CS 59000 CTT Current Topics in Theoretical CS Sep 18, 2012

Lecture 9

Lecturer: Elena Grigorescu *Scribe: Bhaskar Sarma*

1 Introduction

Definition 1 *A k-spanner of a graph G is a spanning subgraph H in which any two vertices are at most k times as far apart in H as they are in G, i.e. for any two vertices u, v we have $d_H(u, v) \leq k \cdot d_G(u, v)$.*

Spanners were introduced by Awerbuch [1] and Peleg and Schäffer [3] in the networks and distributed computing communities.

Definition 2 *Given a directed acyclic graph G, the* transitive closure *of $G = G(V, E)$ is the directed acyclic graph $G'(V, E')$, such that if there exists a path between vertices u and v in G, then there exists an edge between u and v in G'.*

Definition 3 *A k-*transitive closure spanner (k-TC spanner) of graph G, is defined as a k-spanner on the transitive closure of graph G.

Transitive-closure spanners have been used implicitly and explicitly in the CS literature for the past 2 decades, but they were first explicitly formalized and systematically studied in [2]. They have important applications to property testing (in testing monotonicity of functions), local function reconstruction, in building efficient access hierarchy schemes, data structures, etc. For a comprehensive survey, see [4].

We have seen in the last lecture that there exists a 2-TC spanner of the directed line graph on n vertices (where all edges are directed from vertex i to vertex $i+1$), with $O(n \log n)$ edges. The line graph is just a simple example of a graph for which sparse spanners can be constructed.

In this lecture we address the following question - How many edges does one need to add to the line graph L_n to obtain a k-TC spanner, when $k \geq 3$? One can ask the same question for more general families of graphs. It turns out that the techniques useful for line graph TC spanner constructions serve as great starting points for building sparse TC spanners for many other general families, such as directed tree graphs, planar graphs, or graphs with small path separators.

We start of by looking at 3-TC spanners and then present a general solution for k-TC spanners for the directed line graph L_n.

2 3 - TC spanners construction for L_n

Theorem 4 *There exists a 3 - TC spanner with $O(n \log(\log n))$ edges for the directed line L_n.*

Proof We prove this theorem by presenting a way to construct a sparse 3 - TC spanner. As in the construction for 2-TC spanners, we will use recursion, and designate some equally spaced vertices from L_n as *hubs*. We will connect the hubs with all the vertices in the line segments determined by consecutive hubs.

The following are the steps in creating the 3-TC spanner:

1. Designate $h+1$ equally spaced vertices from L_n as the hubs (starting from the first vertex). (For simplicity of presentation but without loss of generality we will assume that the hubs partition the line into n/h segments evenly). Here the number of hubs will be $h = \sqrt{n}$ (again, for simplicity, we will assume this is an integer).
2. For all non-hub vertices v in G, we add directed edges from and to its closest hubs (in the correct direction, i.e. only edges from the transitive closure of L_n).
3. We then connect all hubs in a complete graph (again, where the edges added come from the transitive closure).
4. For each of the h line segments between consecutive hubs we recursively partition them via the same procedure as in the above three steps.

This is a 3-TC spanner because any 2 vertices are connected by a path of length at most 3. To see why, consider two vertices u and v. Suppose $u \in [\text{hub}_i, \text{hub}_{i+1}]$ and $v \in [\text{hub}_j, \text{hub}_{j+1}]$, and $i < j$. Then by construction, $\exists$ a path: $u \to \text{hub}_{i+1} \to \text{hub}_j \to v$, which is of length 3.

If (u, v) both belong to $[\text{hub}_i, \text{hub}_{i+1}]$, then $\exists$ some stage of the recursion, where they belong to different segments and then there is a path of length ≤ 3 by the above proof.

Now we show that the 3 - TC spanner constructed above has $O(\log(\log n))$ edges.

Let $T(n)$ be the number of edges added for L_n. Then $T(n)$ is composed of the following

1. $\sqrt{n} \cdot T(\sqrt{n})$ edges accounting for the edges added for the $\sqrt{n}$ segments of length $\sqrt{n}$ during the recursion step.
2. $2n$ edges that come from connecting each vertex with the closest left and right hubs.
3. $\binom{\sqrt{n}}{2}$ edges between hubs

Therefore, we obtain the following recursion $T(n) = \sqrt{n}T(\sqrt{n}) + 2n + \binom{\sqrt{n}}{2}$ which gives a solution $T(n) = O(n \log(\log n))$.

■

3 k - TC spanner for L_n

In this section we present a proof of the number of edges needed to create a k-TC spanner for a given directed line graph L_n.

Theorem 5 *$\exists$ a k-TC spanner for L_n with $kn\lambda_k(n)$ edges. Here $k \cdot n \cdot \lambda_k(n)$ is the kth- row inverse Ackerman function (defined below).*

Proof As in the proof for a 3-TC spanner we present a construction for a k-TC spanner and then show that the number of edges added by this construction is in accordance to our claim.

Construction:

1. We designate again $h + 1$ equally spaced vertices of L_n as hubs, and partition the graph into h segments. So the size of each segment is n/h

2. As before, we add directed edges for every vertex v of L_n from and to its nearest hubs.
3. We put a $k-2$ TC spanner on the hubs.
4. We recursively construct k-TC spanners on each of the h segments.

