One Line and n Points

Bernd Gärtner* | József Solymosi† | Pavel Valtr
Falk Tschirschnitz* | Emo Welzl

ETH Zürich | Charles University, Prague

{gaertner,solymosi,tschirsc,emo}@inf.ethz.ch | valtr@kam.mff.cuni.cz

ABSTRACT

We analyze a randomized pivoting process involving one line and n points in the plane. The process models the behavior of the RANDOM-EDGE simplex algorithm on simple polytopes with n facets in dimension $n-2$. We obtain a tight $O(\log^2 n)$ bound for the expected number of pivot steps. This is the first nontrivial bound for RANDOM-EDGE which goes beyond bounds for specific polytopes. The process itself can be interpreted as a simple algorithm for certain 2-variable linear programming problems, and we prove a tight $\Theta(n)$ bound for its expected runtime.

The combinatorial structure behind the process is a directed graph over pairs of points, with arc orientations induced by the pivot steps. We characterize the class of graphs arising from one line and n points, up to oriented matroid realizability.

1. INTRODUCTION

Let S be a set of n points in general position in the plane (i.e., no three on a common line), and let ℓ be a vertical line which is disjoint from S and from all intersections of segments connecting pairs of points in S. We use $e \in \binom{S}{2}$ for a pair of points and $\overline{e}$ for the segment $\operatorname{conv} e$. e is called an ℓ-edge if ℓ separates its endpoints.

For a non-vertical line λ, $\operatorname{below}(\lambda)$ denotes the set of points from S below λ. For an edge e, we use $\operatorname{below}(e)$ for the set of points in S that lie below the line carrying $\overline{e}$. In fact, we will frequently use 'below $\overline{e}$' short for 'below the line carrying $\overline{e}$'. Given an ℓ-edge e and $p \in \operatorname{below}(e)$, we use $\operatorname{pivot}(e,p)$ for the unique ℓ-edge $\{p,q\}$, $q \in e$ (see Figure 1(a)).

We are interested in the randomized process given in Figure 2 which we will call *fast process.* It finds the unique ℓ-edge e such that $\operatorname{below}(e) = \emptyset$, given some initial ℓ-edge $\{a,b\}$. Figure 1(b) shows a possible sequence of edges that may arise during the process.

The significance of this process becomes apparent when we interpret it in terms of the simplex method for linear programming. If ℓ is the y-axis, the problem of finding the lowest ℓ-edge determined by points $p_i = (x_i, y_i), i = 1, \dots, n$ can be formulated as the following linear program (LP) in two variables k, d.

$$\begin{array}{ll} \text{maximize} & d \\ \text{subject to} & y_i \geq kx_i + d, \quad i = 1 \dots n. \end{array} \tag{1}$$

If ℓ has points on both sides, this is a bounded linear program, and optimal variable values k^*, d^* correspond to the line $y = k^*x + d^*$ carrying the lowest ℓ-edge. Moreover, the pivot steps performed by the randomized process correspond to pivot steps performed by the dual simplex method on the LP (1). Equivalently, they correspond to pivot steps performed by the regular *primal* simplex method on the *dual* of LP (1), which is a problem in n nonnegative variables $\mu_1, \dots, \mu_n$ and two equality constraints:

$$\begin{array}{ll} \text{minimize} & \sum_{i=1}^n y_i \mu_i \\ \text{subject to} & \sum_{i=1}^n x_i \mu_i = 0, \\ & \sum_{i=1}^n \mu_i = 1, \\ & \mu_i \geq 0, \qquad i = 1, \dots, n. \end{array} \tag{2}$$

In fact, *any* strongly feasible and strongly bounded linear program in n nonnegative variables and two equality constraints can be written in the form of (2), after a suitable coordinate transformation. Here, strongly feasible (bounded, resp.) means that the set of feasible solutions has an interior point (is bounded, resp.).

This means, the randomized process of Figure 2 models the behavior of the RANDOM-EDGE simplex algorithm on any linear pro-

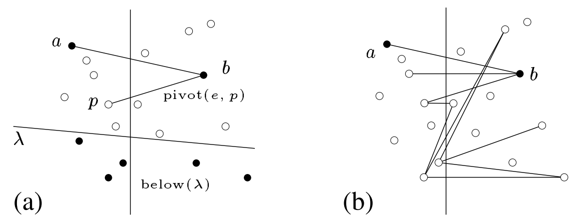

Figure 1: (a) The setup and (b) a pivoting sequence.

$e \leftarrow \{a,b\}$;
while $\operatorname{below}(e) \neq \emptyset$ **do**
  $p \leftarrow_{\text{random}} \operatorname{below}(e)$;
  $e \leftarrow \operatorname{pivot}(e,p)$;

Figure 2: Fast process.

---

*Part of this work was done while the authors were supported by the Swiss Science Foundation (SNF), project No. 21-50647.97.

†Supported by the joint Berlin/Zürich graduate program Combinatorics, Geometry, Computation, financed by German Science Foundation (DFG) and ETH Zürich. On leave from MTA SZTAKI.

STOC'01, July 6-8, 2001, Hersonissos, Crete, Greece.

gram whose feasible region is a simple n-facet polytope in dimension $n-2$. Here, a pivot step corresponds to a move from one vertex to an adjacent vertex along an incident edge chosen at random among the ones that improve the objective function.

Despite the fact that RANDOM-EDGE is the most natural randomized variant of the simplex method, and that it has received considerable attention ([8], Exercise 8.10* in [12], Section 9.10 in [11], [3],[5],[9]), nontrivial bounds for its expected performance could only be derived for specific instances [5]. In particular, the case of general n-facet polytopes in dimension d could only be handled for $n = d+1$ (when the polytope is a simplex—we come back to this case below). If $n = d+2$, as in our scenario, only the trivial $O(n^2)$ bound—obtained from the maximum number of vertices a polytope with this parameters can have—was known.

In this paper, we obtain three new results. First, we prove a tight bound of $O(\log^2 n)$ for the number of pivot steps needed by the randomized process of Figure 2 (equivalently, by RANDOM-EDGE on an n-facet polytope in dimension $n-2$). This is an exponential improvement over the previous bound.

Second, we analyze a variant of this process which we call the *slow process*. Here, each round chooses a random point among *all* points and only performs a pivot step if the chosen point is below the current edge (Figure 3).

$e \leftarrow \{a, b\}$;
while $\text{below}(e) \neq \emptyset$ **do**
  $p \leftarrow_{\text{random}} S$;
  if p below $\overline{e}$ **then**
    $e \leftarrow \text{pivot}(e, p)$;

Figure 3: Slow process.

From the analysis of the fast process, we get an $O(n \log^2 n)$ bound for the expected number of iterations of the slow process. We can prove that the bound is actually $\Theta(n)$, making the slow process an ultimately simple and optimal algorithm for 2-variable LP that can be transformed to (1).

For our third result, we interpret both processes as random walks on the directed graph whose nodes are the ℓ-edges, and whose arcs are the pairs $(e, \text{pivot}(e, p))$, for all ℓ-edges e and points p below e. The underlying undirected graph is what we will call a *grid graph* with n_L rows and n_R columns, where n_L (n_R, resp.) is the number of points to the left (right, resp.) of ℓ (see Figure 4(a)).

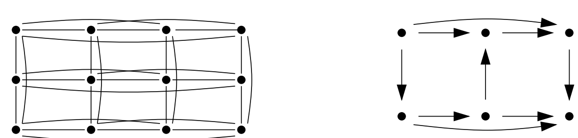
(a) (b)

Figure 4: (a) A grid graph with 3 rows and 4 columns and (b) the forbidden subgraph.

After pairing rows (columns) with points on the left (right) of ℓ in an arbitrary fashion, every node corresponds to some ℓ-edge, and the graph receives its orientation. We observe that if such a grid orientation is induced by a line and n points, then it is acyclic, every subgrid[1] has a unique sink, and no subgrid is isomorphic to the graph shown in Figure 4(b) (see Section 4.)

[1] A graph induced by all the nodes in some of the rows and columns.

