

An Alternate Proof of Near-Optimal Light Spanners*

Greg Bodwin

University of Michigan

Abstract

In 2016, a breakthrough result of Chechik and Wulff-Nilsen [SODA '16] established that every n -node graph G has a $(1 + \varepsilon)(2k - 1)$ -spanner of lightness $O_\varepsilon(n^{1/k})$, and recent followup work by Le and Solomon [STOC '23] generalized the proof strategy and improved the dependence on ε . We give a new proof of this result, with the improved ε -dependence. Our proof is a direct analysis of the often-studied greedy spanner, and can be viewed as an extension of the folklore Moore bounds used to analyze spanner sparsity.

1 Introduction

We study *spanners*, which are a graph-theoretic primitive with applications in graph algorithms, network design, and sketching [1].

Definition 1 (Spanners [13, 14]). *Given a graph G , a t -spanner is an edge-subgraph H that satisfies $\text{dist}_H(u, v) \leq t \cdot \text{dist}_G(u, v)$ for all vertices u, v .*

The usual goal is to design a spanner with a favorable tradeoff between its stretch t and its *size*. There are two different ways that spanner size is commonly measured. One is by the *number of edges* or *sparsity* of the spanner, i.e., the goal is to minimize $|E(H)|$. The stretch/sparsity tradeoff has long been understood, thanks to the following classic theorem by Althöfer et al [3]. Here and throughout the paper, all graphs are undirected and may have arbitrary positive edge weights.

Theorem 1 ([3]). *For all positive integers k, n , every n -node graph G has a $(2k - 1)$ -spanner H on $|E(H)| = O(n^{1+1/k})$ edges. This tradeoff is best possible, assuming the girth conjecture [10].*

The other popular way to measure spanner size is by the *total edge weight* in the spanner $w(H)$. In general, the total edge weight required for a t -spanner might be unbounded, say by scaling up the edge weights of the input graph. So in order to prove stretch/weight tradeoffs, we typically normalize the spanner weight by the weight of a minimum spanning tree (MST) of the original graph. This size measure is called *lightness*:

Definition 2 (Spanner Lightness). *The lightness of a subgraph H of a graph G is the quantity*

$$\ell(H \mid G) := \frac{w(H)}{w(\text{MST}(G))}$$

where $\text{MST}(G)$ is a minimum spanning tree of G .¹ For brevity we also write $\ell(H) := \ell(H \mid H)$.

There has been a long line of work studying the tradeoff between spanner stretch and lightness; see Table 1 for the progression of results. A key work in this sequence was a breakthrough of Chechik and Wulff-Nilsen [8], which established the following analog of Theorem 1:

Theorem 2 ([8]). *For all $\varepsilon > 0$ and positive integers k, n , every n -node graph G has a $(1 + \varepsilon)(2k - 1)$ -spanner H of lightness $\ell(H \mid G) = O_\varepsilon(n^{1/k})$.*

*This work was supported by NSF:AF 2153680. bodwin@umich.edu

¹Throughout this paper, we will assume where convenient that graphs are connected, so that $\text{MST}(G)$ exists. Otherwise, a minimum spanning forest may be used.

Stretch	Lightness	Analyzes Greedy Spanner?	Citation
$2k - 1$	$O(n/k)$	✓	[3]
$(1 + \varepsilon) \cdot (2k - 1)$	$O_\varepsilon(k \cdot n^{1/k})^*$	✓	[7]
$(1 + \varepsilon) \cdot (2k - 1)$	$O_\varepsilon\left(\frac{k}{\log k} \cdot n^{1/k}\right)^*$	✓	[9]
$(1 + \varepsilon) \cdot (2k - 1)$	$O(\varepsilon^{-(3+2/k)} n^{1/k})$		[8]
$(1 + \varepsilon) \cdot (2k - 1)$	$O(\varepsilon^{-1} n^{1/k})^*$		[12]
$(1 + \varepsilon) \cdot (2k - 1)$	$O(\varepsilon^{-1} n^{1/k})$	✓	this paper

Table 1: Work on the stretch/lightness tradeoff for spanners. Lightness bounds marked with a * have an *unconditionally* optimal dependence on n , i.e., they automatically improve if the girth conjecture fails.

The stretch/lightness tradeoff in Theorem 2 is best possible, assuming the girth conjecture [10], and up to its dependence on ε (which could conceivably be improved or even removed). The theorem is proved using an ingenious framework for hierarchical graph clustering. A followup paper by Le and Solomon [12] refined and vastly generalized the scope of this clustering method, improving the hidden ε -dependence in the lightness bound of Theorem 2, and also gaining broad applications to the study of light spanners in various important graph classes (see also [2, 5, 6]).

Perhaps the most popular spanner construction algorithm is the following *greedy algorithm*:

Input: Graph $G = (V, E, w)$, stretch t ;

Let $H = (V, \emptyset, w)$ be the initially-empty spanner;

foreach $(u, v) \in E$ *in order of nondecreasing weight* **do**

| **if** $\text{dist}_H(u, v) > t \cdot w(u, v)$ **then**

| | Add (u, v) to H ;

return H ;

Algorithm 1: The Greedy Spanner Algorithm [3]

The greedy algorithm is well-studied and widely used because it is simple, easy to prove correct, and its stretch/sparsity and stretch/lightness tradeoffs are both known to be *existentially optimal* [3, 11]. That is, the stretch/lightness tradeoff achieved by *any* algorithm – including the clustering method in [8, 12] – is automatically achieved by the greedy algorithm as well. This has motivated interest in spanner size bounds that are proved by directly analyzing the output spanner of the greedy algorithm, rather than turning to alternate constructions. In the context of spanner *sparsity*, there is indeed a simple proof of Theorem 1 that works by directly analyzing the greedy spanner. This proof is called the Moore bounds and it is considered folklore; we recap the proof in Section 2. In the context of spanner *lightness*, there are some arguments that directly analyze the greedy spanner [3, 7, 9], but they all show suboptimal lightness bounds that do not quite match the one in Theorem 2. However, current proofs of near-optimal light spanners in Theorem 2 work by analyzing the alternate clustering-based construction.

The contribution of this paper is a new proof of Theorem 2, with the improved ε -dependence from [12], which directly analyzes the greedy spanner (or, more accurately, which directly analyzes graphs of high weighted girth [9]; see Section 3.1). Our proof also follows the template of the Moore bounds (see also [4]), and so it may also have an advantage in conceptual familiarity to a reader who is primarily comfortable with the literature on spanner sparsity.

Organization. A lot of this paper is optional “warmup” content rather than the main proof, and the enterprising reader can get the full proof by reading Sections 3 and 5 only. Nonetheless, the surrounding warmup proofs, discussions, and puzzles are recommended for intuition.

