

Randomized tree metrics

Lecture 1

Probability

Events, probabilities,
conditional probabilities,
union bound

Random variables,
expected value,
linearity of expectation
Markov's inequality

Algorithms

Approx. 3-SAT,

Big idea: we only want the average

$$\begin{array}{l} \text{Linearity of Expectation} \\ E[X+Y] = E[X] + E[Y] \quad (\text{stochastic variables}) \\ \text{"average sum = sum of averages"} \end{array}$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases} \quad (i=1, \dots, m)$

$$E[Y_i] = 7/8 \quad (1 \text{ clause case})$$

$$E[\text{\# clauses satisfied}] = E\left[\sum_{i=1}^m Y_i\right] = \sum_{i=1}^m E[Y_i] = 7/8 m$$

let $Y_i = \begin{cases} 1 & \text{if } i\text{th clause satisfied} \\ 0 & \text{otherwise} \end{cases}$

$$Z = \sum_{i=1}^m Y_i = \text{\# clauses satisfied}$$

$$E[Z] = \frac{1}{2} E[Z | x_1 = \text{true}] + \frac{1}{2} E[Z | x_1 = \text{false}]$$

$\Rightarrow$ either $E[Z | x_1 = \text{true}]$ or $E[Z | x_1 = \text{false}]$ is $\geq E[Z]$!

QuickSort,

Random divide and conquer

① pick random 'pivot'

② Divide A into two parts



Expected running time of QuickSort

For each $i, j \in [n]$ w/ $i < j$,

$X_{ij} = \begin{cases} 1 & \text{if } i\text{th smallest element is compared to } j\text{th smallest} \\ 0 & \text{otherwise} \end{cases}$

$\sum_{i,j} X_{ij}$ = total # comparisons made by the algorithm
 $\approx$ total running time

$E[\sum_{i,j} X_{ij}]$ = average running time

$$E[\sum_{i,j} X_{ij}] = \sum_{i,j} E[X_{ij}]$$

now analyze $E[X_{ij}]$ in isolation

$$E[X_{ij}] = \Pr[X_{ij} = 1] = \Pr[i\text{th element is compared to } j\text{th}]$$

i th and j th elem selected before any element in between happens w/ probability $\frac{2}{j-i+1}$

$$E[X_{ij}] = \frac{2}{j-i+1}$$

$$E[\sum_{i,j} X_{ij}] = \sum_{i,j} E[X_{ij}] = \sum_{i,j} \frac{2}{j-i+1}$$

$$\leq 2n \left(\frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{n} \right)$$

$$\leq O(n \log n) \quad \text{with harmonic series}$$

Quickselect (HW)

Lecture 2

Probability

Hash Functions,

Universal Hash Functions

$h: [n] \rightarrow k$ s.t. for all pairs $i, j \in [n]$,
 $\Pr[h(i) = h(j)] \leq \frac{1}{k}$ (much weaker than ideal hash)

e.g. $h(x) = (ax + b) \bmod p \bmod k$

where • p prime, $> k$

• a in $\{1, \dots, p-1\}$ unif. random

• b in $\{0, \dots, p-1\}$ unif. random



Markov's Inequality


Markov's Inequality (special case)

Let $X \geq 0$ be nonnegative RV

$$\Pr[X \geq 2E[X]] \leq \frac{1}{2}$$

Algorithms

Frequency Estimation, Heavy Hitters (in streams)



estimate f_i w/
 $\min_{j=1, \dots, d} A_j[h_j(i)]$

Lecture 3 (last Tuesday)

Probability



k-wise independence,
higher order moments
(generalizing universal hash
functions and Markov's inequality)

Algorithms

Hash tables w/
chaining +
linear probing

k-wise independent hash functions
 $h: [U] \rightarrow [m]$
s.t. for any k keys $x_1, x_2, \dots, x_k \in [U]$,
 k values $v_1, v_2, \dots, v_k \in [m]$,
 $P[h(x_1)=v_1, h(x_2)=v_2, \dots, h(x_k)=v_k] = \frac{1}{m^k}$
eg. $h(x) = a_0 + a_1x + a_2x^2 + \dots + a_{k-1}x^{k-1} \pmod p$
where $p \geq U$ prime,
 $a_0, a_1, \dots, a_{k-1} \in [p]$ indep, uniformly random


Important lemma
k-wise independent random variables
 $P[X_1=y_1, X_2=y_2, \dots, X_k=y_k]$
 $= P[X_1=y_1] P[X_2=y_2] \dots P[X_k=y_k]$
for all values $y_1, \dots, y_k$
let $X_1, \dots, X_n \in \{0,1\}$ be 4-wise independent
and $\mu = E[X_1 + \dots + X_n]$. then for $\beta > 0$,
 $P[X_1 + \dots + X_n \geq \mu + \beta] \leq \frac{4 \max\{\mu, \mu^3\}}{\beta^4}$



Lecture 4


Probability

Galton-Watson processes



Algorithms

Random Contractions for min cut



Randomized contractions


Input: $G=(V,E)$, weights $c \in \mathbb{R}_{\geq 0}^E$

① while $|E| > 1$

② sample $e \in E$ in proportion to c

③ "contract" e

④ return set of edges (in original graph)
contracted to remaining edge




Amplification by branching

Instead of restarting at the beginning, restart partway through

High-level algorithm

① randomly contract until prob.
of ruining the min-cut is $\sim 1/2$

② Branch:
make two copies of current graph
recurse on both copies



Lecture 5

Probability

Multiplicative Chernoff Bound

Multiplicative Chernoff bound

let $X_1, \dots, X_n \in [0, 1]$ be independent

$$\mu = E[X_1 + \dots + X_n]$$

$$\epsilon \in (0, 1)$$


- $P[X_1 + \dots + X_n \geq (1 + \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 3}$
- $P[X_1 + \dots + X_n \leq (1 - \epsilon)\mu] \leq e^{-\epsilon^2 \mu / 2}$

Algorithms (/ results)

Random graphs & the giant component

Erdős-Rényi graphs $G(n, p)$

- n vertices
- each edge sampled independently w/ probability p



Theorem 5.3. Consider a random graph $G \sim \mathcal{G}(n, p)$ for $p = c/n$, where c is a constant.