On the other hand, if a grid orientation has these properties, we prove that it is induced by a line ℓ and n points in the sense that we find *pseudolines* through the pairs of points whose relative orders along ℓ determine the orientation, see Figure 5. Moreover, we give an example of such an orientation which is not 'properly' induced by a line and n points.

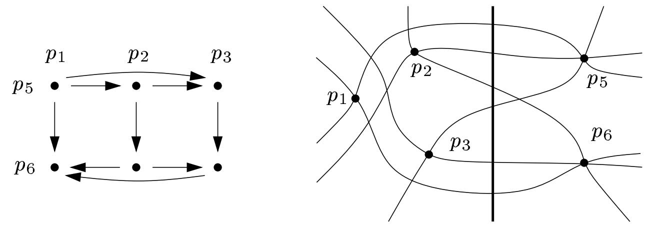

Figure 5: A grid orientation and its realization.

2. UPPER BOUNDS

Here we analyze the expected number of pivots in the process described in Figure 2.

Warm-up.

It is instructive to first have a look at the 1-dimensional counterpart as given in Figure 6. For that we are given a set M of real numbers, and $a \in M$. As above, this process can be interpreted in terms of the simplex algorithm; this time, it models the behavior of the RANDOM-EDGE variant on n-facet polytopes in dimension $n-1$, hence on simplices.

$r \leftarrow a$;
while $r \neq \min(M)$ **do**
  $r \leftarrow_{\text{random}} \{x \in M \mid x < r\}$;

Figure 6: 1-dimensional fast process.

It is not difficult[2] to show that the expected number of while-loops performed is exactly the Harmonic number H_{m-1}, where m is the rank, $1 + \#\{x \in M \mid x < a\}$, of a in M. But we want to provide a rough estimate in the spirit of the later analysis instead. Let X_i, $i \in \mathbf{N}_0$, be the random variable for the number of iterations of the while-loop with

$$2^i \leq \#\{x \in M \mid x < \rho\} < 2^{i+1} \qquad (3)$$

for ρ the value of r at the beginning of the while-loop. For $m \geq 2$, the random variable $Z = \sum_{i=0}^{\lfloor \log_2(m-1) \rfloor} X_i$ gives the overall number of executions of the while-loop. $\mathrm{E}(X_i) \leq 2$ for all $i \in \mathbf{N}_0$, since, whenever (3) holds, we have a chance of at least $\frac{1}{2}$ to choose an element of rank 2^i or smaller.
Hence, $\mathrm{E}(Z) \leq 2(1 + \lfloor \log_2(m-1) \rfloor) = O(\log m)$.

The obvious extension of that analysis to the 2-dimensional process fails, since the number of points below edges appearing in the process oscillates. In fact, the number of points below the current edge is no measure of progress at all. This number may be 1, we pivot, and the number becomes as large as $n-3$ (see Figure 2).

[2] For $m \in \mathbf{N}$, let f_m denote the expected number of iterations when starting with an element of rank m. Then $f_1 = 0$ and $f_{m+1} = 1 + \frac{1}{m} \sum_{i=1}^{m} f_i$. Now $f_m = H_{m-1}$ follows.

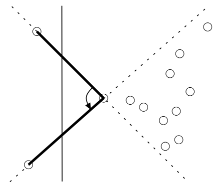

Figure 7: We thought we were so close!

k-Lines as milestones.

Here is the crucial definition that will allow us to measure progress. Given $k \in \mathbf{N}_0$, a non-vertical line λ is called *a k-line of S and ℓ* if on both sides of ℓ there are exactly k points from S below λ. It is easy to see that every point $x \in \ell$ is contained in a k-line for some $k \in \mathbf{N}_0$, as long as x is disjoint from all segments connecting two points in S. Start with a line through x that has large slope so that all points on the right side of ℓ are below, and all on the left side are above. Now rotate the line by decreasing its slope. Eventually, we will reach the situation opposite to what we started with: No points below to the right, all below to the left. All transitions in between change the number of points below on exactly one side by ± 1. Somewhere in between we must have had a transition where the numbers of points below were the same on both sides.

A k-line disjoint from S exists for all k,

$$0 \le k \le m := \min\{n_{\mathrm{L}}, n_{\mathrm{R}}\},$$

n_{L} the number of points in S left of ℓ, and $n_{\mathrm{R}} := n - n_{\mathrm{L}}$. For each i, $0 \le i \le \lfloor \log_2 m \rfloor$, fix some 2^i-line λ_i disjoint from S. Line λ_0 has to be chosen, so that the only edge intersecting ℓ below λ_0 is the edge ε with $\mathrm{below}(\varepsilon) = \emptyset$. Moreover, let $\lambda_{\lfloor \log_2 m \rfloor + 1}$ be some m-line that intersects ℓ above all ℓ-edges (and above[3] $\lambda_{\lfloor \log_2 m \rfloor}$). The line λ_i intersects ℓ below λ_j for $0 \le i < j \le \lfloor \log_2 m \rfloor + 1$.

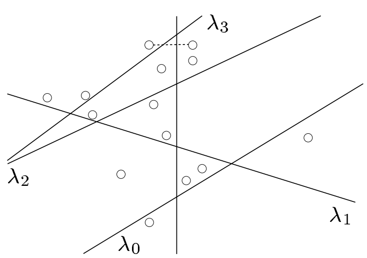

Figure 8: Setting milestones.

We define the random variable X_i, $i = 0, 1, \ldots, \lfloor \log_2 m \rfloor$, as the number of executions of the while-loop (in Figure 2) where the current segment $\overline{\varepsilon}$ intersects ℓ below λ_{i+1} but not below λ_i (ε is the value of e at the beginning of the while-loop). The sequence of these executions we call *phase i of the process*[4]. The careful choice of λ_0 ensures that completion of phase 0 entails completion of the whole process. Hence, $Z = \sum_{i=0}^{\lfloor \log_2 m \rfloor} X_i$ is the random variable whose expectation we want to analyze.
We will show that $\mathrm{E}(X_i) = O(\log n)$ for all i and, hence, $\mathrm{E}(Z) = O((\log n)(1 + \log m)) = O(\log^2 n)$.

Analysis of a single phase.

Fix some i, $0 \le i \le \lfloor \log_2 m \rfloor$, set $k = 2^i$, $\lambda' = \lambda_i$ and $\lambda = \lambda_{i+1}$. So λ' is a k-line, and there are at most $4k$ points below λ (actually exactly, unless $i = \lfloor \log_2 m \rfloor$). We have an edge intersecting ℓ not below λ' but below λ, and the phase starts. The phase ends whenever we reach an edge that intersects ℓ below λ'. Note that for every edge occurring in the phase, one endpoint has to be below λ (since the edge intersects ℓ below λ) and there is an endpoint above λ' (since otherwise, we are already in a new phase).

A few words on what we are heading for. We further split phase i into *strokes*. A stroke starts after we have sampled a point in $\mathrm{below}(\lambda) \cup \mathrm{below}(\lambda')$ (or at the very beginning of the phase) and it finishes, when the next point in $\mathrm{below}(\lambda) \cup \mathrm{below}(\lambda')$ is chosen (or the phase ends; so the last stroke may be empty). If N is the number of strokes, then we can write $X := X_i$ as

$$X = Y_1 + Y_2 + \cdots + Y_N$$

where Y_j is the number of iterations of the jth stroke. Note that N itself is a random variable. We will show that (i) $\mathrm{E}(Y_j \,|\, j \le N) \le O(\log n)$ for all j, and (ii) $\mathrm{E}(N) = O(1)$. It follows that $\mathrm{E}(X) = O(\log n)$.

$$\begin{aligned} \mathrm{E}(X) &= \sum_{j=1}^{\infty} \overbrace{\mathrm{E}(Y_j \,|\, j \le N)}^{O(\log n)} \Pr(j \le N) \\ &= O(\log n) \sum_{j=1}^{\infty} \Pr(j \le N) \\ &= O(\log n)\, \mathrm{E}(N) \;. \end{aligned}$$

As for the points sampled from $\mathrm{below}(\lambda) \cup \mathrm{below}(\lambda')$ we distinguish several cases depending on where the respective new point pivoted into the current edge lies. We will see that each of these situations are more or less promising in our goal to escape this phase.