2 Warmup 1: The Moore Bounds for Spanner Sparsity

In order to demonstrate our proof strategy, we will recap the proof of the stretch/sparsity tradeoff given in Theorem 1. We first observe that the output spanner H of the greedy algorithm with stretch parameter $2k - 1$ has girth (shortest cycle length) $> 2k$ [3]. (We omit this proof, as it is standard, but we note that it is implied by Lemma 8 to follow.) Theorem 1 then follows from the *Moore bounds*, which limit the maximum possible number of edges in a high-girth graph:

Theorem 3 (Moore Bounds). *For any positive integers n, k , every n -node graph H with girth $> 2k$ has $O(n^{1+1/k})$ edges.*

The proof of the Moore bounds is a counting argument over the *edge-simple k -paths* of H . Recall that an edge-simple path is one that does not repeat edges. One part of this counting argument is the following *dispersion lemma*, implying that these paths are “dispersed” around the graph, rather than having several of them concentrated on any given pair of endpoints.

Lemma 4 (Unweighted Dispersion Lemma). *H may not have two distinct edge-simple k -path with the same endpoints s, t .*

Proof. Suppose for contradiction that π_a, π_b are two distinct $s \rightsquigarrow t$ edge-simple k -paths in H . The subgraph $\pi_a \cup \pi_b$ is not a tree, since it contains two distinct $s \rightsquigarrow t$ paths, and so it contains a cycle C . This cycle must have $|C| \leq |\pi_a| + |\pi_b| = 2k$ edges, which contradicts that H has girth $> 2k$. \square

The dispersion lemma implies an absolute upper bound on the number of edge-simple k -paths in H . The other part of the argument is a *counting lemma*, implying a lower bound on the same quantity, which scales with the size $|E(H)|$. Only the last “full” counting lemma in the following sequence is used, but the proof strategy is to bootstrap it by starting with weaker versions.

Lemma 5 (Unweighted Weak Counting Lemma). *If $|E(H)| \geq n$, then H contains an edge-simple k -path.*

Proof. Since $|E(H)| \geq n$, H contains a cycle C . Since H has girth $> 2k$, there are $> 2k$ edges in C . Thus, any subpath of C of length k is an edge-simple k -path. \square

Lemma 6 (Unweighted Medium Counting Lemma). *H contains at least $|E(H)| - n$ edge-simple k -paths.*

Proof. Repeat the following process until no longer possible: find an edge-simple k -path π , record it, and then delete any **one** edge in π from H to ensure that we don’t re-record π in a future round. By the weak counting lemma, we may repeat this process for at least $|E(H)| - n$ rounds. \square

Lemma 7 (Unweighted Full Counting Lemma). *Let $d := |E(H)|/n$. If $d \geq 2$, then H contains $n \cdot \Omega(d)^k$ edge-simple k -paths.*

Proof. Let H' be a random edge-subgraph of H , obtained by keeping each edge independently with probability $2/d$. Let p, p' be the number of edge-simple k -paths in H, H' , respectively. On one hand, for any edge-simple k -path π in H , the probability that π survives in H' is $\Theta(d)^{-k}$, and so $\mathbb{E}[p'] = p \cdot \Theta(1/d)^k$. On the other hand, we have

$$\begin{aligned} \mathbb{E}[p'] &\geq \mathbb{E}[|E(H')| - n] && \text{Medium Counting Lemma} \\ &= |E(H)| \cdot \frac{2}{d} - n \\ &= 2n - n \\ &= n. \end{aligned}$$

Combining these inequalities, we get $n \leq p \cdot \Theta(1/d)^k$, and rearranging gives $p \geq n \cdot \Theta(d)^k$. \square

We are now ready to complete the proof of the Moore bounds. Let $d := |E(H)|/n$. If $d < 2$ then $|E(H)| = O(n)$ and we are done. Otherwise, if $d \geq 2$, then by the full counting lemma H has $n \cdot \Omega(d)^k$ edge-simple k -paths. Meanwhile, the dispersion lemma implies that H has $O(n^2)$ edge-simple k -paths. Comparing these estimates, we get

$$n \cdot \Omega(d)^k \leq O(n^2).$$

Rearranging terms in this inequality, we get $d \leq O(n^{1/k})$, and so $|E(H)| = O(n^{1+1/k})$.

3 Some Ideas About Lightness from Prior Work

We will next recap some helpful reductions from prior work on spanner lightness.

3.1 The Weighted Girth Framework

In our previous proof of the Moore bounds, the first step is to observe that the output spanner of the greedy algorithm has high girth, and then the focus of the proof shifts to bounding sparsity of an arbitrary high-girth graph. For lightness, Elkin, Neiman, and Solomon [9] formalized the analogous method, which we will use in this paper.

Definition 3 (Normalized Weight and Weighted Girth [9]). *For a cycle C in G , we define its normalized weight to be*

$$w^*(C) := \frac{w(C)}{\max_{e \in C} w(e)}.$$

The weighted girth of G is the minimum value of $w^(C)$ over all cycles C in G .*

Lemma 8 ([9]). *The greedy algorithm with parameter t returns a graph H with weighted girth $> t + 1$.*

Proof. Let C be a cycle in G of normalized weight $w^*(C) \leq t + 1$. It suffices to argue that not all edges of C will be added to H . Let (u, v) be the last edge in C considered by the greedy algorithm, and suppose that all previous edges in C were added to H . Then there is a $u \rightsquigarrow v$ path in H , through the other edges in C , of total weight

$$w(C) - w(u, v) \leq (t + 1) \cdot w(u, v) - w(u, v) = t \cdot w(u, v).$$

This inequality means that the greedy algorithm will reject the edge (u, v) , rather than adding it to H , and so the cycle C does not survive in H . \square

We also observe that the greedy algorithm essentially contains a run of Kruskal's algorithm within it, and so the output spanner H contains an MST of G . This means that instead of bounding $\ell(H \mid G)$, we may equivalently bound $\ell(H)$.

3.2 Reduction to Unit-Weight Spanning Cycles

It will be convenient in the main proof to reduce to the setting where the spanner H has a very particular structure for its MST:

Definition 4 (Unit-Weight Spanning Cycles). *We say that a cycle C in H is a unit-weight spanning cycle if C is Hamiltonian (it contains each node exactly once), all edges in C have weight 1, and all edges in $E(H)$ have weight ≥ 1 .*

So, for example, any tree created by deleting any one edge from a unit-weight spanning cycle C is an MST. We will reduce to the case where H has a unit-weight spanning cycle; the parts of this reduction all appear implicitly or explicitly in prior work [8, 9, 12].

Lemma 9. *Let H be a non-forest n -node graph with weighted girth $> t$ and lightness ℓ . Then there exists a graph H' with $O(n)$ nodes, weighted girth $> t$, lightness $\Omega(\ell)$, and a unit weight spanning cycle.*

Proof. We split the reduction into two steps: first we modify to a unit-weight MST, and then we change the MST into a spanning cycle.

Reduction to Unit-Weight MST.

- Rescale the edge weights of G so that the average edge weight in $\text{MST}(H)$ is 1.
- For all edges $e \in \text{MST}(H)$ of weight $w(e) > 1$, add new nodes to H to subdivide e into a path of $\lceil w(e) \rceil$ edges. Each new edge is assigned weight $w(e)/\lceil w(e) \rceil$. Notice that the total weight in $\text{MST}(H)$ is $n - 1$ after the previous rescaling step, and therefore we add at most $n - 1$ new nodes to H in this step. We also notice that weighted girth of H does not change in this step, since the total weight of each cycle C is unchanged, and also C must have a heaviest edge that is not in the MST, which does not change weight.

