Lower Bound for Sparse Euclidean Spanners *

Pankaj K. Agarwal† Yusu Wang† Peng Yin†

Abstract

Given a one-dimensional graph G such that any two consecutive nodes are unit distance away, and such that the minimum number of links between any two nodes (the *diameter* of G) is $O(\log n)$, we prove an $\Omega(n \log n / \log\log n)$ lower bound on the sum of lengths of all the edges (i.e., the *weight* of G). The problem is a variant of the widely studied *partial sum* problem. This in turn provides a lower bound on Euclidean spanner graphs with small diameter and low weight, showing that the upper bound from [1] is almost tight.

1 Introduction

Given a set of n points $V = \{v_1, \ldots, v_n\}$ in $\mathbb{R}^d$ and a geometric graph $G = (V, E)$, define the *weight* of an edge $e = (v_i, v_j)$ as $w(e) = \|v_i v_j\|$, that is, the Euclidean distance between points v_i and v_j. The weight of a subgraph $G' = (V', E')$ is $w(G') = w(E') = \sum_{e \in E'} w(e)$. The *shortest path* between nodes v_i and v_j, denoted by $P_G(v_i, v_j)$, is the smallest-weight path that connects v_i and v_j in G, while the *minimum link path*, denoted by $\pi(v_i, v_j)$, is the one with the smallest number of edges. Define the *diameter* of the graph as $\Delta(G) = \max_{1 \le i,j \le n} |\pi(v_i, v_j)|$.

The problem studied in this paper arises from the study of spanner graphs: a subgraph $G' \subseteq G$ is a *t-spanner* of G if for any $v_i, v_j \in V$,

$$w(P_{G'}(v_i, v_j))/w(P_G(v_i, v_j)) \le t.$$

Ideally, we would like to have a sparse spanner (i.e., with $O(n)$ edges) with low maximum vertex degree, low weight, and small diameter. Arya *et al.* [1] investigated the problem of constructing spanners while optimizing various combinations of the above measures simultaneously. For example, they showed that it is possible to construct a spanner with $O(n)$ edges, $O(\log n)$ diameter, and $O(w(\mathcal{T}) \log n)$ weight, where $\mathcal{T}$ is the minimum spanning tree of G. The remaining question is then whether this combination is optimal. In other words, we wish to know whether there is a graph so that any spanner of it with $O(\log n)$ diameter has $\Omega(w(\mathcal{T}) \log n)$ weight.

To answer this question, we focus on the following problem, which is interesting in its own right. For any graph where every node v_i lies on the x-axis with coordinate i, what is the smallest weight it can have for a given diameter? In particular, we provide an $\Omega(n \log n / \log\log n)$ lower bound on the weight of any such graph with $O(\log n)$ diameter, implying that the result of Arya *et al.* (i.e., $O(n \log n)$) is almost tight. (Note that for the type of graphs that we are inspecting, the weight of its minimum spanning tree is $n - 1$. So from now on, we simply bound the weight of the graph with respect to n.)

Related work. This one-dimensional graph problem is related to the partial sum problem, where given an array of numbers $A[1], \ldots, A[n]$, one would like to construct a data structure of small size so that a partial sum like $S(i,j) = \sum_{i \le k \le j} A[k]$ can be computed efficiently. Roughly speaking, the query time there corresponds to the diameter in our case, while the canonical sets usually constructed for the data structures there correspond to the edge set in our graphs. The partial sum problem is a special case of orthogonal range searching, and has been widely studied. We only give a small sample of results here. For static partial-sum problem, the query time is $\Omega(\alpha(n,m))$ if m units of storage is used [4]. Tight bounds for the partial sum problem in a dynamic setting under various models were provided in [3]. The problem has also been studied for multi-dimensional arrays [2].