Here are the steps in our reasoning: *At any point of the phase*, the following four claims hold.

CLAIM 1. *The expected number of pivots until we sample a new point in*

$$\mathrm{below}(\lambda)$$

is at most $2 \log_2 n$.

PROOF. At least one of the two endpoints of the current edge has to be below λ. So in a subsequence where the new point is contiguously always chosen above λ, the other endpoint below stays the same throughout this sequence. We denote this point by q. If we order the points on the other side of ℓ according to their visibility from q, we get almost the situation as in the one-dimensional process described in Figure 6. In fact, there are two differences which can only improve our expectations: We terminate not only in the lowest point but also in $k - 1$ points. In each step we may also sample on the q's side of ℓ, in which case we immediately terminate (we have surely sampled below λ). Hence, the expected length of such a subsequence is at most[5] $2 \log_2 n$. $\square$

CLAIM 2. *Conditioned on the event that we sample a point in*

$$\mathrm{below}(\lambda) \cup \mathrm{below}(\lambda') \,,$$

the point will be in

$$\mathrm{below}(\lambda')$$

with probability at least $\frac{1}{5}$.

PROOF. Since all edges in this phase intersect ℓ not below λ' it follows: For one side of ℓ, all k points below λ' must also lie below the line through the current edge. That is, *at least* k points below λ' are also below the line through the current edge. On the

[3] This is automatically satisfied, unless m is a power of 2.

[4] Note that phases count *down* during the process.

[5] Even H_{n-m-1} is true.

other hand, *at most* $5k$ points are below λ or λ'. This holds, since $\#\text{below}(\lambda) \leq 4k$, $\#\text{below}(\lambda') = 2k$, and on one side of ℓ, all k points below λ' are also below λ. $\square$

Claims 1 and 2 combined assure that we reach a point below λ' within an expected number of at most $10 \log_2 n$ steps.

So what happens after we see such a point p below λ'? Two cases have to be distinguished, depending on whether p is also below λ or not.

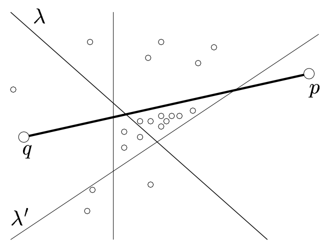

Figure 9: $p \in \text{below}(\lambda') \setminus \text{below}(\lambda)$.

CLAIM 3. *If an endpoint p of the current edge is in*

$$\text{below}(\lambda') \setminus \text{below}(\lambda)\,,$$

then the next point sampled below λ or λ' will be in

$$\text{below}(\lambda) \cap \text{below}(\lambda')$$

with probability at least $\frac{1}{5}$.

PROOF. Two relevant conclusions right away (see Fig. 9): (i) Since p is not below λ and below λ', while λ' intersects ℓ below λ, the lines λ and λ' must intersect on p's side. (ii) Since p is not below λ, the other endpoint q of the current edge has to be.

Before q can be substituted by a point not below λ, the other endpoint has to be below λ. That is, when we first sample a point below λ or λ', point q is still in the edge. But then, on q's side, all k points below λ' are also below the then current edge (that requires a careful argument[6] – for the time being consult Figure 9). Moreover, on q's side, all points below λ' are also below λ (since these lines intersect on the other side). So, summing up, the k points below λ' are both below the current edge and below λ, and they are at disposal, when we sample a point below λ or λ'. The claim follows. $\square$

CLAIM 4. *If an endpoint p of the current edge is in*

$$\text{below}(\lambda) \cap \text{below}(\lambda')\,,$$

then the next point sampled below λ or λ' will be in

$$\text{below}(\lambda') \text{ on the side opposite to } p$$

with probability at least $\frac{1}{5}$.

PROOF. If p is substituted in a pivot, it must be substituted by a point below λ or λ'. This holds, since on p's side of ℓ, everything below the current edge has to be below λ or λ' (again, this needs proof, but see Figure 10). As a consequence, until the first pivot with a point below λ or λ', point p is still an endpoint of the edge. But since p is below λ', on the opposite side everything below λ' is also below the current edge. So there are at least k good choices, and at most $5k$ choices of points below λ or λ'. $\square$

[6]It uses the fact that q was connected to p, a point below λ', and the current edge connects q to a point below the edge $\{p, q\}$. From that it can be shown that the line carrying the current edge must intersect λ' on p's side – thus not on q's.

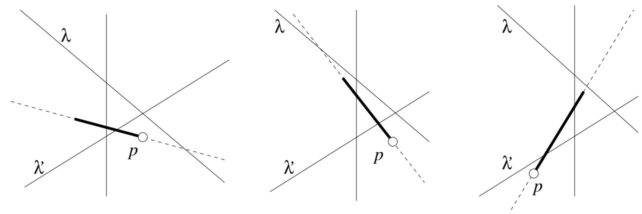

Figure 10: $p \in \text{below}(\lambda') \cap \text{below}(\lambda)$.

Claim 4 entails that if we choose a point below λ' and λ, then – with probability at least $\frac{1}{5}$ the next point chosen below λ or λ' will terminate the phase.

To complete the argument, we look at the sequence of points from $\text{below}(\lambda) \cup \text{below}(\lambda')$ that are pivoted into the current edge. Recall that these are exactly the points that terminate the strokes of a phase (except for the last one). If we can show that the expected length of this sequence is at most some constant c, then the expected length of the whole sequence is at most $2c \log_2 n$ due to Claim 1. Each point in this sequence is classified depending on whether it lies in

$$\begin{array}{ll} \text{Class 0:} & \text{below}(\lambda) \setminus \text{below}(\lambda') \\ \text{Class 1:} & \text{below}(\lambda') \setminus \text{below}(\lambda) \\ \text{Class 2:} & \text{below}(\lambda') \cap \text{below}(\lambda) \end{array}$$

Every point in the sequence considered is in Class 0, 1, or 2. If we have a point in Class 0, the next will be in Class 1 or 2 with probability at least $\frac{1}{5}$. If we have a point in Class 1, the next will be in Class 2 with probability at least $\frac{1}{5}$. (All of this of course conditioned on the event that a next point exists at all, i.e. the phase hasn't stopped already.) Finally, if we are in Class 2, it is the last point in the sequence with probability at least $\frac{1}{5}$.

Now we estimate the expected length of the sequence by the Markov chain[7] depicted in Figure 11, with four states

$$\mathsf{start} = 0,\;\; 1,\;\; 2,\;\; \text{and}\;\; 3 = \mathsf{stop},$$

and the indicated transition probabilities.

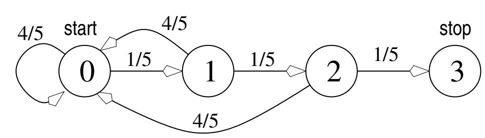

Figure 11: A pessimistic Markov chain.

On one hand, it is easy to calculate that the expected number of steps from start to stop is 155. On the other hand, the chain and our sequence can be coupled so that whenever the chain is in state $s \in \{0, 1, 2, 3\}$, then the corresponding point in the sequence in Class $t \geq s$, or the sequence has ended already. Hence, we have shown that the expected number of pivots in a single phase is bounded by $310 \log_2 n$, and the theorem follows.

THEOREM 2.1. *The expected number of pivots in the process defined in Figure 2 is at most*

$$O((\log n)(1 + \log m)) = O(\log^2 n),$$

where n is the number of points, and m is the smaller of the numbers of points on the two sides of the line.

[7]It simulates a biased coin with success probability $\frac{1}{5}$ and counts the number of experiments until we have three consecutive successes.

The slow process.

If we consider the process in Figure 2 as an algorithm, it remains to specify how to sample from $\mathrm{below}(e)$. It is perhaps worthwhile to mention on the side, that if we can sample efficiently (say in time logarithmic in n), then the process gives a polynomial time algorithm even for exponential size sets. If we sample in the obvious way in $O(n)$ time, then this gives us an $O(n(\log n)^2)$ algorithm. On that we can improve by looking at the alternative slow process[8] in Figure 3. The number of pivots has the same distribution as the process in Figure 2. Here we analyze the number of iterations of the while-loop. To that end, define phases as previously done. Recall that in phase i, the current edge intersects ℓ not below the 2^i-line λ_i.