- Finally, for all edges $e \in E(H)$ of weight < 1 , increase $w(e)$ to 1 (this step is applied to both MST and non-MST edges.) This step increases $w(\text{MST}(H))$ by at most $+(n-1)$, which is a constant factor. Additionally, weighted girth is nondecreasing in this step. This is because (1) for each cycle C , $w(C)$ is increasing, (2) if C has an edge of weight > 1 then its heaviest edge weight does not change, and (3) if all edges in C have weight ≤ 1 , then after this step all edges have weight 1, so its normalized weight is $w^*(C) = |C| > t$.

Reduction to Spanning Cycle. Next, we reduce to the setting where H has a spanning cycle C .




Figure 1: (Left) For the pictured $\text{MST}(H)$, the red line around the outside of the tree traces the tour $T = (a, b, c, b, d, b, a, e, f, e, g, e, a)$. (Right) We construct H' by mapping the tour T to a spanning cycle on $2n - 2$ nodes, making copies of nodes from H as needed.

A tour T of $\text{MST}(H)$ is a circularly-ordered sequence of nodes, with repeats, of the form $T = (v_0, v_1, \dots, v_{2n-2} = v_0)$, which is the node sequence of a closed walk on $\text{MST}(H)$ that visits every edge exactly twice, with opposite orientations. Fix a tour T of $\text{MST}(H)$, and then construct H' as follows:

- A tour of an n -node tree always contains exactly $2n - 2$ nodes. We will take these $2n - 2$ nodes as the vertex set of H' ; that is, some nodes in H have several copies in H' .
- The tour T will be the spanning cycle of H' , meaning that for each pair of adjacent nodes along T , we include the corresponding edge in $E(H')$ with weight 1.
- For each non-spanning-cycle edge $(u, v) \in E(H \setminus T)$, we choose an arbitrary copy $u_i, v_j \in V(H')$ of u, v respectively. Then we include $(u_i, v_j) \in E(H')$ with the same weight as (u, v) .

It is immediate from the construction that H' has $O(n)$ nodes and a unit-weight spanning cycle, and that $w(H') \geq w(H)$, and so lightness only changes by a constant factor. Moreover, we only create one new cycle when we move from H to H' , which is the spanning cycle T itself. Note that $w^*(T) = 2n - 2$, and since we have assumed that H is not a forest, it has a cycle and therefore its weighted girth is $\leq n$. It follows that the weighted girth of H' is at least as large as the weighted girth of H . \square

One of the advantages of reducing to the case where H has a unit-weight spanning cycle is that we can limit its maximum edge weight:

Lemma 10. *Let H be an n -node graph with weighted girth $> t$ and a unit weight spanning cycle C . Then all edges in H have weight $< \frac{n}{2(t-1)}$.*

Proof. Consider an edge (u, v) , and consider the cycle formed by (u, v) and the shorter $u \rightsquigarrow v$ path through the spanning cycle, which uses at most $n/2$ edges. The normalized weight of this cycle is at most

$$\frac{w(u, v) + n/2}{w(u, v)} = 1 + \frac{n}{2w(u, v)}.$$

By the weighted girth of H , this quantity must be $> t$. Rearranging, we get $w(u, v) < \frac{n}{2(t-1)}$. \square

4 Warmup 2: Lightness Bounds via Monotone Paths

We will next prove a weaker version of our main result, with an additional suboptimal k factor, in order to introduce some of our new proof ideas.

Theorem 11 (Warmup). *Let $\varepsilon > 0$, let k, n be positive integers, and let H be an n -node graph with a unit-weight spanning cycle \mathcal{C} and weighted girth $> (1 + 2\varepsilon) \cdot 2k$. Then*

$$w(H) = O\left(\varepsilon^{-1}kn^{1+1/k}\right).$$

The 2ε term in the weighted girth, rather than ε , is purely for convenience in the analysis to follow; by reparametrizing $\varepsilon \leftarrow \varepsilon/2$ it does not affect the theorem statement. We will also assume for convenience that all non-spanning-cycle edges in H have distinct weights; if not, any tiebreaking method will work, e.g. the lexicographically smaller edge is considered lighter. Finally, we arbitrarily choose one direction around the spanning cycle \mathcal{C} to be the **forward** direction, and the reverse to be **backward**.

4.1 Edge-Safe Paths and Warmup Proof Intuition

Our first step is to define the kind of paths that will be the focus of our Moore-bound-like counting argument. The following kind of path is analogous to a *single edge* in the Moore bounds; we will ultimately focus on paths made of k concatenated copies of this basic building block.

Definition 5 (Edge-Safe Paths). *A path π is safe for an edge (u, v) if, for some integer $0 \leq s \leq \varepsilon w(u, v)$, it has the following structure: it starts with a prefix of exactly s **forward** spanning cycle edges, then it uses the edge (u, v) , and then it ends with a suffix of exactly s **backward** spanning cycle edges. We will say that π is *extra-safe* for (u, v) if in fact $s \leq \varepsilon w(u, v)/2$.*




Figure 2: A path that is safe for the edge (u, v) (considering the counterclockwise direction around the spanning cycle to **forward**, and clockwise to be **backward**)

The detail of extra-safety in this definition can be ignored for now; it will become relevant for a minor technical reason in the counting lemma. The requirement that an edge-safe path uses the same number of **forward** and **backward** spanning cycle edges enables the following simple yet important technical claim:

Claim 12. *Let q, q' be paths in H that are safe for edges e, e' respectively, and which share an endpoint node y . If $q \neq q'$, then $e \neq e'$.*

Proof. We prove the contrapositive. Suppose $e = e' =: (u, v)$. By Lemma 10 we have $w(u, v) < n/2$, and so no path safe for (u, v) can use more than n spanning cycle edges. Let $0 \leq s < n$ be the number of **backward** steps along the spanning cycle from v to y . Then both q, q' must end with a suffix of exactly s **backward** steps along the spanning cycle. So they also begin with a prefix of exactly s **forward** steps along the spanning cycle, ending at u , implying equality. \square

In the same way that the Moore bounds focus on paths made up of k edges, a natural proof attempt to bound lightness would be to focus on paths made up of k edge-safe subpaths:

Definition 6 (Safe k -Paths). A path π in H is a safe k -path if it can be partitioned into k subpaths $\pi = q_1 \circ \dots \circ q_k$, where each path q_i is safe for an edge e_i . We say that π is an extra-safe k -path if each path q_i is extra-safe for e_i .

