Sweeping points*

Adrian Dumitrescu† Minghui Jiang‡

August 28, 2009

Abstract

Given a set of points in the plane, and a sweep-line as a tool, what is best way to move the points to a target point using a sequence of sweeps? In a sweep, the sweep-line is placed at a start position somewhere in the plane, then moved orthogonally and continuously to another parallel end position, and then lifted from the plane. The cost of a sequence of sweeps is the total length of the sweeps. Another parameter of interest is the number of sweeps. Four variants are discussed, depending on whether the target is a hole or a pile, and whether the target is specified or freely selected by the algorithm. Here we present a ratio $4/\pi \approx 1.27$ approximation algorithm in the length measure, which performs at most four sweeps. We also prove that, for the two constrained variants, there are sets of n points for which any sequence of minimum cost requires $3n/2 - O(1)$ sweeps.

1 Introduction

Sweeping is a well known and widely used technique in computational geometry. In this paper we make a first study of sweeping as an operation for moving a set of points. The following question was raised by Paweł Żyliński [5]:

> There are n balls on a table. The table has a hole (at a specified point). We want to sweep all balls to the hole with a line. We can move the balls by line sweeping: all balls touched by the line are moved with the line in the direction of the sweep. The problem is to find an optimal sequence of sweeps which minimizes the total sweeping distance covered by the line.

Although the above problem is quite natural, it does not seem to have been studied before. We note an obvious application to robotics, in particular, to the automation of part feeding and to nonprehensile part manipulation [1]. Imagine a manufacturing system that produces a constant stream of small identical parts, which have to be periodically cleared out, or gathered to a collection point by a robotic arm equipped with a segment-shaped sweeping device [1]. Here we study an abstraction of such a scenario, when the small objects and the target are abstracted as points.

*A preliminary version of this paper appeared in the Proceedings of the 11th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2008), MIT, Boston, USA, August 2008; Vol. 5171 of LNCS, Springer Verlag, pp. 63–76.

†Department of Computer Science, University of Wisconsin–Milwaukee, WI 53201-0784, USA. Email: ad@cs.uwm.edu. Supported in part by NSF CAREER grant CCF-0444188.

‡Department of Computer Science, Utah State University, Logan, UT 84322-4205, USA. Email: mjiang@cc.usu.edu. Supported in part by NSF grant DBI-0743670.

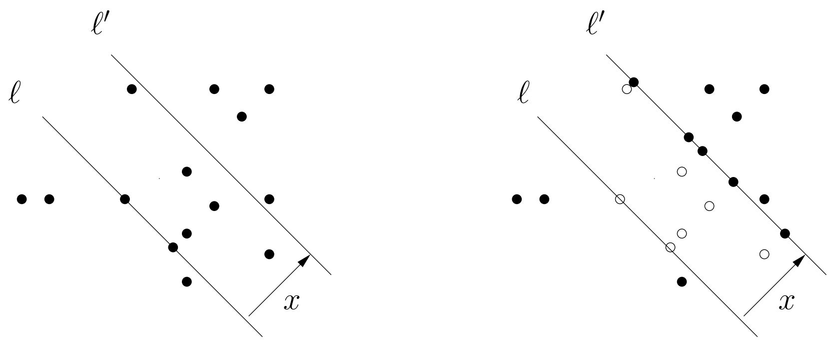

Figure 1: A sweep of length (cost) x involving 7 points. From the initial 7 points swept, only 5 points remain.

We now introduce some definitions to make the problem formulation more precise. We refer to Figure 1. A set S of n points in the plane is given. In a *sweep*, the sweep-line is placed at a start position somewhere in the plane and is moved orthogonally and continuously to another parallel end position. Then the line is lifted from the plane. All points touched by the line are moved with the line in the direction of the sweep. Points can merge during a sweep, and merged points are subsequently treated as one point: a set of collinear points on any segment parallel to the sweep direction and bounded by the initial and the final positions of the line will yield one point after the sweep; see Figure 1 for an example.

A *sweeping sequence* for S is a sequence of sweeps that move all points in S to a target point. The *cost* of a sweeping sequence is the total length of its sweeps. As it will be evident from our Theorem 3, the sweep-line as a tool can be conveniently replaced by a finite sweep-segment of length twice the diameter of the point set.

We consider several variants of the sweeping problem, by making two distinctions on the *target*. First, the target can be either a hole or pile: if the target is a *hole*, then a point stays at the target once it reaches there, i.e., the point drops into the hole; if the target is a *pile*, then a point can still be moved away from the target after it reaches there. While it makes no difference for our algorithms whether the target is a hole or a pile (i.e., our algorithms are applicable to both variants), this subtle difference does matter when deriving lower bounds. Second, the target is either *constrained* to be a specified point or *unconstrained* (an arbitrary point freely selected by the algorithm). The four possible combinations, constrained versus unconstrained (C or U) and hole versus pile (H or P), yield thus four variants of the sweeping problem: CH, CP, UH, and UP.

Our main results are the following: although there exist sets of n points that require $\Omega(n)$ sweeps in any optimal solution (Section 3, Theorem 2), constant-factor approximations which use at most 4 sweeps can be computed in linear or nearly linear time (Section 2, Theorem 1). We also present some initial results and a conjecture for a related combinatorial question (Section 4, Theorem 3), and conclude with two open questions (Section 5).

We now introduce some preliminaries. A sweep is *canonical* if the number of points in contact with the sweep-line remains the same during the sweep. The following lemma is obvious.