CLAIM 5. *For each i, the expected number of iterations in phase i is at most $O(\frac{n}{2^i})$.*

PROOF. Divide phase i into strokes as we did it before. That is, a stroke is ended, whenever we sample a point in

$$(\mathrm{below}(\lambda') \cup \mathrm{below}(\lambda)) \cap \mathrm{below}(\varepsilon),$$

(ε the current value of e), or when the phase ends.

In phase i, there are always at least 2^i points from $\mathrm{below}(\lambda')$ that lie below the current edge (on some side of ℓ, all points below λ_i are also below the current edge). That is, at any point, we sample a point resulting in the termination of the stroke with probability at least $\frac{2^i}{n}$. Therefore, the expected number of iterations in a stroke is at most $\frac{n}{2^i}$. The number of strokes is, of course, the same as in the slow process; its expectation is constant. The claim follows. $\square$

THEOREM 2.2. *The expected number of iterations of the process defined in Figure 3 is*

$$\Theta(n)$$

where n is the number of points, unless the starting edge $\{a, b\}$ is already the lowest ℓ-edge.

PROOF. $\sum_{i=0}^{\lfloor \log_2 m \rfloor} \frac{n}{2^i} < 2n$ and so the upper bound follows from Claim 5. If $\{a, b\}$ is disjoint from the lowest ℓ-edge, a lower bound of $\frac{3}{2}n$ is obvious, since on the average it takes that long until we have sampled both endpoints of the lowest edge at least once. Even if $\{a, b\}$ contains exactly one of the two endpoints of the lowest ℓ-edge, we still need n steps on the average before we meet the other endpoint for the first time. The lower bound follows. $\square$

The coupon collector analysis tells us that it takes $\Theta(n \log n)$ iterations until we expect to have sampled each point at least once. The process analyzed finds the lowest ℓ-edge much before all points have been seen at least once.

3. A LOWER BOUND

Here we prove a lower bound on the expected number of pivot steps in the process of Figure 2. This is our construction: Let ℓ be the y-axis. We place $4n$ points $p_0, \ldots, p_{2n-1}, s_0, \ldots, s_{2n-1}$ onto the lines $x = \pm 1$ and $x = \pm n$ in such a way that $\{p_j \mid j \geq n\} \in \mathrm{below}(\{p_i, s_j\})$ whenever $j < i < n$, and $\{s_j \mid j \geq n\} \in \mathrm{below}(\{p_i, s_j\})$, whenever $i < j < n$; cf. Figure 12. Finally, to get general position, we slightly perturb the points.

[8]The '$\mathrm{below}(e) \neq \emptyset$'-test can be made once in n rounds only, thus causing amortized constant cost. Or, after every pivot, we can go through all points in random order (without replacement) until we find the first point in $\mathrm{below}(e)$; if no such point is found, we are done. Compared to the 'pure version', this can only speed up the procedure.

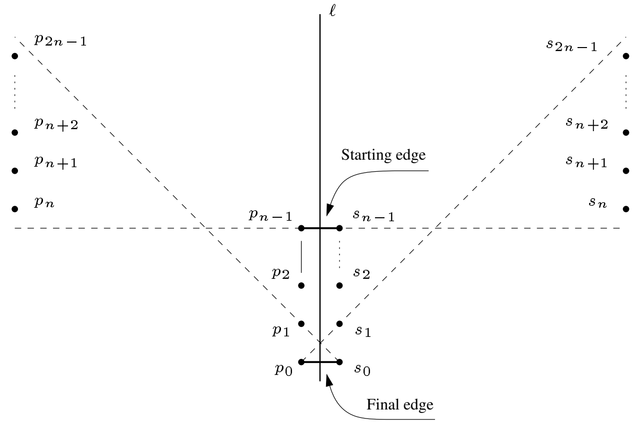

Figure 12: Instance for the lower bound.

Again, we divide the sequence of pivots into distinct phases. We are in phase a as long as $\min(i, j) = a$, where i, j are the indices of the points defining the current edge e.[9]

We define the random variables $X_i, i = 0, \ldots, n-1$, as the number of pivots during the phase i, Z shall denote the total number of pivots under the assumption that we start with the edge $\{p_{n-1}, s_{n-1}\}$. We aim to compute the expected value of Z, but the approach $\mathrm{E}(Z) = \sum_{i<n} \mathrm{E}(X_i)$ turns out to be unsuitable. We therefore introduce the auxiliary random variables Y_i, $Y_i = X_i +$ (# of pivots in the consecutive phase), $i = 1, \ldots, n-1$. $\overline{X_i}, \overline{Y_i}$ shall denote the random variables X_i, Y_i, respectively, under the condition that phase i is actually entered. Since

$$\Pr(\text{Phase } i \text{ is entered}) = \begin{cases} \frac{1}{i+1} & \text{for all } i, 0 \leq i < n-1 \\ 1 & \text{for } i = n-1 \end{cases}$$

we get

$$\begin{aligned} 2\,\mathrm{E}(Z) &= 2\sum_{i<n} \mathrm{E}\big(\overline{X_i}\big) \Pr(\text{Phase } i \text{ is entered}) \\ &= \mathrm{E}\big(\overline{X_{n-1}}\big) + \mathrm{E}\big(\overline{Y_{n-1}}\big) + \sum_{i<n-1} \frac{\mathrm{E}\big(\overline{Y_i}\big)}{i+1} + \mathrm{E}\big(\overline{X_0}\big) \\ 2\,\mathrm{E}(Z) &\geq \sum_{i=1}^{n-2} \frac{\mathrm{E}\big(\overline{Y_i}\big)}{i+1}. \end{aligned} \tag{4}$$

To estimate the expectations of the $\overline{Y_i}$ we first compute $t_{r,s}$, the expected number of flips until we leave the current phase $a = \min(r, s)$, where r and s denote the number of points on the left and on the right of ℓ below the current edge. W.l.o.g. $a = s < r$, say. The simple recursion

$$t_{r,s} = \begin{cases} 0 & \text{whenever } r = s = 0 \\ 1 + \frac{1}{r+s} \sum_{r'=s}^{r-1} t_{r',s} & \text{otherwise} \end{cases}$$

gives us

$$t_{r,s} = \begin{cases} \mathrm{H}_r & \text{whenever } s = 0 \\ \mathrm{H}_{r+s} - \mathrm{H}_{2s} + 1 & \text{otherwise.} \end{cases}$$

We have to assume the worst: possibly the first phase is over after just one flip! But for the second phase with probability $p = \frac{1}{2}$ the minimum will lie on the other side of ℓ than before – which puts n points ($p_n, \ldots, p_{2n-1}$ or $s_n, \ldots, s_{2n-1}$) back into the game.

[9]Note that we follow the general trend: Phases count *down* during the process.

