Uniform Buy-at-Bulk Network Design[1]

R. Inkulu
http://www.iitg.ac.in/rinkulu/

[1] slides last updated in 2013

Outline

1. Buy-at-Bulk Network Design
2. Tree Metrics
3. Conclusions

Description

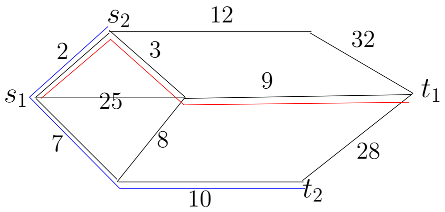

$d_1 = 300 \ d_2 = 100$

$C(u)$ per unit length

$C(300) + C(100) \geq C(400)$

$C(100) \leq C(300) \leq C(400)$

Given an undirected graph $G(V, E)$ with the following:

length $L : E \rightarrow R^+$,

k source-sink pairs; each associated with a demand $d_i > 0$.

capacity u can be purchased at $C(u)$ per unit length on any edge, and

function $C : R^+ \rightarrow R^+$ is both *non-decreasing* and *subadditive*

Find a minimum cost multicommodity unsplittable flow c

(i.e., $\min \sum_{(u,v)\in E} C(c_{uv}) l_{uv}$)

that meets the demands of all the source-sink pairs.

Lower Bounds

- NP-hard problem —no proof provided
- No $O(\lg^{\frac{1}{2}-\epsilon} n)$ apprx algo possible unless $NP \subseteq ZPTIME(n^{polylog(n)})$ —not proof provided

Herewith, we design a randomized apprx algo with $O(\lg n)$ apprx factor in expectation.

Optimal algo when G is a tree

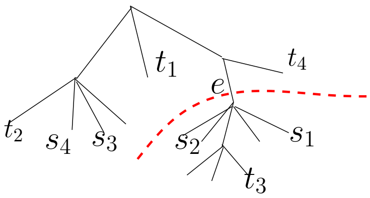

1. for each s_i-t_i
   (a) for each edge $e(u, v) \in dist_G(u, v)$, mark with e that additional d_i units of flow required to be pushed along e
2. for each edge e in G that is marked, purchase the corresp. capacity at bulk

in any optimal solution, capacity c_e for any edge e must be at least equal to the summation of demands that e separates

Algorithm when G is not a tree

$$
\begin{tikzcd}
G \arrow[r, "?"'] & T' \\
G_{sol} & T'_{opt} \arrow[l, "?"']
\end{tikzcd}
$$

our algo

Algorithm when G is not a tree (cont)

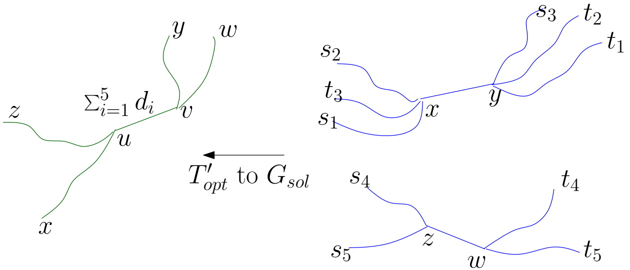

1. probabilistically approximate metric completion of $G(V, E)$ (w.r.t. L) by a spanning tree metric $T'(V, E')$

   i.e., T' such that for any $u, v \in V$,

   $dist_G(u, v) \leq T'(u, v) \leq O(\lg n) dist_G(u, v)$ (latter in expectation)

Algorithm when G is not a tree (cont)

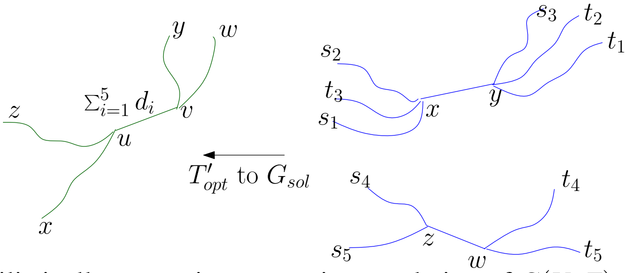

1. probabilistically approximate metric completion of $G(V, E)$ (w.r.t. L) by a spanning tree metric $T'(V, E')$

   i.e., T' such that for any $u, v \in V$,

   $dist_G(u, v) \leq T'(u, v) \leq O(\lg n) dist_G(u, v)$ (latter in expectation)
2. solve the problem on T'