- If $c > 1$, then with high probability, there is exactly one connected component of G with $\Omega(n)$ vertices, and all other components have size $\leq O(\log n)$.
- For $c < 1$, then with high probability, all connected components of G has size $< O(\log n)$.


Lecture 6

Probability (& techniques)

Approximations via linear programming & randomized rounding

Linear programs (LPs)
Similar to IPs, except no integrality constraints

min/max $\sum_{j=1}^n c_j x_j$ linear objective
over $x \in \mathbb{R}^n$ linear integral variables
st. $A_1 x \leq b_1$ linear constraints
 $A_2 x = b_2$
 $A_3 x \geq b_3$



Continuous optimization vs Discrete optimization

randomization

Algorithms


Max Sat Set Cover Covering Integer Programs

For each variable x_i , independently, randomly set $x_i = t$ with probability g_i , and $x_i = f$ otherwise.

$$E[\# \text{ sat. clauses}] \geq (1 - 1/e) \text{OPT}$$

k	oblivious	LP	average
1	1/2	$z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
2	3/4	$(3/4)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
3	7/8	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
4	15/16	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
5	31/32	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$
...	...	...	...
k	$1 - 1/2^k$	$(1 - 1/e)z_i \geq (3/4)z_i$	$\geq (3/4)z_i$

Can we get $\frac{3}{4}$ -APX?



Rounding LP solution
 $x \in [0,1]^n \rightsquigarrow F \subseteq \{S_1, \dots, S_n\}$

For each set S_j , independently, take S_j with probability x_j .


$$E[|F|] = \sum x_j \leq \text{OPT}$$

Fix a point i .

$$E[\# \text{ sets in } F \text{ covering } i] = \sum_{j \in S_i} x_j \geq z_i$$

$$P[F \text{ not cover } i] = \prod_{j \in S_i} (1 - x_j) \leq e^{-\sum_{j \in S_i} x_j} \leq e^{-z_i}$$

$$P[F \text{ cover } i] \geq 1 - e^{-1}$$



Lecture 7

Probability

Continuous Random Variables

Variance, Chebyshev's Variance Reduction

Median trick

Variance

$$\text{Var}(X) = E[(X - E[X])^2] = E[X^2] - (E[X])^2$$

Chebyshev's inequality

$$P[|X - E[X]| \geq d] \leq \frac{\text{Var}(X)}{d^2}$$

"median trick", "median of means"

- ① take mean to reduce variance, get constant prob. of error
- ② take median to amplify success prob. exponentially

Algorithms

Distinct Elements



Goal: estimate # of distinct elements in the stream

Suppose you had an ideal hash function

$$h: [m] \rightarrow [0,1]$$

and the following facts:

let $Y_1, \dots, Y_k \in [0,1]$ independent and uniformly random

$$X = \min \{Y_1, \dots, Y_k\}$$

$$\text{then } E[X] = \frac{1}{k+1}$$

$$\textcircled{1} X \leftarrow 1$$

② for each element e

$$\textcircled{a} X \leftarrow \min\{X, h(e)\}$$

$$/* E[X] = \frac{1}{k+1}, k = (\text{\# distinct elements}) */$$

$$\textcircled{3} \text{ return } \frac{1}{X} - 1$$

$$\text{recall } E[X] = \frac{1}{k+1}$$

$$\text{goal: show } |X - E[X]| \leq \frac{\epsilon}{k+1}$$

w/ good prob.

$$P[|X - E[X]| \geq \frac{\epsilon}{k+1}]$$

$$\leq \text{Var}(X) \cdot \left(\frac{k+1}{\epsilon}\right)^2 \leftarrow \text{Chebyshev's inequality}$$

where

$$\text{Var}(X) = \frac{k}{(k+1)^2(k+2)} \approx \frac{1}{k^2}$$

Q: how to reduce $\text{Var}(X)$?

we can get $(1 \pm \epsilon)$ -APX

w/ 10% error prob. in $O(\frac{1}{\epsilon^2})$ space

what about high probability?

① make q copies $Y_1, \dots, Y_q$ of

② return median of $Y_1, \dots, Y_q$

Algo correct $\Leftrightarrow \geq \frac{1}{2}$ of Y_i 's correct

let $Z = \# \text{correct } Y_i$'s $E[Z] = \frac{q}{2}$

$$P[Z \leq (1 - \frac{\epsilon}{4})E[Z]] \leq e^{-\left(\frac{\epsilon}{4}\right)^2 E[Z]/2} \leq \frac{1}{n^{O(1)}} \text{ for } q = O(\log n)$$

ideal hash $h: [m] \rightarrow [0,1]$

let $Y_1, \dots, Y_k \in [0,1]$

indep. and unifs. rand.

$$X = \min \{Y_1, \dots, Y_k\}$$

$$\text{then } E[X] = \frac{1}{k+1}$$

$$\text{and } E[X^2] = \frac{2}{(k+1)(k+2)}$$

$$\text{Var}(X) = E[(X - E[X])^2]$$

$$X_1, \dots, X_k \text{ indep. copies of } X$$

$$\bar{X} = \frac{1}{k}(X_1 + \dots + X_k)$$

$$\text{Var}(\bar{X}) = \frac{1}{k} \text{Var}(X)$$

$$P[X_1 + \dots + X_k \leq (1 - \epsilon)k] \leq e^{-\epsilon^2 k/4}$$

Lecture 8

Probability

Gaussian random variables

Gaussian random variables

$$x \sim \mathcal{N}(\mu, \sigma^2)$$

mean variance

Def: the nicest, most pleasant, convenient random variable ever

Fact let $x \sim \mathcal{N}(\mu, \sigma^2)$, $d \in \mathbb{R}$.

$$\text{then } x+d \sim \mathcal{N}(\mu+d, \sigma^2)$$

$$dx \sim \mathcal{N}(d\mu, d^2\sigma^2)$$

Fact let $x_1 \sim \mathcal{N}(\mu_1, \sigma_1^2)$, $x_2 \sim \mathcal{N}(\mu_2, \sigma_2^2)$

$$\text{then } x_1 + x_2 \sim \mathcal{N}(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$$