Unfortunately, this natural attempt breaks. Specifically, the dispersion lemma fails: it is possible to have two edge-simple safe k -paths that share endpoints, without implying that H has a cycle of small normalized weight. See Figure 3 for a diagram and discussion of this failure case:




Figure 3: In this picture the solid blue lines represent non-spanning-cycle edges, and we imagine that (a, b) is much heavier than all others. The dashed lines represent spanning cycle subpaths of length $\epsilon w(a, b)$ each. This graph contains two different safe 3-paths with endpoints s, t . One is (s, a, b, t) , which uses spanning cycle subpaths of length 0 between its non-spanning-cycle edges. The other has the form $(s, c, \dots, a, b, \dots, d, t)$, which uses long spanning cycle subpaths before and after its middle edge (a, b) . However, neither of the two cycles (s, a, c, \dots, s) nor (t, b, \dots, d, t) have small normalized weight.

The point made by [7, 9] is essentially that the counterexample in Figure 3 can be avoided if we narrow our focus to paths that do not have one edge (a, b) that is much heavier than the others. Indeed, the dispersion lemma holds for safe k -paths whose non-spanning-cycle edge weights differ by at most a factor of 2 [7], or even a factor of k [9]. The point made in our warmup proof is that this counterexample can also be avoided if we narrow our focus to “monotonic” paths, in which heavy edges (a, b) are allowed, but they must occur at the *end* of the path rather than somewhere in the middle.

4.2 Monotone Safe Paths and the Dispersion Lemma

Recall in the following definition that we have assumed for convenience that the non-spanning-cycle edges of H have distinct weights.

Definition 7 (Monotone Safe k -Paths). Let π be a safe k -path, which can be partitioned into subpaths $\pi = q_1 \circ \dots \circ q_k$ where each q_i is safe for an edge e_i . We say that π is monotone if these edges are increasing in weight, that is, $w(e_1) < \dots < w(e_k)$.

Lemma 13 (Monotone Dispersion Lemma). H may not have two distinct monotone safe k -paths with the same endpoints s, t .

Proof. Seeking contradiction, let π^a, π^b be two distinct monotone safe k -paths with the same endpoints s, t . Let the decomposition of π^a into safe paths be $\pi^a = q_1^a \circ \dots \circ q_k^a$, where each subpath q_i^a is safe for the edge e_i^a . We use similar notation for π^b .

Let j be the largest index for which $q_j^a \neq q_j^b$. Since the subpaths following q_j^a, q_j^b are equal, q_j^a, q_j^b must share an endpoint node, which we will call y . By Claim 12, q_j^a, q_j^b are safe for distinct edges $e_j^a \neq e_j^b$. Assume without loss of generality that $w(e_j^a) > w(e_j^b)$. By monotonicity, it follows that $e_j^a \notin \pi_b[s \rightsquigarrow y]$. We can therefore find a cycle

$$C \subseteq \pi^a[s \rightsquigarrow y] \cup \pi^b[s \rightsquigarrow y]$$

in which e_j^a is the heaviest edge. We can bound the normalized weight of C as:




Figure 4: The setup in our proof of the monotone dispersion lemma; each colorful segment is an edge-safe path. The proof in this case takes $j = 3$ and argues that $w^*(C) \leq (1 + 2\varepsilon) \cdot 6$.

$$\begin{aligned}
 w^*(C) &\leq \frac{\sum_{i=1}^j w(q_i^a) + w(q_i^b)}{w(e_j^a)} \\
 &\leq \frac{\sum_{i=1}^j (1 + 2\varepsilon) (w(e_i^a) + w(e_i^b))}{w(e_j^a)} \\
 &\leq \frac{(j)(1 + 2\varepsilon) (2w(e_j^a))}{w(e_j^a)} \\
 &\leq (1 + 2\varepsilon) \cdot 2j.
 \end{aligned}$$

Since $j \leq k$, this contradicts that H has weighted girth $> (1 + 2\varepsilon) \cdot 2k$, completing the proof. \square

4.3 Hiker Paths and the Counting Lemma

The following is a famous puzzle in graph theory:


A HIKER PUZZLE

We are vacationing in Graph National Park, which has n landmarks (nodes) and m trails (undirected edges) connecting pairs of landmarks. Each trail has a difficulty rating (weight). We would like to hike as many trails as possible, without repeating any trails. However, we will get increasingly tired as we hike, and so we are only willing to hike trails in nonincreasing order of difficulty. That is, after we hike a trail t_i , our next trail t_{i+1} must depart from the endpoint of t_i and its difficulty rating may not be higher than t_i 's difficulty rating. **Prove:** If we start at the right landmark, then we can hike at least $2m/n$ trails.

A valid hiker path, 4 trails long.

SOLUTION

This problem, like so many others in life, is solved by asking our friends for help. We invite friends to join our vacation until we have gathered a group of n total hikers, and then we start with one hiker standing at each landmark. For each trail $t = (u, v)$, considered in descending order of difficulty, we ask the hiker currently standing at u and the hiker currently standing at v to hike the trail, switching places with each other.



In total, our n hikers will hike $2m$ trails, and so there must exist a hiker who hiked a path of length at least $2m/n$. This hiker hiked their trails in descending order of difficulty, and so their path satisfies the puzzle.

We will now prove counting lemmas for monotone safe k -paths. The medium and full counting lemmas from the Moore bounds generalize easily, but a new idea is needed for the weak counting lemma. Our proof will take direct inspiration from the hiker puzzle.

Lemma 14 (Warmup Weak Counting Lemma). *If $w(H \setminus \mathcal{C}) \geq \varepsilon^{-1}kn$, then H contains a monotone extra-safe k -path.*

Proof. Start by placing a hiker at each node of H . Then, consider the non-spanning-cycle edges of H in order of increasing weight. When an edge (u, v) is considered, for each path π that is extra-safe for (u, v) , we ask the two hikers at either endpoint of π to hike π , thus switching places with each other. We note that no hiker walks (u, v) more than once: the set of hikers within $\varepsilon w(u, v)/2$ forward steps of u on the spanning cycle and the set of hikers within $\varepsilon w(u, v)/2$ forward steps of v on the spanning cycle must be disjoint. Otherwise, we would have a cycle of normalized weight $\leq 1 + \varepsilon$ formed by the edge (u, v) together with a $u \rightsquigarrow v$ path in the spanning cycle of total length at most $\varepsilon w(u, v)$, which would contradict the weighted girth of H .

There are at least $\varepsilon w(u, v)/2$ paths that are extra-safe for each edge (u, v) , and two hikers hike each such path (one in each direction). Thus, after all non-spanning-cycle edges of H are considered, in total our n hikers have hiked at least $\varepsilon w(H \setminus \mathcal{C}) \geq kn$ extra-safe paths, and the path walked by each hiker is a monotone extra-safe path. Thus, there exists a hiker who hiked a monotone extra-safe path of length at least k . \square

The bootstrapping process from the weak to the medium and full counting lemmas essentially works exactly as in the Moore bounds, with minor adjustments.

