Traversing a set of points with a minimum number of turns

Sergey Bereg* Prosenjit Bose† Adrian Dumitrescu ‡ Ferran Hurtado § Pavel Valtr¶

December 5, 2008

Abstract

Given a finite set of points S in $\mathbb{R}^d$, consider visiting the points in S with a polygonal path which makes a minimum number of turns, or equivalently, has the the minimum number of segments (links). We call this minimization problem the *minimum link spanning path* problem. This natural problem has appeared several times in the literature under different variants. The simplest one is that in which the allowed paths are axis-aligned. Let $L(S)$ be the minimum number of links of an axis-aligned path for S and let G_n^d be a $n \times \cdots \times n$ grid in $\mathbb{Z}^d$. Kranakis *et al.* [10] showed that $L(G_n^2) = 2n - 1$ and $\frac{4}{3}n^2 - O(n) \leq L(G_n^3) \leq \frac{3}{2}n^2 + O(n)$ and conjectured that, for all $d \geq 3$, $L(G_n^d) = \frac{d}{d-1}n^{d-1} \pm O(n^{d-2})$. We prove the conjecture for $d = 3$ by showing the lower bound for $L(G_n^3)$. For $d = 4$, we prove that $L(G_n^4) = \frac{4}{3}n^3 \pm O(n^{5/2})$.

For general d, we give new estimates on $L(G_n^d)$ that are very close to the conjectured value. The new lower bound of $(1 + \frac{1}{d})n^{d-1} - O(n^{d-2})$ improves previous result by Collins and Moret [5], while the new upper bound of $(1 + \frac{1}{d-1})n^{d-1} + O(n^{d-3/2})$ differs from the conjectured value only in the lower order terms.

For arbitrary point sets, we give an exact bound on the minimum number of links needed in an axis-aligned path traversing any planar n-point set. We obtain similar tight estimates (within 1) in any number of dimensions d. For the general problem of traversing an arbitrary set of points in $\mathbb{R}^d$ with an axis-aligned spanning path having a minimum number of links, we present a constant ratio (depending on the dimension d) approximation algorithm.

Keywords: Computational geometry, minimum link spanning path, approximation algorithms.

*Dept. of Computer Science, University of Texas at Dallas, Richardson, TX 75083, USA. Email: besp@utdallas.edu

†School of Computer Science, Carleton University, Ottawa, Ontario, Canada, K1S 5B6. Email: jit@scs.carleton.ca

‡Dept. of Computer Science, University of Wisconsin–Milwaukee, WI 53211, USA; Partially supported by NSF CAREER grant CCF-0444188. Email: ad@cs.uwm.edu

§Dept. de Matem. Aplicada II, Universitat Politecnica de Catalunya, Barcelona, Spain; Partially supported by projects MECMTM2006-01267 and Gen. Cat. 2005SGR00692. Email: ferran.hurtado@upc.edu

¶Department of Applied Mathematics and Institute for Theoretical Computer Science, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic. Email: valtr@kam.mff.cuni.cz Work by P. Valtr was partially supported by the project 1M0545 of the Ministry of Education of the Czech Republic.

1 Introduction

Let S be a finite set of points in $\mathbb{R}^d$. A polygonal path P is *axis-aligned* if every link of P is parallel to one of the coordinate axes. An axis-aligned *spanning path* of S is an axis-aligned polygonal path that passes through all the points of S. We only consider axis-aligned paths in this paper. The *link length* — or just *length* for short — of a spanning path P, denoted $L(P)$, is the number of line segments (called *links*) of the path. In many cases, it is desirable to have a path traversing a given point set, which has a small (possibly minimum) number of links. We call this minimization problem the *minimum link spanning path* problem. This problem was introduced and first studied by Kranakis et al. [10]. As noted there, without the restriction to axis-aligned paths, the problem of computing a spanning path with the minimum number of links is NP-complete, since the NP-complete problem of *edge embedding on a grid* [8] can be reduced to it. An alternate proof of NP-completeness for the minimum link spanning path problem was given by Arkin et al. [3] in the context of minimum link watchman tours. It is not known (to the best of our knowledge) whether the problem remains NP-complete for axis-aligned paths, although we suspect it does since Hassin and Megiddo [9] showed that it is NP-complete to determine the minimum number of axis-parallel lines required to cover a set of points in 3 (or higher) dimensions. Recently, Gaur and Bhattacharya [2] found a $(d-1)$-approximation algorithm for the latter problem.

Indeed, the problem we study appears to have connections with the problem of hitting objects (points, in particular) by straight lines as studied by Megiddo and Tamir [12] as well as Hassin and Megiddo [9]; the problem of hitting a set of points with straight lines has applications in medicine (radiotherapy) as well as military applications (in destroying a set of targets on the ground with a bomber, see [9] for details). Some of these applications carry over to the spanning path problem. One advantage of studying axis-aligned spanning paths with a minimum number of turns instead of arbitrary (polygonal) paths under the same optimization criterion is that by definition, axis-aligned paths provide large turning angles, which are often desirable. It is worth mentioning that for arbitrary spanning paths, and every $n \geq 4$, there exist point sets of size n that require a turning angle as low as 30° as shown by Fekete and Woeginger [7]. It is conjectured [7] that no point set requires a smaller turning angle; see [7] for other challenging angle problems for paths and tours.

The *link length of* S, denoted $L(S)$, is the minimum number of links in an axis-aligned spanning path of S. Alternatively, we can think of the link length of a path as being one plus the number of 90°-turns of the path. Consequently a 180°-turn is counted as two turns: if a path goes (say in $\mathbb{R}^3$) from (x, y, z) to (x, y, z') and then back to (x, y, z'') (assuming that z' is not between z and z''), we say that there is a link of length zero at (x, y, z') moving in the xy-plane [6]. We note that all our results hold in the case when 180°-turns are counted as single turns. Note that these definitions do not require that the spanning path avoid self-intersections, nor that the path stay inside the grid. The *size* of a link ℓ of a path P is the number of points in S covered by ℓ and not covered by any previous link of P, cf. [6]. In a link from (x, y, z_1) to (x, y, z_2), z is the *moving* coordinate, while x and y are the *constant* coordinates. Analogously, in $\mathbb{R}^d$, a link has one moving coordinate and $d-1$ constant coordinates.

A wide variety of *covering tour* or *covering path* problems have been recently investigated by Arkin et al. [1], where one has to find a polygonal tour for a *cutter* so that it sweeps out (*mills*) a specified region, in order to minimize a cost that depends mainly on the number of turns. These problems arise naturally in manufacturing applications, automatic tool path generation, automatic inspection systems, robotic exploration, and other areas. Many of these milling problems (in both — tour and path formulations) are NP-hard even restricted to orthogonal polygons and axis-aligned motion of the cutting tool: for instance, discrete milling, orthogonal milling, and integral orthogonal milling fall in this category, see [1] for details. In our paper the region to be traversed is: (i) a cube or a box in $\mathbb{R}^d$, or (ii) an arbitrary set of points — and we look at both variants in higher dimensions. In the first part, we study what are the best ways to traverse a cube or a box in $\mathbb{R}^d$, in minimizing the number of turns of an axis-aligned spanning path. In the final part we study the same question when traversing arbitrary sets of points in $\mathbb{R}^d$.

