600.469 / 600.669 Approximation Algorithms **Lecturer:** Michael Dinitz
Topic: Tree Embeddings **Date:** 03/26/15
Scribe: Alexander Schiffhauer

7.1 Background

Metric: A pair (V, d) such that $\forall\, u, v, w \in V$

1. $d_{uv} = 0 \iff u = v$
2. $d_{uv} = d_{vu}$
3. $d_{uv} \leq d_{uw} + d_{wv}$

Note that it is common to simply refer to the metric as d instead of the pair (V, d).

7.2 Tree Embeddings

Tree Metric: A pair (V', T) over some metric (V, d) such that

1. T is a tree on V' that has only non-negative edge lengths
2. $V' \supseteq V$

Similarly as above, it is common to simply refer to the tree metric as T instead of the pair (V', T).

The distance between any two vertices $u, v \in V'$ is denoted T_{uv}. And because T is a tree, the path from u to v is unique, which implies T_{uv} is the distance of the shortest u-v path in T.

Distortion: A value α such that $d_{uv} \leq T_{uv} \leq \alpha \cdot d_{uv} \;\forall\, u, v \in V$.

Tree Embedding: A tree metric (V', T) that approximates a metric (V, d) with distortion α.

Note that it is common to say a metric (V, d) embeds into (V', T) with distortion α.

7.3 FRT

Fakcharoenphol, Rao and Talwar Algorithm (FRT): A randomized, polynomial-time algorithm that embeds (V, d) into (V', T) such that $\forall\, u, v \in V$

1. $d_{uv} \leq T_{uv}$ (proved below)
2. $\mathrm{E}[T_{uv}] \leq \alpha \cdot d_{uv}$, where $\alpha = \mathcal{O}(\log n)$ (proved next class)

The only constraint on the original metric is that $\forall\, u, v \in V$, $d_{uv} \geq 1$.

Note that there exist metrics such that their embedding into any probabilistic-based tree metric must have distortion $\Omega(\log n)$. So, for some metrics, we know this bound is asymptotically tight.

7.4 Hierarchical Cut Decomposition

Hierarchical Cut Decomposition: A tree embedding (V', T) over some metric (V, d) such that T is a rooted tree with $\log \Delta + 1$ levels (the root node is at level $\log \Delta$ and leaf nodes at level 0), where Δ is the smallest power of 2 greater than $2 \cdot \max_{u,v \in V} d_{uv}$, and $V' \supseteq V$ such that

1. The representative of T's root node is V.
2. There is a bijection between the representatives of T's leaf nodes and the vertices of V.
3. The representative of a node in T at level i is a subset S of V such that its vertices are enclosed by a ball (centered on some node in S) with radius r, where $2^{i-1} \leq r < 2^i$.
4. The representatives of a node's children in T create a partition of their parent's representative.
5. the length of an edge between a level i node and a level $i+1$ node is 2^{i+1}.

Consider the hierarchical cut decomposition (V', T) of some metric (V, d).

Lemma: *If the least common ancestor of two leaf nodes u and v in T is at level i then $T_{uv} \leq 2^{i+2}$. Furthermore, $T_{uv} \geq d_{uv} \; \forall \, u, v \in V$*

Proof: Let u and v be leaf nodes in T, and let w be u and v's least common ancestor such that it is at level i. Then by construction $2^i \leq T_{uw} = \sum_{i=1}^{i} 2^i \leq 2^{i+1}$ and similarly $2^i \leq T_{wv} \leq 2^{i+1}$, and hence $2^{i+1} \leq T_{uv} \leq 2^{i+2}$. Also, since w is at level i, by the definition of a hierarchical cut decomposition u and v are both in a ball of radius 2^i, so by the triangle inequality $d_{uv} \leq 2^{i+1}$ and so $T_{uv} > d_{uv}$.

7.5 Constructing a Hierarchical Cut Decomposition

Constructing a Hierarchical Cut Decomposition

```
Let π be a permutation of V, chosen uniformly at random
Let r_0 be a value in [1/2, 1), chosen uniformly at random
Let r_i = r_0 · 2^i for all i such that 1 ≤ i ≤ log Δ
Let T be a tree with only a root node (at level log Δ) whose representative is V
for i ← log Δ to 1 do
    Let 𝒞 be the set of nodes at level i
    for C ∈ 𝒞 do
        S ← C
        for j ← 1 to n do
            P ← the nodes in S enclosed by a circle centered on π(j) with radius r_{i−1}
            if P ≠ ∅ then
                S ← S \ P
                Add P to T as a child of C at level i − 1
return T
```

We will analyze algorithm (due to FRT) in the next class. For now, let's see an example of why tree embeddings are useful.