Notation. Assume from now on that $V = \{v_1, \ldots, v_n\}$ is a set of n ordered points in $\mathbb{R}^1$ such that any two consecutive points are unit distance apart. A *block* of nodes $[i : j]$ is defined as $\{v_i, v_{i+1}, \ldots, v_j\}$, and v_i and v_j are referred to as *endpoints* of the block. Let $\chi(v_k)$ be the *covering* of a node v_k, defined as the number of edges that span over v_k, i.e., the number of edges (v_i, v_j) such that $i < k < j$. set $\chi(G) = \max_{v \in V} \chi(v)$. Two edges (v_i, v_j) and (v_k, v_l) *intersect* if $i < k < j < l$. A graph is called a *stack* if it only contains non-intersecting edges. A *cluster* in a stack graph G is a maximal subgraph $G' = (V', E')$ induced by $V' = [i : j]$ such that edge $(v_i, v_j) \in E'$, and no edge in $E \setminus E'$ spans over any point in V'.

2 Diameter and Weight

Weight and covering. The following lemma converts the problem of relating $w(G)$ and $\Delta(G)$ into the problem of relating $\chi(G)$ and $\Delta(G)$. Let $\chi(n, \delta)$ denote the smallest covering for a graph with n nodes and diameter at most δ; let

*P.A. is supported by by NSF grants EIA-98-70724, EIA-99-72879, ITR-333-1050, CCR-97-32787, and CCR-00-86013, and by a grant from the U.S.-Israeli Binational Science Foundation. Y.W. is supported by NSF grants ITR-333-1050 and CCR-02-04118. P.Y. is supported by NSF grant CCR-03-26157.

†Department of Computer Science, Duke University, Durham, NC 27708-0129, U.S.A. Email: {pankaj, wys, py}@cs.duke.edu.

$w(n, \delta)$ denote the weight of such a graph.

LEMMA 2.1. *If $\chi(n, \delta) \geq g(n)$ where $g(n)$ is a concave function, then $w(n, \delta) = \Omega(ng(n))$.*

Proof. Let $G = (V, E)$ be a graph with diameter δ, covering $\chi(n, \delta)$, and weight $w(n, \delta)$. A vertex v is *heavy* if $\chi(v) \geq g(n)/6$; otherwise, v is *light*. Let V_h be the set of heavy nodes. We claim that $|V_h| \geq n/2$, which implies the lemma.

Suppose to the contrary that $|V_h| < n/2$. Decompose V_h into a set of disjoint maximal blocks $B = \{B_1, \ldots, B_k\}$, that is, no larger block $B' \subseteq V_h$ can contain B_i, for any $1 \leq i \leq k$. By contracting the induced subgraph of each B_i into a single node v_i, for $1 \leq i \leq k$, we obtain a new graph $G' = (V', E')$. Obviously, $|V'| \geq n - |V_h| \geq n/2$, and $\Delta(G') \leq \Delta(G) = \delta$. Furthermore, we have that $\chi(G') \leq g(n)/3$. To see this, first note that all light vertices from V remain light in V'. For a contracted heavy vertex v_i, $\chi(v_i) \leq \chi(v_i^-) + \chi(v_i^+)$, where v_i^- (v_i^+) is the vertex in V' to the immediate left (right) of v_i along the horizontal line. Since v_i^- and v_i^+ are light, $\chi(v_i) \leq g(n)/3$. We further reduce the size of V' to $n/2$ by contracting the subgraph induced by the first $|V'| - n/2 + 1$ vertices into a single node. This produces a graph with $n/2$ nodes and diameter at most δ, and its covering is at most $C = g(n)/3$. Since $g(n)$ is a concave function, this leads to a contradiction, as $\chi(n/2, \delta) \geq g(n/2) \geq g(n)/2 > C$.

Stack graphs. The next lemma allows us to focus only on the covering and diameter of a stack graph.

LEMMA 2.2. *For any graph $G = (V, E)$, there is a stack graph $\mathcal{S} = (V, \mathcal{E})$ such that $\chi(\mathcal{S}) \leq \chi(G)$ and $\Delta(\mathcal{S}) \leq (\chi(G) + 1)\Delta(G)$.*

Proof. Intuitively, in order to obtain a stack graph, we wish to split an edge e if it intersects other edges. However, we have to do it in a way such that we do not introduce new intersections while removing an old one. In particular, given a graph $G' = (V', E')$, where $V' = [i : j]$ is a block of nodes from V, let $\bar{E}' = E' \setminus \{(v_i, v_j)\}$ if edge (v_i, v_j) exists. We now find edge $\bar{e}$ whose left endpoint is leftmost in $\bar{E}'$. If there are more than one of such edges, choose the one with the rightmost right endpoint. Let v_s be the right endpoint of $\bar{e}$. For each edge $(v, v') \in \bar{E}'$ that covers v_s, we split it at v_s into two new edges (v, v_s) and (v_s, v'). After splitting all edges covering v_s, we obtain a new edge set $\bar{E}^*$ where no edge from it covers node v_s (see Figure 1). We repeat this process recursively in the two subgraphs induced by nodes $\{v_i, \ldots, v_s\}$ and by $\{v_s, \ldots, v_j\}$.