Lemma 15 (Warmup Medium Counting Lemma). *H contains at least $\Theta(\varepsilon) \cdot (w(H \setminus \mathcal{C}) - \varepsilon^{-1}kn)$ monotone safe k -paths.*

Proof. Repeat the following process until no longer possible. Find a monotone extra-safe k -path π , with decomposition $\pi = q_1 \circ \dots \circ q_k$. Let e_1 be the edge for which q_1 is extra-safe. Notice that, for any integer $0 \leq s \leq \varepsilon w(e_1)/2$, we can record a (not-necessarily-extra-)safe k -path by modifying π by adding s additional forward spanning cycle edges to the start of every path q_i , and also adding s additional backward spanning cycle edges to the end of every path q_i .

We record $\Theta(\varepsilon w(e_1))$ monotone safe k -paths in this way. We then delete the edge e_1 , to ensure that we do not re-record any of these paths in a future round. By the weak counting lemma, we

may repeat this process at least until $w(H \setminus \mathcal{C}) < \varepsilon^{-1}kn$. It follows that we will record at least $\Theta(\varepsilon) \cdot (w(H \setminus \mathcal{C}) - \varepsilon^{-1}kn)$ monotone safe k -paths before halting. \square

Lemma 16 (Warmup Full Counting Lemma). *Let $d := w(H \setminus \mathcal{C})/n$. If $d \geq 2k\varepsilon^{-1}$, then H contains at least $kn \cdot \Omega(\varepsilon d/k)^k$ monotone safe k -paths.*

Proof. Let H' be a random edge-subgraph of H obtained by keeping the spanning cycle \mathcal{C} deterministically, and keeping each non-spanning-cycle edge independently with probability $2k\varepsilon^{-1}/d$. Let p, p' be the number of monotone safe k -paths in H, H' , respectively. On one hand, monotonicity implies that any monotone safe k -path π uses k *distinct* non-spanning-cycle edges. Thus, the probability that π survives in H' is $\Theta(k\varepsilon^{-1}/d)^k$, and so

$$\mathbb{E}[p'] = p \cdot \Theta\left(\frac{k\varepsilon^{-1}}{d}\right)^k.$$

On the other hand, we have

$$\begin{aligned} \mathbb{E}[p'] &\geq \mathbb{E}\left[\Theta(\varepsilon) \cdot (w(H' \setminus \mathcal{C}) - \varepsilon^{-1}kn)\right] && \text{Medium Counting Lemma} \\ &= \Theta(\varepsilon) \cdot (\mathbb{E}[w(H' \setminus \mathcal{C})] - \varepsilon^{-1}kn) \\ &= \Theta(\varepsilon) \cdot \left(w(H \setminus \mathcal{C}) \cdot \frac{2k\varepsilon^{-1}}{d} - \varepsilon^{-1}kn\right) \\ &= \Theta(\varepsilon) \cdot (2kn\varepsilon^{-1} - \varepsilon^{-1}kn) \\ &= \Theta(\varepsilon) \cdot (\varepsilon^{-1}kn) \\ &= \Theta(kn). \end{aligned}$$

Comparing the two previous bounds on $\mathbb{E}[p']$, we get

$$\Theta(kn) \leq p \cdot \Theta\left(\frac{k\varepsilon^{-1}}{d}\right)^k.$$

Rearranging this inequality gives our desired inequality of

$$p \geq kn \cdot \Theta\left(\frac{d}{k\varepsilon^{-1}}\right)^k. \quad \square$$

We are now ready to complete the proof of Theorem 11, which is essentially the same as in the Moore bounds. Let $d := w(H)/n$. If $d < 2k\varepsilon^{-1}$, then we have $w(H) = O(kn\varepsilon^{-1})$ and we are done. Otherwise, if $d \geq 2k\varepsilon^{-1}$, then we may apply the full counting lemma to say that H has $kn \cdot \Omega(\varepsilon d/k)^k$ monotone safe k -paths. Meanwhile, the dispersion lemma implies that H has $O(n^2)$ such paths. Comparing these estimates, we get

$$kn \cdot \Omega\left(\frac{d}{k\varepsilon^{-1}}\right)^k \leq O(n^2).$$

Rearranging terms in this inequality to isolate d , we get

$$d \leq O\left(\varepsilon^{-1}k^{(k-1)/k}n^{1/k}\right) = O\left(\varepsilon^{-1}kn^{1/k}\right),$$

and thus $w(H) = w(H \setminus \mathcal{C}) + w(\mathcal{C}) = nd + n = O(\varepsilon^{-1}kn^{1+1/k})$.

5 Full Proof: Light Spanners via Bucket-Monotone Paths

We are now ready to prove our main theorem:

Theorem 17 (Main Theorem). *Let $\varepsilon > 0$, let k, n be positive integers, and let H be an n -node graph with a unit-weight spanning cycle \mathcal{C} and weighted girth $> (1 + 4\varepsilon) \cdot 2k$. Then*

$$w(H) = O\left(\varepsilon^{-1}n^{1+1/k}\right).$$

By the reductions in Section 3, this implies Theorem 2. As in the warmup proof, we use 4ε rather than ε in the weighted girth purely for convenience, and we will arbitrarily define a **forward** and **backward** direction around the spanning cycle.

5.1 Bucket-Safe Paths and Full Proof Intuition

Our full proof will save a factor of k in lightness over the previous warmup proof. Before proceeding to this main proof, let us give some intuition on where exactly the previous proof is losing this factor of k . We will consider two extreme opposing cases for the structure of H , and see how our previous proof loses this k -factor in either one. In the following we define the *bucket* B_i to be the set of edges in $H \setminus \mathcal{C}$ with weights in the range $[2^i, 2^{i+1})$.

Special Case 1: Well-Separated Edge Weights. Let us imagine that we are in an extreme special case in which no two edges in $H \setminus \mathcal{C}$ lie in the same bucket; in particular, any two edges e_1, e_2 have weights $w(e_1), w(e_2)$ that differ by a factor of at least 2. In this case, the k -factor loss from our previous argument is because we are too stingy in our definition of edge-safe paths. We allowed only $s \leq \varepsilon w(e)$ spanning cycle edges before and after each spanning cycle edge e , but in fact our dispersion lemma would still hold if we allowed approximately $\varepsilon k w(e)$ such edges. This is because the separation of edge weights implies that the total number of non-spanning-cycle edges will *telescope* over the monotone path; e.g., revisiting the proof of Lemma 13 there will still be only $O(\varepsilon k w(e_j^a))$ in the cycle C that we analyze. These ultimately cost only a negligible $+O(\varepsilon k)$ term in the bound on weighted girth. This leads to our first piece of intuition on how to improve the argument:

We can budget $O(\varepsilon k 2^i)$ non-spanning-cycle edges per bucket B_i , rather than $O(\varepsilon w(e))$ non-spanning-cycle edges per edge e .

Special Case 2: Bucketed Edge Weights. Let us now imagine the opposite special case, in which all edges in $H \setminus \mathcal{C}$ lie in the same bucket B_i . In this case, as discussed earlier, *monotonicity* is not needed in the proof; we can use safe k -paths that use edges from B_i in any order. For these (not-necessarily-monotone) safe k -paths, the weak counting lemma can be improved by a factor of k : they are guaranteed to exist once $w(H \setminus \mathcal{C}) \geq \varepsilon^{-1}n$, rather than $w(H \setminus \mathcal{C}) \geq \varepsilon^{-1}kn$. This leads to our second piece of intuition on how to gain a k -factor:

The non-spanning-cycle edges within any bucket B_i do not need to occur in monotone order.

