Distributed Spanner Approximation

Michal Dory, with: Keren Censor-Hillel (PODC 2018)

The Distributed Models

- Vertices exchange messages in **synchronous** rounds

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

k-spanners

- Sparse subgraphs that **preserve distances** up to a multiplicative factor of k

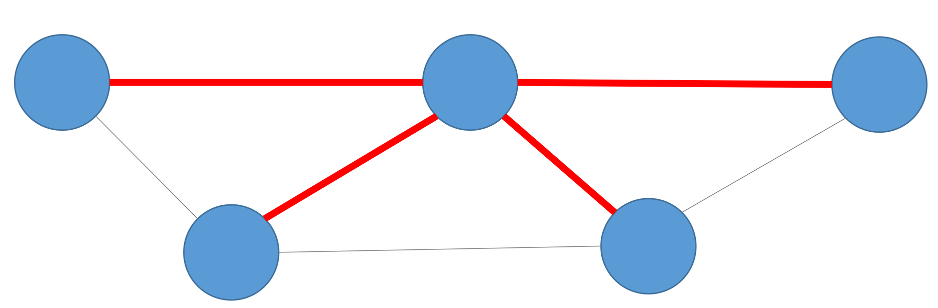

- Many **global** constructions:
$(2k-1)$-spanners: $O(n^{1+1/k})$ edges

An α-approximation

Number of edges $\leq \alpha \cdot OPT$

Spanner approximation

- Many **sequential** algorithms and hardness results
- In **LOCAL**: $(1+\epsilon)$-approximation in $O(poly(\log n/\epsilon))$ rounds

Can we give efficient CONGEST algorithms?
No! at least in directed or weighted graphs.