Algorithm when G is not a tree (cont)

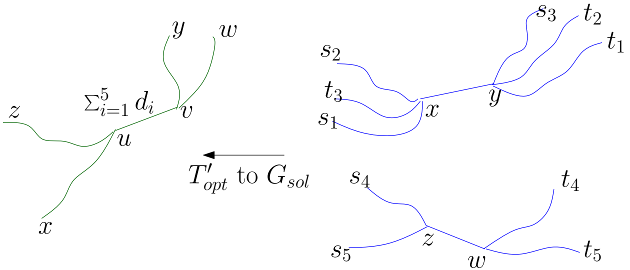

1. probabilistically approximate metric completion of $G(V, E)$ (w.r.t. L) by a spanning tree metric $T'(V, E')$

   i.e., T' such that for any $u, v \in V$,

   $dist_G(u, v) \leq T'(u, v) \leq O(\lg n) dist_G(u, v)$ (latter in expectation)
2. solve the problem on T'
3. for each s_i-t_i
   - (a) map each edge in $(x, y) \in T'(s_i, t_i)$ to a shortest path P between x and y in G
   - (b) for each edge $e \in P$, mark with e that additional d_i units of flow required to be pushed along e

Algorithm when G is not a tree (cont)

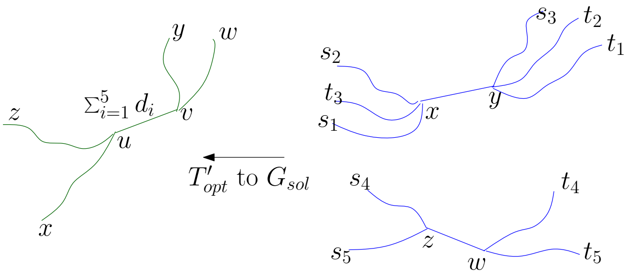

1. probabilistically approximate metric completion of $G(V,E)$ (w.r.t. L) by a spanning tree metric $T'(V,E')$

   i.e., T' such that for any $u, v \in V$,

   $dist_G(u,v) \leq T'(u,v) \leq O(\lg n) dist_G(u,v)$ (latter in expectation)

2. solve the problem on T'
3. for each s_i-t_i
   a. map each edge in $(x,y) \in T'(s_i,t_i)$ to a shortest path P between x and y in G
   b. for each edge $e \in P$, mark with e that additional d_i units of flow required to be pushed along e
4. for each edge e in G that was marked, purchase capacity at bulk

Overview of the Analysis

$$\begin{tikzcd} G \arrow[r] & T' \\ G_{sol} & T'_{opt} \arrow[l] \end{tikzcd}$$

our algo

$$\begin{tikzcd} G_{opt} \arrow[d] \\ T'_{sol} \end{tikzcd}$$

$$cost(our\ algo) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt}) \text{ in exp}$$

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n)cost(G_{opt})$ in exp

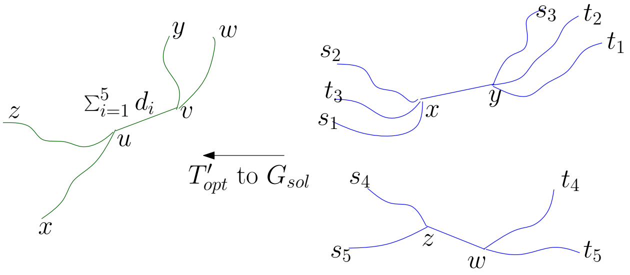

- T' is chosen s.t. for any $u, v \in V$ in G, $dist_G(u, v) \leq T'(u, v)$
- SPs corresp. to two edges in T' may pass through the same edge in G: *subadditive C*

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt})$ in exp

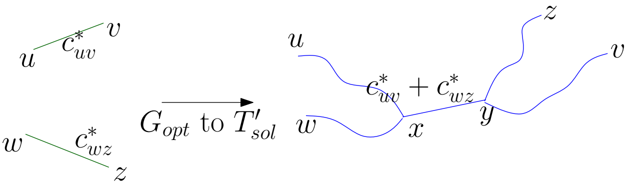

procedure to obtain T'_{sol} from G_{opt}:

- for each edge $e(u, v) \in G$ with opt capacity c^*_e, mark for all edges along the unique path between u and v that additional c^*_e units of capacity required to be pushed along e
- buy-at-bulk

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt})$ in exp (cont)

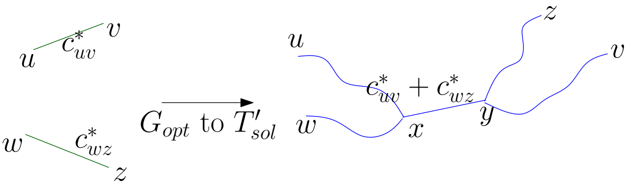

- for every edge xy in T', if the capacity xy uses in T'_{opt} is c'_{xy}, then the capacity xy uses in T'_{sol} is at least c'_{xy} — this is due to xy in T'_{opt} is separating the same demand as xy in T'_{sol}
- *non-decreasing C*

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt})$ in exp: express capacities in T'_{sol} in terms of G_{opt}

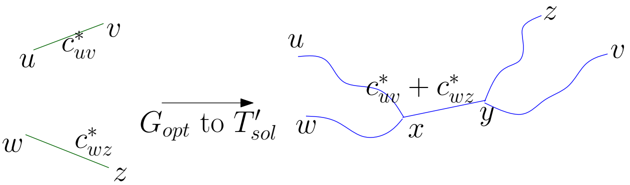

$cost_{T'_{sol}}$

$$= \sum_{(x,y)\in T'} T'(x,y) C\Big(\sum_{(u,v)\in E:(x,y)\in u \sim v \text{ in } T'} c^*_{uv}\Big)$$

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n)cost(G_{opt})$ **in**

exp: express capacities in T'_{sol} in terms of G_{opt}

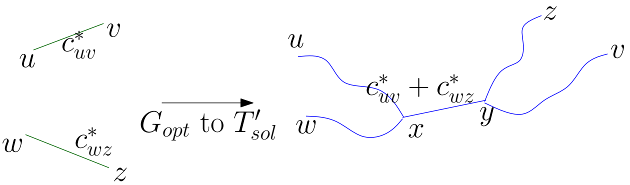

$cost_{T'_{sol}}$

$$= \sum_{(x,y)\in T'} T'(x,y) C\left(\sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} c^*_{uv}\right)$$

$$\leq \sum_{(x,y)\in T'} T'(x,y) \sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} C(c^*_{uv}) \leftarrow \text{subadditivity of } C$$

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt})$ in exp: express capacities in T'_{sol} in terms of G_{opt}

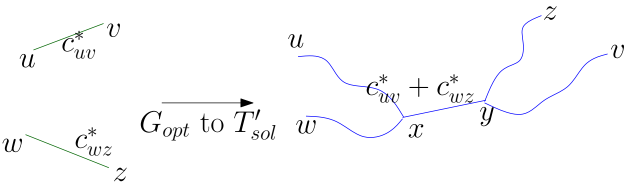

$cost_{T'_{sol}}$

$$= \sum_{(x,y)\in T'} T'(x,y) C\left(\sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} c^*_{uv}\right)$$

$$\leq \sum_{(x,y)\in T'} T'(x,y) \sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} C(c^*_{uv}) \leftarrow \text{subadditivity of } C$$

$$= \sum_{(u,v)\in E} C(c^*_{uv}) \sum_{(x,y)\in u\sim v \text{ in } T'} T'(x,y)$$

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n) cost(G_{opt})$ in exp: express capacities in T'_{sol} in terms of G_{opt}

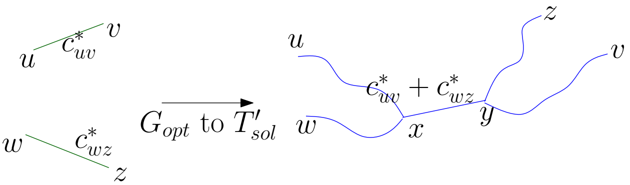

$cost_{T'_{sol}}$

$$= \sum_{(x,y)\in T'} T'(x,y) C\left(\sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} c^*_{uv}\right)$$

$$\leq \sum_{(x,y)\in T'} T'(x,y) \sum_{(u,v)\in E:(x,y)\in u\sim v \text{ in } T'} C(c^*_{uv}) \leftarrow \text{subadditivity of } C$$

$$= \sum_{(u,v)\in E} C(c^*_{uv}) \sum_{(x,y)\in u\sim v \text{ in } T'} T'(x,y)$$

$$= \sum_{(u,v)\in E} C(c^*_{uv}) T'(u,v)$$

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n)cost(G_{opt})$ **in exp: express T' distances in terms of distances in G**