We recently learned of the work by Stein and Wagner [13, 14] on the Minimum Bends Traveling Salesman problem for planar point sets, in which they study several variants of these problems and obtain various algorithms and complexity results. Although their problem is slightly different from ours since they study Hamiltonian cycles whereas we study Hamiltonian paths, their proof techniques and algorithms clearly provide a proof of Theorem 5 and the two-dimensional part of Theorem 7. We were unaware of these results when we initially wrote this article and fully acknowledge that Stein and Wagner's result [13, 14] predates our work. To the best of our knowledge, our extensions of these algorithms to higher dimensions remain new contributions. We decided to keep our descriptions of the two-dimensional versions of our algorithms for finding Hamiltonian paths in Sections 5.1 and 5.3, as our algorithms are slightly different and provide the reader with the necessary intuition behind the generalization of our algorithms to higher dimensions.

Let $G_{n_1,\ldots,n_d}$ denote the grid points (points with integer coordinates) in $[1, n_1] \times \cdots \times [1, n_d]$. For simplicity we write G_n^d for $G_{n,\ldots,n} \subset \mathbb{Z}^d$. For the square grid in the plane Kranakis et al. [10] have obtained an exact bound: $L(G_n^2) = 2n - 1$. We provide an alternative proof of this result that extends to higher dimensions and yields better bounds.

For any d, trivially we have $L(G_n^d) \geq n^{d-1}$, since there are n^d points and each link in the path can cover at most n points [5, 10]. So the main difficulty lies in establishing the constant c_d in front of the leading term. In particular, it is not clear apriori whether c_d tends to 1 as d tends to infinity. However, this has been shown to be true by Kranakis et al. [10], see equations (1 and 2) below. The problem of estimating $L(G_n^d)$ (for fixed d, with n going to infinity) appears in the recent collection of research problems by Braß, Moser, and Pach [4], and also in the survey article by Maheshwari et al. [11].

The best known upper bound for the cube [10] (see also [5]) is $L(G_n^3) \leq \frac{3}{2}n^2 + O(n)$. It has been conjectured [10] that this bound is tight. The first non-trivial lower bound in three dimensions was proved by Kranakis et al. [10], namely that $L(G_n^3) \geq 1.023\ldots n^2$, via an involved calculation. This was further improved by Collins and Moret [5] to $L(G_n^3) \geq \frac{7}{6}n^2 - O(n)$. The best result up to this date, is the following recent one due to Collins [6] which makes use of a proof technique that we refer to as the *nested cubes argument*: $L(G_n^3) \geq \frac{4}{3}n^2 - O(n)$. Here we close this gap and thereby prove the conjecture, by using a different proof technique that we call the *hyperplane argument*. We summarize our results below.

Theorem 1 $\frac{3}{2}n^2 - O(n) \leq L(G_n^3) \leq \frac{3}{2}n^2 + O(n)$.

More generally, we obtain almost tight bounds on the minimum number of links needed to traverse an $a \times b \times c$ grid:

Theorem 2 *Let $a \leq b \leq c$ be three positive integers. If $c < a + b - 1$ then $L(G_{a,b,c}) = ab + bc + ac - (a^2 + b^2 + c^2)/2 \pm O(c)$. If $c \geq a + b - 1$ then $L(G_{a,b,c}) = 2ab \pm O(b)$.*

The hyperplane argument gives also a tight lower bound for the cube in $\mathbb{R}^4$, so the conjecture is (essentially) verified also for $d = 4$:

Theorem 3

$$\frac{4}{3}n^3 - O(n^2) \leq L(G_n^4) \leq \frac{4}{3}n^3 + O(n^{5/2}).$$

Kranakis et al. [10] prove the following bound: For all $0 < \epsilon < 1$ the following inequalities holds for any sufficiently large $d = d(\epsilon)$

$$\frac{L(G_n^d)}{n^{d-1}} \geq 1 + \frac{1}{2}\left(1 - \exp\left(\frac{-1}{d(d-1)}\right)\right) \tag{1}$$

and

$$\frac{L(G_n^d)}{n^{d-1}} \leq 1 + \frac{1}{2(d-3)^{1-\epsilon}} + \exp\left(-(d-3)^{\epsilon}\right). \tag{2}$$

As noted in [5], the upper bound implies the existence of an absolute constant $C > 0$, such that

$$L(G_n^d) \leq \left(1 + \frac{C + \log d}{2(d-3)}\right) n^{d-1} + O(n^{d-2}), \tag{3}$$

The lower bound was further improved by Collins and Moret [5]:

$$L(G_n^d) \geq \left(1 + \frac{1}{2d}\right) n^{d-1} - O(n^{d-2}). \tag{4}$$

Here we give bounds that are very close to the conjectured value of $(1 + \frac{1}{d-1})n^{d-1} + O(n^{d-2})$ [10]:

Theorem 4 *If $d \geq 4$ then:* $\left(1 + \frac{1}{d}\right) n^{d-1} - O(n^{d-2}) \leq L(G_n^d) \leq \left(1 + \frac{1}{d-1}\right) n^{d-1} + O(n^{d-3/2})$.

In the final part of our paper (Section 5), we deal with arbitrary point sets and provide an exact bound on the minimum number of links needed for an axis-aligned path to traverse any planar point set of size n. Let $f(n)$ be the minimum positive integer N such that every set of n points in the plane admits an axis-aligned spanning path with at most N links. We show the following:

Theorem 5

$$f(n) = \begin{cases} n, & \text{if } n = 1, 2, 3 \\ n + 1, & \text{if } n \geq 4. \end{cases}$$

We also obtain similar tight estimates (within 1) in any number of dimensions d. Let $f_d(n)$ be the minimum positive integer N such that every set of n points in $\mathbb{R}^d$ admits an axis-aligned spanning path with at most N links. Note that $f(n) = f_2(n)$.

Theorem 6 *Let $d \geq 3$. If $n \leq 3$ then $f_d(n) = (d-1)n + 2 - d$. For $n \geq 4$ we have $(d-1)n + 2 - d \leq f_d(n) \leq (d-1)n + 3 - d$.*

Finally, we present a simple constant ratio (depending on the dimension d) approximation algorithm for the problem of traversing an arbitrary set of points in $\mathbb{R}^d$ by an axis-aligned spanning path with a minimum number of links.

Theorem 7 *There exists a ratio 2 polynomial-time approximation algorithm for the problem of traversing an arbitrary set of points in the plane by an axis-aligned spanning path with a minimum number of links. Furthermore, in $d \geq 3$ dimensions, there is an approximation algorithm with ratio d^2 for this problem.*

2 Traversing the square

In this section we present our *hyperplane argument*. We first use this argument to obtain the planar result of [10] for traversing the square. We then use the hyperplane argument in Section 3 and Section 4 to obtain a better bound for $d = 3$. The key intuition behind the hyperplane argument for bounding paths on the general d-dimensional cube is the following: we find a small but carefully selected subset of the cube whose traversal by every axis-aligned spanning path requires (roughly) as many links as the traversal of the whole cube. The points in the subset are the grid points contained in a finite set of suitably chosen hyperplanes intersecting the cube. Each link in the spanning path covers at most a constant number of points (which depends on the dimension d). The ability to apply this technique in order to obtain lower bounds increases in complexity as the dimension increases. In order to highlight the salient points behind our technique, we review it in its simplest form where the input is a 2-dimensional square.

By the planar result of [10], the link length of a square grid G_n^2 is $L(G_n^2) = 2n - 1$. This means that every axis-aligned spanning path of the $n \times n$ grid contains at least $2n - 1$ links, and this bound can be achieved. We show a stronger claim, namely that a certain *small* subset X_n of G_n^2 requires the same number of links in every spanning path. Let $X_n = \{(i, i), (i, n - i + 1) \mid i = 1, \ldots, n\}$, see Figure 1. Note that $|X_n| = 2n - 1$ for n odd, and $|X_n| = 2n$ for n even.