Lemma 1. *Any sweep sequence can be decomposed into a sweep sequence of the same cost, consisting of only canonical sweeps. In particular, for any point set S, there is an optimal sweep sequence consisting of only canonical sweeps.*

Proof. Let $|S| = n$. A non-canonical sweep can be decomposed into a sequence of at most n canonical sweeps in the same direction and of the same total cost. □

Throughout the paper, we use the following convention: if A and B are two points, $\overline{AB}$ denotes the line through A and B, $\overrightarrow{AB}$ denotes the ray starting from A and going through B, AB denotes the line segment with endpoints A and B, and $|AB|$ denotes the length of the segment AB.

2 A four-sweep algorithm

In this section, we present a four-sweep algorithm applicable to all four variants CH, CP, UH, and UP.

Theorem 1. *For any of the four variants* CH*,* CP*,* UH*, and* UP *of the sweeping problem (with n points in the plane),*

(I) *A ratio $\sqrt{2}$ approximation that uses at most 4 sweeps can be computed in $O(n)$ time;*

(II) *A ratio $4/\pi \approx 1.27$ approximation that uses at most 4 sweeps can be computed in $O(n \log n)$ time.*

Proof. We consider first the constrained variant, with a specified target o. Let S be the set of n points, and let $S' = S \cup \{o\}$. We next present two algorithms.

(I) **Algorithm A1.** Choose a rectilinear coordinate system xoy whose origin is o (of arbitrary orientation). Compute a minimal axis-parallel rectangle Q containing S'. Denote by w and h its width and height respectively, and assume w.l.o.g. that $h \leq w$. Perform the following (at most four) sweeps: (i) sweep from the top side of Q to the x-axis; (ii) sweep from the bottom side of Q to the x-axis; (iii) sweep from the left side of Q to the y-axis; (iv) sweep from the right side of Q to the y-axis. Figure 2 illustrates the execution of the algorithm on a small example.

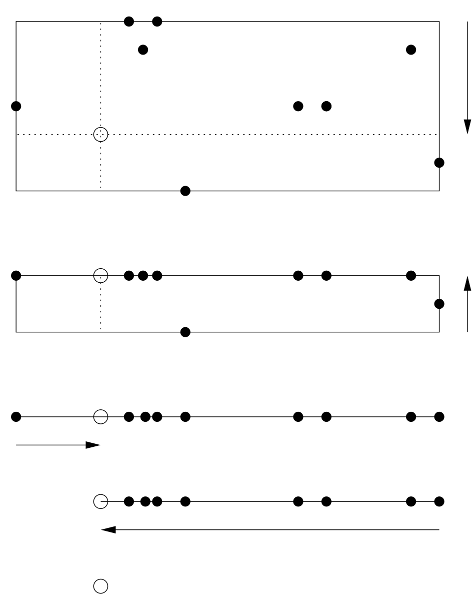

Figure 2: Running the four-sweep algorithm.

Analysis. Clearly, the algorithm gives a valid solution, whose total cost is ALG $= w + h$. Let OPT be the cost of an optimal solution. We first argue that the approximation ratio of our algorithm is at most 2; we then improve this bound to $\sqrt{2}$.

We first show that OPT $\geq w$. Let p and q be the two extreme points of S' with minimum and maximum x-coordinates. Assume first that $p, q \in S$. Let p' and q' be the projection points of p and q on the x-axis throughout the execution of the sweep sequence. Put $w_1 = |p'o|$, and $w_2 = |oq'|$. Note that after the sweep sequence is complete, p' and q' coincide with the origin o. Further note that every sweep brings either p' or q' closer to o, but not both. Finally, observe that to bring p' to o requires a total sweep cost of at least w_1, and similarly, to bring q' to o requires a total sweep cost of at least w_2. Therefore the total sweep cost is at least $w_1 + w_2 = w$, thus OPT $\geq w$. Since the total sweep cost is

$$\text{ALG} = w + h \leq 2w \leq 2 \cdot \text{OPT},$$

the ratio 2 follows when $p, q \in S$. The case when o is one of the two extreme points p and q is completely analogous.

We now argue that OPT $\geq (w+h)/\sqrt{2}$. Let X be an arbitrary sequence consisting of k sweeps which solves the given instance S. For $i = 1, \ldots, k$ let x_i be the cost of the ith sweep, and $\alpha_i \in [0, 2\pi)$ be its direction. Write $x = \sum_{i=1}^{k} x_i$. Indeed, the ith sweep of cost x_i can reduce the current semi-perimeter of Q (the width plus the height of Q) by at most $\sqrt{2}\, x_i$; see also (1) below. Here the points in S are considered moving, so S', and its enclosing rectangle Q change continuously as an effect of the sweeps. Since the semi-perimeter of Q drops from $w + h$ to 0, by summing over all sweeps, we get that in any sweep sequence for S of total cost x,

$$\sqrt{2}\, x = \sqrt{2} \sum_{i=1}^{k} x_i \geq w + h,$$

thus

$$\text{ALG} = w + h \leq \sqrt{2} \cdot \text{OPT},$$

and the approximation ratio $\sqrt{2}$ follows.

(II) **Algorithm A2.** First compute a minimum perimeter rectangle Q_0 containing S'. This takes $O(n \log n)$ using the rotating calipers algorithm of Toussaint [3]. Let now xoy be a rectilinear coordinate system in which Q_0 is axis-aligned. Let w and h be its width and height respectively. Then perform the four sweeps as in Algorithm A1.