We can therefore deduce that

$$\begin{aligned}
&\mathrm{E}\left(\overline{Y}_i\right) \\
&\geq 1 + \frac{1}{2i}\left[\underbrace{\sum_{0\leq i'<i} 1}_{\text{new } i'} + \underbrace{\sum_{0<j'<i}\left(1+\mathrm{H}_{(n+i)+j'} - \mathrm{H}_{2j'}\right) + \mathrm{H}_{n+i}}_{\text{new } j'}\right] \\
&\geq \frac{1}{2}\left[\mathrm{H}_{n-i+1} - \mathrm{H}_{2i}\right] \\
&\quad + \frac{1}{2i}\left[\mathrm{H}_{n+i} - \mathrm{H}_{n-i+1} + n\left(\mathrm{H}_n - \mathrm{H}_{n-i+1}\right)\right] \\
&\geq \frac{1}{2}\left[\mathrm{H}_{n-i+1} - \mathrm{H}_{2i}\right]
\end{aligned}$$

Using this in (4) we get:

$$\begin{aligned}
\mathrm{E}(Z) &\geq \frac{1}{2}\sum_{i<n-1}\frac{\mathrm{E}\left(\overline{Y_i}\right)}{i+1} \\
&\geq \frac{1}{4}\sum_{i<n-1}\frac{\mathrm{H}_{n-i+1} - \mathrm{H}_{2i}}{i+1} \\
&\geq \frac{1}{4}\left[\sum_{0<i<n}\left(\frac{\mathrm{H}_n}{i+1} - \frac{1}{i+1}\left(\frac{1}{n} + \cdots + \frac{1}{n-i+2}\right)\right)\right. \\
&\qquad \left. - \sum_{0<i<n}\frac{\mathrm{H}_{2i} + \mathrm{H}_{2i+1}}{2(i+1)}\right] \\
&\geq \frac{1}{4}\left[\mathrm{H}_n^2 - 2\mathrm{H}_n - \frac{1}{2}\left(\mathrm{H}_{2n}^2 + \sum_{i=1}^{2n-1}\frac{1}{i^2}\right)\right] \\
&\geq \frac{1}{4}\left(\mathrm{H}_n^2 - \frac{1}{2}\mathrm{H}_{2n}^2\right) + O(\log n).
\end{aligned}$$

THEOREM 3.1. *There exist instances of one line and n points in general position in the plane such that the expected number of pivot steps for the fast process described in Fig. 2 satisfies the following bound:*

$$\mathrm{E}(Z) \geq \frac{1}{8}\log^2 n + O(\log n).$$

Before we move on to the next section we would like to describe another noteworthy instance for which this lower bound holds: Let ℓ be the y-axis. We place $2n$ points $p_0, \ldots, p_{n-1}, s_0, \ldots, s_{n-1}$ onto the graph of the function $y = f(x) = \log|x|$ such that $x_{p_j} < x_{p_i} < 0 < x_{s_i} < x_{s_j}$ for any $i < j < n$. Furthermore, and most importantly, $\{p_j | j \neq i\} \in \mathrm{below}(\{p_i, s_i\})$, and $\{s_j \mid j \neq i\} \in \mathrm{below}(\{p_{i+1}, s_i\})$, cf. Figure 13.

Note that there exists a sequence of pivots which visits all possible ℓ-edges! That is, there are simple $(n-2)$-polytopes with n facets where we can pivot through all vertices in a monotone fashion (w.r.t. some linear function).

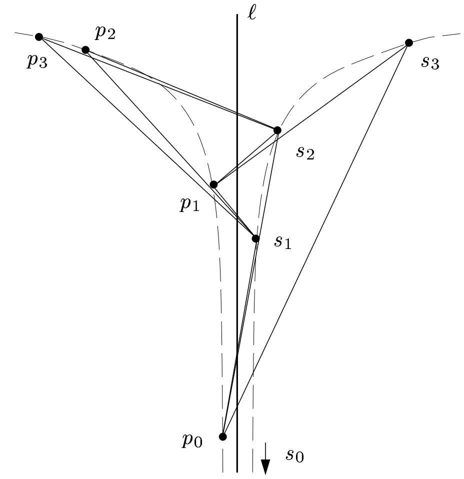

Figure 13: Another instance for the lower bound.

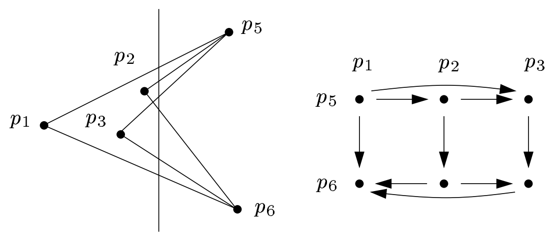

Figure 14: Grid orientation.

4. ADMISSIBLE GRID ORIENTATIONS

Every instance of one line and n points gives us a *grid orientation*—the directed graph behind the randomized processes we have considered in the previous sections, see Figure 14.

Furthermore, the following three properties are necessary:

1. The orientation is acyclic.
2. Every nonempty subgrid has a unique sink.
3. No subgrid is isomorphic to the 'forbidden subgrid' depicted in Figure 4(b).

Properties 1 and 2 are fairly easy to see; in fact, by dualizing our scenario, a grid orientation can also be used to describe the edge orientations that a simple n-facet polytope in dimension $n-2$ may assume under a linear objective function. In this view, properties 1 and 2 boil down to *abstract objective functions* (AOF)([7], [13]).

Property 3 can be obtained either by inspection of all possible classes of instances with 5 points, or by observing that, in the dual view, this condition is equivalent to the so-called *path condition* found by Holt and Klee [6]: Source and sink of the 1-skeleton of a polytope oriented by a generic linear function can be connected by d vertex-disjoint paths. This condition is necessary for an AOF (a grid orientation, in our view) to be induced by a linear objective function (an instance of one line and n points, in our view).

Let us call a grid orientation with properties 1–3 *admissible*. A natural question is whether any admissible grid orientation is actually induced by one line and n points. (In this case, we call the orientation *realizable*.) As we show in this section, this is not the case. Invoking duality again, this also shows that Holt and Klee's path condition is not sufficient for an AOF to be induced by a linear objective function—a fact that is not surprising, but that has to our knowledge not been worked out so far. (For 3-polytopes, Klee and Mihalisin have proved that the path condition *is* sufficient [10].)

Hyperline sequences.

Let $\mathcal{C}$ be a planar configuration of n distinct points in general position, indexed by $E_n = \{1, \ldots, n\}$. Adding to E_n the new elements $\{\bar{i} | i \in E_n\}$, we get the *signed index set* $\overline{E}_n$ (with $\bar{\bar{s}} = s, \forall s \in \overline{E}_n$). $\operatorname{conv}\mathcal{C}$ shall denote the convex hull of $\mathcal{C}$.

Rotating an oriented line in counterclockwise order around p_i, $i \in E_n$, and looking at the successive positions where it coincides with lines defined by pairs of points (p_i, p_j) defines the *hyperline sequence* π_i over $\overline{E}_n$, cf. [2]. If a point p_j is encountered by the rotating line in positive direction from p_i, it will be recorded as a positive index j, otherwise as a negative index $\overline{j}$. The whole sequence is recorded in the order induced by the rotating line, and an arbitrary half-period is chosen to represent it, cf. Figure 15 for a simple example.

Note that the hyperline sequence of a point which belongs to the convex hull of the point configuration has a half-period consisting only of positive indices. In other words, it is represented by some permutation of the other points.

Similarly, we can define for any additional point s the hyperline sequence of s with respect to $\mathcal{C}$. If s is outside $\operatorname{conv}\mathcal{C}$, and not on any of the lines defined by two points of $\mathcal{C}$, this sequence is represented by a permutation $\pi \in \mathcal{S}_n$.

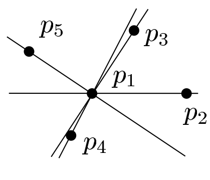

Figure 15: $\pi_1 = 23\overline{4}5$.

Suppose we are traveling on a directed curve which does not intersect $\operatorname{conv}\mathcal{C}$, but which intersects every line spanned by two points of $\mathcal{C}$ exactly once. We call such a curve a *directed pseudoline passing* $\mathcal{C}$.

Along this pseudoline, we get a sequence Σ of hyperline sequences w.r.t. $\mathcal{C}$, represented by permutations over the n points.[10] When we cross the line defined by the pair of points $\{p_i, p_j\}$, we get a new permutation, arising from the previous one by transposing the adjacent elements i and j. In particular, the last permutation is the reversal of the first one. We say that Σ is witnessed by the pseudoline.

We are, in fact, dealing with simple *allowable sequences* of permutations, also known as *wiring diagrams* [1]; this is illustrated in the example in Figure 16. Note that an allowable sequence can compactly be encoded by a sequence of transpositions $ij := (i, j)$; we write $ij < kl$ if the former transposition happens earlier in the sequence.

We call an allowable sequence *realizable* if there is a point configuration $\mathcal{C}$ such that the sequence can be witnessed by some directed pseudoline $\mathcal{L}$ passing $\mathcal{C}$. This notion is illustrated in Figure 17. W.l.o.g. we shall assume that the configuration of points lies *to the right* of the directed pseudoline.