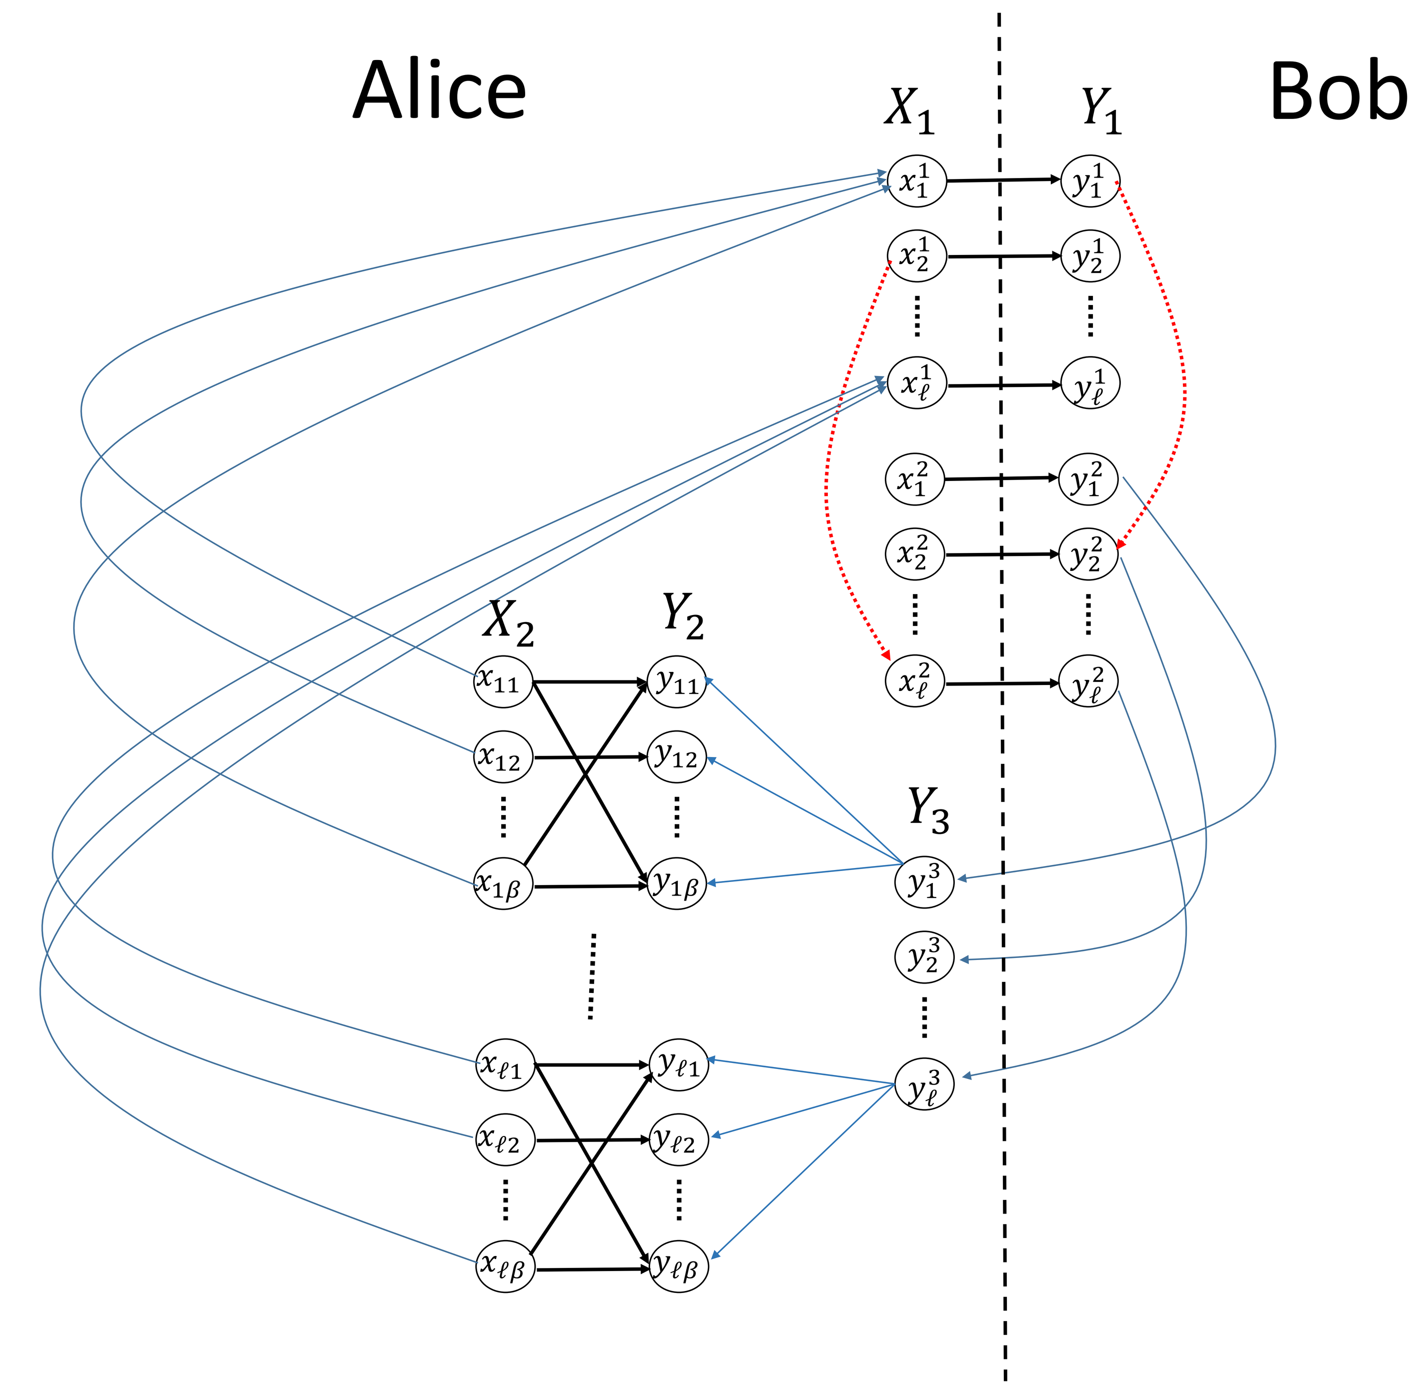

Directed k-spanner for $k \geq 5$:

Randomized - $\widetilde{\Omega}(\sqrt{n/\alpha})$ rounds, Deterministic - $\widetilde{\Omega}(n/\sqrt{\alpha})$

$\Rightarrow$ A strict **separation** between the **LOCAL** and **CONGEST** models

Computer Science Department, Technion - Israel Institute of Technology