$\sum_{(u,v)\in E} C(c^*_{uv})T'(u,v)$

$\leq O(\lg n)\sum_{(u,v)\in E} C(c^*_{uv})dist_G(u,v)$ in expectation

(as T' is chosen s.t. for any $u, v \in V$ in G, $T'(u,v) \leq O(\lg n)dist_G(u,v)$ in expectation)

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n)cost(G_{opt})$ in exp: express T' distances in terms of distances in G

$\sum_{(u,v)\in E} C(c^*_{uv})T'(u,v)$

$\leq O(\lg n) \sum_{(u,v)\in E} C(c^*_{uv}) dist_G(u,v)$ in expectation

(as T' is chosen s.t. for any $u, v \in V$ in G, $T'(u,v) \leq O(\lg n) dist_G(u,v)$ in expectation)

$\leq O(\lg n) \sum_{(u,v)\in E} C(c^*_{uv}) l_{uv}$ in expectation

$cost(G_{sol}) \leq cost(T'_{opt}) \leq cost(T'_{sol}) \leq O(\lg n)cost(G_{opt})$ in exp: express T' distances in terms of distances in G

$\sum_{(u,v)\in E} C(c^*_{uv})T'(u,v)$

$\leq O(\lg n) \sum_{(u,v)\in E} C(c^*_{uv})dist_G(u,v)$ in expectation

(as T' is chosen s.t. for any $u, v \in V$ in G, $T'(u,v) \leq O(\lg n)dist_G(u,v)$ in expectation)

$\leq O(\lg n) \sum_{(u,v)\in E} C(c^*_{uv})l_{uv}$ in expectation

$= O(\lg n)cost_{G_{opt}}$ in expectation

Apprx factor in expectation

suggested randomized algorithm outputs a solution with apprx factor $O(\lg n)$ in expectation:

$$
\begin{aligned}
&cost_{\text{our algo}} \\
&\qquad \le cost(T'_{opt}) \\
&\qquad \le cost(T'_{sol}) \\
&\qquad = O(\lg n) cost(G_{opt}) \text{ in expectation}
\end{aligned}
$$

Done except for devising a randomized algo to construct a T' wherein for any $u, v \in V$, $dist_G(u, v) \le T'(u, v) \le O(\lg n) dist_G(u, v)$ (latter in expectation)

Outline

Metric Space

A *metric space* is a pair (X, D) where X is a set and $D : X \times X \to [0, \infty)$ is a metric satisfying:

- $D(x, y) \geq 0$
- $D(x, y) = 0$ iff $x = y$
- $D(x, y) = D(y, x)$
- $D(x, y) + D(y, z) \geq D(x, z)$ (triangle inequality)

ex. $(\mathcal{R}^d, L_2^d)$

A metric space (X, D) is a *finite metric space* if $|X|$ is finite.

Metric completion of a graph (a.k.a. graph metric)

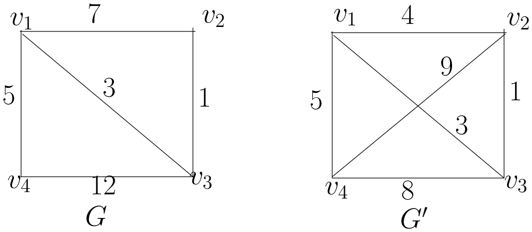

- All pair shortest distance graph G' corresp. to input graph G is a finite metric

Embedding

Let (X, D') and (Y, D'') be two (finite) metric spaces. Any one-to-one map $f : X \to Y$ is termed as an *embedding*.

