


Solution of in-class exercise 1: Number of Cells in Arrangements

(Solution only for the upper bound on the number of cells.)

We use induction on n , the number of lines in the arrangement.

Induction Basis ($n = 1$). With one line the number of cells is 2, which equals $\binom{1+1}{2} + 1$.

Induction step ($n - 1 \Rightarrow n$). By the induction hypothesis we know that for *any* arrangement S' of $n - 1$ lines, the number of cells in the arrangement is at most $\binom{n}{2} + 1$. Let us look now at an arbitrary arrangement S of n lines and delete one of these lines ℓ (the dashed line in the figure). This yields an arrangement S' of $n - 1$ lines and we lose at most n additional cells: Everytime ℓ crosses a line of the arrangement S' , it splits a cell of S' into two. Since there are at most $n - 1$ lines to cross, this splitting happens at most n times.



Thus the number of cells formed by S is at most n more than the number of cells formed by S' . The number of cells formed by S is at most

$$\binom{n}{2} + 1 + n = \binom{n}{2} + \binom{n}{1} + 1 = \binom{n+1}{2} + 1,$$

thus proving the claim (recall $\binom{n}{k} + \binom{n}{k-1} = \binom{n+1}{k}$ for $1 \leq k \leq n$).


Solution of in-class exercise 2: Locating a Point in a Line Arrangement

The lemma contains a chain of two inequalities:

$$\sum_{v \text{ inner node}} |\bar{S}_v| \leq 2 \cdot \sum_{v \text{ inner node}} |S_v| \leq 2n^2.$$

The second inequality is easy to derive: The number $|S_v|$ of coordinates for level v can be bounded by the number of edges that belong to level v , and the total number of edges does not exceed n^2 (cf. exercise 1).

So it remains to look at the first inequality. Let us first look at an example, so that we may believe the statement. Below we have depicted a tree to store $n = 7$ levels: On the left-hand side, the tree for the original sets S_v , and on the right-hand side, the tree with the 'enhanced' sets \bar{S}_v . We have annotated the nodes with the number ($|S_v|$ and $|\bar{S}_v|$, respectively) of x -coordinates stored in that node.



Indeed our example has

$$\begin{aligned} \sum_{v \text{ inner node}} |\bar{S}_v| &= \left(1 + \frac{1}{2} + \frac{1}{4}\right) (n_1 + n_2 + n_3 + n_4) + \left(1 + \frac{1}{2}\right) (n_5 + n_6) + n_7 \\ &\leq 2(n_1 + n_2 + n_3 + n_4 + n_5 + n_6 + n_7) \\ &= 2 \cdot \sum_{v \text{ inner node}} |S_v|. \end{aligned}$$

The example suggests that we should group the nodes of the tree by their depth. (We refrain to speak of the 'levels' of the tree, because that will almost certainly cause confusion vis-à-vis the levels v of the line arrangement). Thus, let h denote the height of the tree and for $0 \leq i \leq h$ let

$$m_i := \sum_{\substack{v \text{ inner node} \\ \text{at depth } i}} |S_v|, \quad \bar{m}_i := \sum_{\substack{v \text{ inner node} \\ \text{at depth } i}} |\bar{S}_v|.$$

Now

$$\sum_{v \text{ inner node}} |\bar{S}_v| = \sum_{k=0}^h \bar{m}_k \leq m_h + \sum_{k=0}^{h-1} \left(m_k + \frac{\bar{m}_{k+1}}{2}\right) = \sum_{v \text{ inner node}} |S_v| + \frac{1}{2} \cdot \sum_{v \text{ inner node}} |\bar{S}_v| - \frac{\bar{m}_0}{2}$$

which yields the claim.

Solution of in-class exercise 3: Nearest Neighbor Changes

This is the classical use case for backward analysis. Instead of considering the points as being inserted one by one and counting the number of nearest neighbors occurring, we get the very same number if we start with the whole set $P_n := P$, and then for each $i = n - 1, n - 2, \dots, 1$, obtain P_i from P_{i+1} by removing a point $p \in P_i$ chosen uniformly at random.