The above process produces a stack graph $\mathcal{S}$ without increasing the covering for any vertex $v \in V$. Furthermore, a split on an edge $e \in E$ happens only if e intersects some edge that covers its left endpoint, and each split removes at least one such intersection. This implies that $\Delta(\mathcal{S}) \leq (\chi(G) + 1)\Delta(G)$. We omit the proof from here.

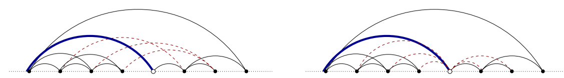

Figure 1: Block $[i : j]$ with $\bar{e}$ (the thick edge) and v_s (hollow point). Dotted edges intersect $\bar{e}$ in (a) and are split in (b).

We now present the relation between the diameter and covering of stack graphs. It is natural to consider a stack graph $\mathcal{S} = (V, \mathcal{E})$ as a tree T: Each node τ represents an edge $e_\tau \in \mathcal{E}$, and the subtree rooted at τ consists of all the edges from $\mathcal{E}$ contained in e_τ. An edge $e_1 \in \mathcal{E}$ is a child of edge e if $e_1 \subset e$ (i.e, e contains e_1) and there is no other edge $e' \in \mathcal{E}$ such that $e_1 \subset e' \subset e$. We add an artificial edge connecting nodes v_1 and v_n (if it does not exist), which corresponds to the root of T. The depth of T is at most $\chi(\mathcal{S}) + 1$, and every edge can have at most $2\Delta(\mathcal{S})$ children. Therefore,

$$n = |V| \leq (2 \cdot \Delta(\mathcal{S}))^{\chi(\mathcal{S})+2}. \tag{2.1}$$

Putting everything together. By (2.1) and Lemma 2.2, for any graph G with n nodes, there is a corresponding stack graph $\mathcal{S}$ also with n nodes:

$$n \leq (2 \cdot \Delta(\mathcal{S}))^{\chi(\mathcal{S})+2} \leq [\, 2 \cdot \Delta(G) \cdot (\chi(G) + 1) \,]^{\chi(G)+1} .$$

Take logarithms on both sides,

$$\chi(G) + 1 \geq \frac{\ln n}{\ln \Delta(G) + \ln(\chi(G) + 1) + \ln 2}.$$

Substituting $\Delta(G) = O(\log n)$ in the above inequality we obtain $\chi(G) = \Omega(\log n / \log\log n)$: Lemma 2.1 now implies the main result:

THEOREM 2.1. *Given any 1-dimensional graph G with unit distance between consecutive nodes and $\Delta(G) = O(\log n)$, $w(G) = \Omega(n \log n / \log\log n)$.*

COROLLARY 2.1. *There is a graph that any of its t-spanners with diameter $O(\log n)$ has $\Omega(w(\mathcal{T}) \log n / \log\log n)$ weight.*

References

[1] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid. Euclidean spanners: Short, thin and lanky. *Proc. 27th Ann. ACM Sympos. Theo. Comput.*, 1995, 489–498.

[2] B. Chazelle and B. Rosenberg. The complexity of computing partial sums off-line. *Int. J. Comput. Geom. Appl.*, 1(1991), 33–45.

[3] M. Pătraşcu and E. D. Demaine. Tight bounds for the partial-sums problem. *Proc. 15th Ann. ACM-SIAM Sympos. Discrete Algorithms*, 2004, 20–29.

[4] A. C. Yao. Space-time tradeoff for answering range queries. *Proc. 14th Ann. ACM Sympos. Theo. Comput.*, 1982, 128–136.