Distributed Spanner Approximation

Michal Dory, Technion

Joint work with: Keren Censor-Hillel, Technion

Spanners

A ***k*-spanner** of a graph G is a subgraph of G that preserves distances up to a multiplicative factor of k.

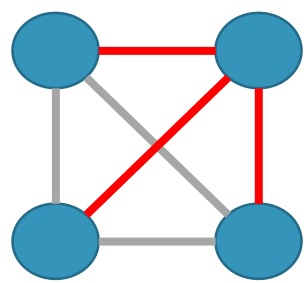

Spanners

- There are many constructions which give a global guarantee on the size of the spanner:

$$(2k-1)\text{-spanners with } O(n^{1+1/k}) \text{ edges}$$

- This is optimal in the worst case assuming Erdős's girth conjecture.

Spanner Approximation

- What about approximating the minimum k-spanner?
- There are graphs where any 2-spanner has $\Omega(n^2)$ edges, this is also true for k-spanners in directed graphs.

In the sequential setting:

- 2-spanner: $O\left(\log\frac{|E|}{|V|}\right)$-approximation [Kortsarz and Peleg 1994]
- Directed k-spanner: $O(\sqrt{n}\ \log n)$-approximation [Berman et al. 2013]

<u>Hardness Results:</u>

- 2-spanner: $\Omega(\log n)$ [Kortsarz 2001]
- Directed k-spanner: $\Omega(2^{(\log^{1-\varepsilon} n)})$[Elkin and Peleg 2007]
- Undirected k-spanner: $\Omega(2^{(\log^{1-\varepsilon} n)/k})$ [Dinitz, Kortsarz and Raz 2016]

The Distributed Models

Vertices exchange messages in **synchronous** rounds

The model	Message size
LOCAL	unbounded
CONGEST	$\theta(\log n)$ bits

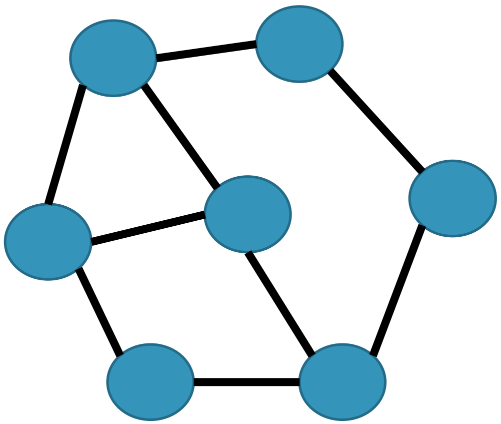

In the LOCAL model

Directed k-spanners:

Approximation	Number of rounds	
$O(\sqrt{n}\ \log n)$	$O(k \log n)$	[Dinitz and Nazari, 2017]
$O(n^{\epsilon})$	constant	[Barenboim, Elkin and Gavoille, 2016]
$(1+\epsilon)$	$O(poly(\log n/\epsilon))$	Our Results

Spanner Approximation

- Can we give efficient approximations also in the **CONGEST** model?

Spanner Approximation

- Can we give efficient approximations also in the **CONGEST** model?

Approximating k-spanners in **directed** or **weighted** graphs is hard in the CONGEST model.

Spanner Approximation

- Can we give efficient approximations also in the **CONGEST** model?

Approximating k-spanners in **directed** or **weighted** graphs is hard in the CONGEST model.

This gives a strict **separation** between the **LOCAL** and **CONGEST** models.

Our Results

Directed k-spanner for $k \geq 5$:

- Randomized algorithms - $\widetilde{\Omega}(\sqrt{n/\alpha})$ rounds for an α-approximation.
- Deterministic algorithms - $\widetilde{\Omega}(n/\sqrt{\alpha})$

Weighted k-spanner for $k \geq 4$:

- Directed graphs - $\widetilde{\Omega}(n)$
- Undirected graphs - $\widetilde{\Omega}(n/k)$