Fact let $x \sim \mathcal{N}(0, 1)$, $t < \frac{1}{2}$

$$\text{then } E[e^{tx^2}] = \frac{1}{\sqrt{1-2t}}$$

Algorithms

Dimensionality reduction

Bag-of-words

literally any sentence is fine

any ↑ is literally sentence English
five encodes text in high dimensions

Lemma 8.8. Let $A \sim \mathcal{N}^{k \times d}$ be a random matrix where each coordinate A_{ij} is an independently drawn sample from $\mathcal{N}$. Let $\epsilon > 0$ be sufficiently small. For any vector x ,

$$\mathbb{P}\left[(1-\epsilon)\|x\|^2 \leq \frac{1}{k}\|Ax\|^2 \leq (1+\epsilon)\|x\|^2\right] \geq 1 - 2e^{-k/8}$$

$$k \left\{ \left[\frac{1}{k} \|Ax\|^2 \right] \leftarrow \alpha_i \right.$$

$$E[\|Ax\|^2] = \sum_{i=1}^k E[(Ax)_i^2] = k\|x\|^2$$

$$E[(\sum_{i=1}^k A_{ij} x_j)^2] = \|x\|^2$$

$$A_{ij} x_j \sim \mathcal{N}(0, x_j^2)$$

$$\sum_{j=1}^n A_{ij} x_j \sim \mathcal{N}(0, \|x\|^2)$$

$$(Ax) \sim \mathcal{N}^k(0, \|x\|^2)$$

$$\Pr[\|Ax\|^2 \geq k\|x\|^2 \cdot (1+\epsilon)] \leq e^{-\epsilon^2 k/8}$$

Dimensionality Reduction - high level

$O(\log n)$ dimensions suffice

where $n = \# \text{ points}$

$$\mathbb{R}^{\text{very large}} \xrightarrow{f(x)} \mathbb{R}^{O(\log n)}$$

One can map n points into $\mathbb{R}^{O(\log n/\epsilon^2)}$ while preserving all distances up to $(1 \pm \epsilon)$ -APX

Lecture 9

Probability

Locality-sensitive hashing

Fix $\theta \in [0, \pi]$. given q , either:

(a) output $p \in P$ w/ $\angle(p, q) \leq 4\theta$

(b) declare no $p \in P$ w/ $\angle(p, q) \leq \theta$

define $h: S^{d-1} \rightarrow \{-1, 1\}^k$ by

$$h_i(x) = \text{sign}(\langle g_i, x \rangle) \text{ w/ } g_i \sim N^d$$

when querying y :

• if there is $x \in P$ w/ $\angle(x, y) \leq \theta$, then it collides w/ probability $\geq 1/\sqrt{n}$

• expect to collide w/ at most one point $x \in P$ w/ $\angle(x, y) > 2\theta$

Lemma let $\bullet x, y \in S^{d-1}$
 $\bullet g \sim N^d$
 then

$$P[\text{sign}(\langle g, x \rangle) \neq \text{sign}(\langle g, y \rangle)] = \frac{\angle(x, y)}{\pi}$$


$$k = \pi(\log(n)/2\theta)$$



d-dimensional sphere

Random line embedding & Buckets

$$h(x) = \lfloor \langle g, x \rangle + d \rfloor \text{ w/ } g \sim N^d, d \in [0, 1] \text{ unifs. rand.}$$




Algorithms

Approximate nearest neighbors

Nearest neighbor search

Input: n points $P \subset \mathbb{R}^d$




Query: given $q \in \mathbb{R}^d$, find closest point in P to q

Types of distances

• Euclidean distance

$$\|x - z\|_2 = \sqrt{\sum_i (x_i - z_i)^2}$$

• Angle



Lecture 9

Probability (and techniques)

LP duality

L_1 -metric embeddings

Algorithms

Sparsest Cut

Lecture 9

Probability (and techniques)

Algorithms

Randomized
Tree metrics

Metrics $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$ s.t.


(a) Symmetry:

$$d(u,v) = d(v,u)$$

(b) Positivity:

$$d(u,v) \geq 0,$$

$$d(u,v) = 0 \Leftrightarrow u = v$$



(c) Triangle inequality:

$$d(u,v) \leq d(v,w) + d(w,v) \quad \forall u,v,w$$

eg. Euclidean metrics, shortest path metrics, cut metrics

Problems about metrics

Buy-at-bulk network design

Input: • undirected graph $G=(V,E)$

• edge lengths $l: E \rightarrow \mathbb{R}_{>0}$


• demands $b(s,t) \geq 0 \quad \forall s,t \in V$

• for each edge $e \in E$,

subadditive cost function $f_e(x) > 0$

$$f_e(x+y) \leq f_e(x) + f_e(y)$$

(per unit length of unit capacity)




Output: (s,t) -path $P_{st} \quad \forall s,t \in V$

Goal: minimize $\sum_{e \in E} l(e) f_e(u_e)$ w/ $u_e = \sum_{s,t: e \in P_{st}} b(s,t)$

k-median clustering

Input: n points V , metric $d: V \times V \rightarrow \mathbb{R}_{\geq 0}$

Output: k "centers" $c_1, \dots, c_k$



Goal: minimize $\sum_{v \in V} \min_{i=1, \dots, k} d(v, c_i)$

Metric labeling

Input: graph $G=(V,E)$, edge weights $w:E\rightarrow\mathbb{R}_{>0}$,
 k labels $\mathcal{L}$, label costs $c:\mathcal{L}\times\mathcal{L}\rightarrow\mathbb{R}_{>0}$, metric $d:V\times V\rightarrow\mathbb{R}_{\geq 0}$

Output: labeling $f:V\rightarrow\mathcal{L}$



Goal: minimize $\sum_{v \in V} c(v, f(v)) + \sum_{\{u,v\} \in E} w(u,v) d(f(u), f(v))$

Other metric-based problems

(sometimes via correspondance between cuts $\leftrightarrow$ metrics)

