d_i}.$$

In a weird sense of the word "average", the random cut is good on "average". We want to argue that one of the cuts S_i is good. Consider the following elementary fact.

Lemma 3. *Let $a_1,\ldots,a_h,b_1,\ldots,b_h > 0$. Then*

$$\min_i \frac{a_i}{b_i} \leq \frac{\sum_i a_i}{\sum_i b_i} \leq \max_i \frac{a_i}{b_i}.$$

Lemma 4. *Let $y : \binom{V}{2} \to \mathbb{R}$ be a metric induced by a line embedding. Then one can partition the vertices into two sets $(S,\bar{S})$ such that*

$$\frac{\sum_{e\in\partial(S)} c(e)}{\sum_{|\{s_i,t_i\}\cap S|=1} d_i} \leq \frac{\sum_{\{u,v\}\in E} c(u,v)y(u,v)}{\sum_{(s_i,t_i)} d_i y(s_i,t_i)}.$$

So much for metrics given by line embeddings. How about a metric obtained as a sum of line metrics? Recall that the L_1 metric on $\mathbb{R}^d$ is defined by

$$\|x-y\|_1 = \sum_{i=1}^{n} |x_i - y_i|.$$

Suppose y was the L_1-metric of an embedding $f : V \to \mathbb{R}^h$. We can think of this as the sum of h line metrics $f_1,\ldots,f^h$. Applying Lemma **??** again, we see that one of these line embeddings, say f_j, has sparsity no worse then y. From the line embedding $f_j : V \to \mathbb{R}$, we can extract a cut with sparsity at most that of f_j.

Lemma 5. *Let $y : \binom{V}{2} \to \mathbb{R}$ be the L_1-metric over an explicit embedding of X. Then one can partition the vertices into two sets $(S,\bar{S})$ such that*

$$\frac{\sum_{e\in\partial(S)} c(e)}{\sum_{|\{s_i,t_i\}\cap S|=1} d_i} \leq \frac{\sum_{\{u,v\}\in E} c(u,v)y(u,v)}{\sum_{(s_i,t_i)} d_i y(s_i,t_i)}.$$

To sum up: L_1 metrics can be rounded without loss. We can find an L_1-metric with sparsity within a factor α of the sparsest metric y, then we can covert that into an α-approximate minimum cut.

random-Frechet($d : V \times V \to \mathbb{R}_{\geq 0}$)

1. for $i = 1, \ldots, \lceil \log |V| \rceil$
   - A. let S_i sample each $v \in V$ independently with probability 2^{-i}
   - B. for each $v \in V$
     - 1. $v_i \leftarrow \min_{s \in S_i} y(s, v)$

Figure 1: A $O(\log n)$ dimension, randomized Frechét embedding with $O(\log n)$ distortion in expectation

4.3 Bourgain's embedding theorem

Theorem 6. *Let $y : \binom{V}{2} \to \mathbb{R}_{\geq 0}$ be a metric and $\delta \in (0, 1)$. For $h = O(\log(n) \log(1/\delta))$, one can construct a randomized embedding $f : V \to \mathbb{R}^h$ such that for all $u, v \in V$,*

$$f(u) - f(v) \leq O(\log(1/\delta))y(u, v)$$

deterministically, and

$$\mathrm{P}[f(u) - f(v) \geq y(u, v)] \geq 1 - \delta.$$

Corollary 7. *Let $y : \binom{V}{2} \to \mathbb{R}_{\geq 0}$ be a metric and $\delta \in (0, 1)$. For $h = O\big(\log^2(n)\big)$, one can construct a randomized embedding $f : V \to \mathbb{R}^h$ such that for all $u, v \in V$,*

$$y(u, v) \leq f(u) - f(v) \leq O(\log n)y(u, v)$$

deterministically, and

$$\mathrm{P}[f(u) - f(v) \geq y(u, v)] \geq 1 - \delta.$$

In this section, we consider the randomized algorithm of Linial, London, and Rabinovich (1995), which makes algorithmic the embedding of Bourgain (1985). The pseudocode is given in Figure 1. The algorithm is extremely simple. We generate $\lceil \log n \rceil$ coordinates. For $i = 1, \ldots, n$, we sample a set S_i where each point is sampled independently with probability $1/2^i$. For each vertex v, we find the distance between v and (the closest point in) S_i. This gives the ith coordinate of v.

4.3.1 Low-distortion in expectation

Abusing notation, for a vertex v and coordinate i, we let v_i denote the ith coordinate of the embedding of v.