Bucket-Safe Paths. Our full proof will use *bucket-safe paths*, which replace the *edge-safe paths* from the previous warmup proof:

Definition 8 (Bucket-Safe Paths). *A path π in H is safe for bucket B_i if it is non-backtracking (meaning that it does not repeat any edge twice in a row), all of its non-spanning-cycle edges are in B_i , and for some integer $0 \leq s \leq \varepsilon k 2^i$ it contains exactly $2s$ spanning cycle edges, where the first s are in the **forward** direction and the last s are in the **backward** direction. We say that π is **extra-safe** if in fact $s \leq \varepsilon k 2^{i-1}$.*

We note that any empty (single-node) path is considered to be safe for every bucket. Bucket-safe paths incorporate both pieces of intuition: they allow the edges within bucket B_i to occur in any order, and they budget $\varepsilon k 2^i$ non-spanning-cycle edges for this bucket. The following technical claim extends Claim 12 in the natural way to bucket-safe paths, and it has essentially the same proof:

Claim 18. *Let q, q' be bucket-safe paths in H , let σ, σ' be the sequences of non-spanning-cycle edges used by q, q' respectively, and suppose that q, q' share an endpoint node y . If $q \neq q'$, then $\sigma \neq \sigma'$.*

Proof. We will prove the contrapositive. Suppose that $\sigma = \sigma' =: ((u_1, v_1), \dots, (u_j, v_j))$. For all pairs of nodes v_i, u_{i+1} , and also for v_j, y , there are two possible spanning cycle paths between these nodes (going around \mathcal{C} in either direction). One of these two spanning cycle paths must use $\geq n/2$ spanning cycle edges. It follows from Lemma 10 that the maximum edge weight in H is $W < \frac{n}{4\varepsilon k}$, and thus any bucket-safe path uses at most $2\varepsilon k W < n/2$ spanning cycle edges. Thus q, q' must choose the same spanning cycle paths between each of these pairs of nodes. This implies that q, q' are identical on their suffix following the node u_1 . By structure of bucket-safe paths, this then implies that q, q' must use prefixes of the same number of **forward** steps before u_1 , giving equality. \square

Our plan is to use bucket-safe paths as building blocks to construct *bucket-monotone paths*, which are formed by concatenating bucket-safe paths in monotonic order of bucket weight, exactly like we previously constructed monotone safe paths by concatenating edge-safe paths.

5.2 Bucket-Monotone Paths and the Dispersion Lemma

We focus on the following paths in our proof:

Definition 9 (Bucket-Monotone Safe k -Paths). *A path π in H is a bucket-monotone safe k -path if it has exactly k non-spanning-cycle edges in total, and it can be partitioned into (possibly empty) subpaths $\pi = q_0 \circ \dots \circ q_j$, where each subpath q_i is safe for bucket B_i (and so these bucket weights are increasing along π). We say that π is extra-safe if each subpath q_i is extra-safe for B_i .*

We note that, although each individual bucket-safe path q_i is non-backtracking, a bucket-monotone safe k -path may backtrack, e.g. when edges from q_{i+1} backtrack edges from q_i . This will be used in the counting lemma. The dispersion lemma for bucket-monotone safe k -paths is similar to the one from the warmup proof, and does not contain any major technical departures.

Lemma 19 (Dispersion Lemma). *H may not have two distinct bucket-monotone safe k -paths with the same endpoints s, t .*

Proof. Seeking contradiction, let π^a, π^b be two distinct bucket-monotone safe k -paths with the same endpoints s, t . Let the decomposition of π^a into bucket-safe paths be $\pi^a = q_0^a \circ q_1^a \circ \dots$, and let σ_i^a be the (possibly empty) sequence of non-spanning-cycle edges used in q_i^a . We use similar notation for π^b .

Let j be the last index on which $q_j^a \neq q_j^b$. Note that q_j^a, q_j^b share an endpoint node, which we will call y . Then by Claim 18, we have $\sigma_j^a \neq \sigma_j^b$. Thus the $s \rightsquigarrow y$ prefixes of π^a, π^b are distinct, and in particular there exists a cycle

$$C \subseteq \pi^a[s \rightsquigarrow y] \cup \pi^b[s \rightsquigarrow y]$$

where C contains at least one edge from $\sigma_j^a \cup \sigma_j^b$. Let e^* be the heaviest edge in C . Our goal is now to bound $w(C)$, and it will be helpful to separately count the contribution of the spanning cycle and non-spanning-cycle edges.

Non-Spanning-Cycle Edges: Since π^a, π^b are bucket-safe k -paths, together they contain at most $2k$ non-spanning-cycle edges, each of weight $\leq w(e^*)$. So these contribute at most $2kw(e^*)$ to $w(C)$.

Spanning Cycle Edges: Since each subpath q_i^a, q_i^b is safe for bucket B_i , it contains at most $\varepsilon k 2^{i+1}$ spanning cycle edges. So the total number of spanning cycle edges in $\pi^a[s \rightsquigarrow y] \cup \pi^b[s \rightsquigarrow y]$ is at most

$$2 \cdot \sum_{i=1}^j \varepsilon k 2^{i+1} < \varepsilon k 2^{j+3}.$$

Finally, since e^* is in bucket B_j , we have $w(e^*) \geq 2^j$. Putting the parts together, we have

$$w^*(C) = \frac{w(C)}{w(e^*)} < \frac{2kw(e^*) + \varepsilon k 2^{j+3}}{w(e^*)} \leq 2k + 8\varepsilon k = (1 + 4\varepsilon) \cdot 2k.$$

This contradicts that H has weighted girth $> (1 + 4\varepsilon) \cdot 2k$, which completes the proof. \square

5.3 New Hiker Paths and the Counting Lemma

Much like the warmup proof, the medium and full counting lemmas generalize easily from the Moore bounds, but a new conceptual idea is needed for the weak counting lemma. For intuition, and for fun, we will make up an extension of the previous hiker puzzle that captures the gist of how our new weak counting lemma extends the one from the warmup proof.


ANOTHER HIKER PUZZLE

Graph National Park once again has n landmarks (nodes), m trails (undirected edges), and each trail has a difficulty rating (weight). It is possible for many trails to receive the same difficulty rating. The park has also installed a shuttle system, with a fleet of shuttles that drive in both the **forward** and **backward** direction in a loop (Hamiltonian/spanning cycle) around the landmarks. We have a Visitor's Pass that lets us ride the shuttle for up to $2t$ stops per day, which can be split across several trips if we like.

We are planning a multi-day backpacking trip to the park. At any time we may camp overnight at a landmark. We have the following constraints:

- To avoid boredom, we must hike at least one trail per day. We may also never *backtrack* a trail, meaning that after we hike a trail, our next action cannot be to immediately hike the same trail again in the reverse direction. (But we can otherwise repeat trails, e.g., by riding the shuttle back to the start of the trail and hiking it again.)
- Over the entire trip, we must hike all trails in nonincreasing order of difficulty. Additionally, each time we camp we wake up sore, and so the trails we hike on the following day must be *strictly* easier than all trails hiked on the previous day.