Analysis. Assume w.l.o.g. that $w + h = 1$. For $\beta \in [0, \pi/2)$, let $Q(\beta)$ denote the minimum perimeter rectangle of orientation β containing S'; i.e., one of the sides of $Q(\beta)$ makes an angle β with the positive direction of the x-axis. Let $w(\beta)$ and $h(\beta)$ denote the initial values of the width and height of $Q(\beta)$ respectively. Note that $[0, \pi/2)$ covers all possible orientations β of rectangles enclosing S'.

As in the proof of the ratio $\sqrt{2}$ approximation ratio, recall that for any $i \in \{1, \ldots, k\}$, the ith sweep of cost x_i can reduce the current semi-perimeter of $Q(\beta)$ by at most $x_i\sqrt{2}$. In fact we can be more precise by taking into account the direction of the sweep: the reduction is at most

$$x_i \left(|\cos(\alpha_i - \beta)| + |\sin(\alpha_i - \beta)| \right) \leq x_i \sqrt{2}. \tag{1}$$

Since X solves S, by adding up the reductions over all sweeps $i \in \{1, \ldots, k\}$, we must have—since $w(\beta) + h(\beta) \geq 1$, for every $\beta \in [0, \pi/2)$:

$$\sum_{i=1}^{k} x_i \left(|\cos(\alpha_i - \beta)| + |\sin(\alpha_i - \beta)| \right) \geq 1. \tag{2}$$

We integrate this inequality over the β-interval $[0, \pi/2]$; x_i and α_i are fixed, and each term is dealt with independently. Fix $i \in \{1, \ldots, k\}$, and write $\alpha = \alpha_i$ for simplicity. Assume first that $\alpha \in [0, \pi/2)$. A change of variables yields

$$\begin{aligned} & \int_0^{\pi/2} (|\cos(\alpha - \beta)| + |\sin(\alpha - \beta)|) \, \mathrm{d}\beta \\ = & \int_\alpha^{\alpha+\pi/2} (|\cos\beta| + |\sin\beta|) \, \mathrm{d}\beta \\ = & \int_\alpha^{\pi/2} (\cos\beta + \sin\beta) \, \mathrm{d}\beta + \int_{\pi/2}^{\alpha+\pi/2} (-\cos\beta + \sin\beta) \, \mathrm{d}\beta \\ = & (\sin\beta - \cos\beta)\Big|_\alpha^{\pi/2} + (-\sin\beta - \cos\beta)\Big|_{\pi/2}^{\alpha+\pi/2} \\ = & (1 - \sin\alpha + \cos\alpha) + (-\cos\alpha + \sin\alpha + 1) = 2. \end{aligned}$$

Let

$$G(\alpha) = \int_\alpha^{\alpha+\pi/2} (|\cos\beta| + |\sin\beta|) \, \mathrm{d}\beta.$$

It is easy to verify that $G(\alpha) = G(\alpha + \pi/2)$ for any $\alpha \in [0, 2\pi)$, hence the integration gives the same result, 2, for any $\alpha_i \in [0, 2\pi)$, and for any $i \in \{1, \ldots, k\}$. Hence by integrating (2) over $[0, \pi/2]$ yields

$$2\left(\sum_{i=0}^{k} x_i\right) \geq \frac{\pi}{2}, \quad \text{or} \quad x \geq \frac{\pi}{4}.$$

Since this holds for any valid sequence, we also have $\text{OPT} \geq \frac{\pi}{4}$. Recall that $\text{ALG} = w + h = 1$, and the approximation ratio $4/\pi$ follows.

To extend our results to the unconstrained variant requires only small changes in the proof. Instead of the minimum semi-perimeter rectangle(s) enclosing $S' = S \cup \{o\}$, consider the minimum semi-perimeter rectangle(s) enclosing S. All inequalities used in the proof of Theorem 1 remain valid. We also remark that the resulting algorithms execute only two sweeps (rather than four): from top to bottom, and left to right, with the target being the lower-right corner of the enclosing rectangle. □

2.1 A lower bound on the approximation ratio of Algorithm A2

It is likely that the approximation ratio of our four-sweep algorithm is slightly better than what we have proved: we noticed that for both cases, when h is large and when h is small relative to w, our estimates on the reduction are slightly optimistic. However, the construction we describe next, shows that the ratio of our four-sweep algorithm cannot be reduced below 1.1784 (for either variant).

Perhaps the simplest example to check first is the following. Take the three vertices of a unit (side) equilateral triangle as our point set. For the constrained variant, place the target at the triangle center: the optimal cost is at most $\sqrt{3}$ by 3 sweeps (in fact, equality holds, as shown in the proof of Theorem 3), while the four-sweep algorithm uses $1 + \sqrt{3}/2$. The ratio is about 1.077.

We now describe a better construction that gives a lower bound of about 1.1784; see Figure 3. Place n points uniformly (dense) on the thick curve $\mathcal{C}$ connecting B and C. For the constrained variant, place the target at point B. The curve is made from the two equal sides of an obtuse isosceles triangle with sharp angles $\alpha = \arctan(1/2) \approx 26.565°$, then "smoothed" around the obtuse triangle corner. $\triangle ABC$ is an isosceles triangle with sides $AB = AC = \sqrt{5}$ and $BC = 4$, with height $AD = 1$, and with angles $\angle ABC = \angle ACB = \alpha = \arctan(1/2)$. E and F are two points on AB and AC, respectively, such that $DE \perp AB$ and $DF \perp AC$.