Define X_i to be the nearest neighbor of q in P_i . We are interested in the random variable $X := \{X_i | i \in \{1..n\}\}$. To this end we just calculate the probability $x_i = \Pr[X_i \neq X_{i+1}]$ for $1 \leq i \leq n - 1$. Clearly, then,

$$\mathbf{E}[X] = 1 + \sum_{i=1}^{n-1} x_i.$$

We now note that $X_i \neq X_{i+1}$ occurs if and only if $P_i = P_{i+1} \setminus \{X_{i+1}\}$. This yields that

$$x_i = \Pr[X_i \neq X_{i+1}] = \Pr[P_i = P_{i+1} \setminus \{X_{i+1}\}] = \frac{1}{i+1},$$

where the second equality follows from the fact that P_i is generated from P_{i+1} by removing a point uniformly at random. As a very important remark: please note that this simple argument only works because **no matter** what P_{i+1} is, P_i is always generated by removing one of its elements uniformly at random and the random variable X_i is **completely determined** by P_i (what the currently nearest point is does not depend on the insertion history). You have to verify this before you apply a backward analysis of this type.

Finally,

$$\mathbf{E}[X] = 1 + \sum_{i=1}^{n-1} x_i = H_n.$$

Solution of in-class exercise 4: Linear Separability in Linear Time

- (a) Suppose $R(\mathcal{S}) \neq \emptyset$. Clearly, this implies that $R(\mathcal{S} \setminus \{H\}) \neq \emptyset$. Therefore and since $H_1, H_2 \in \mathcal{S} \setminus \{H\}$, both $x := c(\mathcal{S})$ and $x' := c(\mathcal{S} \setminus \{H\})$ are well-defined points. Now there are two cases. Either $x' \in H$ or $x' \notin H$.

Suppose $x' \in H$. In that case, we readily have $x' \in R(\mathcal{S})$. And since x' was the point of smallest x -coordinate in $R(\mathcal{S} \setminus \{H\}) \supseteq R(\mathcal{S})$, it will a fortiori be the point of smallest x -coordinate in $R(\mathcal{S})$. Therefore $x' = x$ in that case.

Now suppose $x' \notin H$. In that case, we have $x' \notin R(\mathcal{S})$. Now consider the line segment s connecting x' and x . Since $x \in H$ but $x' \notin H$, s must cross ℓ_H at an intersection point p . Moreover, since $x', x \in R(\mathcal{S} \setminus \{H\})$ and since $R(\mathcal{S} \setminus \{H\})$ is a convex set, the whole line segment s runs inside $R(\mathcal{S} \setminus \{H\})$. Since $p \in H$ by construction, it follows $p \in R(\mathcal{S})$. Clearly, x' must have a strictly smaller x -coordinate than x , otherwise x would have been a possible choice for $c(\mathcal{S} \setminus \{H\})$ in the first place. Therefore, going further away from x' along s than p would only make x -coordinates larger, from which we conclude $p = x$ and thus $x \in \ell_H$.

- (b) For any halfplane $G \in \mathcal{S} \setminus \{H\}$, let us say that G *points right* (with respect to H) if the intersection of G and ℓ_H is a ray that extends infinitely to the right. Similarly, let us say that G *points left* (with respect to H) if the intersection of G and ℓ_H is a ray that extends infinitely to the left. There is no third option because we assumed not to have vertical or parallel lines.

We can now proceed in the following way:

SolveInc. Input: $\mathcal{S}, H, x' = c(\mathcal{S} \setminus \{H\})$.

- (1) Check whether $x' \in H$. If so, return x' .
- (2) Else, initialize the (point) variables $h_l := (-\infty, 0)$ and $h_r := (+\infty, 0)$.
- (3) Now for all $G \in \mathcal{S} \setminus \{H\}$, do:
 - (3.1) Let y be the intersection point of the lines ℓ_G and ℓ_H
 - (3.2) If G points right and y has larger x -coordinate than h_l , update $h_l := y$
 - (3.3) If G points left and y has smaller x -coordinate than h_r , update $h_r := y$
- (4) If h_l has strictly larger x -coordinate than h_r , return 'empty', else return h_l .

That the running time is linear is trivial.

Now we argue for correctness. Suppose $R(\mathcal{S}) \neq \emptyset$. Since the conditions of (a) are then satisfied, we know that $c(\mathcal{S})$ is either equal to x' or on ℓ_H . In the first case, the algorithm returns on line (1). In the second case, the algorithm maintains two boundary points h_l and h_r such that on ℓ_H , exactly the points between h_l and h_r are admissible. Since we know $c(\mathcal{S}) \in \ell_H$ and $c(\mathcal{S}) \in R(\mathcal{S})$, we must have $c(\mathcal{S}) = h_l$ because this is the point of smallest x -coordinate within that set.