Buy-at-bulk network design

k-median clustering

metric labeling

group Steiner tree

metrical task systems



vehicle routing

hierarchical placement


distributed k-server

covering Steiner tree

6



Tree metrics: shortest path metric on a tree



Shortest paths

(or sum/min over multiple such trees)

- $\tilde{O}(mn)$ time in general graphs
 - super easy on trees!
 - LCA queries
 - path-to-root queries
- } $\log(n)$ time w/ link-cut trees, etc

Other metric-based problems

(sometimes via correspondence between cuts $\leftrightarrow$ metrics)

Buy-at-bulk network design

k-median clustering

metric labeling

group Steiner tree

metrical task systems



vehicle routing

hierarchical placement

distributed k-server

covering Steiner tree

⌊



these problems are easy for tree metrics

Approximation by tree metrics

Input: undirected $G=(V,E)$ w/ edge weights $w:E \rightarrow \mathbb{R}_{>0}$

(High-level)

Goal: compute tree T (containing V as nodes)
so that shortest paths on T are close to G

$$(\text{stretch } e) = \frac{\overset{\text{(distance in } T\text{)}}{d_T(u,v)}}{\underset{\text{(distance in } G\text{)}}{d_G(u,v)}} \quad \text{for } e=\{u,v\}$$

ideally $(\text{stretch } e) = \text{polylog}(n) \quad \forall e$

Approximation by tree metrics

Input: undirected $G=(V,E)$ w/ edge weights $w:E \rightarrow \mathbb{R}_{\geq 0}$

(High-level)


Goal: compute tree T (containing V as nodes)
so that shortest paths on T are close to T

for $e=\{u,v\}$

$$(\text{stretch } e) = \frac{\overset{\text{(distance in } T\text{)}}{d_T(u,v)}}{\underset{\text{(distance in } G\text{)}}{d_G(u,v)}}$$

ideally

$$(\text{stretch } e) = \text{polylog}(n) \quad \forall e$$



Approximation by tree metrics

Input: undirected $G=(V,E)$ w/ edge weights $w:E \rightarrow \mathbb{R}_{\geq 0}$

(High-level)

Goal: compute tree T (containing V as nodes)
so that shortest paths on T are close to G

① deterministic low-stretch spanning trees [AKPW95]

spanning tree T s.t.

$$\frac{1}{|E|} \sum_{e \in E} (\text{stretch } e) \leq n^{o(1)}$$

Approximation by tree metrics

Input: undirected $G=(V,E)$ w/ edge weights $w:E \rightarrow \mathbb{R}_{>0}$

(High-level)

Goal: compute tree T (containing V as nodes)
so that shortest paths on T are close to G

① deterministic low-stretch spanning trees [AKPW95]


② Randomized dominating tree metric [Bartal 96, 98; FRT04]

Auxiliary tree T w/ V as leaves s.t.


("dominating")

$$\bullet d_T(u,v) \geq d_G(u,v) \quad \forall u,v$$

$$\bullet E[\text{stretch } e] \text{ small } \forall e \in E$$



Hierarchical Tree Metric



edge from height $i+1$ to i has weight d^i for
fixed d . (we will have $d=4$)

On Approximating Arbitrary Metrics by Tree Metrics

Yair Bartal *



Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Journal of Computer and System Sciences 69 (2004) 485–497

JOURNAL OF
COMPUTER
AND SYSTEM
SCIENCES

<http://www.elsevier.com/locate/jcss>

Abstract

We improve the result of [Bart96] on probabilistic approximation of metric spaces by “hierarchically well-separated tree” metric spaces. We obtain an approximation factor of $O(\log n \log \log n)$ getting the gap to the lower bound within lower order factors. We also give a deterministic version of the result which gives a tree with low average distortion of distances.

The results have applications in a variety of areas including approximation algorithms, on-line algorithms and distributed computation and hence we obtain new approximation bounds for these applications.

1 Introduction

1.1 Probabilistic Metric Approximations and HSTs

We study the problem of approximating arbitrary metric spaces by distances in more simple metric spaces. This allows reducing the solution of problems to the more simple spaces.

Approximations of metric spaces by more simple metric spaces has been intensively studied in various areas of mathematics and computer science. We mention but a few. Johnson and Lindenstrauss [JL84] and Bourgain [Bour85] study embeddings in Hilbert spaces (from a perspective of functional analysis). Graham and Winkler [GW85] study embeddings in $\mathbb{Z}^d$ (from a graph theoretic motivation). Algorithmic applications in distributed computation and graph algorithms, have led to the notion of graph spanners, introduced by Peleg and Ullman [PU89] and later studied in many papers includ-

ing [PS89, ADDJS90, CDNS92], and to low-distortion embeddings in low-dimensional real normed spaces by Linial, London and Rabinovich [LLR94].

In [Bart96] we define *probabilistic approximation* of metric spaces by a set of “simpler” metric spaces $\mathcal{S}$. Given a metric space M over a finite set V , we consider sets $\mathcal{S}$ of “simple” metric spaces over V such that distances in M are dominated by the corresponding distances in metric spaces in $\mathcal{S}$. M is said to be α -*probabilistically-approximated* by $\mathcal{S}$, if there exists a probability distribution over $\mathcal{S}$ such that for every pair of nodes in V , the expected ratio between the distance in a metric space in $\mathcal{S}$ chosen from the probability distribution, and the distance in M is at most α .

Regarding a randomized algorithm as a distribution over deterministic algorithms, its performance ratio is the expected performance ratio of these algorithms over its own coin tosses. The probabilistic metric approximation approach provides a *novel technique* for bounding the performance ratio of randomized algorithms for optimization problems on metric spaces where a solution can be expressed as a linear combination of distances in the metric space. If every metric space is α -probabilistically-approximated by $\mathcal{S}$, and the performance ratio of randomized algorithms for $\mathcal{S}$ is at most β , then the performance ratio of randomized algorithms for any metric space is at most $\alpha\beta$.

The advantage of using probabilistic metric space approximation is that certain classes of simple metric spaces which cannot provide a good deterministic approximation of arbitrary metric spaces may provide a good probabilistic approximation.