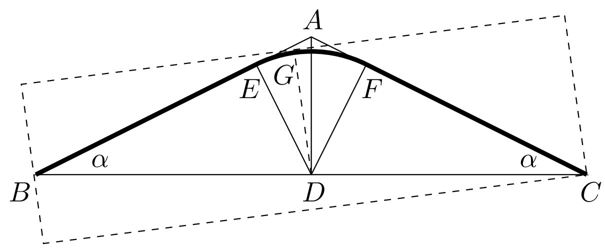

Figure 3: A continuous convex curve $\mathcal{C}$ that gives a lower bound of about 1.1784 on the approximation ratio of Algorithm A2.

The curve $\mathcal{C}$ consists of the two segments BE and CF and a curve $\mathcal{C}_0$ connecting E and F, defined as follows. For an arbitrary point G on $\mathcal{C}_0$, $\angle ADG = \beta \le \alpha$, the length of the segment DG is

$$|DG| = f(\beta) = 4\cos\alpha + 4\sin\alpha - 4\cos\beta - 2\sin\beta.$$

Observe that $4\sin\alpha - 2\cos\alpha = 0$ holds by the definition of α, hence

$$\frac{\mathrm{d}}{\mathrm{d}\beta} f(\beta) = 4\sin\beta - 2\cos\beta \le 4\sin\alpha - 2\cos\alpha = 0,$$

where the derivative reaches zero at E and F (when $\beta = \alpha$). So $\mathcal{C}$ is a continuous convex curve. For a rectangle that circumscribes the curve $\mathcal{C}$ with one side tangent to $\mathcal{C}_0$ at point G, its width and height are $|BC|\cos\beta$ and $|DG| + |CD|\sin\beta$, respectively. Hence its semi-perimeter is

$$\begin{aligned} & |BC|\cos\beta + |DG| + |CD|\sin\beta \\ = \;& 4\cos\beta + (4\cos\alpha + 4\sin\alpha - 4\cos\beta - 2\sin\beta) + 2\sin\beta \\ = \;& 4\cos\alpha + 4\sin\alpha. \end{aligned}$$

Therefore the semi-perimeter of a minimum rectangle with orientation β, where $0 \le \beta \le \alpha$, that encloses $\mathcal{C}$ is a constant: $4\cos\alpha + 4\sin\alpha$. Since the length of $\mathcal{C}_0$ is

$$\begin{aligned} 2\int_{\beta=0}^{\alpha} f(\beta)\mathrm{d}\beta &= 2\int_{\beta=0}^{\alpha} (4\cos\alpha + 4\sin\alpha - 4\cos\beta - 2\sin\beta)\mathrm{d}\beta \\ &= 8(\cos\alpha + \sin\alpha)\alpha + 2(-4\sin\beta + 2\cos\beta)\Big|_0^\alpha \\ &= 8(\cos\alpha + \sin\alpha)\alpha + 2(-4\sin\alpha + 2\cos\alpha - 2) \\ &= 8(\cos\alpha + \sin\alpha)\alpha - 4, \end{aligned}$$

and since $|BE| = |CF| = 2\cos\alpha$, the length of $\mathcal{C}$ is $8(\cos\alpha + \sin\alpha)\alpha - 4 + 4\cos\alpha$. The ratio of the minimum semi-perimeter and the curve length is (after simplification by 4, and using the values $\cos\alpha = 2/\sqrt{5}$, $\sin\alpha = 1/\sqrt{5}$, $\alpha = \arctan(1/2)$)

$$\frac{4\cos\alpha + 4\sin\alpha}{8(\cos\alpha + \sin\alpha)\alpha - 4 + 4\cos\alpha} = \frac{3}{6\arctan(1/2) - \sqrt{5} + 2} = 1.1784\ldots.$$

Finally, observe that OPT is at most the length of $\mathcal{C}$. This gives a lower bound of 1.1784 on the approximation ratio of Algorithm A2, which holds for all four variants.

3 Point sets for the constrained variants that require many sweeps

In this section we show that some point sets require many sweeps in an optimal solution, i.e., the number of sweeps is not just a constant. In what follows, the target is constrained to a specified point, and may be either a hole or a pile, i.e., we refer to both constrained variants CH and CP.

Theorem 2. *For the two constrained variants* CH *and* CP*, and for any* $n \geq 3$*, there are sets of* n *points for which any optimal sweeping sequence consists of at least* $3n/2 - O(1)$ *sweeps.*