We now show that this is a k-TC spanner, namely for any two vertices u, v we have a path of length at most k connecting u and v. Suppose $u \in [\text{hub}_i, \text{hub}_{i+1}]$ and $v \in [\text{hub}_j, \text{hub}_{j+1}]$, and $i < j$. Then, by construction, $\exists$ a path as follows: u is adjacent to hub_{i+1}, hub_{i+1} is connected by a path of length $k-2$ with $\rightarrow \text{hub}_j$ (since the hubs belong to a $k-2$ TC spanner) and hub_j is adjacent to vertex v. Clearly this gives a path of length $\leq k$ between u and v.

If u and v belong to the same segment (i.e. $i = j$), then $\exists$ a level of recursion such that they belong to different segments and thus by the above proof they are at a distance of at most k.

As before we can compute the number of edges added in this recursion as by setting up a recursion. Let $T_k(n)$ be the number of edges in the resulting k-TC spanner for L_n. Then $T_k(n) = h \cdot T(n/h) + 2n + T_{k-2}(h)$.

In the expression above the first term comes from the top level recursion, the second from connecting each vertex to its nearest left/right hubs and the last term comes from placing a $(k-2)$ spanner on the hubs.

We now look at what the value of h should be and for that we need to define the inverse Ackerman function. We will then finish the proof by using induction on k.

Inverse Ackerman function

Definition 6 *Given, $f : R^+ \rightarrow R^+$ such that $f(n) < n$, we define $f^*(n)$ to be the* $\min k \in Z^+$ *such that $f^{(k)} = f(f(f(\ldots)) < 2$ (where f is composed with itself k times)*

Examples:

1. If $f(n) = n/2$, then $f^*(n) = \log n$
2. If $f(n) = \sqrt{n}$, then $f^*(n) = \log \log n$

Definition 7 *The k^{th} row inverse Ackerman function is defined recursively as follows*

$\lambda_0(n) = n/2$
$\lambda_1(n) = \sqrt{n}$
$\lambda_k(n) = \lambda^*_{n-2}(n)$, *for $k \geq 2$.*

We can now return to the proof of the bound on the number of edges of a k TC spanner.

Let $h = n/f(n)$, where $f(n) = \lambda_{k-2}(n)$

We will prove the bound by induction on k. For $k = 3$ we have already checked that the bound holds. Assume that the size of a $(k-2)$-TC spanner on the line L_h is $\leq (k-2)h\lambda_{k-2}(h)$. We have $(k-2)h\lambda_{k-2}(h) \leq (k-2)\frac{n}{f(n)} f(n/f(n) \leq (k-2)\frac{n}{f(n)} f(n) \leq (k-2)n$, (where we used that f is an increasing function.)

Plugging this in the expression for number of edges, we get $T(n) \leq 2n + (k-2)n + \frac{n}{f(n)} T(n/f(n))$, whose solution is $T(n) = k \cdot n \cdot f^*(n) = k \cdot n \cdot \lambda^*_{k-2}(n) = k \cdot n \cdot \lambda_k(n)$ ■

4 Generalizations

This theorem generalizes to directed rooted trees and directed planar graphs as well. To see a flavor of such generalizations we state a bound for k-TC spanners on the directed grid graph $[\sqrt{n}] \times [\sqrt{n}]$ where all edges are directed from left to right and upwards.

Theorem 8 *([2]) There $\exists$ a k-TC spanner of size $O(n \cdot \log n \cdot \lambda_k(n))$ for the directed grid $[\sqrt{n}] \times [\sqrt{n}]$.*

We mention the main ideas of the proof when $k = 2$. The proof also follows by recursion, where this time we partition the grid into 2 equally-sized grids (the lower and the upper grids each of size $[\frac{\sqrt{n}}{2} \times \sqrt{n}]$) and will recurse on these halves. The middle row, say $L_{\sqrt{n}}$ is called a path separator, and first one builds a 2-TC spanner H of this line. To connect a vertex u from the lower sub-grid to a vertex v in the upper sub-grid, one needs to perform the following steps: (1) 'project' these vertices to the middle row to obtain vertices u' and v' (2) connect u to u' and to all the vertices that u' is connected to in the 2-TC spanner H (3) similarly, connect v to v' and to all the vertices that v' is connected to in the 2-TC spanner H.

One can check that this gives a 2-TC spanner for the grid with the claimed number of edges.

One application of this bound for $k = 2$ is for testing monotonicity of a function defined on the directed grid, i.e. say $f : [\sqrt{n}] \times [\sqrt{n}] \to \{0, 1\}$. The above theorem implies a test with $\log^2 n << n$ many queries.

References

[1] Baruch Awerbuch. Communication-time trade-offs in network synchronization. pages 272–276, 1985.

[2] Arnab Bhattacharyya, Elena Grigorescu, Kyomin Jung, Sofya Raskhodnikova, and David P. Woodruff. Transitive-closure spanners. In *SODA*, pages 932–941, 2009.

[3] David Peleg and Alejandro A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.

[4] Sofya Raskhodnikova. Transitive-closure spanners: A survey. In *Property Testing*, pages 167–196, 2010.