Lemma 8. *For $u, v \in V$ and $i \in \mathbb{N}$, $|u_i - v_i| \leq y(u, v)$.*

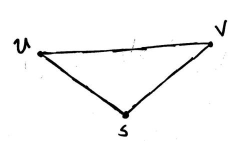

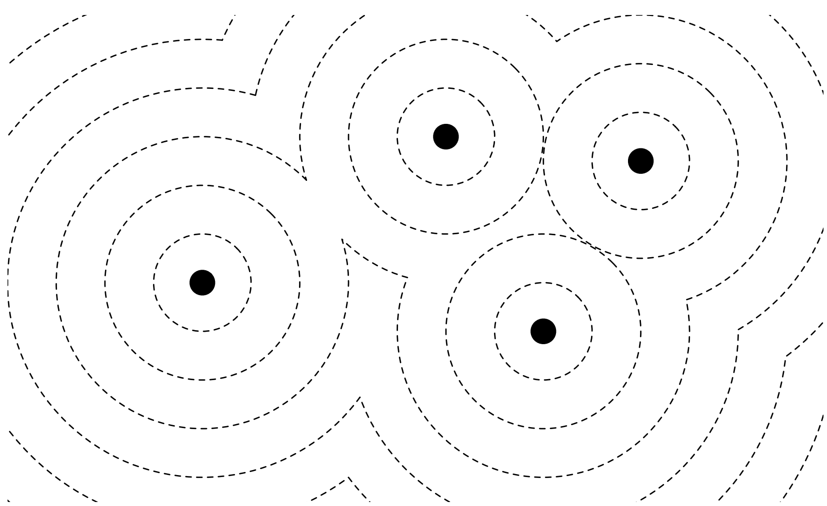

Figure 2: Level sets by distance from a set of points S, encoding one coordinate of a Frechét embedding.

Proof. By the triangle inequality, we have both

$$y(s \mid u) - y(s \mid v) \leq y(u \mid v) \text{ and } y(s \mid v) - y(s \mid u) \leq y(u \mid v).$$

for all $s \in S_i$. ■

Lemma 8 immediate implies the $\|u - v\|_1 \leq O(\log n) y(u, v)$, since there are $O(\log n)$ dimensions and each can contribute at most $y(u, v)$. The harder part is showing the lower bound: informally, we want to show that $\|u - v\|_1 \geq y(u, v)$, up to constant factors. This lower bound is too strong; instead, we settle for the same inequality but only in expectation.

Lemma 9. *For $u, v \in V$, we have $\mathrm{E}[\|u - v\|_1] \geq c y(u, v)$ for some constant c.*

Proof. For ease of notation, let $\delta = y(u, v)$. For each i, let r_i be the minimum length r such that there are at least 2^i points at distance $\leq r$ from u, and 2^i points at distance $\leq r$ from v; i.e.,

$$r_i = \arg\min_{r>0} \left\{ |\{x : y(u, x) \leq r\}| \geq 2^t, |\{x : y(v, x) \leq r\}| \geq 2^t \right\}$$

We claim that

> *For each index i, we have*
>
> $$|u_{i+1} - v_{i+1}| \geq (\min\{r_i, \delta/2\} - \min\{r_{i-1}, \delta/2\})$$
>
> *with constant probability $c > 0$.*

Suppose the above claim holds. Let k be the largest index such that $r_{k-1} \leq \delta/2$ (for which the claim applies). We have

$$\mathrm{E}[\|u - v\|_1] \geq \sum_{i=0}^{k} \mathrm{E}[|u_{i+1} - v_{i+1}|] \overset{\text{(a)}}{\geq} c \sum_{i=0}^{k} (\min\{r_i, \delta/2\} - \min\{r_{i-1}, \delta/2\}) \overset{\text{(b)}}{\geq} \frac{c\delta}{4},$$

where (a) applies our claim and (b) is by telescoping sums and recalling that $r_{k+1} > \delta/2$.

It remains to prove the claim. We have two cases: *(a)* $r_i \leq \delta/2$, and *(b)* $r_{i-1} < \delta/2 \leq r_i$. We assume without loss of generality that r_i is defined by u; i.e., $|\{x : y(u, c) < r_i\}| < 2^i$.

Case 1: $r_i \le \delta/2$. Let $U = \{x : y(u,x) < r_i\}$, and let $V = \{x : y(v,x) \le r_{i-1}\}$. We have $|U| < 2^i$ and $|V| \ge 2^{i-1}$. Since $r_{i-1} < r_i \le \delta/2$, U and V are disjoint.

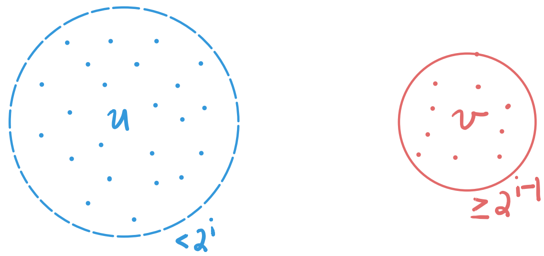

S_{i+1} samples each point with probability 2^{-i-1}. By direct calculation, S_{i+1} samples no points from U with constant probability, and at least one point from V with constant probability. Since U and V are disjoint, whether any point from U is sampled and whether any point from V is sampled is independent. Thus S_{i+1} samples a point from V and no points from U simultaneously with some constant probability $c > 0$. In this event, we have $u_{i+1} \ge r_i$ and $v_{i+1} \le r_{i-1}$, so $u_{i+1} - v_{i+1} \ge r_i - r_{i-1}$. In expectation, we have

$$\mathrm{E}[|u_{i+1} - v_{i+1}|] \ge c(r_i - r_{i-1}),$$

as desired.