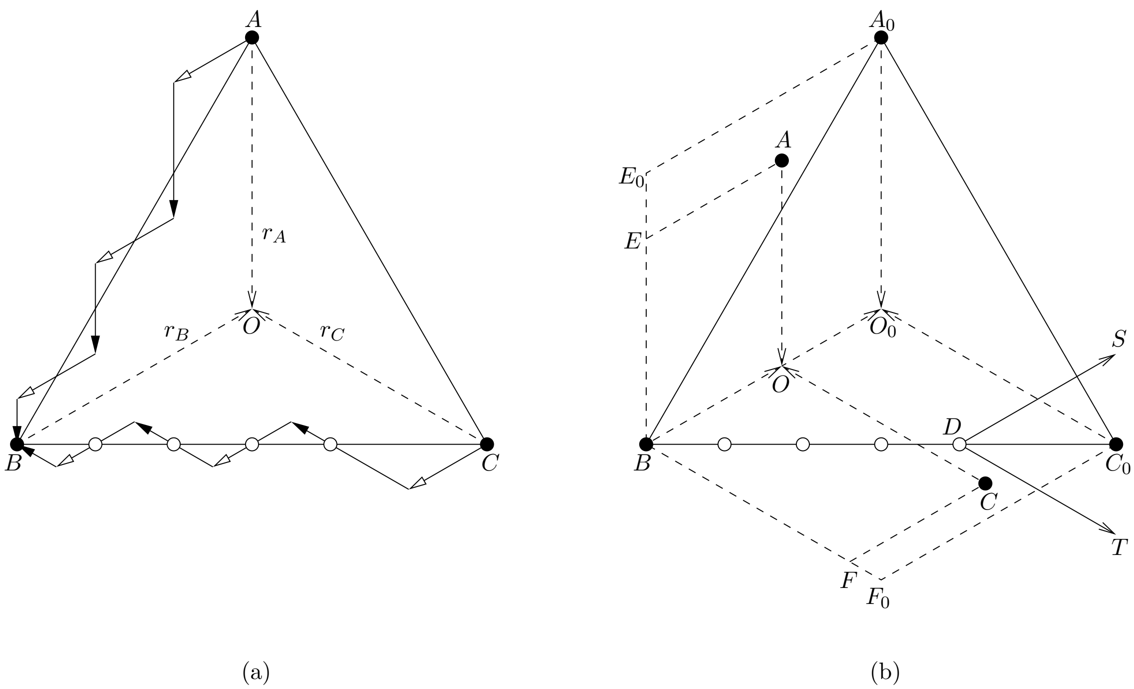

Figure 4: A construction with three points A, B, and C (black points) forming a unit equilateral triangle and $n-3$ arbitrary points (white points) on the edge BC. The target is at the point B. Initially: $A = A_0$, $B = B_0$, $C = C_0$. (a) An optimal sweeping sequence. (b) Some properties of optimal sweeping sequences are illustrated.

We now proceed with the proof of Theorem 2. We refer to Figure 4(a). Our set S consists of three points A, B, and C (black points) forming a unit equilateral triangle and $n-3$ points (white points) arbitrary placed on the edge BC. The target is at the point B. For convenience, we place $\triangle ABC$ initially with B at the origin and $\overrightarrow{BC}$ along the x axis. In what follows, we refer to the intermediate positions of the moving points: input points from the set S (such as A, B, C, D, etc.) or other auxiliary points (such as E and F) during a sequence of sweeps. When the intermediate position of a point does not coincide with its original position, we avoid the possible ambiguity by adding a subscript 0 to the label of the original position. For example, the two labels A and A_0 in the figure refer to the intermediate and the original positions, respectively, of the same point A. Initially, we have $A = A_0$, $B = B_0$, $C = C_0$. We will show in Lemma 2 that $B = B_0$ (that is, B remains stationary) during any optimal sequence; this is evident for the CH variant, but not so for the CP variant.

Define three rays: a ray r_A from A in the $3\pi/2$ direction, a ray r_B from B in the $\pi/6$ direction, and a ray r_C from C in the $5\pi/6$ direction. The three rays from A, B, and C initially intersect at a single point $O = O_0$, the center of $\triangle A_0B_0C_0$. We will show below that this concurrency property is maintained throughout any optimal sweeping sequence for S. We now define six special types of sweeps:

Type A: A is moved in the direction $\overrightarrow{AO}$. B and C are not moved.

Type BC: B and C are moved together in the direction $\overrightarrow{OA}$. A is not moved.

Type B: B is moved in the direction $\overrightarrow{BO}$. A and C are not moved.

Type AC: A and C are moved together in the direction $\overrightarrow{OB}$. B is not moved.

Type C: C is moved in the direction $\overrightarrow{CO}$. A and B are not moved.

Type AB: A and B are moved together in the direction $\overrightarrow{OC}$. C is not moved.

We note that, for the CH variant, the three types involving B, namely types BC, B, and AB are in fact not used, since point B will remain at the hole throughout any sweeping sequence.

For each of the six types, each moved point (among A, B, and C) is moved for a distance equal to the sweep length, that is, the moved point is on the sweep-line during the sweep. If a sweeping sequence consists of only sweeps of the six special types, then it can be easily verified (by induction) that the three rays from A, B, and C still intersect at a single point O after each sweep; see Figure 4(b).

The three segments A_0O_0, B_0O_0, and C_0O_0 determine two parallelograms $A_0O_0B_0E_0$ and $C_0O_0B_0F_0$ (each is a rhombus with two $60°$ angles), as shown in Figure 4(b). We now observe some properties of sweeps of the three types A, C, and AC. Consider how a sweep changes the two parallelograms $AOBE$ and $COBF$, initially $A_0O_0B_0E_0$ and $C_0O_0B_0F_0$: a sweep of type A reduces the two sides AO and BE; a sweep of type C reduces the two sides CO and BF; a sweep of type AC reduces the three sides AE, CF, and OB (note that the side OB is shared by the two parallelograms). During any sweeping sequence of the three types A, C, and AC, the point A always remains inside the rhombus $A_0O_0B_0E_0$, and point C inside the rhombus $C_0O_0B_0F_0$.