For any realizable sequence Σ and a realizing point configuration $\mathcal{C}$, the following facts are easily seen to hold.

1. Point p_i is extreme in $\mathcal{C}$, that is, $\operatorname{conv}\mathcal{C} \neq \operatorname{conv}(\mathcal{C} - \{i\})$, if and only if there is a permutation of Σ that has i in its first or last position.

2. Two points p_i, p_j with $i \neq j$ are endpoints of a common edge of $\operatorname{conv}\mathcal{C}$ if and only if there are consecutive permutations σ and σ' in Σ such that σ has i in its first (or its last) and σ' has j in its first (or its last, resp.) position.

[10] We only record the hyperline sequences obtained from points not on any of the lines determined by $\mathcal{C}$.

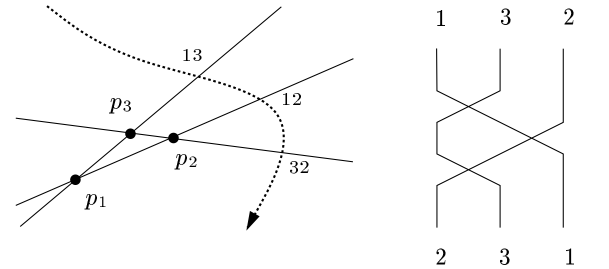

Figure 16: Allowable sequence witnessed by a pseudoline; represented by a sequence of transpositions (left) and a wiring diagram (right).

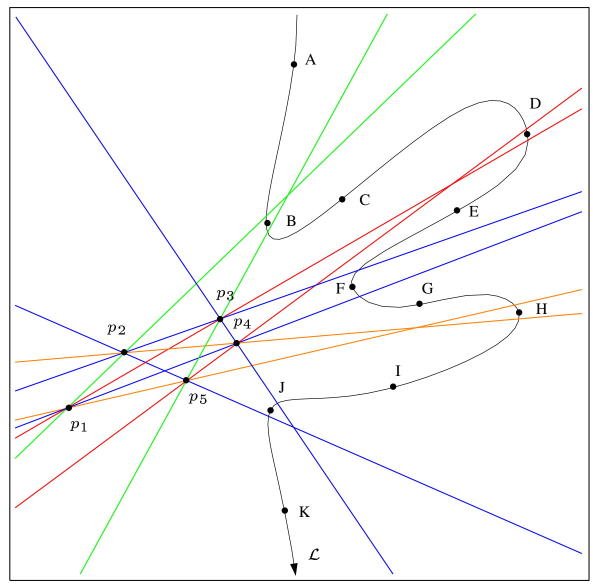

Figure 17: The sequence given by the moves $12, 53, 54, 13, 23, 14, 15, 24, 34, 25$ **is witnessed by the pseudoline** $\mathcal{L}$ **passing through the points** $A, B, \ldots, K$**, hence realizable.**[11]

3. If Σ contains the transposition ij, the ray r_{ij}—starting from p_i and going through p_j—intersects the witnessing pseudoline.

4. Assume that p_i, p_j, p_k, p_l are the vertices of a convex quadrilateral in counterclockwise orientation, and that $kl < ji$ in Σ. Then the rays r_{ji} and r_{kl} intersect. (In particular, they intersect before they meet the witnessing pseudoline.)

Allowable sequences and admissible grid orientations.

Given an allowable sequence Σ, there is a natural way to get an AOF-p-graph with $\binom{n}{2}+1$ rows and n columns: we order the nodes in row i according to the i-th permutation in Σ, while the nodes in column j are ordered from top to bottom, see Figure 18. Let us denote the resulting graph by $G(\Sigma)$.

[11] This illustrates that our notion of *realizability* deviates from e.g. the realizability of circular sequences as in [4].

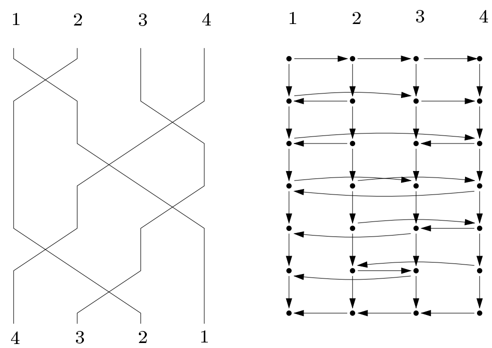

Figure 18: Each allowable sequence defines an admissible grid orientation.[12]

LEMMA 4.1. *For any allowable sequence Σ, $G(\Sigma)$ is an admissible grid orientation.*

PROOF. It is obvious that $G(\Sigma)$ is acyclic, and that every subgrid has a unique sink (namely the unique sink of its bottommost row). Moreover, as we witness the reversal of a pair $\{i, j\}$ only once in Σ, the forbidden 2×3-subgrid can clearly not occur. $\square$

A non realizable admissible orientation.

Consider the graph $G(\Sigma)$ with $\Sigma = 651423 \xrightarrow{65} 561423 \xrightarrow{23} 561432 \xrightarrow{14} 564132 \xrightarrow{64} 546132 \xrightarrow{13} 546312 \xrightarrow{54} 456312 \xrightarrow{12} 456321 \xrightarrow{63} 453621 \xrightarrow{62} 453261 \xrightarrow{53} 435261 \xrightarrow{43} 345261 \xrightarrow{61} 345216 \xrightarrow{52} 342516 \xrightarrow{51} 342156 \xrightarrow{42} 324156$. Assume there is a point configuration $\mathcal{C} = \{p_1, \ldots, p_6\}$ and a directed pseudoline $\mathcal{L}$ passing $\mathcal{C}$ such that $\mathcal{L}$ witnesses Σ.

By Facts 1 and 2 above, each p_i is extreme, and $p_1, p_2, p_3, p_4, p_5, p_6$ is the order of the points along $\operatorname{conv} \mathcal{C}$. We can even deduce from Fact 3 that we must have them in counterclockwise order, because e.g. ray r_{65} intersects $\mathcal{L}$ before ray r_{54} does.

Then, by Fact 4, the rays r_{14} and r_{23}, r_{63} and r_{12}, as well as r_{52} and r_{61} intersect in points M, N and P (before they meet $\mathcal{L}$).

The combinatorial structure of the point configuration is, therefore, as in Figure 19.

We want to argue that it is impossible to draw this picture with straight lines connecting the points. For this, suppose that each point p_i is given in normalized homogeneous coordinates, i.e. in the form $p_i = (x_i\, y_i\, 1)^T$.

We use the bracket notation $[\,P_i\, P_j\, P_k]$ as shorthand for the orientation determinant $\det |P_i P_j P_k|$ which equals the (oriented) area of the triangle spanned by the three points; its sign indicates the orientation of the triple, which is positive if and only if they appear in counterclockwise order. (For this to hold, the homogenizing coordinates of all three points must be positive.)

It is easy to see that M can be expressed as

$$M = [p_2\, p_3\, p_1]\, p_4 - [p_2\, p_3\, p_4]\, p_1.$$

We even know that M has positive homogenizing coordinate in this representation, because $[p_2\, p_3\, p_1] > [p_2\, p_3\, p_4]$. This again follows from the fact that p_1, p_4, M appear in this order along $\overline{p_1 p_4}$. By the same reasoning, the homogenizing coordinates of

$$N = [p_1\, p_2\, p_6]\, p_3 - [p_1\, p_2\, p_3]\, p_6$$
$$\text{and} \quad P = [p_6\, p_1\, p_5]\, p_2 - [p_6\, p_1\, p_2]\, p_5$$

[12]For better visibility we shall omit in our drawings all directed edges whose direction is non ambiguous.