7.6 Group Steiner Tree

Recall the Group Steiner Tree problem:

Input:, A graph $G = (V, E)$, edge costs $c_e \geq 0 \,\forall\, e \in E$, a root vertex $r \in V$, and Steiner groups $g_1, \ldots, g_k$ such that each group is a subset of V.

Feasibles: A tree T such that $\forall\, i \in [k]$, $\exists\, v \in g_i$ such that T has a path between r and v.

Objective: $\min \sum\limits_{e \in T} c_e$

It is NP-hard to approximate the Group Steiner Tree problem to a factor better than $\Omega(\log n)$ for all graphs (see last class). However, Garg, Konjevod, and Ravi devised an algorithm (GKR) that approximates the Group Steiner Tree problem to a factor of $\Omega(\log n \cdot \log k)$ on trees.

Theorem: Using FRT to embed the input for the Group Steiner Tree problem into a tree metric and using the GKR algorithm over this embedding yields an expected $\mathcal{O}(\log^2 n \cdot \log k)$ approximation to the Group Steiner Tree problem.

Proof: Without loss of generality, assume the input to the Group Steiner Tree problem is metric distances, d (this is the metric completion of the edge costs and the input graph). Using FRT over this input yields a random tree T. Using GKR over T yields a $\mathcal{O}(\log n \cdot \log k)$ approximation T' for T. Shortcutting T' yields a cycle C on the group nodes.

In the following series of inequalities, let X_e denote the cost of an edge (or set of edges) over a graph (or tree) X, let S be the group nodes connected by OPT, let C_S be the cycle on S achieved from shortcutting S, and let T_S be the subtree of T over S.

$$
\begin{aligned}
\mathrm{E}[C] &\leq \mathrm{E}[T_C] && \text{distances in } T \text{ are nondecreasing} \\
&\leq \mathrm{E}[2 \cdot T_{T'}] && \text{shortcutting costs at most a factor of 2} \\
&\leq 2 \cdot \mathrm{E}[T_{T'}] && \text{linearity of expectation} \\
&\leq 2\mathrm{E}[\mathcal{O}(\log n \cdot \log k) \cdot \mathrm{OPT}(T)] && \text{GKR} \\
&\leq 2\mathcal{O}(\log n \cdot \log k) \cdot \mathrm{E}[\mathrm{OPT}(T)] && \text{linearity of expectations} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \mathrm{E}[\mathrm{OPT}(T)] && \text{asymptotic notation} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \mathrm{E}[T_{T_S}] && \text{by construction} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \mathrm{E}[T_{C_S}] && \text{by construction} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \mathrm{E}\Big[\sum\nolimits_{(u,v) \in C_S} T_{uv}\Big] && \text{by construction} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \sum\nolimits_{(u,v) \in C_S} \mathrm{E}\big[T_{uv}\big] && \text{linearity of expectations} \\
&\leq \mathcal{O}(\log n \cdot \log k) \cdot \sum\nolimits_{(u,v) \in C_S} \mathrm{E}\big[\log n \cdot d_{uv}\big] && \text{FRT}
\end{aligned}
$$

$$
\begin{aligned}
&\leq \mathcal{O}(\log^2 n \cdot \log k) \cdot \sum\nolimits_{(u,v)\in C_S} d_{uv} && \text{linearity of expectations} \\
&\leq \mathcal{O}(\log^2 n \cdot \log k) \cdot 2\mathrm{OPT} && \text{shortcutting} \\
\mathrm{E}[C] &\leq \mathcal{O}(\log^2 n \cdot \log k) \cdot \mathrm{OPT} && \text{asymptotic notation}
\end{aligned}
$$

7.7 Metric Embeddings

Metric Embedding: A metric (V, d) that embeds into another metric (V, d') with distortion α such that, $\forall\, u, v \in V, d_{uv} \leq d'_{uv} \leq \alpha \cdot d_{uv}$.

Given a metric (V, d) with a known β-approximation ALG over some metric space d', embedding (V, d) into (V, d') and solving with ALG yields an $\alpha \cdot \beta$ approximation for the original problem.