Incidentally we remark that there are two completely different ways to achieve the optimal value for G_n^2, which are illustrated in Figure 1(c,d): a spiral path where the size of successive links gradually decreases from n to 1, and a zig-zag path that alternates between long and short links. Both of these traversal modes are used in the best known spanning path of the 3-dimensional cube, see Section 3.

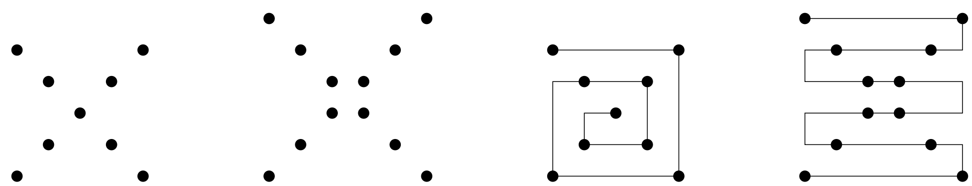

Figure 1: The hyperplane argument in two dimensions: two subsets of points of G_n^2 (for odd n and even n) which require spanning paths with many links: (a) X_5; (b) X_6; (c) An optimal spanning path of X_5 whose corresponding sequence $(v_i)_{i\geq 1}$ is $2, 1, 1, 0, 2, 1, 1, 0, 1$. (d) An optimal spanning path of X_6 whose corresponding sequence $(v_i)_{i\geq 1}$ is $2, 0, 2, 0, 2, 0, 2, 0, 2, 1, 1$.

We now show that the link length of the $n \times n$ square is identical to the link length of X_n.

Lemma 8 $L(X_n) = 2n - 1$.

Proof. Clearly $L(X_n) \leq 2n - 1$ since $X_n \subset G_n^2$, and $L(X_n) \leq L(G_n^2) \leq 2n - 1$. We now show that $L(X_n) \geq 2n - 1$: consider the sequence v_i, for $i \geq 1$, where v_i is the size of the i-th link in the path, namely the number of new points visited by the path using that link. By construction of the set X_n, we have $v_i \in \{0, 1, 2\}$ for any i.

We first consider the case where n is even, and let $v_i = v_j = 2$ be two consecutive 2's in the sequence, where $i < j$. We claim that there must exist a zero in between them, that is, there is an element $v_k = 0$ in the sequence for some $i < k < j$. This is can be easily seen from the construction of X_n in Figure 1. Consider the 4 regions formed by the two lines containing the points of X_n. Let us label these 4 regions as R_1, R_2, R_3, R_4 in clockwise order, and look at the possible ways a path can cross from one region to another. For a segment in the path to have its corresponding element in the sequence to be of value 2, its endpoints must be in opposite regions, that is in R_1 and R_3 (or R_2 and R_4 resp.), depending on whether the segment is horizontal or vertical. A simple inductive argument shows that any path joining two points properly contained in a given region must have at least one segment having both endpoints in one region, which implies a value of zero. One can now charge the second element of 2 (v_j) to the the zero element v_k, and the claim follows from the equality $|X_n| = 2n$.

In the case where n is odd, notice that the above claim still holds. In addition, we need the following fact. Let ℓ be a link of the path which visits the center grid-point $(\frac{n+1}{2}, \frac{n+1}{2})$ of X_n: this gives an element $v_j = 1$ in the sequence. If there exists an element of 2 preceding this 1 in the sequence, that is, $v_i = 2$ for some $i < j$, then there must be a zero element in between them, i.e., $v_k = 0$ for some $i < k < j$. Similarly, if there exists an element of 2 which comes after this 1 in the sequence, that is, $v_i = 2$ for some $j < i$, then there must be a zero element in between them, i.e., $v_k = 0$ for some $j < k < i$. The claim now follows from the equality $|X_n| = 2n - 1$, which concludes the proof of the theorem. (Alternatively, for n odd, it suffices to take the central point with multiplicity 2 and then the proof works in the same way as for n even.) $\square$

3 Traversing the cube

The best known upper bound for traversing the cube $L(G_n^3) \leq \frac{3}{2}n^2 + O(n)$, comes from the following spanning path construction given in [10], see [10, 6] for an illustration: Fix a direction, say the z-axis, and assume for simplicity that n is even. The path spirals around the outer part of the cube in each xy-plane until a box of size $\frac{n}{2} \times \frac{n}{2} \times n$ remains in the center. The grid points in this remaining box are then covered by moving up and down along the vertical (z) axis. Note that the average link size decreases from n to $n/2$ in the first part, while staying at about $n/2$ in the second part; the average link size for the whole path is about $2n/3$, and the number of links in the path is $3n^2/2 + O(n)$, see [10] for details. We now show the optimality of this path (modulo lower order terms), as conjectured by Kranakis et al. [10].

3.1 Proof of Theorem 1

We extend our approach for the plane (the hyperplane argument) to three-dimensions. For simplicity we assume that n is odd, i.e. $n = 2k + 1$. Consider the cube $Q_n = [-k, k]^3 \cap \mathbb{Z}^3$ and the four planes $\Pi_1 : x + y + z = 0$, $\Pi_2 : x = y + z$, $\Pi_3 : y = x + z$, $\Pi_4 : z = x + y$. Let $\Pi = \cup_{i=1}^{4} \Pi_i$ and $X := X_n = \Pi \cap Q_n$. We refer to the four planes Π_i as the planes of X.

It suffices to prove:

Lemma 9 $L(X) \geq \frac{3}{2}n^2 - O(n)$.

Proof. First, we show that for every $i \in \{1, 2, 3, 4\}$, $\Pi_i \cap Q_n$ contains $3n^2/4 + 1/4$ points. By symmetry it suffices to count the points in $\Pi_1 \cap Q_n$. We project these points onto the plane $z = 0$. Note that every point has a unique projection. The projection of $\Pi_1 \cap [-k, k]^3$ is $\{(x, y) \mid -k \leq x, y \leq k, -k \leq x + y \leq k\}$, see Figure 2. It is easily verified that $|\Pi_1 \cap Q_n| = (2k + 1)^2 - k(k + 1) = 3n^2/4 + 1/4$.

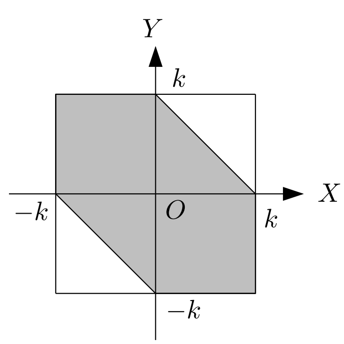

Figure 2: Counting points of $\Pi_1 \cap Q_n$. The shaded area is the projection of $\Pi_1 \cap [-k, k]^3$ onto the xy-plane.

The set X_n contains $4|\Pi_1 \cap Q_n| = 3n^2 - O(n)$ points since any two planes Π_i and Π_j intersect in at most n grid points. Let P be a spanning path of X_n. It suffices to prove that the average size of a link in P is at most $2 - O(1/n)$.