On the other hand if $R(\mathcal{S}) = \emptyset$, then we *must* end up in the situation that h_1 lies right of h_r as otherwise h_1 , for instance, would be contained in all halfspaces contrary to the assumption.

(c) Consider the following algorithm. **Solve**. Input: \mathcal{S}

- (1) If $\mathcal{S} = \{H_1, H_2\}$, return g as the certificate.
- (2) Else, select $H \in \mathcal{S} \setminus \{H_1, H_2\}$ uniformly at random.
- (3) Recursively call **Solve**($\mathcal{S} \setminus \{H\}$).
- (4) If the result is 'empty', return 'empty'
- (5) Else, if the result is a point x' , call **SolveInc**(\mathcal{S}, H, x') and return the result.

The base case is clearly fine. General correctness then follows by induction from the correctness of **SolveInc**.

The more difficult part will be to determine the running time. At first sight, it looks like this algorithm might require quadratic time; after all, we are calling **SolveInc** linearly many times and we know that **SolveInc** uses up to linear time in the worst case. The crucial observation is that **SolveInc**, on input $(\mathcal{S}, H, c(\mathcal{S} \setminus \{H\}))$, requires linear time only if $c(\mathcal{S}) \neq c(\mathcal{S} \setminus \{H\})$. If on the other hand those certificates are equal, then **SolveInc** can return on line (1) and takes only constant time to terminate.

Let us first only consider sets \mathcal{S} such that $R(\mathcal{S}) \neq \emptyset$. Let t_n be expected running time for $n = |\mathcal{S}|$ such halfplanes. We have $t_2 = \mathcal{O}(1)$. Now for $n \geq 3$, t_n is bounded according to the above observation as

$$t_n \leq t_{n-1} + \Pr [c(\mathcal{S}) \neq c(\mathcal{S} \setminus \{H\})] \cdot \mathcal{O}(n) + \Pr [c(\mathcal{S}) = c(\mathcal{S} \setminus \{H\})] \cdot \mathcal{O}(1).$$

But what is the probability $\Pr [c(\mathcal{S}) \neq c(\mathcal{S} \setminus \{H\})]$? According to (a), this can happen only if $\ell_H \ni c(\mathcal{S})$. But since H is chosen uniformly at random from all halfplanes (except for H_1, H_2) and since $c(\mathcal{S})$ cannot lie on any more than two boundary lines, the probability that this happens is bounded by $\frac{2}{n-2}$. We therefore obtain that

$$t_n \leq t_{n-1} + \frac{2}{n-2} \cdot \mathcal{O}(n) + \mathcal{O}(1) = t_{n-1} + \mathcal{O}(1).$$

From this it readily follows that $t_n = \mathcal{O}(n)$ as desired.

The only thing we have left to do is to consider the case when the result will be 'empty'. However, consider that as soon as a recursive step returns 'empty', all higher recursion levels return in constant time. So the transition from a certificate point to the result 'empty' happens at most once in the course of the algorithm. Before that, the bound we just derived applies. Therefore the total time is no bigger than linear in that case either.

Solution of in-class exercise 5: Minimum Cut

- (a) We observe that $\mu(G) = 2$ and that the depicted multigraph G has two disjoint minimum cuts $\{(3, 4), (3, 6)\}$ and $\{(6, 9), (7, 10)\}$. Therefore, after contracting any edge, at least one of those cuts must still be present, and hence $\Pr[\mu(G) = \mu(G/e)] = 1$.
- (b) The algorithm \mathcal{B} simply returns the maximum $B := \max\{A_1, \dots, A_k\}$. This implies that $B = \mu(G)$ if and only if $A_i = \mu(G)$ for some index i . The “only if” is obvious, and the “if” follows from the assumption that $A_i \leq \mu(G)$ holds for all indices i . Therefore, the probability for $B = \mu(G)$ is equal to the probability that for at least one i we have $A_i = \mu(G)$. This gives us, by using the inequality $1 - x \leq e^{-x}$,

$$\begin{aligned}\Pr[B = \mu(G)] &= \Pr[\exists i: A_i = \mu(G)] = 1 - \Pr[\forall i: A_i < \mu(G)] \\ &= 1 - \left(1 - \frac{1}{n}\right)^k \\ &\geq 1 - \exp\left(-\frac{1}{n} \cdot k\right) \geq 1 - n^{-42}\end{aligned}$$

if we choose $k := \lceil 42n \cdot \ln(n) \rceil$.