Case 2: $r_i > \delta/2 > r_{i-1}$. Let $U = \{x : y(u,x) \le \delta/2\}$ (with $\delta/2$ in place of r_i) and let $V = \{x : y(u,v) \le r_{i-1}/2\}$.

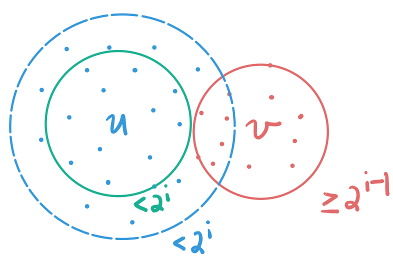

By the same argument as above, we have that S_{i+1} samples a point from V and no points in U with some constant probability $c > 0$. In this event, $u_{i+1} - v_{i+1} \ge \delta/2 - r_{i-1}$. ■

Theorem 10. `randomized-Frechet` *embeds V into $\mathbb{R}^{\lceil \log n \rceil}$ such that*

$$\|u - v\|_1 \le \lceil \log n \rceil y(u,v) \text{ and } \mathrm{E}[\|u - v\|_1] \ge c y(u,v) \text{ for all } u, v \in V,$$

for some absolute constant $c > 0$.

Amplification

Theorem 11. *With probability of error $1/\operatorname{poly}(n)$, the average $O(\log n)$ embeddings produced by* `randomized-Frechet` *is an embedding V into $\mathbb{R}^{O(\log^2 n)}$ such that*

$$c y(u,v) \le \|u - v\|_1 \le C \log(n) y(u,v) \text{ for all } u, v \in V$$

for absolute constants $c, C > 0$.

Proof sketch. Fix $u, v \in V$. We treat each coordinate difference $|u_i - v_i|$ over the $O(\log n)$ `random-Frechet` embeddings as a random variable bounded above by $y(u, v)$ and, whose total expectation $\geq c \log(n) y(u, v)$. By Chernoff inequalities, the sum is strongly concentrated at the mean; scaling down by $\log n$ (from averaging) gives the desired result. ∎

Theorem 12. *With probability of error* $1/\operatorname{poly}(n)$*, the average* $O(\log k)$ *embeddings produced by* `randomized-Frechet` *is an embedding* V *into* $\mathbb{R}^{O(\log^2 n)}$ *such that*

$$cy(u, v) \leq \|u - v\|_1 \leq C \log(n) y(u, v) \text{ for all } u, v \in V$$

for absolute constants $c, C > 0$.

Proof sketch. Fix $u, v \in V$. We treat each coordinate difference $|u_i - v_i|$ over the $O(\log n)$ `random-Frechet` embeddings as a random variable bounded above by $y(u, v)$ and, whose total expectation $\geq cO(\log(k)) y(u, v)$. By Chernoff inequalities, the probability of deviating by more than a constant factor is $< 1/k$. Scaling down by $\log(k)$ gives us the desired result. ∎

5 Exercises

Exercise 1. Consider the random cut, which was based on a threshold $\theta \in (0, 1)$ chosen uniformly at random. Show that for *all* $\theta \in (0, 1)$, the corresponding cut is a minimum $\{s, t\}$-cut.

Exercise 2. Using only the fact thet the maximum (s, t)-flow equals the minimum (discrete) cut, prove that if the capacities are integral, then the maximum flow can always be taken to be integral.

References

Menger, Karl (1927). "Zur allgemeinen Kurventheorie". In: *Fundamenta Mathematicae* 10.1, pp. 96–115.

Ford, Jr., L. R. and D. R. Fulkerson (1956). "Maximal flow through a network". In: *Canadian Journal of Mathematics* 8, pp. 399–404.

Bourgain, Jean (1985). "On Lipschitz embedding of finite metric spaces in Hilbert space". In: *Israel Journal of Mathematics* 52.1-2, pp. 46–52.

Linial, Nathan, Eran London, and Yuri Rabinovich (1995). "The Geometry of Graphs and Some of its Algorithmic Applications". In: *Combinatorica* 15.2. Preliminary version in FOCS 1994, pp. 215–245.

Leighton, Frank Thomson and Satish Rao (1999). "Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms". In: *J. ACM* 46.6, pp. 787–832. Preliminary version in FOCS, 1988.