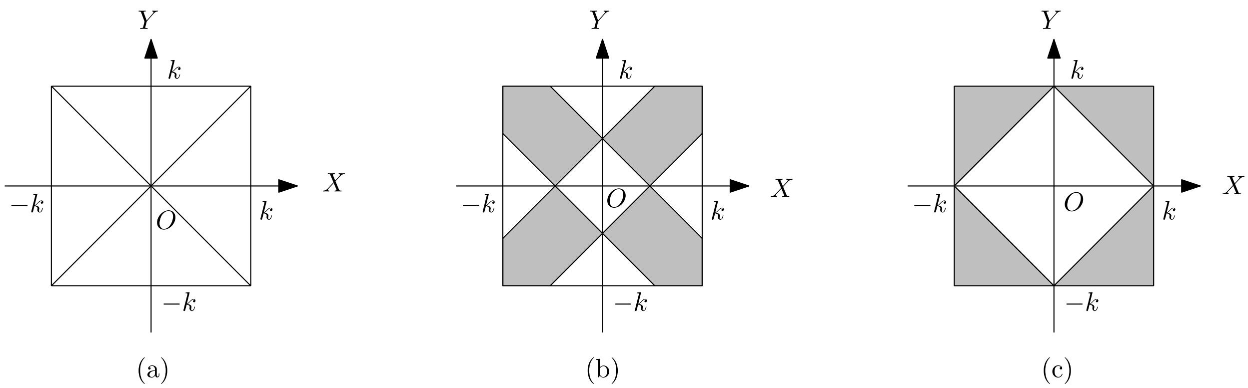

Figure 3: The slices of the cube $[-k,k]^3$ by the plane (a) $z=0$, (b) $z=c, 0<c<k$, and (c) $z=k$.

Note that a link of P can intersect all four planes of X, for example a line $x=z=c$ where $0<c<k/3$. We consider the cross-cut of the cube $[-k,k]^3$ by a plane $z=c, |c| \leq k$, see Figure 3. The planes Π_i intersect it in four line segments. We color the regions containing the vertices of the square $[-k,k]^2$ in grey and the other regions in white. If a link has endpoints in a grey area then it crosses at most two planes. We consider similar cross-cuts of the cube for the other two coordinate axes and employ the same coloring scheme; by symmetry, all colorings (in all three axes directions) look the same!

We decompose the path into subpaths, each lying in a square-shaped cross-cut of the cube and also analyze what happens when the path moves from one plane to another orthogonal plane. For each subpath we use a charging scheme similar to the one used in the planar case (proof of Lemma 8). The corresponding v-sequence has elements in $\{0,1,2,3,4\}$. We obtain another sequence v' derived from v by replacing elements equal to 1 or 3 with elements equal to 0, 2 or 4 as follows. When a link of the path crosses from an unshaded region to a shaded region and crosses 3 (resp. 1) of the four planes in X, we charge 2 to each crossing into a shaded region and 0 to each crossing out of the shaded region. Finally, by dividing each element in v' by 2 we obtain a sequence v'', whose elements are in $\{0,1,2\}$ as in the planar case.

We also make the observation that when, say a vertical link crossing all four planes in X is the last one in a vertical plane, the link lands in the center unshaded region of the next horizontal plane (cross-cut section) where the path lies. The argument for the planar case in combination with the above observation yields that $L(X) \geq (|X| - O(1))/2 = 3n^2/2 - O(n)$, and thus completes the proof. □

Next we show how to generalize the above arguments to obtain almost tight bounds on the minimum number of links needed to traverse an $a \times b \times c$ grid in three-space.

3.2 Proof of Theorem 2

The lower bound is obtained by the hyperplane argument. For simplicity we assume that a, b, and c are odd, that is $a = 2k_a+1, b = 2k_b+1$, and $c = 2k_c+1$. We consider the box $B = [-k_a, k_a] \times [-k_b, k_b] \times [-k_c, k_c]$ and the same four planes as for the 3-cube. They all intersect in the center of the box. In the case where $c \leq a+b-1$, the number of grid points in Π_1 is

$$n_1 = ab - (k_a + k_b - k_c)^2 \pm O(c),$$

see Figure 4. By symmetry, the other planes contain the same number of grid points. The same argument as above shows that the average link size is $2 - O(1/n)$. Therefore, the lower bound follows from

$$2n_1 = ab + bc + ac - (a^2 + b^2 + c^2)/2 \pm O(c).$$

In the case where $c > a + b - 1$, we truncate the input by applying the above argument to the $a \times b \times b$ grid to obtain that $L(G_{a,b,c}) = 2ab \pm O(b)$.

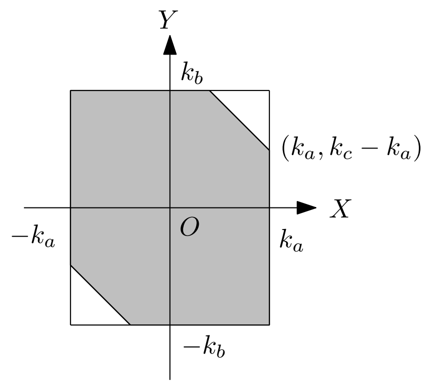

Figure 4: Counting points of $\Pi_1 \cap B$. The shaded area is the projection of $\Pi_1 \cap B$ onto the xy-plane.

The upper bound uses a construction similar to the path construction for the cube. If $c \geq a + b - 1$, we take the ab vertical segments of (Euclidean) length $c - 1$, each covering c points, and connect them by $ab - 1$ short segments. The resulting path has $2ab - 1$ links. If $c < a + b - 1$, we first take a spiral in each of the c rectangles of size $a \times b$, until rectangles of size about $(a - b + c - 2)/2 \times (b - a + c - 2)/2$ remain. Put $a' = (a - b + c - 2)/2$ and $b' = (b - a + c - 2)/2$ for the side lengths of the remaining rectangles. The number of complete rotations in each spiral is $x = (a + b - c + 2)/4$. It is verified that $a' + 2x = a$, $b' + 2x = b$, and $x \geq 1$, $a', b' \geq 0$. We connect the spirals by short segments and traverse the remaining box by vertical segments of (Euclidean) length $c - 1$, connected by short segments. The total number of links in the path is

$$\begin{aligned} & 4xc + 2a'b' + O(c) \\ = \; & c(a + b - c + 2) + \frac{(c-2)^2 - (a-b)^2}{2} + O(c) \\ = \; & ab + bc + ca - \frac{a^2 + b^2 + c^2}{2} + O(c), \end{aligned}$$

as required. □

In [6], Collins considered traversals of $G_{\gamma n,\gamma n,n}$, for $\gamma \in (0, 1)$, and showed that for $\gamma > 1/2$, an upper bound $L(G_{\gamma n,\gamma n,n}) \leq (2\gamma - 1/2)n^2$ can be obtained by a hybrid path spiraling in each horizontal plane until the size of each link gets down to $n/2$ and then switching to vertical mode. He conjectured that these paths are optimal. Using our tight bound for traversing the cube (Theorem 1) we can easily confirm this fact (modulo a linear term): it follows from the equality $4(1 - \gamma)/2 + (2\gamma - 1/2) = 3/2$ (or even directly, with no calculation). Embed the box $G_{\gamma n,\gamma n,n}$ into the cube G_n^3 and assume that there is a better way to traverse the box. Then using a hybrid path spiraling in each horizontal plane until the size of each link gets down to γn and then continuing with the supposedly better path for the box $G_{\gamma n,\gamma n,n}$, one would get a path traversing the cube having fewer than $3n^2/2$ links (modulo a linear term), which is a contradiction.