A natural candidate for a class of “simple” metric spaces is the class of tree metric spaces. Indeed, intensive recent research has been concerned with approximating metric spaces by tree metrics, including work on *numerical taxonomy* [BG92, ABFNP96] and *tree spanners* [Cai92, CC95].

However, there exist simple metric spaces for which every specific tree dominating the distances of the metric space has distortion $\Omega(n)$ (trivially matched by the

A tight bound on approximating arbitrary metrics by tree metrics

Jittat Fakcharoenphol,^{a,1} Satish Rao,^{b,2} and Kunal Talwar^{c,*,3}

^aKasetsart University, Bangkok, Thailand

^bComputer Science Division, University of California, Berkeley, CA 94720, USA

^cComputer Science Division, University of California, Berkeley, CA 94720, USA

Received 3 September 2003; revised 12 April 2004

Available online 24 June 2004

Abstract

In this paper, we show that any n point metric space can be embedded into a distribution over dominating tree metrics such that the expected stretch of any edge is $O(\log n)$. This improves upon the result of Bartal who gave a bound of $O(\log n \log \log n)$. Moreover, our result is existentially tight; there exist metric spaces where any tree embedding must have distortion $\Omega(\log n)$ -distortion. This problem lies at the heart of numerous approximation and online algorithms including ones for group Steiner tree, metric labeling, buy-at-bulk network design and metrical task system. Our result improves the performance guarantees for all of these problems.

© 2004 Elsevier Inc. All rights reserved.

Keywords: Metrics; Embeddings; Tree Metrics

1. Introduction

1.1. Metric approximations

The problem of approximating a given graph metric by a “simpler” metric has been a subject of extensive research, motivated from several different perspectives. A particularly simple metric of

*Corresponding author.

E-mail address: kunal@cs.berkeley.edu (K. Talwar).

¹Supported in part by a DPST scholarship and NSF Grant CCR-0105533. Work done while the author was at University of California, Berkeley.

²Supported by NSF Grant CCR-0105533.



³Supported by NSF Grants CCR-0105533 and CCR-9820897.

*Bell-Labs, Lucent Technologies, 600 Mountain Avenue, Murray Hill, NJ 07974. E-mail: yair@research.bell-labs.com


Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.


STOC '98 Dallas Texas USA
Copyright ACM 1998 0-89791-962-9/98 \$5.00


Theorem In randomized polynomial time,
 randomized hierarchical dominating tree metric w/
 $E[\text{stretch } e] \leq O(\log n) \quad \forall e \in E$



Scoopedy-whoop-de-scoop







Let D = diameter, $L = \lceil \log_4 D \rceil$



1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ units. at random

2. for $i = L, L-1, \dots, 0$

a. for each v_j in order

i. $C_{ij} \leftarrow B(v_j, d4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$

3. Use $C_{i,j}$'s to partition V in a tree from top down.





(each node associated w/ subset of V , starting w/ root = V
 node w/ set S at height i has children correspond.
 to (nonempty) sets of the form $S \cap C_{i,j}$)

Analysis

Fix $u, v \in V$.

need to show

$$E[d_T(u, v)] \leq O(\log n) d(u, v)$$



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. For $i = L, L-1, \dots, 0$
 - a. For each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

need to show

$$E[d_T(u,v)] \leq O(\log n) d(u,v)$$

let k = height where u and v are split


$$\text{then } d_T(u,v) \approx 4^k$$

u, v split by cluster centered at some $w \in V$

then we say

w "contributes" 4^k to $d_T(u,v)$

$$d_T(u,v) = O\left(\sum_w (\text{contrib. of } w \text{ to } d_T(u,v))\right)$$



Let D = diameter, $L = \lceil \log_4 D \rceil$



1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

need to show $E[d_T(u,v)] \leq O(\log n) d(u,v)$

$$d_T(u,v) = O\left(\sum_w (\text{contrib. of } w \text{ to } d_T(u,v))\right)$$

Claim: let w be ℓ th closest vertex to u or v . then

$$E[\text{contrib. from } w] \leq \frac{4 d(u,v)}{\ell}$$



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u,v)$



need to show $E[d_T(u,v)] \leq O(\log n) d(u,v)$

$$d_T(u,v) = O\left(\sum_w (\text{contrib. of } w \text{ to } d_T(u,v))\right)$$

Claim: let w be l th closest vertex to u or v . then

$$E[\text{contrib. from } w] \leq \frac{4 d(u,v)}{l}$$

if so:



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u,v)$

Claim: let w be l th closest vertex to u or v . then $E[\text{contrib. from } w] \leq \frac{4 d(u,v)}{l}$

Proof wlog $d(w,u) \leq d(w,v)$


define two necessary events

$E_1: d(w,u) \leq d 4^k \leq d(w,v)$

$E_2: w$ ordered before any of the $l-1$ vertices closer to u or v

E_1 and E_2 are independent (!)

$\Pr[E_2] = 1/l$



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unis. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.

$k = \text{height where } u \text{ and } v \text{ are split}$
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$


$$E_1: d(w, u) \leq d 4^k \leq d(w, v)$$

2 cases.

case 1: $d(u, v) \geq 4d(w, u)$

case 2: $d(u, v) \leq d(w, u)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)




Let D = diameter, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.

k = height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

$$E_1: d(w, u) \leq d 4^k \leq d(w, v)$$

case 1: $d(u, v) \geq 4d(w, u)$



Let D = diameter, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. For $i = L, L-1, \dots, 0$
 - a. For each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.


k = height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

$$E_1: d(w, u) \leq d 4^k \leq d(w, v)$$

case 2: $d(u, v) \leq d(w, v)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)

claim: choice of k is unique



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

$$E_k: d(w, u) \leq d 4^k \leq d(w, v)$$


case 2: $d(u, v) \leq d(w, v)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)

claim: choice of k is unique

let $d \in [1, 2]$, k satisfy E_k

for $k+1$:



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

$$E_k: d(w, u) \leq d 4^k \leq d(w, v)$$


case 2: $d(u, v) \leq d(w, v)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)

claim: choice of k is unique

let $d \in [1, 2]$, k satisfy E_k

for $k-1$:



Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$ s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$


$$E_1: d(w, u) \leq d 4^k \leq d(w, v)$$

case 2: $d(u, v) \leq d(w, v)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)

✓ choice of k is unique

$$P[E_1] \leq$$



Let D = diameter, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.

k = height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

$$E_1: d(w, u) \leq d 4^k \leq d(w, v)$$


case 2: $d(u, v) \leq d(w, v)$

($d(w, u)$ might be much bigger than $d(u, v)$, which is risky)

✓ choice of k is unique

$$✓ P[E_1] \leq \frac{d(u, v)}{4^k}$$

$$\begin{aligned} E[\text{contrib. from } w \sim] \\ &\leq 4^{k+1} \cdot P[E_1] \cdot P[E_2] \\ &\leq 4 d(u, v) / \epsilon \end{aligned}$$





Let $D = \text{diameter}$, $L = \lceil \log_4 D \rceil$

1. let $v_1, \dots, v_n$ order V randomly
 $d \in [1, 2]$ unifs. at random
2. for $i = L, L-1, \dots, 0$
 - a. for each v_j in order
 - i. $C_{i,j} \leftarrow B(v_j, d 4^{i-1}) \setminus \bigcup_{k < j} C_{i,k}$
3. Use $C_{i,j}$'s to partition V in a tree from top down.

$k =$ height where u and v are split
 u, v split by cluster centered at some $w \in V$
 then w "contributes" 4^k to $d_T(u, v)$

Theorem In randomized polynomial time,
 randomized hierarchical dominating tree metric w/
 $E[\text{stretch } e] \leq O(\log n) \quad \forall e \in E$



Other metric-based problems

(sometimes via correspondence between cuts $\leftrightarrow$ metrics)

Buy-at-bulk network design

k-median clustering

metric labeling

group Steiner tree

metrical task systems



vehicle routing

hierarchical placement

distributed k-server

covering Steiner tree

⌊



these problems are easy for tree metrics

Chapter 11

Tree Metrics

11.1 Introduction

Many flow and cut problems are ultimately about paths, whether packing paths in flow, cutting paths in cuts, evaluating shortest paths in fractional cuts, etc. Obviously there are many paths between any two points in a graph and this makes all these problems nontrivial. An extremely simple setting, then, would be a graph where there is a unique path between any two vertices: by definition, a *tree*. Many graphs problems become trivial in a tree. Here we will study a bold approach to graph algorithms based on this idea: process the input graph G to produce a tree T that preserves its salient properties, solve the problem on T , and lift the solution back to G . Of course a tree T cannot preserve all of G , so we will only preserve specific properties and only approximately at that, in such a way that is appropriate to the problem at hand.

In this discussion, we will focus on preserving the shortest path metric of a graph. Let G be an undirected graph, and let d_G denote the shortest path metric in G . Let T be a spanning tree of T (with the same edge weights), and let d_T denote the shortest path metric in T . (Of course, the shortest path in T is also the only path.) Ideally, we want $d_T(u, v)$ to resemble $d_G(u, v)$ as much as possible for each edge $e = \{u, v\}$. In general, we have

$$d_G(u, v) \leq d_T(u, v) \text{ for all } u, v \in V$$

simply because T is a subgraph of G . For an edge $e = \{u, v\}$, we say that the *stretch* of e is defined as the ratio

$$(\text{stretch } e) \stackrel{\text{def}}{=} \frac{d_T(u, v)}{d_G(u, v)}.$$

We say that d_T has *uniform stretch* (at most) α , for $\alpha \geq 1$, if every edge has stretch at most α .

A natural goal is to obtain a spanning tree with small uniform stretch. However the n -vertex cycle C_n presents a lower bound of $n - 1$. Indeed, any spanning tree T of C_n is obtained by dropping one edge e ; this edge e is stretched around the cycle, so to speak, and has stretch $n - 1$.

The n -vertex cycle C_n indicates that we cannot, in the worst-case, find spanning trees with uniform stretch better than $n - 1$. (We leave it as exercise 11.1 to obtain a matching upper bound.) The rest of this chapter discusses two different approaches that obtain better bounds for relaxations of this problem. (In the lecture, we only discuss the second, randomized one in detail.)

Low-stretch spanning trees. Alon, Karp, Peleg, and West [AKPW95] showed how to compute a spanning tree T where the average stretch among all edges is $n^{o(1)}$, in the sense that

$$\frac{1}{|E|} \sum_{e \in E} (\text{stretch } e) \leq n^{o(1)}$$

Dominating tree metrics. Bartal [Bar96; Bar98] ignored the requirement that T is a spanning tree of G ; more generally he sought auxiliary trees T where the vertex of G correspond to the leaves of T , while retaining the property T that $d_T \geq d_G$. He produced *randomized* trees where *for each edge* the average stretch was $\text{polylog}(n)$:

$$\mathbf{E}_T[(\text{stretch } e)] \leq O(\text{polylog}(n)) \text{ for all } e \in E.$$

Note that this is a different sense of “average stretch” than above. We will present an algorithm of [FRT04] building on [Bar98] to obtain (per-edge) average stretch of $O(\log n)$.

11.2 Low-Stretch Spanning Trees

We first present the low-stretch spanning trees of [AKPW95]. Here we recall that we want to compute a spanning tree with low stretch on average over all the edges.

Suppose our goal was to obtain average stretch (roughly) D for a parameter $D > 0$. Let us partition the graph in vertex-disjoint subgraphs each with radius at most $D/2$ from some center vertex, and compute a shortest path tree from each center. Then every edge *within* a neighborhood has stretch at most D ; it remains to address the edges that are cut by the partition. We can try to address these edges recursively by contracting every subgraph/subtree into a single vertex, leaving a multigraph G' consisting of the cut edges, and recursing on this graph. This produces a tree T' on the contracted multigraph G' ; expanding out the vertices of T' by the underlying

[AKPW95] algorithm for $m^{o(1)}$ -average stretch spanning trees:

1. *Compute a low diameter decomposition:* repeatedly, until no vertices remain:
 - A. Select a remaining vertex v and compute a shortest path tree of v in the remaining graph.
 - B. Remove the set C_v all vertices (remaining) in V within some distance $R' \leq R$ such that

$$|\delta(C_v)| \leq O(\log(m))(1 + |E[C_v]|)$$

where $E[C_v]$ denotes the set of all (remaining) edges incident to some vertex in v .