Lemma 2. *The optimal cost for S is $\sqrt{3}$. Moreover, any optimal sequence for S consists of only sweeps of the three special types A, C, and AC, with a subtotal cost of $\sqrt{3}/3$ for each type.*

Proof. We first show that the optimal cost for S is at most $\sqrt{3}$. We refer to Figure 4(a) for a sweeping sequence of $n-1$ alternating steps: (i) one sweep of type AC (the white arrow); (ii) two sweeps, one of type A and the other of type C (the black arrows). Each step, except the first and the last, merges C with a white point, in sequential order from right to left. The total number of sweeps in this sequence is $(3n-3)/2$ when n is odd, and is $(3n-4)/2$ when n is even. The total cost of this sequence is $|A_0O_0| + |B_0O_0| + |C_0O_0| = 3 \cdot \sqrt{3}/3 = \sqrt{3}$.

We next show that the optimal cost for S is at least $\sqrt{3}$. Consider an optimal sequence for S. Assume w.l.o.g. that the sequence is canonical. We construct three paths, from the three points A_0, B_0, and C_0 to a single point, such that their total length is at most the cost of the sequence. Each sweep in the sequence that moves one or two of the three points A, B, and C corresponds to an edge in one of the three paths, with the sweep length equal to the edge length: (i) if a sweep moves only one of the three points, then the corresponding edge extends the path from that point, along the sweep direction; (ii) if a sweep moves two of the three points, then the corresponding edge extends the path from the third point, along the opposite sweep direction. We note that, for the three points A, B, and C, each three-point sweep is useless, and each two-point sweep is equivalent to a one-point sweep in the opposite direction, in the sense that the resulting triangles $\triangle ABC$ are congruent. When the three points finally meet at the target, the three paths also end at a single point (which could be different from the target).

The total length of the three paths is at least the total length of a Steiner tree for the three points A_0, B_0, and C_0. It is well known [2] that the minimum Steiner tree for the three points A_0, B_0, and C_0 is unique, and consists of exactly three edges of equal length $\sqrt{3}/3$, from the three points to the center O_0 of $\triangle A_0B_0C_0$. It follows that the optimal cost for S is at least $\sqrt{3}$. Together with the matching upper bound achieved by the sequence illustrated in Figure 4(a), we have shown that the optimal cost for S is exactly $\sqrt{3}$.

The uniqueness of the minimum Steiner tree for the three points A_0, B_0, and C_0 implies that every sweep in the optimal sequence must be of one of the six special types, with a subtotal cost of $|A_0O_0| = |B_0O_0| = |C_0O_0| = \sqrt{3}/3$ for each of the three groups: A and BC, B and AC, and C and AB. To complete the proof, we next show that sweeps of the three types B, AB, and BC never appear in the optimal sequence. Consider the two possible cases for the target:

1. The target is a hole, that is, a point stays at the target once it reaches there. Since B is already at the target, it must stay there. So this case is obvious, as noted after our definition of the six types.
2. The target is a pile, that is, a point can be moved away from the target after it reaches there. Although B is already at the target, it can still be moved away. The only sweeps that move B are of the three types B, AB, and BC. Such sweeps all have a positive projection in the direction $\overrightarrow{BO}$, and can only move B away from the target (and cannot move it back); therefore they cannot appear in the optimal sequence.

This completes the proof of Lemma 2. □

Let D be the rightmost white point. Figure 4(b) shows the initial position of D. Later in Lemma 4, we will prove that D remains at its initial position until it is merged with C. In Lemma 3 however, we don't make any assumption of D being at its original position. Let $\overrightarrow{DS}$ and $\overrightarrow{DT}$ be two rays from D with directions $\pi/6$ and $-\pi/6$, respectively.

Lemma 3. *Consider an optimal sweeping sequence. If C is moved above the line $\overline{DS}$ or below the line $\overline{DT}$, then C remains either above $\overline{DS}$ or below $\overline{DT}$ until either C or D coincides with the target.*

Proof. We refer to Figure 4(b). Assume w.l.o.g. that the sweeping sequence is canonical. Consider each remaining sweep in the sequence after C is at a position above $\overline{DS}$ or below $\overline{DT}$:

Type C. Consider two cases: C is above $\overline{DS}$ or below $\overline{DT}$.

1. C is above $\overline{DS}$. If D is not moved, then C is moved further above $\overline{DS}$. If both C and D are moved (when $CD \perp CO$), then they are moved for the same distance in the same direction, and C remains above $\overline{DS}$.
2. C is below $\overline{DT}$. Since $\overline{DT}$ is parallel to the sweep direction $\overrightarrow{CO}$, C remains below $\overline{DT}$,

Type AC. Consider two cases: C is above $\overline{DS}$ or below $\overline{DT}$.

1. C is above $\overline{DS}$. Since DS is parallel to the sweep direction $\overrightarrow{OB}$, C remains above $\overline{DS}$.
2. C is below $\overline{DT}$. If D is not moved, then C is moved further below $\overline{DT}$. If both C and D are moved, then they are moved for the same distance in the same direction, and C remains below $\overline{DT}$.

Type A. Note that C may be both above $\overline{DS}$ and below $\overline{DT}$. We divide the two cases in an alternative way without overlap: C is either (i) above $\overline{DS}$ and not below (i.e., above or on) $\overline{DT}$ or (ii) below $\overline{DT}$.