4 Traversing the cube in higher dimensions

4.1 $d = 4$: proof of Theorem 3

The upper bound is a special case of that in Theorem 4, whose proof appears in the next section. We prove the lower bound. Let the 4-cube be centered at the origin and consider eight hyperplanes with equations

$$x_1 \pm x_2 \pm x_3 \pm x_4 = 0.$$

The argument is essentially the same as for $d = 3$. Let $p_1 p_2 .... p_k$ be a path traversing the cube. For each i, let s_i be the number of hyperplanes intersecting the segment $p_i p_{i+1}$ and let t_i be the number of hyperplanes intersecting the halfline $p_i p_{i+1}$ (intersections in p_i are not counted, unless $i = 1$). Then $t_{i+1} \leq 4 + (t_i - s_i)$. Summing up all these inequalities and the "initial" inequality $t_1 \leq 8$ gives $s_1 + s_2 + \cdots + s_k \leq 4(k + 1)$. The number of links is bounded by $k \geq s/4 - 1$ where s is the total number of points in the hyperplanes.

The number of points on the hyperplane $x_1 + x_2 + x_3 + x_4 = 0$ is the same as the number of points in 3-cube $[-k, k]^3$ satisfying $-k \leq x_1 + x_2 + x_3 \leq k$. These points reside in the truncated cube which is obtained by removing two tetrahedra, see Fig. 5. The volume of the truncated cube is $V - 2v = V - V/3 = 2V/3$ where $V = (2k)^3$ is the volume of the cube and v is the volume of one tetrahedron. Therefore the hyperplane contains at least $2n^3/3 - O(n^2)$ points. Eight hyperplanes contain at least $16n^3/3 - O(n^2)$ points which require at least $4n^3/3 - O(n^2)$ links to traverse.

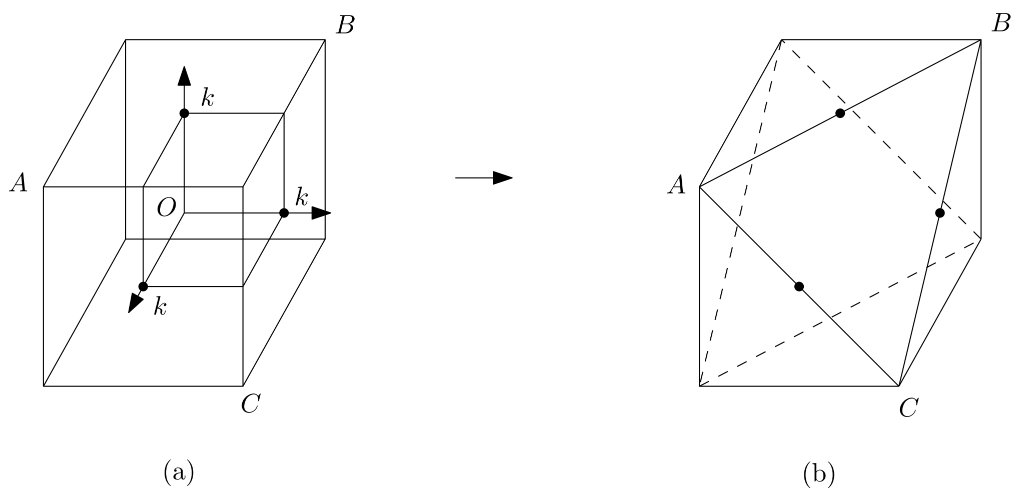

Figure 5: (a) The cube $[-k, k]^3$. (b) The volume between two planes $x_1 + x_2 + x_3 = \pm k$.

Remark. It is worth to note that the hyperplane argument is getting weaker in dimensions higher than 4 (quite rapidly with increasing dimension).

4.2 Proof of Theorem 4

4.2.1 Lower bound

We revisit the *nested cubes argument* introduced in [5, 6], and rederive the main results in these papers with a simplified and more concise reasoning. Section 10.2 (pp. 423) of [4] mentions an irreparable error in the proof of one of the results from [5]; our derivation is self-contained and does not make use of any of the previous results (lemmas) from [5, 6], except two basic observations to start with, which are given below.

The idea of the argument is that any spanning path must visit points close to the center of the cube (say in a smaller concentric cube) using multiple links; while doing so, each link visiting the smaller cube (or whose supporting line intersects the smaller cube) must be followed by a link significantly shorter than the cube size n, and thus forces more links in the path. We follow with the precise details.

For simplicity, suppose that n is odd, that is $n = 2k+1$, and that Q_n is centered at the origin, that is $Q_n = [-k, k]^d \cap \mathbb{Z}^d$ — the proof for n even is similar. Consider the sequence $Q_0, Q_1, \ldots, Q_k$ of $k+1$ cubes with sides $1, 3, 5, 7, \ldots, n$ nested one inside the other, and centered at the origin. We have that Q_i is the set of $(2i+1)^d$ grid points in the i-th cube, $i = 0, 1, \ldots, k$. Let $P = \ell_1, \ldots, \ell_m$ be a spanning path of G_n^d that consists of m links. The *distance* of a link $\ell \in P$ from the center of G_n^d is defined to be the maximum of the $d-1$ absolute values of the constant coordinates in ℓ. We partition the links of P into $k+1$ disjoint sets: $A_0, \ldots, A_k$, where A_i contains the links of the path P at distance i from the origin. Write $a_i = |A_i|$. Define the *distance label* function $\beta \colon \{1, 2, \ldots, m\} \to \{0, 1, \ldots, k\}$, where $\beta(h) = i$ if and only if $\ell_h \in A_i$.

Two important observations from [5, 6] are:

Observation 10 *If a link of P belongs to A_i, it must be followed immediately by a* short *link, i.e., one whose size is at most $k+i$.*

Proof. Let the extent of the moving coordinate in the link be $[x_1, x_2]$. The claim follows from the inequality $|x_1| \le i$. □

Observation 11 *Let $2 \le h \le m-1$. If $\beta(h-1) = i$ and $\beta(h+1) = j$, for some $i, j \in \{0, 1, \ldots, k\}$, then the size of ℓ_h is at most $i+j$.*

Proof. The extent of the moving coordinate in the link is $[x_1, x_2]$, where $|x_1| \le i$ and $|x_2| \le j$, whence $|x_2 - x_1| \le i + j$ (here one counts only new points visited by the link). □

Since the grid points of Q_i can be covered only by links in $A_0 \cup \ldots \cup A_i$, and each such link can cover at most $2i+1$ points of Q_i, we have the following system of $k+1$ inequalities:

$$\sum_{j=0}^{i} a_j \ge (2i+1)^{d-1}, \;\; i = 0, \ldots, k. \tag{5}$$

Consider now an optimal spanning path of Q_n with m links: $P = \ell_1, \ldots, \ell_m$, where $m = L(G_n^d)$. By using Observation 10 for each link in P, and the fact that the first link ℓ_1 covers at most n points, we have:

$$n + \sum_{i=0}^{k} a_i(k+i) \ge n^d. \tag{6}$$

We are interested in minimizing the length of the path $m = \sum_{i=0}^{k} a_i$ subject to (5) and (6). Write $S_i = \sum_{j=0}^{i} a_j$ for $i = 0, \ldots, k$, and $S = S_k$ (so $m = S$), and observe that

$$\sum_{i=0}^{k} i a_i = kS - \sum_{i=0}^{k-1} S_i. \tag{7}$$

The system (5) says that $S_i \ge (2i+1)^{d-1}$, for $i = 0, \ldots, k$. We substitute (7) in (6) and obtain

$$n + kS + \left(kS - \sum_{i=0}^{k-1} S_i \right) \ge n^d. \tag{8}$$

Using the well-known estimate (for positive integers s, t with t fixed)

$$\sum_{i=1}^{s} i^t = \frac{s^{t+1}}{t+1} + O(s^t),$$

we obtain a lower bound on $\sum_{i=0}^{k-1} S_i$:

$$\begin{aligned}\sum_{i=0}^{k-1} S_i &\geq \sum_{i=0}^{k-1} (2i)^{d-1} \\ &\geq 2^{d-1}\frac{k^d}{d} - O\left(n^{d-1}\right) \\ &= \frac{n^d}{2d} - O\left(n^{d-1}\right).\end{aligned}$$