Prove: If we start at the right landmark, we can hike at least $2mt/n$ trails.




A valid hiker path, 4 trails long and split over 3 days, assuming $s_1 + s_2 \leq 2t$ and $s_3 \leq 2t$.

(this space intentionally left blank to avoid spoilers)

SOLUTION

As in our previous hiker puzzle, we assemble a squad of n hikers, and we start with one hiker standing at each landmark. We plan out hiker paths by the following process.

- First, we group the trails by difficulty rating. Let B_0 be all the trails that are tied for highest difficulty rating, let B_1 be the trails that are tied for second-highest difficulty rating, etc. We will plan hiker paths for each day i using the following steps:
 - **(Dawn)** At dawn, each hiker plans a path that consists of riding the shuttle t stops in the **forward** direction from their current position (but they do not yet walk their planned path). Notice an **invariant**: for each node v , and for each integer $1 \leq s \leq t$, there exists a hiker who reaches v as the endpoint of their s^{th} shuttle stop. As we modify paths in the following step, this invariant will remain.
 - **(Morning)** Next, the hikers collectively consider the trails $(u, v) \in B_i$ one at a time, in arbitrary order. For each integer $1 \leq s \leq t$, consider the two hikers h_u, h_v who currently plan to reach the nodes u, v as their s^{th} shuttle stop. We insert (u, v) into the paths of these two hikers right after their s^{th} shuttle stops. Following (u, v) , the two hikers swap paths: h_u takes h_v 's previous path from v onward, and h_v takes h_u 's previous path from u onward.
 - **(Afternoon)** Finally, in the afternoon, the hikers all travel the paths they have planned out. Then they end the day by riding the shuttle t steps **backward** and camping at the landmark they reach.



Each trail (u, v) is inserted into the paths of $2t$ hikers (perhaps counting with multiplicity). So our n hikers hike $2mt$ trails in total, and so there is a hiker who hikes at least $2mt/n$ trails. There is one last detail: we are supposed to ensure that this hiker hikes at least one trail per day. Notice that, if a hiker does not hike any trails on day i , then their planned path for the day consists of riding the shuttle t stops **forward**, riding the shuttle t stops **backward**, and camping back at the same landmark where their day began. The hiker can therefore skip this day entirely in their itinerary, and so they hike at least one trail in each un-skipped day.

We now give our weak counting lemma, which is more or less a repeat of the hiker path protocol from the above puzzle, with some adaptations and slightly more formal exposition of some steps.

Lemma 20 (Weak Counting Lemma). *If $w(H \setminus \mathcal{C}) \geq 4\epsilon^{-1}n$, then H contains a bucket-monotone extra-safe k' -path, for some $k' \geq k$.*

Proof. Start by placing a hiker at each node of H . We will define a protocol to generate paths for these hikers, which loops on three phases; each loop is called a **day**.

- **(Dawn)** On day i at dawn, each hiker plans a path that consists of walking $t_i := \lfloor \varepsilon k 2^{i-1} \rfloor$ steps **forward** on the spanning cycle (but they do not walk these edges yet). We notice that the following invariant holds:

Invariant: for each node v , for each integer $1 \leq s \leq t_i$ there exists a hiker who plans to reach v as the endpoint of the s^{th} spanning cycle edge that they walk, and also there exists a hiker who plans to end their journey at v .

- **(Morning)** On day i in the morning, the hikers collectively consider the edges $(u, v) \in B_i$ one at a time, in arbitrary order. For each integer $1 \leq s \leq t_i$, by the invariant, we may let h_u, h_v be the two hikers who respectively reach u, v at the end of their s^{th} **forward** spanning cycle steps. We insert the edge (u, v) into the planned paths of these two hikers immediately after their s^{th} spanning cycle steps. Following this edge, these hikers plan to swap paths; h_u takes the path suffix previously used by h_v from v onward, and h_v takes the path suffix previously used by h_u from u onward.

Note that each suffix swap preserves the invariant, since each $(s+1), \dots, t_i$ spanning cycle edge used by h_u becomes the $(s+1), \dots, t_i$ spanning cycle edge used by h_v (and vice versa), and the node at which h_u plans to end their journey becomes the node at which h_v plans to end their journey (and vice versa).

- **(Afternoon)** Finally, for each hiker h , let f_h be the number of contiguous **forward** spanning cycle edges that form a suffix of their planned path. We replace this suffix with $t_i - f_h$ **backward** steps on the spanning cycle. (We can think of this step as adding a suffix of t_i **backward** steps on the spanning cycle to the end of each hiker's path, and then repeatedly canceling adjacent **forward** and **backward** steps that backtrack each other.)

Each hiker then walks their planned path, and ends the day.

For correctness of this hiker protocol, we note that by the invariant, it remains true that one hiker stands at each node at the end of the day and thus the invariant holds again in the following day at dawn. This construction implies that the path hiked by each hiker on day i is extra-safe for B_i . Thus, each hiker's overall journey (concatenating their path across each day) forms a bucket-monotone extra-safe path. The number of hikers who traverse an edge $(u, v) \in B_i$ is

$$2t = 2\lfloor \varepsilon k 2^{i-1} \rfloor \geq \frac{\varepsilon k 2^{i+1}}{4} \geq \varepsilon k \cdot \frac{w(u, v)}{4}.$$

Summing over all edges in $H \setminus \mathcal{C}$, the total number of non-spanning-cycle edges traversed by our n hikers is therefore at least

$$\sum_{(u, v) \in E(H \setminus \mathcal{C})} \varepsilon k \cdot \frac{w(u, v)}{4} = \varepsilon k \cdot \frac{w(H \setminus \mathcal{C})}{4} \geq kn.$$

So there exists a hiker who hikes a bucket-monotone extra-safe k' -path, for some $k' \geq k$. \square

The medium and full counting lemmas can now be bootstrapped from the weak counting lemma by the usual method, with a few minor extra details.

Lemma 21 (Medium Counting Lemma). *H contains at least $\Theta(\varepsilon) \cdot (w(H \setminus \mathcal{C}) - 4\varepsilon^{-1}n)$ bucket-monotone safe k -paths.*

Proof. Repeat the following process until no longer possible. Find a bucket-monotone extra-safe k' -path π , for some $k' \geq k$. Let e_1, \dots, e_k be the first k non-spanning-cycle edges used by π . Let B_i be the bucket that contains e_1 , and let B_j be the bucket that contains e_k . Thus, omitting empty paths at the beginning and end, we may write the decomposition of π into bucket-extra-safe paths as $\pi = q_i \circ \dots \circ q_j$, where each subpath q_x is extra-safe for bucket B_x . We can use π to record a family of bucket-monotone (not-necessarily-extra-)safe k -paths by the following process:

- Truncate π immediately after e_k . After this truncation, the last path q_j in the decomposition might not end with a suffix of the appropriate number of **backward** edges, and so we add the appropriate number of **backward** edges as a suffix to π to restore the fact that q_j is extra-safe for B_j .