2. *Recurse:* Let G' be the multi-graph obtained from the input graph by contracting each C_v to a single vertex. Recurse on G' to obtain a spanning tree T' , and return the tree T obtained by replacing each C_v with the corresponding shortest path tree.

shortest path trees gives a spanning tree T . Recursively, we might expect that every edge cut edge e has stretch $O(R)$ in T' , but this expands out to stretch $O(R^2)$ with respect to T because passing through a vertex in T' actually corresponds to traversing a path of length $O(R)$ in the underlying tree.

So we have a problem where each step of the recursion induces an additional factor of R . Now, recall that we want to preserve *average* stretch. Let $f(m)$ be the stretch obtained by the recursive approach. We have

$$f(m) \leq D(\# \text{ internal edges}) + (D + 1)f(\# \text{ external edges}).$$

Studying this recursion, we can see that if the number of external edges was extremely small, then we might hope that it can offset the extra factor of D . How can we minimize the number of external, or cut, edges? Region growing! Low diameter decompositions!

If we partition the graph by region growing techniques, we can ensure that

$$(\# \text{ external edges}) \leq \frac{O(\log m)}{R}m.$$

This revises the recursive bound as

$$f(m) \leq Dm + (D + 1)f(c \log(m)m/D)$$

for a constant $c > 0$. For $D = e^{\sqrt{\log(m)/\log \log m}}$ the recursion is bounded by $f(m) = e^{O(\sqrt{\log m \log \log m})}$. (The calculations are given below.)

Theorem 11.1. *The AKPW algorithm returns a spanning tree with average stretch $e^{O(\sqrt{\log m \log \log m})}$.*

Solving the recurrence. We have

$$f(m) \leq Dm + (D + 1)f(c_0 \log(m)m/D)$$

where $\epsilon = c_0 \log(m)/D$ for a constant c_0 . The height of the recursion is

$$h = O(\log_{D/c_0 \log(m)} m) = O(\log(m)/\log(D/c_0 \log(m))) = O\left(\frac{\log(m)}{\log(D)}\right)$$

assuming $D = \Omega(\log m)$. Unrolling the recursion gives

$$f(m) \leq O(Dm) \sum_{i=0}^h ((1 + 1/D)c_0 \log(m))^i \leq Dme^{O(h \log \log m)}.$$

To minimize the RHS, we can instead minimize the logarithm of the RHS, $\log(m) + \log(D) + O(h \log \log m)$. Choosing D to make the last two terms (roughly) equal, we have

$$\log(D) = \frac{\log(m) \log \log(m)}{\log(D)},$$

hence

$$\log(D) = \sqrt{\log(m) \log \log(m)}.$$

Then $D = e^{\sqrt{\log(m) \log \log(m)}}$ gives

$$f(m) \leq me^{O(\sqrt{\log(m) \log \log(m)})},$$

as desired.

11.3 Hierarchical Tree Metrics

The [AKPW95] algorithm was able to obtain low stretch *in total*. This means that a few unlucky edges might have extremely high stretch. In this section we want every edge to have low stretch, in some sense. Unfortunately we already know that it is impossible to guarantee $o(n)$ stretch for every edge simultaneously. But a different possibility opens up when we allow for randomization. Perhaps we can output a *randomized* tree T , such that for each edge e , the expected stretch of e is $o(n)$. Such a claim does *not* contradict the lower bound for the n -vertex cycle; in fact, one can get constant expected stretch for the cycle which we leave as exercise 11.2.

To build some intuition, let us start from the [AKPW95] algorithm for inspiration (even though ultimately we will not produce a spanning tree). A high level goal is to inject randomization so that every edge has a decent chance at having low stretch. To this end there are at least two natural ways to introduce randomization into [AKPW95].

1. We can make the radii of the clusters randomized, instead of a deterministic function of the total number of edges cut.
2. Second, the order of vertices that center the clusters can be randomized, which would seem most equitable.

Both of these ideas will be reflected at a high level in the following algorithm which we now present.

The algorithm we present will produce a randomized *hierarchical tree metric* over V . This means that the tree T will be rooted, with V at the leaves, and the edges between height i and height $i - 1$ have length α^i for a fixed constant α . Here we choose $\alpha = 4$ to simplify calculations, though we note that $\alpha = 2$ is more common, and the analysis can be adjusted to accommodate any fixed constant. The convenience of a hierarchical tree T is that the tree distance $d_T(u, v)$ is entirely determined by the height of their least common ancestor, and within a constant factor of the biggest edges at the top of the corresponding subtree. This additional structure turns out to be useful for several other problems.

We now present [FRT04]’s randomized algorithm. We assume the input is an edge-weighted graph where the minimum edge length is normalized to 1. We let $D = \max_{u,v} d(u, v)$ denote the diameter of the graph. Below we describe the algorithm in detail and first we give a high level description. For every radius of the form $\alpha 4^i$, we are randomly scooping out balls of size $\alpha 4^i$, where $\alpha \in [1, 2]$ is drawn uniformly at random, and the vertices at the center of the balls are in random order. A key and subtle point is that we use the *same* (random) α and ordering for every i . the

intersections of these balls (across i 's) induce a laminar family of sets over V which are arranged as a tree.

[FRT04]'s algorithm producing a randomized hierarchical tree metric.