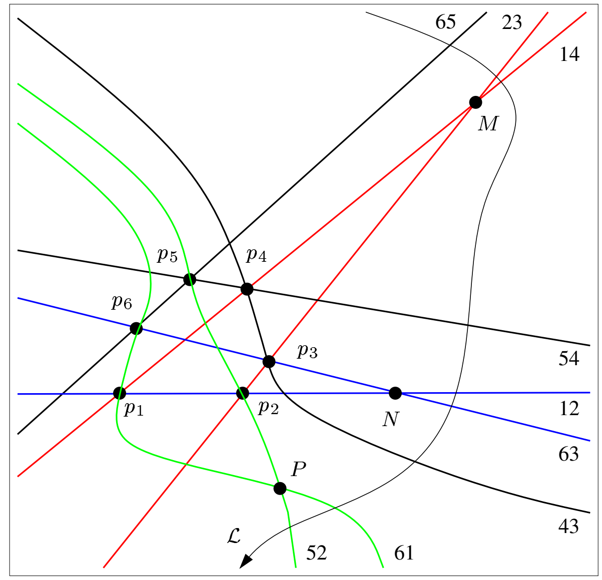

Figure 19: Constraints on the location of the 6 points.

are positive, too. As $\mathcal{L}$ intersects the ray r_{65} before it intersects any of r_{14}, r_{23}, M lies to the right of r_{65}; equivalently, $[\,M\, p_5\, p_6] > 0$ must hold. This can be rewritten as

$$\begin{aligned} 0 &< [\,M\, p_5\, p_6] \\ &= [(\,[p_2\, p_3\, p_1]\, p_4 - [p_2\, p_3\, p_4]\, p_1)\, p_5\, p_6] \\ &= [p_2\, p_3\, p_1]\,[p_4\, p_5\, p_6] - [p_2\, p_3\, p_4]\,[p_1\, p_5\, p_6]\,. \end{aligned} \tag{5}$$

Analogously, we see that $[\,N\, p_4\, p_5] > 0$, which can be rewritten as

$$[p_1\, p_2\, p_6]\,[p_3\, p_4\, p_5] - [p_1\, p_2\, p_3]\,[p_6\, p_4\, p_5] > 0, \tag{6}$$

and $[\,P\, p_3\, p_4] > 0$, which yields

$$[p_6\, p_1\, p_5]\,[p_2\, p_3\, p_4] - [p_6\, p_1\, p_2]\,[p_5\, p_3\, p_4] > 0. \tag{7}$$

Adding (6) and (7), we get

$$[p_6\, p_1\, p_5]\,[p_2\, p_3\, p_4] - [p_1\, p_2\, p_3]\,[p_6\, p_4\, p_5] > 0,$$

a contradiction to (5).

5. PSEUDO-REALIZABILITY

We have seen that not all admissible grid orientations are induced by one line and n points. However, in this section, we show that this is 'almost' true. To simplify notation, we will assume that the gridgraph has L rows, associated with indices $\mathcal{L} = \{i, j, k, \ldots\}$, and R columns, indexed with capital letters $\mathcal{R} = \{I, J, K, \ldots\}$. We use Greek letters $\alpha, \beta, \gamma, \ldots$ for any other indices. Thus, nodes correspond to pairs (i, I), and we write $v \to w$ to indicate the fact that there is a directed edge from node v to node w.

For a given admissible grid orientation G, we construct point sets $S_L = \{p_i, p_j, p_k, \ldots\}$ and $S_R = \{p_I, p_J, p_K, \ldots\}$ ($S_L \cup S_R = S$) of sizes L and R ($L + R = n$), respectively, separated by a vertical line ℓ. Along with this, we will have an arrangement $\mathcal{A}$ of $\binom{n}{2}$ oriented pseudolines λ_{pq} through all pairs of points $p, q \in S$.

The conditions are the following:

- If $(i, I) \to (i, J)$ in G, then p_J is to the right of λ_{p_i, p_I}, and
- If $(i, I) \to (j, I)$ in G, then p_j is to the right of λ_{p_i, p_I}.

This means, G is *pseudo-realizable*, leading to 'realizations' as the one in Figure 5. In the following, we will not construct S and $\mathcal{A}$ directly; rather, we construct hyperline sequences, from which we can deduce the existence of suitable S and $\mathcal{A}$ by using a connection between hyperline sequences and oriented matroids.

Hyperline sequences revisited.

The rule to create a set of hyperline sequences for a point configuration can be carried over to the case where the points are joined by pseudolines: The hyperline sequence π_i for a point $p_i, i \in E_n$ describes the order in which we encounter the pseudolines passing through p_i and some other point p_j in counterclockwise order around p_i.

Given an abstract set of hyperline sequences (not necessarily coming from a point configuration with pseudolines), we want to deduce orientations for triples α, β, γ. For this, we define $\chi : \overline{E}_n^3 \to \{-1, 1\}$ (partially) by setting $\chi(\alpha\beta\gamma) := 1$ if β and γ occur in this order in π_α, and $\chi(\alpha\beta\gamma) := -1$ if β and γ occur in reverse order in π_α. We can write this as

$$\chi(\alpha\beta\gamma) := \begin{cases} 1, & \text{if } \alpha : ..\beta..\gamma.., \\ -1, & \text{if } \alpha : ..\gamma..\beta... \end{cases}$$

We also define $\chi(\alpha\gamma\beta) = -\chi(\alpha\beta\gamma)$. Extending this partial definition by using the rules

$$\chi(\alpha\overline{\beta}\gamma) = \chi(\alpha\beta\overline{\gamma}) = -\chi(\alpha\overline{\beta\gamma}) = -\chi(\alpha\beta\gamma),$$

we obtain values $\chi(\alpha\beta\gamma)$ for all distinct $\alpha, \beta, \gamma \in E_n$. We call χ consistent if $\chi(\alpha\beta\gamma) = \chi(\sigma(\alpha)\sigma(\beta)\sigma(\gamma))$ exactly for the even permutations $\sigma \in S_3$. In this case, we call χ an *abstract determinant function* [2], and we say that the hyperline sequence admits an abstract determinant function. It has been shown in [2] that an abstract determinant function fulfills the axioms of a *rank-3-chirotope*. Oriented matroid theory [1] then guarantees that this chirotope has a representation as a pseudoconfiguration of points which is consistent with the chirotope information (and hence with the hyperline sequence). That is, there are points $S = \{p_\alpha, p_\beta, p_\gamma, \dots\}$ and oriented pseudolines $\mathcal{A}$ through pairs of points such that $\chi(\alpha\beta\gamma)$ is positive (negative) if and only if p_γ is to the left (to the right) of the pseudoline through p_α and p_β.

Therefore, the goal now is to derive from a given admissible grid orientation G a hyperline sequence which admits an abstract determinant function whose values on 'mixed' triples are consistent with the grid orientation, meaning that

$$(i, I) \to (i, J) \iff \chi(iJI) > 0, \tag{8}$$
$$(i, I) \to (j, I) \iff \chi(Iij) > 0. \tag{9}$$

This finally gives us the desired relation between the original grid orientation and the pair $(S, \mathcal{A})$ constructed from the chirotope. The condition that S subdivides into S_L and S_R which can be separated by a line, can then be derived easily from the fact that the grid orientation is acyclic.

The requirements (8) and (9) already determine the order in which the indices of $\mathcal{R}$ must appear in the sequences $\pi_i, i \in \mathcal{L}$ (and vice versa).

Namely, for every index $i \in \mathcal{L}$, the graph gives us an ordering $I_1, \dots, I_t$ of $\mathcal{R}$ such that

$$(i, I_1) \to (i, I_2) \to \cdots \to (i, I_R).$$

Restricted to $\mathcal{R}$, the sequence π_i must then be representable by the following half-period:

$$i : I_R I_{R-1} \dots I_1.$$

Similarly, for an index $I \in \mathcal{R}$ such that

$$(i_1, I) \to (i_2, I) \to \cdots \to (i_L, I),$$

we obtain the restriction

$$I : i_1 i_2 \dots i_L.$$

It remains to insert the remaining indices into the sequences in a consistent way. We describe how to do this for sequences $\pi_i, i \in \mathcal{L}$, the other case is symmetric.