- An embedding in which no distances shrink is termed as an *expansive embedding*.
- distortion of an expansive embedding f is $\max_{x,y \in X} \frac{D''(f(x), f(y))}{D'(x,y)}$

We intend to construct an expansive ($dist_G(u, v) \leq T(u, v)$) *tree metric* $(V \subseteq V', T)$ corresp. to the graph metric $(V, dist_G)$ such that $T(u, v) \leq O(\lg n) dist_G(u, v)$ in expectation. Further, we enforce $V = V'$ by building a *spanning tree metric* (V, T').

Hierarchical cut decomposition of the metric $(V, dist_G)$

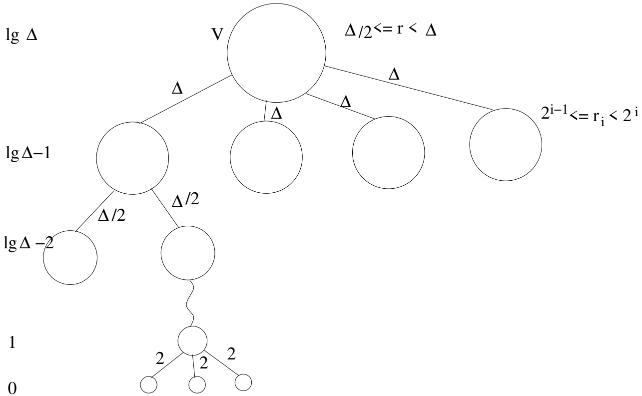

here, Δ is the smallest power of 2 greater than 2 $max_{u,v \in V} dist_G(u, v)$

- root has the entire V
- each leaf node corresp. to a unique point in V for convenience, let $dist_G(u, v) \geq 1$
- nodes in each level together partition V

let us refer vertices of T as nodes while the vertices of V as points

Hierarchical cut decomposition is a tree metric

- $V \subseteq V'$
- positive edge lengths
- (V', T) is an expansive metric

  lowest level at which u and v belong to the same is $\lfloor \lg_2 dist_G(u, v) \rfloor$
- what about the distortion?

Randomized Algorithm to construct (V', T)

1. pick a permutation π of V
2. pick a random number r_0 in $[1/2, 1)$; set radius $r_i = 2^i r_0$ for all balls at each level i
3. root is associated with points in ball $B(\text{any point}, \Delta)$ i.e., V itself
4. for each node v in each level i ($i > 0$)

   let S be the set of points associated with v

   (a) for every j from 1 to n

      (i) if $S' = B(\pi(j), r_{i-1}) \cap S \neq \phi$ then create a child node to v and associate points in S' to it

      (ii) $S = S - S'$

- takes polynomial time

Algorithm in execution

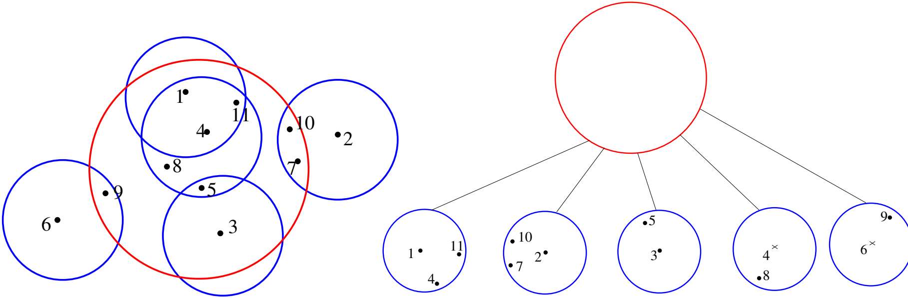

points belonging to a tree node are
shown with filled circles

Bounding the expected distortion

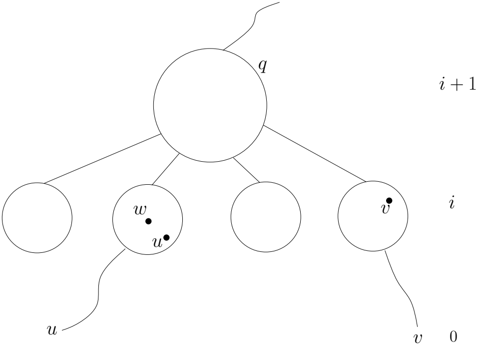

- If LCA of u and v is at level i, then $T(u,v) \leq 2^{i+2}$.
- $E[T(u,v)]$

  $= \sum_{w \in V} \sum_{i=0}^{\lg \Delta - 1}$ (prob. $u \in B(w, r_i)$, $v \notin B(w, r_i)$ and $B(w, r_i)$ is a child of q)$*(2^{i+3}) \leq O(\lg n) dist_G(u, v)$

  w.l.o.g. suppose u is nearer to w than v

  — only intuition behind the proof is given

Transforming Tree Metric (V', T) to a Spanning Tree Metric (V, T')