1. C is above $\overline{DS}$ and not below $\overline{DT}$. Then C is above D. Since C is not moved, D is not moved either. So C remains above $\overline{DS}$ and not below $\overline{DT}$.
2. C is below $\overline{DT}$. Since $\overline{DT}$ is parallel to $\overline{CO}$, D is above $\overline{CO}$. The sweep may move A down to O and correspondingly move D down until it is on the horizontal line through O, but no further. So D remains above $\overline{CO}$, and C remains below $\overline{DT}$.

The proof of Lemma 3 is now complete. □

Lemma 4. *In any optimal sequence, each white point is not moved until it is merged with C, in sequential order from right to left.*

Proof. Assume w.l.o.g. that the sweeping sequence is canonical. Lemma 2 shows that the sweeps in any optimal sequence are of the three types A, C, and AC. Let σ_1 be the first sweep that moves a white point, and let D_1 be the first white point moved. If the sweep σ_1 is of type A, then A would be moved below the x axis (recall that in a sweep of type A the sweep-line always goes through A), and any subsequent sweep that moves A, of type A or AC, would move A further below the x axis and never to B. This contradicts the validity of the sequence. Therefore σ_1 must be of type C or AC.

We claim that C must be merged with the rightmost white point D before the sweep σ_1. We will prove the claim by contradiction. Suppose the contrary.

Our proof by contradiction is in two steps: In the first step, we will show that C is either above $\overline{DS}$ or below $\overline{DT}$ at the beginning of sweep σ_1. In the second step, we will show that the assumed optimal sequence is not valid.

First step. The sweep-line of σ_1 goes through D_1 during the sweep. Since σ_1 is of type C or AC, C is also on the sweep-line of σ_1. Consider two cases for the relation between D_1 and D:

1. $D_1 \neq D$ (D_1 is to the left of D on the x axis). Then every point on the sweep-line, including C, is either above $\overline{DS}$ or below $\overline{DT}$.

2. $D_1 = D$. Then every point on the sweep-line, except D, is either above $\overline{DS}$ or below $\overline{DT}$. Since C is not merged with D before σ_1, C is either above $\overline{DS}$ or below $\overline{DT}$.

In either case, C is either above $\overline{DS}$ or below $\overline{DT}$.

Second step. From Lemma 3, C remains either above $\overline{DS}$ or below $\overline{DT}$ until either C or D coincides with the target. This, as we will show in the following, implies that the sweeping sequence is not valid. Consider the two possible cases for the target as either a pile or a hole:

1. The target is a pile, that is, a point can be moved away from the target after it reaches there. Then C remains either above $\overline{DS}$ or below $\overline{DT}$ even after either C or D reaches the target. It follows that C and D never merge, and hence cannot end up together at the target. Therefore the sweeping sequence is not valid.

2. The target is a hole, that is, a point stays at the target once it reaches there. Let σ_2 be a sweep in the sequence that moves D to the target. We consider the three possible cases for the type of σ_2:

Type AC. The sweep-line of σ_2 goes through the two points A and C. As D is moved to the point B by σ_2, both parallelograms $AOBE$ and $COBF$ shrink to the point B, that is, both A and C are moved to the target together with D. Then A, C, and D must have been merged even before the sweep σ_2. This is impossible because C is above $\overline{DS}$ or below $\overline{DT}$ until either C or D reaches the target.

Type C. It follows by the same argument (the parallelogram $COBF$ shrinks to the point B) that C and D are merged before the sweep σ_2, which is again impossible.

Type A. It follows by the same argument (the parallelogram $AOBE$ shrinks to the point B) that A and D are merged before the sweep σ_2, above the line $\overline{BO}$. This is impossible because D cannot be moved above $\overline{BO}$: a sweep of type AC does not change the distance from D to $\overline{BO}$; a sweep of type A can only move D further below $\overline{BO}$; a sweep of type C can move D to BO but not above $\overline{BO}$, since C itself cannot be moved above $\overline{BO}$.

In each case, D cannot be moved to the target. Therefore the sweeping sequence is not valid.

We have shown that the sequence is not valid with the target as either a pile or a hole. By contradiction, this proves our original claim that C must be merged with D before the sweep σ_1.

As soon as C is merged with D, we can consider D as deleted. The point set now reaches a configuration similar to the original configuration: the two points B and C are on the x axis with all the (unmoved) white points between them, and A alone is above the x axis. But now we have one less white point. Repeating the argument in the preceding paragraphs inductively completes the proof of Lemma 4. □

We are now in position to finalize the proof of Theorem 2. We have shown that in an optimal sequence, C must be merged with the white points one by one from right to left. Since the sweeps are not along the x axis, each of the $n-3$ white point requires at least one sweep to be merged. The total number of sweeps in the sequence is at least $n - O(1)$. We obtain a tighter estimate (that matches our previous sweep sequence for S) as follows. Between two consecutive merges, C has to be moved to the left by alternating sweeps of types AC and C. Between two sweeps of type AC, since C is moved by a sweep of type C, A must also be moved by a sweep of type A, to make $AC \perp OB$ for the next sweep of type AC. Therefore each merge requires either one sweep of type AC or two sweeps of types A and C, in an alternating pattern as shown in Figure 4(a). The total number of sweeps in the sequence is at least $3n/2 - O(1)$. This completes the proof of Theorem 2.