We substitute this bound in (8) and derive

$$n + 2kS \geq \left(1 + \frac{1}{2d}\right) n^d - O\left(n^{d-1}\right), \tag{9}$$

which implies the result in [5]:

$$S \geq \left(1 + \frac{1}{2d}\right) n^{d-1} - O\left(n^{d-2}\right). \tag{10}$$

Using Observation 11 for each of the links $h = 2, \ldots, m-1$ (instead of Observation 10) yields a better bound. Inequality (6) can be replaced by the sharper inequality

$$2n + 2\sum_{i=0}^{k} i a_i \geq n^d, \tag{11}$$

since each link ℓ_h can contribute at most two terms equal to $\beta(h)$ to the sum above, and there are $a_{\beta(h)}$ such terms (and the first and last link cover each at most n points). Inequality (11) yields

$$\sum_{i=0}^{k} i a_i \geq \frac{n^d}{2} - n, \tag{12}$$

and further (doing the same substitutions as before) we get

$$kS - \sum_{i=0}^{k-1} S_i \geq \frac{n^d}{2} - n. \tag{13}$$

Finally,

$$kS \geq \frac{n^d}{2} + \frac{n^d}{2d} - O\left(n^{d-1}\right) \tag{14}$$

yields:

$$S \geq \left(1 + \frac{1}{d}\right) n^{d-1} - O(n^{d-2}). \tag{15}$$

This completes the proof of the lower bound in Theorem 4.

4.2.2 Upper bound

The path for traversing G_n^d is obtained by generalizing the planar *spiral* construction. Let $s := 2\lfloor\sqrt{n}/4\rfloor$, $[n] := \{1, 2, \ldots, n\}$ and $S := \{s/2+1, s/2+2, \ldots, n-s/2\}$. The grid G_n^d can be partitioned into n^{d-2} subgrids of size $n \times n$ as follows. For each point $b \in G_n^{d-2}$, we consider the subgrid

$$G(b) = \{x \in G_n^d \mid x_1 = b_1, \ldots, x_{d-2} = b_{d-2}\}.$$

It can be viewed as the 2-dimensional grid $\{x_{d-1}, x_d\} \in [n]^2$ of n^2 points and we cover its $n^2-(n-s)^2$ points by $2s$ links of the spiral as in Figure 1 (c) leaving $(n-s)^2$ uncovered points of the grid $\{x_{d-1}, x_d\} \in S^2$. We connect the spirals at their ends using $n^{d-2}-1$ short segments in a fashion corresponding to any fixed axis-parallel path of Euclidean length $n^{d-2}-1$ traversing the points of G_n^{d-2}. The set of points which remain to be covered is $[n]^{d-2} \times S^2$. The number of links in the spanning path so far is $n^{d-2}(2s+1)-1$.

For each point $b \in G_n^{d-4}$ and each point $c \in S^2$, we consider the subgrid $G(b,c) = \{x \in G_n^d \mid x_1 = b_1, \ldots, x_{d-4} = b_{d-4}, x_{d-1} = c_1, x_d = c_2\}$. It can be viewed as the 2-dimensional grid $\{x_{d-3}, x_{d-2}\} \in [n]^2$ of n^2 points and we cover its $n^2-(n-s)^2$ points by $2s$ links of the spiral as before, leaving $(n-s)^2$ uncovered points of the grid $\{x_{d-3}, x_{d-2}\} \in S^2$. We connect the spirals at their ends by $n^{d-4}(n-s)^2-1$ short segments. The set of points which remain to cover is $[n]^{d-4} \times S^4$. The number of links in the spanning path in this step is $n^{d-4}(n-s)^2(2s+1)-1$.

Suppose that d is even. After $d/2$ steps, the uncovered space is S^d and the number of links in the path is

$$(2s+1)\sum_{i=1}^{d/2} n^{d-2i}(n-s)^{2i-2} + O(1) = dsn^{d-2} + O(n^{d-2}).$$

The cube S^d can be traversed recursively. Let $l(n)$ be the length of the resulting path. Then

$$l(n) = l(n-s) + dsn^{d-2} + h(n) \tag{16}$$

where $h(n) = O(n^{d-2})$. We show by induction that $l(n) \leq dn^{d-1}/(d-1) + cn^{d-3/2}$ for some constant $c > 0$. It follows from

$$\begin{aligned} & l(n) \\ = \; & l(n-s) + dsn^{d-2} + h(n) \\ \leq \; & \frac{d}{d-1}(n-s)^{d-1} + c(n-s)^{d-3/2} + dsn^{d-2} + h(n) \\ \leq \; & \frac{d}{d-1}n^{d-1} + cn^{d-3/2} \end{aligned}$$

since

$$\begin{aligned} & n^{d-1} - (n-s)^{d-1} \\ = \; & s(n^{d-2} + n^{d-3}(n-s) + \cdots + (n-s)^{d-2}) \\ \geq \; & (d-1)s(n-s)^{d-2} \\ = \; & (d-1)sn^{d-2} - sO(sn^{d-3}) \\ = \; & (d-1)sn^{d-2} - O(n^{d-2}) \end{aligned}$$

and similarly

$$n^{d-3/2} - (n-s)^{d-3/2} = \Omega(sn^{d-5/2}) = \Omega(n^{d-2}).$$

For odd d, the argument is similar. Let S_2 be $\{s+1, s+2, \ldots, n-s\}$. First, using $(d-1)/2$ above steps, we reduce the d-cube along the dimensions $2, 3, \ldots, d$. The uncovered space is $[n] \times S^{d-1}$. Then we reduce the space along the dimensions $1, 2, \ldots, d-1$ using $(d-1)/2$ steps. The uncovered space is $S \times S_2^{d-2} \times S$. Then we apply the reduction step for dimensions 1 and d. The uncovered space is S_2^d. There are $sn^{d-2} + O(n^{d-2})$ new links at every step. The total number of links in the spanning path is $2sdn^{d-2} + O(n^{d-2})$. The recurrence is now

$$l(n) = l(n-2s) + 2dsn^{d-2} + O(n^{d-2}). \tag{17}$$

It has the same solution $l(n) = dn^{d-1}/(d-1) + O(n^{d-3/2})$. This completes the proof of the upper bound in Theorem 4. □

5 Arbitrary point sets: combinatorial bounds and approximation algorithms

In this section, we study the situation where the point set is arbitrary but the path is still axis-aligned. We say that a point set is in *general position* if no two x-coordinates and no two y-coordinates of the points are the same.

5.1 Proof of Theorem 5

We first show the lower bound. Consider the set of n points: $\{(1,1), \ldots, (n-3, n-3), (n-2, n-1), (n-1, n-2), (n, n)\}$ shown in Figure 6. We leave it to the reader to verify that any axis-aligned spanning path requires at least $n+1$ links.