(1) repeat until there does not exist a vertex pair u, w such that $u \in V, w \notin V$ and w is the parent of u
   (a) contract edge uw
   (b) identify merged node with $u (\in V)$
(2) multiply the length of every remaining edge by four

— intutively explained why it won't change the bounds on $T'(u, v)$

Outline

Example specializations of uBatB

- *Steiner tree problem*: Given a non-negative edge-weighted connected undirected graph $G(V,E)$ together with a set $S \subseteq V$, find a minimum cost tree in G that spans all vertices in S and any subset of *Steiner* vertices $V - S$.

- *Generalized Steiner forest problem*: Given an undirected graph $G(V,E)$, $\forall_{e \in E}\, w_e \geq 0$, and k pairs of vertices $s_i, t_i \in V$, find a minimum-cost subset of edges $F \subseteq E$ such that every s_i-t_i pair is connected in (V,F).

Open problems around uBatB

- reducing the gap between lower and upper bounds
- algorithm with apprx factor $O(\lg k)$
- devising algorithms for special graphs

Beyond uBatB

- Non-uniform BatB

Other example applications of tree embeddings

- *Group Steiner tree*: Given an undirected graph $G(V, E)$ with $c : E \to \mathcal{R}^+$, and groups of vertices $V_1, V_2, \ldots, V_k \subseteq V$ find a minimum cost subtree of G that contains at least one vertex from each group.

- *Communication spanning trees*: Given an undirected graph $G(V, E)$ with nonnegative costs on edges, requirement value r_{ij} for every pair i and j, and the communication cost of a spanning tree T is defined as $\sum_{ij} r_{ij} * SP_{T(i,j)}$, find a spanning tree T of minimum communication cost.

And, several networking heuristics that use tree metrics are awaiting Algorithms.

Other popular metric embeddings

Let $|X| = n$, D' is an arbitrary metric, d denotes the number of dimensions, and X' is finite. Then,

- Bourgain's Theorem:

  existence of $(X, D') \hookrightarrow^{O(\lg n)} (\mathcal{R}^{O(\lg^2 n)}, L_p)$

- dimension reduction due to the Johnson-Lindenstrauss Lemma:

  existence of $(X, L_2^d) \hookrightarrow^{(1+\epsilon)} (\mathcal{R}^{O(\epsilon^{-2} \lg n)}, L_2)$

- Feige's volume respecting embeddings:

  $Vol(X) = sup_{f:X \to l_2} Evol(f(X))$ (f requires to be a contraction)

  k-distortion of f is $sup_{P \subset X, |P|=k} (\frac{Vol(P)}{Evol(f(P))})^{\frac{1}{k-1}}$

Significant theoretical ideas explored

- extended a trivial solution for trees to graphs
- interpreted a graph theoretic problem as a geometric problem
- embedded a finite metric space into a tree metric
- randomized + approximation algortihm

  hence, the apprx factor in expectation
- designed an algorithm for a network design problem

  uBatB is a generalization of several network design problems

References

- B. Awerbuch and Y. Azar. Buy-at-bulk network design. FOCS 1997, pg 542-547.
- J. Fakcharoenphol, S. Rao and K. Talwar A tight bound on approximating arbitrary metrics by tree metrics. Journal of Computer and System Sciences, 69:485-497, 2004.
- Y. Bartal. Probabilistic Approximation of Metric Spaces and its Algorithmic Applications. FOCS 1996, 184-193.
- P. Indyk. Algorithms applications of low-distortion geometric embeddings. FOCS 2001.
- David P. Williamson and David B. Shmoys. The Design of Approximation Algorithms. Cambridge University Press, 2011.
- Jiri Matousek. Lectures on Discrete Geometry. Springer, 2002.