4 A combinatorial question for the unconstrained variants

The following related question suggests itself: What is the maximum cost required for sweeping a planar point set of unit diameter to a single (unspecified) target point? Depending on whether the target point is a hole or a pile, define

$$\rho_H = \sup_S \inf_X \mathrm{cost}(X), \quad \text{for the variant UH,}$$

and

$$\rho_P = \sup_S \inf_X \mathrm{cost}(X), \quad \text{for the variant UP,}$$

where S ranges over all finite planar point sets of unit diameter, and X ranges over all sweeping sequences for S. We give estimates on the two numbers ρ_H and ρ_P in the following theorem:

Theorem 3. $1.73 \approx \sqrt{3} \leq \rho_H \leq \rho_P \leq 2$.

Proof. Any sweeping sequence for the UP variant is also a sweeping sequence for the UH variant, so we have $\rho_H \leq \rho_P$. We first prove the upper bound $\rho_P \leq 2$. Let S be an arbitrary finite planar set with unit diameter. Let p and q be two points in S at unit distance. Then S is contained in a rectangle with width 1 (parallel to the line pq) and height at most 1. A sweep along the width and a sweep along the height reduce the rectangle to a single point (the pile), at a cost of at most 2.

We next prove the lower bound $\rho_H \geq \sqrt{3}$. Let S be the set of three vertices of an equilateral triangle T of unit side. Let X^* be an optimal sequence of canonical sweeps for S. Using the same idea as in the proof of Lemma 2, we construct three paths, from the three vertices of T to a common point, such that their total length is at most the cost of X^*. It follows that the cost of X^* is at least the total length of a minimum Steiner tree for the three vertices, which is $\sqrt{3}$ [2]. Note that our analysis for this example is tight: three sweeps along the edges of the minimum Steiner tree clearly move the three points of S to a single point (the center of T). □

The reader can observe that a weaker lower bound $\rho_H \geq \pi/2 \approx 1.57$ follows from our result in Theorem 1 applied to a set of n points uniformly distributed on a circle (for large n). We think the upper bound in Theorem 3 is best possible, for instance, in the same case of n points uniformly distributed on a circle of unit diameter, for n going to infinity:

Conjecture 1. $\rho_H = \rho_P = 2$.

5 Concluding remarks

It is natural to consider the sweeping problem for infinite (bounded) sets as well. We implicitly have done so in Section 2.1, with the curve $\mathcal{C}$. Let S be an infinite and bounded set in the plane. Then the optimum cost of sweeping S (to a pile or a hole) is the infimum cost over all finite sweeping sequences for S (of the same type). In particular, for the (points in the) disk of unit diameter, we conjecture that both these optimal costs are equal to 2 in the unconstrained variants.

Other interesting examples for the sweeping problem are open curves. The analysis of Algorithm A2, in particular the lower bound construction, is related to the following question: Which curve of length 1 maximizes the minimum perimeter of an enclosing rectangle? According to a result of Welzl [4], every closed curve of unit length can be enclosed in a rectangle of perimeter $\frac{4}{\pi}$. Consider an open curve C of unit length. By doubling it, or by adding to it a symmetric copy connecting its endpoints, we get a closed curve C' of length 2. By the result we just mentioned, C' can be enclosed in a rectangle of perimeter $\frac{8}{\pi}$. It follows that any open curve of unit length can be enclosed in a rectangle of perimeter $\frac{8}{\pi}$. Equivalently, any open curve of unit length can be enclosed in a rectangle of semi-perimeter $\frac{4}{\pi} = 1.2732\ldots$. On the other hand, our lower bound construction gives an open curve of unit length for which the semi-perimeter of any enclosing rectangle is at least $1.1784\ldots$.

Interestingly enough, the example in Theorem 2 can be extended to an infinite set, where the length of an optimal sweeping sequence is finite, $\sqrt{3}$, but an infinite number of sweeps is required. For instance, consider the (infinite) sequence of points on a unit segment BC, whose distance from B is $\frac{1}{i}$, for $i = 1, 2, \ldots$, and a point A placed at the vertex of an equilateral triangle ΔABC.

Besides Conjecture 1, two interesting questions (for any of the four variants) remain open:

(1) What is the computational complexity of the sweeping problem for n points? It was not obvious to us whether there is any (computable, exponential) upper bound on the number of sweeps required. Is the sweeping problem even decidable?

(2) If it is, does there exist a polynomial time algorithm for generating an optimal sweeping sequence, given n points? Can the number of sweeps in an optimal solution always be bounded by a polynomial in n? i.e., is there always an optimal solution with a polynomial number of sweeps?

Acknowledgment. We are grateful to Paweł Żyliński for sharing his dream problem with us. We would also like to thank the anonymous reviewers for their insightful comments.

References

[1] D. Halperin, L. Kavraki, and J.-C. Latombe: Robotics, in J. Goodman and J. O'Rourke (editors), *Handbook of Discrete and Computational Geometry*, pages 1065–1093, Chapman & Hall, second edition, 2004.

[2] F.K. Hwang, D.S. Richards, and P. Winter: *The Steiner Tree Problem*, vol. 53 of *Annals of Discrete Mathematics*, North-Holland, Amsterdam, 1992.

[3] G. Toussaint: Solving geometric problems with the rotating calipers, *Proceedings of Mediterranean Electrotechnical Conference* (MELECON '83), Athens.

[4] E. Welzl: The smallest rectangle enclosing a closed curve of length π, manuscript, 1993. Available at `http://www.inf.ethz.ch/personal/emo/SmallPieces.html`.

[5] P. Żyliński, personal communication, June 2007.