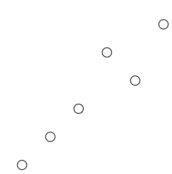

Figure 6: A set of n points which requires $n+1$ links in a spanning path.

We now prove the upper bound. A set of points ordered from left to right is said to form an *ascending chain* if their y-coordinates form a non-decreasing sequence. Similarly, a *descending chain* of points is a sequence of points with non-increasing y-coordinates. The first step is to observe that an ascending or descending chain with k points can be traversed by a path with at most k links. Moreover, the direction of the last (or first) link in the chain can be specified in advance to horizontal or vertical (e.g., by an easy inductive path construction). An axis-parallel path is said to be an *ascending path* if it can traversed so that each link is oriented rightwards or upwards (in the positive direction of the x or y-axis). Similarly, a *descending path* is one that can traversed so that each link is oriented to the right or downwards.

Two axis-parallel rectangles r' and r'', so that $r' \subset r''$, are said to be *properly nested* is no side of r' overlaps with any side of r''. Let $\mathcal{R} = r_1 \subset r_2 \subset \ldots \subset r_k$ be a sequence of properly nested axis-aligned rectangles in the plane. A set of points R is said to form a *rectangle set* if each point in R is contained in the interior of one of the sides of a rectangle in $\mathcal{R}$, and if each rectangle boundary contains at least four points with at least one point in the interior of each side of the rectangle.

We first show how to partition a given set S of n points into an ascending chain A, a descending chain D and a rectangle set R using the following iterative process. Initiate this process by creating five empty lists L_{SW}, L_{NE}, L_{NW}, L_{SE} and L_R. The construction proceeds as follows:

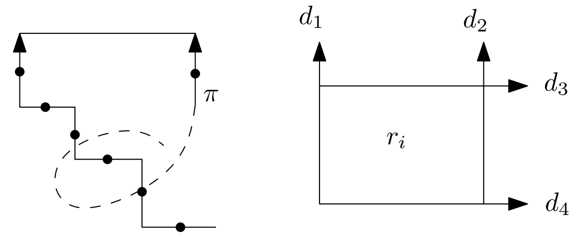

(a) (b)

Figure 7: (a) Combining two paths whose end links have the same direction. (b) The directions d_1, d_2, d_3 and d_4.

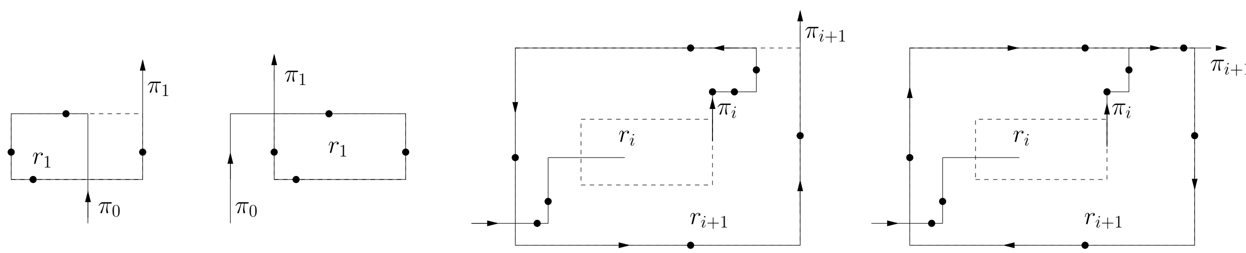

(a) (b) (c) (d)

Figure 8: (a) and (b): Illustrations for the induction base. (c) Extending π_i by first traversing $A_{i+1} \setminus A_i$, and then the points on r_{i+1} using four links. (d) Extending π_i by first traversing $A_{i+1} \setminus A_i$, and then the points on r_{i+1} using five links.

Current step: Find the smallest axis-aligned rectangle r that encloses the current set S. We denote by $SW(r)$ (resp., $SE(r)$, $NE(r)$, $NW(r)$) the SW corner (resp., SE corner, NE corner, NW corner) of r. If any of the corners of r belongs to S, then (i) we append it to the corresponding list L_{SW}, L_{NE}, L_{NW} or L_{SE}, (ii) remove it from S, and (iii) repeat the iterative step. If no corner of r belongs to (the current) S, then (i) append r to a list of nested rectangles L_R, (ii) include in R all points in S on the boundary of r, and remove these points from S, and (iii) repeat the iterative step. We iterate this step as long as S is non-empty.

Concatenate the list L_{SW} and the reversed list L_{NE} into an ascending chain of points A. Concatenate the list L_{NW} and the reversed list L_{SE} into an descending chain of points D. Reverse the list of rectangles L_R and call $\mathcal{R}$ the resulting sequence of rectangles; note that $\mathcal{R}$ is a sequence of properly nested rectangles. We have partitioned S into three sets: $S = A \cup D \cup R$, as claimed.

We show that $A \cup R$ can be traversed by a path π with at most $|A| + |R|$ links. Since (by an earlier observation) D can be traversed with at most $|D|$ links so that the last link has the same direction as the last link of π, the two paths can be combined (by adding an extra link) in a path with at most $|A|+|R|+|D|+1 = n+1$ links, and the result will follow; see Figure 7 (a). The case when $R = \emptyset$ is immediate, so assume that R is not empty.

It remains to show how to construct π, which is done inductively. As a general idea, π is formed by concatenating (extending) an ascending path traversing L_{SW} with a spiral path traversing $R \cup L_{NE}$ that starts in the innermost rectangle r_1 and goes outward from there. Let $r_1 \subset r_2 \subset \cdots \subset r_k$ be the rectangles in $\mathcal{R}$ in nested order. Let A_i be the set of points of A whose x-coordinates do not exceed the x-coordinate of the right side of r_i, i.e., $A_i = L_{SW} \cup (L_{NE} \cap r_i)$. Let R_i be the set of points of R contained in r_i, i.e., $R_i = R \cap r_i$. Let d_1, d_2, d_3, d_4 be the rays (directions) along the sides of r_i oriented as in Figure 7 (b). We will prove by induction on i the following claim.

Claim. $A_i \cup R_i$ *can be traversed by a path* π_i *with at most* $|A_i| + |R_i|$ *links and terminating direction* d_1, d_2, d_3 *or* d_4.

The induction base: $i = 1$. Let π_0 be an ascending path traversing the points in A_1, where the last link of π_0 points upwards (w.l.o.g. by a previous observation). Refer to Figure 8 (a,b). Extend π_0 by traversing r_1 clockwise or counter-clockwise using four or five links, depending on whether the extension of the last link of π_0 intersects r_1, and on whether the top side of r_1 is empty of points strictly to the right (resp. left) of the last link of π_0. It can be checked that the resulting path π_1 has terminating direction d_1, d_2, d_3 or d_4, and that the number of extra links used does not exceed the number of points on r_1 that are covered.

The induction step. Let π_i be the path traversing $A_i \cup R_i$, with terminating direction d_1, d_2, d_3 or d_4 (corresponding to r_i). Refer to Figure 8 (c,d). Extend π_i by first traversing $A_{i+1} \setminus A_i$, and then the points on r_{i+1} in a manner similar to the induction base.

Assuming now the claim, extend π_k to traverse the points in $A \setminus A_k$, if any. Set π to be the resulting path. This completes the proof of the theorem. □

Note that $A \cup D$ cannot be always traversed by a path of length $|A \cup D|$, as our lower bound example shows; and this is the reason for considering $A \cup R$ instead.