1. Let $v_1, \dots, v_n$ be a uniformly random ordering of V . Let $L = \lceil \log_4 D \rceil$. Let $\alpha \in [1, 2]$ be drawn uniformly at random.
2. For i from L down to 0,
 - (a) For each vertex v_j in order,
 - i. Let $C_{i,j}$ be the set of vertices at distance at most $\alpha 4^{i-1}$ from v_j , excluding any vertex already included by the cluster of a previous $C_{i,j}$.
3. We use the $C_{i,j}$'s to arrange the vertices as leaves in a tree T hierarchically as follows. For each intermediate node x at height i , the leaves in the subtree rooted at x corresponds to a set of vertices with diameter at most 4^i . The root at height $L + 1$ corresponds to V . The nodes at height L correspond to the clusters $C_{L,j}$. In general, for a node x at height i corresponding to a set $S \subseteq V$, its children correspond to the (nonempty) intersections of S with clusters $C_{i,j}$. (Here a cluster center v_j may not be in S .) Observe the leaves (at height 0) each correspond to a single vertex V because $C_{0,j} = \{v_j\}$ for all j . Each edge descending from height i is given weight 4^i .

So much for the algorithm. Here, then, is the key claim.

Theorem 11.2. *[FRT04]'s algorithm produces a randomized hierarchical tree T such that for each edge e , $\mathbf{E}[(\text{stretch } e)] \leq O(\log n)$.*

To prove the theorem, fix an edge $e = \{u, v\}$. Recall that $d(u, v)$ is decided, up to a constant factor, by the height k where u and v are first separated. (Then $d(u, v) = O(4^k)$.) When this occurs, u and v are separated in particular by a cluster of radius $\alpha 4^k$ centered at some vertex w ; in this event, we say that w “contributes” 4^k to $d(u, v)$ (which upper bounds the diameter of the remaining vertices). By this terminology, we have

$$d_T(u, v) \leq \sum_w O(1)(\text{contrib. of } w \text{ to } d_T(u, v)). \quad (11.1)$$

Now, fix a vertex w . Suppose that w was the ℓ th closest vertex to u or v (i.e., with respect to $\min\{d(w, u), d(w, v)\}$). The key lemma, which we analyze below, is that

$$\mathbf{E}[\text{contrib. from } w \text{ to } d_T(u, v)] \leq O(1/\ell)d(u, v).$$

Taking expectations of eq. (11.1) and applying the bound above to each w gives $O(\log n)$ stretch, as desired.

It remains to prove the key lemma, as follows.

Lemma 11.3. *Let w be the ℓ th closest vertex to u or v . Then*

$$\mathbf{E}[\text{contrib. from } w \text{ to } d_T(u, v)] \leq 4d(u, v)/\ell.$$

Proof. We first observe that w contributes to $d_T(u, v)$ only if the following two events both occur.

E_1 . $d(w, u) \leq \alpha 4^k \leq d(w, v)$ for some k .

E_2 . w is ordered before any of the $\ell - 1$ vertices that are closer to u or v .

Indeed, the necessity of the first condition is clear. The second is necessary because if any closer vertex would otherwise cluster either u or v (or both) before w .

We also observe that the above conditions are *independent*, since the first event depends (only) on α and the second event depends on the random ordering, which are independent. It is also clear that the second event E_2 occurs with probability $1/\ell$. It remains to analyze E_1 . We have two cases.

Case 1: $d(u, v) \geq 4d(w, u)$. Then for any k satisfying the inequality in E_1 , this inequality and the triangle inequality imply that

$$4^{k+1} \leq 2d(w, v) \leq 2(d(w, u) + d(u, v)) \leq (5/2)d(u, v).$$

Thus the contribution from w is at most $O(d(u, v))$. It follows that

$$\mathbf{E}[\text{contrib. from } w \text{ to } d_T(u, v)] \leq 2.5d(u, v) \mathbf{P}[E_2] = 2.5d(u, v)/\ell,$$

as desired.

Case 2: $d(u, v) \leq d(w, u)$. We first note that there may not be any k in item E_1 satisfying inequality E_1 ; if not, then the claim is immediate. Henceforth we assume such a k exists. We claim that the choice of k is unique. Indeed, suppose there exists some choice of k and $\alpha \in [1, 2]$ such that the inequality in E_1 holds. We have

$$d(w, u) \stackrel{(a)}{\geq} d(w, v) - d(u, v) \stackrel{(b)}{>} 4^k - 4^k/2 > 2 \cdot 4^{k-1}$$

which rules out smaller values of k . Here (a) is by the triangle inequality and (b) is by assumption on α . To rule out larger values of k , we have

$$d(w, v) \leq d(w, u) + d(u, v) < 4^{k+1}$$

by similar reasoning.

Thus the choice of k in E_1 is unique; fix k as such. Now we have

$$\mathbf{P}[E_1] \leq \frac{|[d(w, u), d(w, v)]|}{|[4^k, 2 \cdot 4^k]|} \stackrel{(c)}{\leq} \frac{d(u, v)}{4^k}$$

by (c) the triangle inequality. Thus

$$\mathbf{E}[\text{contrib. from } w \dots] \leq 4^{k+1} \mathbf{P}[E_1] \mathbf{P}[E_2] \leq 4d(u, v)/\ell,$$

as desired. □

11.4 Exercises

Exercise 11.1. Design and analyze an algorithm that computes a spanning tree with uniform stretch $n - 1$ (matching the lower bound induced by the cycle).

Exercise 11.2. For the n -vertex cycle C_n , describe a randomized tree metric where each edge has expected stretch $O(1)$.

Exercise 11.3. Recall that the low-stretch spanning tree of [AKPW95] obtained average stretch $n^{o(1)}$. We consider extensions to the weighted average. Let $w(e) \in \mathbb{R}_{>0}$ be a positive weight for every edge, and let $W = \sum_{e \in E} w(e)$ be the total weight. We assume for simplicity that the weights are between 1 and $\text{poly}(n)$. We still treat the edges as unit length edges. Design and analyze an algorithm to compute a spanning tree T such that

$$\frac{1}{W} \sum_{e \in E} w(e)(\text{stretch } e) \leq n^{o(1)}.$$

Exercise 11.4. Show how to use the randomized tree metric to randomly round the sparsest cut LP and obtain $O(\log n)$ -approximation for sparsest cut.

Exercise 11.5. Prove that the $O(\log n)$ bound is tight for tree metrics (up to constants).