Randomized Algorithms 2025A – Lecture 7*
Probabilistic Embedding into Dominating Trees

Robert Krauthgamer

1 Metric Embeddings

Definition (metric space): We say that (X, d) is a *metric space*, if X is a set (of points), and $d : X \times X \to \mathbb{R}_+$ (a distance function) is symmetric, non-negative (with 0 only between a point and itself), and satisfies the triangle inequality.

Prime examples: A simple example is the Euclidean space $\mathbb{R}^d$. Or one can take a subset of its points.

Given a graph with positive (or non-negative) edge weights $G = (V, E, w)$, its shortest-path metric d_G is a metric on the vertex set V. Or one can take a subset $V' \subset V$.

Optimization problems: Many optimization problems are naturally defined on metric spaces, for example TSP and k-median. (The input may specify a subset of the points to be visited, clustered, potential centers, etc.)

Definition (embedding): An *embedding* of a metric space (X, d_X) into a metric space (Y, d_Y) is a map $f : X \to Y$. Its *distortion* is the least $D = D_1 D_2 \geq 1$ such that

$$\forall x, x' \in X, \quad \frac{1}{D_1} d_X(x, x') \leq d_Y(f(x), f(x')) \leq D_2 \cdot d_Y(d_X(x, x')).$$

Remark: In many cases, we can scale distances in Y and thus assume WLOG that $D_1 = 1$ (or alternatively $D_2 = 1$).

Definition (tree metric): A metric space (X, d) is called a *tree metric* if there exists a tree $T = (V, E, w)$ with $X \subset V$ and its shortest-path metric d_T satisfies

$$\forall x, x' \in X, \qquad d(x, x') = d_T(x, x').$$

Exer: Show that a metric space (X, d) is a tree metric if and only if it satisfies the following (called 4-point condition)

$$\forall x, y, z, w \in X, \quad d(w, x) + d(y, z) \leq \max\{(d(w, y) + d(x, z), d(w, z) + d(x, y)\}.$$

*These notes summarize the material covered in class, usually skipping proofs, details, examples and so forth, and possibly adding some remarks, or pointers. The exercises are for self-practice and need not be handed in. In the interest of brevity, most references and credits were omitted.

Many optimization problems can be solved in polynomial time in tree metrics, including TSP and k-median (hint: use dynamic programming).

Observation: Given a metric space (X, d_X) and a distortion-D embedding of it into a tree metric (Y, d_Y), one can compute a D-approximate solution for TSP and k-median.

This promising approach has the following serious obstacle, which we will bypass using randomization.

Theorem 1 [Rabinovich and Raz, 1998]: Every embedding of the shortest-path metric of C_n, an unweighted n-cycle, into a tree metric has distortion $\Omega(n)$.

Remark: The special case where the tree is a spanning tree of C_n is easy, but the general case requires a proof.

Example [Karp]: Let T be a spanning tree of C_n obtained by removing one edge uniformly at random. Then for all $x, y \in C_n$,

$$d_T(x, y) \geq d_{C_n}(x, y).$$
$$\mathbb{E}[d_T(x, y)] \leq 2d_{C_n}(x, y).$$

Exer: Extend this to a cycle with edge lengths (hint: sample proportionally to the edge lengths).

1.1 Probabilistic Embedding

Probabilistic embedding into trees: A *probabilistic embedding* of a metric (X, d) into trees is a probability distribution over mappings $f : X \to T$ and tree metrics (T, d_T).

The tree T is called *dominating* if

$$\forall x, y \in X, \qquad d_T(f(x), f(y)) \geq d(x, y).$$

The probabilistic embedding has *distortion* $D \geq 1$ if

$$\forall x, y \in X, \qquad d_T(f(x), f(y)) \leq D \cdot d(x, y).$$

Example: We saw above that the n-cycle C_n admits a probabilistic embedding into dominating trees with distortion 2.

Remark: T is random (not fixed) and may contain Steiner points (points that are not images under f).

1.2 Probabilistic Embedding into Dominating Trees

Theorem 2 [Bartal'96, Fakcharoenphol-Rao-Talwar'03]:

Every n-point metric admits a probabilistic embedding into trees with distortion $O(\log n)$.

The main usage of this theorem is that it "reduces" problems about X to problems about a tree (metric), which is usually easier.

Example application: Metric TSP:

Given a TSP instance which is an n-point metric space (X, d), apply the theorem to randomly construct a tree T with metric d_T. Now solve TSP on this tree optimally by going around the tree twice (assuming all leaves are point in X, otherwise we can prune such vertices). Finally, output the same tour (same permutation of points) as a solution to TSP on (X, d).

Analysis: First bound the algorithm's performance

$$ALG(X, d) \le ALG(X, d_T) = TSP(X, d_T),$$

then bound the expectation of the optimum in the tree

$$\mathbb{E}[TSP(X, d_T)] \le O(\log n) TSP(X, d).$$

Key property: the objective is linear in the distances.

Remark: It works similarly even we only have an $O(1)$-approximation for TSP in trees.

Remark: There is a much better algorithm for metric TSP (approximation 2 by twice MST, and even 3/2 by Christofides), but this approach works also for generalizations like vehicle routing.

Proof of Theorem 2:

Assume WLOG that the minimum interpoint distance in X is 2, and denote the maximum as $\Delta = \text{diam}(X)$, and $L = \lceil \log_2 \Delta \rceil$.

We may refer to X as a complete graph, to every pair of points (x, y) as an edge.

Definition (hierarchical decomposition): A *hierarchical decomposition* of X is a sequence $P_L, \dots, P_1, P_0$ of partitions of X, such that

a) $P_L = \{X\}$ (the trivial partition)

b) every element of P_i, referred to as a *cluster* $S \subseteq X$, has diameter at most 2^i. Thus, $P_0 = \{X\}$ (all clusters are singletons).

c) each P_i is a refinement of P_{i+1}, i.e., each element of P_i, referred to as a *cluster* $S \subseteq X$, is contained entirely in some cluster of P_{i+1}.

Building a tree: Given a hierarchical decomposition, we build a tree metric T with $L+1$ levels, where the vertices at level i are the clusters of P_i. Start with a root that corresponds to the single cluster X of P_L. Let each cluster of P_i be the child of the cluster in P_{i+1} that contains it, and let the edge between them have length 2^i. The leaves correspond to clusters that are singletons, and we can thus let the embedding f map each $x \in X$ to the leaf which is the singleton cluster $\{x\}$.

Exer: Extend the proof below to obtain a tree T' whose vertex set is exactly X (without additional vertices).

Hint: Get rid of non-leaf vertices in T by "mapping" them to leaves.

Lemma 3: For every two points $x, y \in X$ there is a unique integer i such that x, y are in the same cluster of P_{i+1} but not of P_i. Moreover, $d_T(x,y) \in [2 \cdot 2^i, 4 \cdot 2^i)$.

Proof: immediate.

Lemma 4: This (hierarchical) tree metric d_T dominates (X, d).

Proof: immediate (and uses Lemma 3).

Lemma 5: Suppose the hierarchical decomposition is randomized and guarantees, for a certain $\alpha > 0$, that

$$\forall x, y \in X, \forall i, \qquad \Pr[x, y \text{ are in different clusters of } P_i] \leq \alpha \tfrac{d(x,y)}{2^i}.$$

Then the embedding has distortion $O(\alpha \log \Delta)$, i.e., $\mathbb{E}[d_T(x,y)] \leq O(\alpha \log \Delta) d(x,y)$.

Remark: This is often weaker than Theorem 2, and we will later show a stronger bound.

Proof: Was seen in class.