5.2 Higher dimensions: proof of Theorem 6

We first show the lower bound. Let $S = \{p_1, \ldots, p_n\}$ be a set of points in $\mathbb{R}^d$ in general position. Any rectilinear path from a point p_i to a point p_j has at least $d - 1$ turns. Thus a path traversing S has at least $(d-1)(n-1)$ turns and $(d-1)(n-1)+1 = (d-1)n + 2 - d$ links.

We now prove the upper bound. Let $S = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$. The bound is trivial for $n = 1, 2$. For $n = 3$, there is an order of the points p_1, p_2, p_3 and $1 \le k \le d$ such that $x_k(p_2) \in [x_k(p_1), x_k(p_3)]$. The segment along k-th axis can be used to traverse p_2. The total number of turns is $2(d-1)$ as required.

Suppose now that $n \ge 4$. We project the points onto the (x_1, x_2)-plane. By Theorem 5 there is a traversal π_2 of the projected points with at most $n + 1$ links. We lift the path π_2 into $\mathbb{R}^d$ in the following way. Let $p_1, \ldots, p_n$ be the order of the points along the path π_2. We lift the first segment using the last $d - 2$ coordinates of p_1. There is a turn point on the path π_2 between p_1 and p_2. We insert $d - 2$ segments at the turn point and change the last $d - 2$ coordinates so that they correspond to the ones of p_2. We continue to use lifting in the same way for the remaining turns of π_2. The total number of segments is $n + 1 + (d-2)(n-1) = (d-1)n + 3 - d$, and the theorem follows. □

5.3 Approximation algorithms

As mentioned in the introduction, the *minimum link spanning path* problem appears to be related to the problem of covering a set of points by minimum number of lines — *point covering with lines* — (in both the general, and in the axis-parallel case). Megiddo and Tamir have shown that point covering with lines is NP-complete in general (for lines of arbitrary orientation) for any $d \ge 2$ [12]. For the axis aligned version, the planar case of point covering with lines can be solved exactly in polynomial time, while the higher dimensional variants ($d \ge 3$) are again NP-complete [9]. Hassin and Megiddo have given a linear-time d-approximation algorithm for point covering in $\mathbb{R}^d$ by axis-parallel lines [9].

Proof of Theorem 7. Consider first the planar case $d = 2$: consider an arbitrary (axis-aligned) spanning path, say with k links, traversing a planar point set S; it immediately gives k axis-parallel lines covering S.

Conversely, given ℓ lines covering S, they can be connected in a spanning path with at most 2ℓ axis-parallel links as follows: assume that there are h horizontal lines and v vertical lines in the cover (where $h + v = \ell$). Place the point set in a (sufficiently large) axis-aligned box B containing all the points. The k lines are clipped by the box in h horizontal segments and v vertical segments. Link the h horizontal segments (in any order) by adding $h - 1$ suitable vertical segments from the boundary of B and obtain a subpath with $2h - 1$ links. Proceed similarly with the vertical segments and obtain another subpath with $2v - 1$ links. Combine the two subpaths by adding two more links (one horizontal and one vertical) in between: the final path has $2(h + v) = 2\ell$ links and spans all points in S.

The above argument gives us a ratio 2 approximation algorithm in the plane: by the result of Hassin and Megiddo [9], the problem of computing a set of axis-parallel lines of minimum cardinality covering a given set of points in the plane can be solved exactly in polynomial time (by formulating it as a vertex cover problem in a bipartite graph). If an optimal line cover having ℓ lines is used, the ratio 2 is implied.

We now consider the case $d \geq 3$. Assume that the approximation algorithm of Hassin and Megiddo for covering S yields $\ell = \ell_1 + \ldots + \ell_d$ lines, where L_i denotes the set of lines parallel to the i-th coordinate ($\ell_i = |L_i|$). We clip the lines with the box B, and then, for each $i = 1, \ldots, d$, connect (in any order) the ℓ_i segments parallel to the i-th coordinate in a subpath by adding $d - 1$ segments on the boundary of B in between any two consecutive segments in L_i. The resulting ith subpath has $d\ell_i - (d - 1)$ links. Combine the $d - 1$ subpaths (in any order) by adding d links in between any two consecutive subpaths. The final path has $d(\sum_{i=1}^{d} \ell_i) - d(d-1) + (d-1)d = d\ell$ links and spans all points in S. Since the line cover obtained from the algorithm of Hassin and Megiddo is by itself a d-approximate solution, it implies a $d \times d = d^2$ approximation for our spanning path covering problem. □

6 Conclusions

We have studied several variants of the following problem: given a finite set of points S in $\mathbb{R}^d$, visit the points in S with an axis-aligned polygonal path that makes a minimum number of turns, or equivalently, has the minimum number of links. We conclude with the following open problems:

- Can the gap between the bounds on path length for traversing the cube in higher dimensions be further reduced?
- Is the *minimum link spanning path* problem still NP-complete for axis-aligned paths?
- Can the approximation ratio in Theorem 7 be improved? Note that in the plane, point covering with lines can be solved exactly.

References

[1] E. M. Arkin, M. A. Bender, E. D. Demaine, S. P. Fekete, J. S. B. Mitchell, and S. Sethia, Optimal covering tours with turn costs, *SIAM Journal on Computing*, **35(3)** (2005), 531–566.

[2] D. R. Gaur, B. Bhattacharya, Covering points by axis parallel lines, *Proc. 23rd European Workshop on Computational Geometry*, (2007), 42–45.

[3] E. M. Arkin, J. S. B. Mitchell, and C. D. Piatko, Minimum-link watchman tours, *Information Processing Letters*, **86** (2003), 203–207.

[4] P. Braß, W. Moser, and J. Pach, *Research Problems in Discrete Geometry*, Springer, New York, 2005.

[5] M. J. Collins and M. E. Moret, Improved lower bounds for the link length of rectilinear spanning paths in grids, *Information Processing Letters*, **68** (1998), 317–319.

[6] M. J. Collins, Covering a set of points with a minimum number of turns, *International Journal of Computational Geometry & Applications*, **14(1-2)** (2004), 105–114.

[7] S. P. Fekete and G. J. Woeginger, Angle-restricted tours in the plane, *Computational Geometry: Theory and Applications*, **8(4)** (1997), 195–218.

[8] F. Gavril, Some NP-complete problems on graphs, *Proc. 11th Conference on Information Sciences and Systems*, 1977, 91–95.

[9] R. Hassin and N. Megiddo, Approximation algorithms for hitting objects with straight lines, *Discrete Applied Mathematics* **30(1)** (1991), 29–42.

[10] E. Kranakis, D. Krizanc and L. Meertens, Link length of rectilinear Hamiltonian tours in grids, *Ars Combinatoria*, **38** (1994), 177–192.

[11] A. Maheshwari, J-R. Sack and H. N. Djidjev, Link distance problems, in *Handbook of Computational Geometry (J-R. Sack and J. Urrutia, eds.)*, chap. 12, Elsevier, 2000, pp. 519–558.

[12] N. Megiddo and A. Tamir, On the complexity of locating linear facilities in the plane, *Operations Research Letters*, **1(5)** (1982), 194–197.

[13] C. Stein and D. P. Wagner, Approximation algorithms for the minimum bends traveling salesman problem, Proc. 8th Internat. Conf. on Integer Programming and Combinatorial Optimization, LNCS 2081, pp. 406-421, 2001.

[14] D. P. Wagner, Path planning algorithms under the link-distance metric, PhD thesis, Darmouth College, 2006.