Advanced Algorithms - Lecture	January 26, 2005
Lecture 14	
Lecturer: Yair Bartal	*Scribe: Or Sattath*

1 Probabilistic Embedding of Metric Spaces

1.1 The Steinter Tree Problem and Metrical Embedding

This problem is a min problem.

Input: A graph G, and a weight function: $W : E \longrightarrow \mathbf{R}^+$, and a subset of the vertices, S.

Output: A "minimum weighted" spanning tree on S (the spanning tree can have vertices from all of G). The idea of "minimum weighted" is that the sum of the weights of all the edges in the spanning tree will be minimized.

This problem has an online version - All the graph is known in stage zero. In every stage we add a vertex to the wanted span tree, and the algorithm should add edges that will cause that vertex to be in the spanning tree.(once an edge is added, it can't be removed in a later stage). If the graph was a tree, then the solution is optimal (the approximation factor is 1). There is a solution that is $\theta(log(n))$ approximation.

Lemma 1 *Assume there is a metric space (X, d_X), for which there is an embedding by a bijective function (HAD HAD ERKI AND AL in Hebrew), f, into another metric space (Y, d_Y) , where (Y, d_Y) is a tree, and the distortion factor is α. (Embedding - SHIKUN in Hebrew, is just a function between one metric space to another metric space). So, if there is an algorithm for solving a problem on Y, A_Y, which is β approximation to the offline algorithm on Y, then X has an $\alpha\beta$ approximation.*

Proof: Assume w.l.o.g. that f isn't shrinking $d_X(u, v) \leq d_Y(f_Y(u), f_Y(v))$. For any order σ over X, f induces an order σ_Y over Y.
Run A_Y on σ_Y, and simulate it on X.

$$A(\sigma) \leq A_Y(\sigma_Y)$$

Assume there is an adversary ADV_X on X. We'll define a new adversary A_Y that simulates ADV_X on Y. We get:

$$ADV_Y(\sigma_Y) \leq \alpha * ADV_X(\sigma)$$

$$A(\sigma) \leq A_Y(\sigma_Y) \leq \beta * ADV_Y(\sigma_Y) \leq \alpha\beta * ADV_X(\sigma)$$

Figure 1:

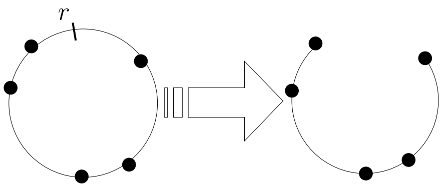

■

The problem with this idea is that the distortion factor is $\Omega(n)$ - for example the circle graph.

1.2 Probabilistic Embedding

Definition 2 (probabilistic embedding) *Given a metric space (X, d_X), and a set S of all the metic spaces of trees with size $|X|$, and for each $Y \in S$ there is a function $f_Y : X \to Y$ that is not shrinking. If there is a distribution over S s.t. $\forall u, v \in X, E[d_Y(f_Y(u), f_Y(v))] \leq \alpha d_X(u, v)$ Then it is called a probabilistic embedding with distortion factor α.*

Lemma 3 *if (X, d_X) has probabilistic embedding to S, and $\forall Y \in S$ exists an algorithm A_Y that is β competitive, then X has a random algroithm that is $\alpha\beta$- competitive.*

Proof: Very similar to lemma 1.

$$E[ADV_Y(\sigma_Y)] \leq \alpha * ADV_X(\sigma)$$

$$E[A(\sigma)] \leq E[A_Y(\sigma_Y)] \leq E[\beta * ADV_Y(\sigma_Y)] \leq \alpha\beta * ADV_X(\sigma)$$

■

Lemma 4 *a wighted circle graph over n verices has probabilistic embedding in trees (or more specifically in a line metric) with a distortion factor 2.*

Proof: Randomally choose a point from the circle r, where the length of the circle is the sum of all the weights. Create the line graph, by removing the edge where the point was picked up from. See figure 1 for visualization. Assuming that the length of the circle is 1, we get:

$$E[d_{T_{X(u,v)}}] = Pr[x \notin (u, v)] * [1 - d(u, v)] + Pr[x \in (u, v)] * d(u, v) \leq$$

$$Pr[x \notin (u, v)] * 1 + Pr[x \in (u, v)] * d(u, v) \leq 2d(u, v)$$

■

Theorem 5 (Bartal and FRT) *evey metric space over n points has probabilistic embedding in trees with distortion factor $O(log(n))$.*

Definition 6 (k-HST(hierarchial seperated tree)) *A tree, where for each vertex u, there is a label $\Delta(u) \in \mathbf{R}^+$. The mertic space is only on the leaves, and for every 2 leaves, u, v , $d(u,v) = \Delta(lca(u,v))$, and for every junction u, and it's descendant v $\Delta(v) \leq \frac{\Delta(u)}{k}$ and for every leaf u, $\Delta(u) = 0$.*

Remark: HST is a generalization of an ultrametric tree - (1-HST is an ultrametric tree).

The previous theorem applies for all k with $O(k * log_k(n))$.

Proof: We'll not do all the proof in this lecture. We'll need the following definition:

Definition 7 ((λ, Δ) probabilistic partition) *$M = (V, d)$ - a metric space. A partition with diameter Δ is a partition of V to clusters $C_1, C_2, ..., C_t$ s.t. $\bigcup_i C_i = V$, $\forall i, j\ i \neq j$, $C_i \cap C_j = \emptyset$,and $\forall i\ diam(C_i) \leq \Delta$.*
The diameter of a set is the maximum distance between two elements in the group: $diam(A) = max_{a,b \in A} d(a,b)$.
A (λ, Δ) probabilistic partition is a distribution over partitions with a diamater Δ , s.t $\forall u, v\ Pr[u, v\ isn't\ in\ the\ same\ cluster] \leq \frac{\lambda}{\Delta} d(u,v)$, where $\lambda \geq 1$.

Lemma 8 *Let M be a metric space s.t. $\forall \Delta\ \exists (\lambda, \Delta)$ probabilistic partition. Then M has a probabilistic embedding of an k-HST with a distortion factor of $O(\lambda k * log_k D)$, where $D = diam(M)$.*

■