LEMMA 5.1. *Assume that $i : I_R I_{R-1} \dots I_1$ is the partial list for i, and for some $j \in \mathcal{L}, j \neq i$ we have $(i, I_R) \to (j, I_R)$. Then there is a unique index $\tau_{ij} \in \{1, \dots, R\}$ such that*

$$\begin{aligned} (i, I_t) &\to (j, I_t), & t &\geq \tau_{ij}, \\ (i, I_t) &\leftarrow (j, I_t), & t &< \tau_{ij}. \end{aligned}$$

Similarly, if $(i, I_R) \leftarrow (j, I_R)$,
there is a unique index $\tau_{ij} \in \{1, \dots, R\}$ such that

$$\begin{aligned} (i, I_t) &\leftarrow (j, I_t), & t &\geq \tau_{ij}, \\ (i, I_t) &\to (j, I_t), & t &< \tau_{ij}. \end{aligned}$$

The lemma easily follows from the fact that no forbidden subgrid (as described in Fig. 4(b)) exists and all 2×2 subgrids have unique sinks. This implies that the orientation of the edge connecting (i, I_t) and (j, I_t) can change at most once as we let t decrease from R to 1.

The lemma points out a canonical way to insert j into π_i: in the first case (i.e. $(i, I_R) \to (j, I_R)$), we get

$$i : I_R \dots I_\tau \overline{j} I_{\tau-1} \dots I_1,$$

in the second case we obtain

$$i : I_R \dots I_\tau j I_{\tau-1} \dots I_1.$$

Doing this for all j, we obtain a sequence π_i which is complete up to the order of elements j, k that give rise to the same value of $\tau = \tau_{ij} = \tau_{ik}$ in the lemma.

Recall that $\chi(ijk)$ can be read off from the order of the elements j, k in the sequence π_i. To prove that the partial map χ is extendable to an abstract determinant function we therefore need to show that

1. whenever the values are determined by the construction above, the orders of elements j, k in π_i, i, k in π_j, and i, j in π_k are consistent, and
2. whenever elements j, k give rise to the same value of τ, $\tau = \tau_{ij} = \tau_{ik}$, their order in π_i can be determined in a consistent way.

We only include the proof for the first fact.

Observe, that $\tau_{ij} = \tau_{ji}, \forall i, j$. Furthermore, whenever an element j is included into the hyperline sequence of some element i as $\overline{j}$, the hyperline sequence of j will contain i and vice versa. More precisely, j is represented by a positive index in i's hyperline sequence if and only if the sink of the $2 \times R$ subgrid consisting of the two rows i and j lies in row i. Extending this argument we can deduce the following lemma.

LEMMA 5.2. *Consider the hyperline sequences of the indices in $\mathcal{L}$. For any $\alpha, 0 \leq \alpha < |\mathcal{L}|$, there exists a unique index $i_\alpha \in \mathcal{L}$ whose hyperline sequence contains exactly α positive indices of $\mathcal{L}$. (The analog property holds for the sequences of the indices $\mathcal{R}$.)*

PROOF. Consider the global sink of the graph. All edges point to it, in particular the ones connecting it to the nodes of the same column. Thus, the other indices in $\mathcal{L}$ are represented by positive indices in the sink row's hyperline sequence. Deleting this row we get an induced subgraph which again has a global sink. In the hyperline sequence of the index representing the according row we find all the other indices of $\mathcal{L}$ represented by positive indices – with the exception of the one corresponding to the original global sink. Reapplying this argument the lemma becomes obvious. $\square$

As all other cases can be dealt with in very much the same way, we assume that j and k are represented by positive indices in the hyperline sequence of i and k appear as positive index in π_j. (So negative indices represent i and j in the hyperline sequence of k and k in j's sequence.) Furthermore assume $\tau_{ij} > \tau_{ik}$, i.e. $\chi(ijk) = 1$. The hyperline sequence for i is then of the form: $i : A_i j B_i k C_i$, where

$$\left\{\begin{array}{l} A_i = \{I \in S_R | (j,I) \to (i,I), (k,I) \to (i,I)\}, \\ B_i = \{I \in S_R | (k,I) \to (i,I) \to (j,I)\}, \\ C_i = \{I \in S_R | (i,I) \to (j,I), (i,I) \to (k,I)\}, \end{array}\right. \qquad (10)$$

and $S_R = A_i \cup B_i \cup C_i$. By our assumption, B_i cannot be empty. So let J be an index in B_i. We will show that whenever the order of i, k in π_j and i, j in π_k is given, it is consistent with the order of j, k in π_i. This means, $\chi(jik) \neq 1, \chi(kji) \neq 1$.

So assume for a contradiction $\chi(jik) = 1$, $j : A_j k B_j \bar{i} C_j$, and $S_R = A_j \cup B_j \cup C_j$. Hence, index J from above must be in one of the sets A_j, B_j or C_j. Since for any index $I \in A_j \cup B_j$ $(j,I) \to (i,I)$ (implying $I \in A_i$), and for any index $I \in C_j$ $(j,I) \to (k,I)$ (implying $I \notin B_i$) must hold, this gives a contradiction to B_i being non-empty.

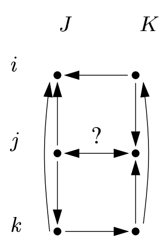

Figure 20: The case $\chi(kji) = 1$.

Assuming $\chi(kji) = 1$ also leads us to a contradiction: The hyperline sequence for k would be of the form $k : A_k \bar{j} B_k \bar{i} C_k$. $S_R = A_k \cup B_k \cup C_k$ and

$$\left\{\begin{array}{l} A_k = \{I \in S_R | (k,I) \to (i,I), (k,I) \to (j,I)\}, \\ B_k = \{I \in S_R | (j,I) \to (k,I) \to (i,I)\}, \\ C_k = \{I \in S_R | (i,I) \to (k,I), (j,I) \to (k,I)\}. \end{array}\right.$$

Considering Eq. 10, we see that $A_i \supseteq B_k = \{I \mid (j,I) \to (k,I) \to (i,I)\}$ and $A_k \supseteq B_i = \{I \mid (k,I) \to (i,I) \to (j,I)\}$. Taking a representative from each set, J from $B_k \subseteq A_i$ and K from $B_i \subseteq A_k$ say, we see – as shown in Fig. 20 – that however the edge $(j,J) - (j,K)$ is directed, we always have a subgraph which does not have a unique sink. Contradiction.

So we showed that up to this point, the defined values of $\chi(ijk)$, $\chi(jki)$, $\chi(kij)$ are equal. We can now determine the so far undefined values without inconsistencies; we omit the proof.

THEOREM 5.1. *Every admissible grid-orientation is 'pseudo-realizable', i.e. induced by a line ℓ, n points, and pseudolines through the pairs of points whose relative orders along ℓ determine the orientation.*

6. REFERENCES

[1] A. Björner, M. Las Vergnas, N. White, B. Sturmfels, and G. M. Ziegler. *Oriented Matroids*. Cambridge University Press, Cambridge, 1993.

[2] J. Bokowski, S. Mock, and I. Streinu. The Folkman-Lawrence Topological Representation Theorem for Oriented Matroids – An Elementary Proof in Rank 3, 1999.

[3] A. Z. Broder, M. E. Dyer, A. M. Frieze, P. Raghavan, and E. Upfal. The worst-case running time of the random simplex algorithm is exponential in the height. *Inform. Proc. Letters*, 56(2):79–81, 1995.

[4] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, Heidelberg, West Germany, 1987.

[5] B. Gärtner, M. Henk, and G. M. Ziegler. Randomized simplex algorithms on Klee-Minty cubes. *Combinatorica*, 18(3):349–372, 1998.

[6] F. Holt and V. Klee. A proof of the strict monotone 4-step conjecture. In J. Chazelle, J. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational geometry*, Contemporary Mathematics. Amer. Math. Soc., 1998.

[7] G. Kalai. A subexponential randomized simplex algorithm. In *Proc. 24th Annu. ACM Sympos. Theory Comput.*, pages 475–482, 1992.

[8] G. Kalai. *Math. Programming*, 79:217–233, 1997.

[9] D. G. Kelly. Some results on random linear programs. *Methods of Operations Research*, 40:351–355, 1981.

[10] J. Mihalisin and V. Klee. Convex and linear orientations of polytopal graphs.

[11] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, New York, NY, 1995.

[12] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice Hall, Englewood Cliffs, NJ, 1982.

[13] K. Williamson Hoke. Completely unimodal numberings of a simple polytope. *Discrete Applied Math.*