- For each integer $0 \leq s \leq \lfloor \varepsilon k 2^{i-1} \rfloor$, notice that we can generate a bucket-monotone safe k -path by modifying π by adding a prefix of s additional **forward** spanning cycle edges, and a suffix of s additional **backward** spanning cycle edges, to each nonempty path q_x in the decomposition of π .

We record $\lfloor \varepsilon k 2^{i-1} \rfloor = \Theta(\varepsilon k w(e_1))$ bucket-monotone safe k -paths in this way. We then delete e_1 to ensure that we do not re-record these paths in a future round. By the weak counting lemma, we can repeat this process at least until $w(H \setminus \mathcal{C}) < 4\varepsilon^{-1}n$. So we record at least $\Theta(\varepsilon) \cdot (w(H \setminus \mathcal{C}) - 4\varepsilon^{-1}n)$ paths in total. \square

The following claim will help with a technical detail in the full counting lemma:

Claim 22. *Every bucket-monotone safe k -path π uses k distinct non-spanning-cycle edges.*

Proof. Suppose for contradiction that there is a repeated non-spanning-cycle edge (u, v) in π . Let B_i be the bucket that contains (u, v) , and let q_i be the subpath of π that is safe for B_i . Notice that q_i must contain both copies of (u, v) , and so it has a cycle C as a subpath between these copies. Let e^* be the heaviest edge in C . We may bound the cycle weight $w(C)$ as follows. First, there are at most k non-spanning-cycle edges in C , and so the total contribution of these edges to $w(C)$ is at most

$$\sum_{(u,v) \text{ non-sp-cyc edge in } C} w(u, v) \leq k \cdot w(e^*).$$

Meanwhile, the number of spanning cycle edges in C is at most $\varepsilon k 2^i$. Putting these together, we have

$$w^*(C) \leq \frac{w(C)}{w(e^*)} \leq \frac{k \cdot w(e^*) + \varepsilon k 2^i}{w(e^*)} \leq k + 2\varepsilon k = (1 + 2\varepsilon)k.$$

This contradicts the weighted girth of H . \square

Lemma 23 (Full Counting Lemma). *Let $d := w(H \setminus \mathcal{C})/n$. If $d \geq 5\varepsilon^{-1}$, then H contains at least $n \cdot \Omega(\varepsilon d)^k$ bucket-monotone safe k -paths.*

Proof. Let H' be a random edge-subgraph of H obtained by keeping the spanning cycle \mathcal{C} deterministically, and keeping each non-spanning-cycle edge independently with probability $5\varepsilon^{-1}/d$. Let p, p' be the number of bucket-monotone safe k -paths in H, H' respectively.

On one hand, by Claim 22, every such path π uses k *distinct* non-spanning-cycle edges. Thus, the probability that π survives in H' is $\Theta(\varepsilon^{-1}/d)^k$, and so

$$\mathbb{E}[p'] = p \cdot \Theta\left(\frac{\varepsilon^{-1}}{d}\right)^k.$$

On the other hand, we have

$$\begin{aligned} \mathbb{E}[p'] &\geq \mathbb{E}[\Theta(\varepsilon) \cdot (w(H' \setminus \mathcal{C}) - 4\varepsilon^{-1}n)] && \text{Medium Counting Lemma} \\ &= \Theta(\varepsilon) \cdot (\mathbb{E}[w(H' \setminus \mathcal{C})] - 4\varepsilon^{-1}n) \\ &= \Theta(\varepsilon) \cdot \left(w(H \setminus \mathcal{C}) \cdot \frac{5\varepsilon^{-1}}{d} - 4\varepsilon^{-1}n \right) \\ &= \Theta(\varepsilon) \cdot (5\varepsilon^{-1}n - 4\varepsilon^{-1}n) \\ &= \Theta(\varepsilon) \cdot (\varepsilon^{-1}n) \\ &= \Theta(n). \end{aligned}$$

Comparing the two previous bounds on $\mathbb{E}[p']$, we get

$$\Theta(n) \leq p \cdot \Theta\left(\frac{\varepsilon^{-1}}{d}\right)^k.$$

Rearranging now gives our desired inequality of $p \geq n \cdot \Theta(\varepsilon d)^k$. \square

We now complete the proof of Theorem 17 in the usual way. Let $d := w(H \setminus \mathcal{C})/n$. If $d < 5\varepsilon^{-1}$, then we have $w(H) = O(\varepsilon^{-1}n)$ and we are done. Otherwise, by the dispersion and full counting lemmas, the number of bucket-monotone safe k -paths in H is at least $n \cdot \Omega(\varepsilon d)^k$, and at most $O(n^2)$. So we have

$$\begin{aligned} n \cdot \Omega(\varepsilon d)^k &\leq O(n^2) \\ d &\leq O\left(\varepsilon^{-1}n^{1/k}\right) \end{aligned}$$

and thus $w(H) = w(H \setminus \mathcal{C}) + w(\mathcal{C}) = nd + n = O(\varepsilon^{-1}n^{1+1/k})$.

Acknowledgments

I am grateful to Michael Elkin for helpful technical comments on an earlier draft of this paper.

References

- [1] Reyhan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020. 1
- [2] Stephen Alstrup, Søren Dahlgaard, Arnold Filtser, Morten Stöckel, and Christian Wulff-Nilsen. Constructing light spanners deterministically in near-linear time. *Theoretical Computer Science*, 907:82–112, 2022. 2
- [3] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993. 1, 2, 3
- [4] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Partially optimal edge fault-tolerant spanners. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3272–3286. SIAM, 2022. 2
- [5] Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen. Minor-free graphs have light spanners. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 767–778. IEEE, 2017. 2
- [6] Glencora Borradaile, Hung Le, and Christian Wulff-Nilsen. Greedy spanners are optimal in doubling metrics. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2371–2379. SIAM, 2019. 2
- [7] Barun Chandra, Gautam Das, Giri Narasimhan, and José Soares. New sparseness results on graph spanners. In *Proceedings of the eighth annual symposium on Computational geometry*, pages 192–201. ACM, 1992. 2, 7
- [8] Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 883–892. Society for Industrial and Applied Mathematics, 2016. 1, 2, 4
- [9] Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 442–452. Springer, 2014. 2, 4, 7
- [10] Paul Erdős. Extremal problems in graph theory. In *Proceedings of the Symposium on Theory of Graphs and its Applications*, page 2936, 1963. 1, 2
- [11] Arnold Filtser and Shay Solomon. The greedy spanner is existentially optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 9–17, 2016. 2
- [12] Hung Le and Shay Solomon. A unified framework for light spanners. In *Proceedings of the 55th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*. ACM, 2023. 2, 4

- [13] David Peleg and Jeffrey Ullman. An optimal synchronizer for the hypercube. *SIAM Journal on Computing (SICOMP)*, 18(4):740—747, 1989. [1](#)
- [14] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)*, 36(3):510–530, 1989. [1](#)