

2 Linear solution

Remove from the graph considered before all edges between cells with different topmost tracks. Now we can contract all those components to single nodes and add the deleted edges thereby obtaining a new graph H . For simplicity number the nodes from 1 to k where node 1 represents the component containing the upper left corner.

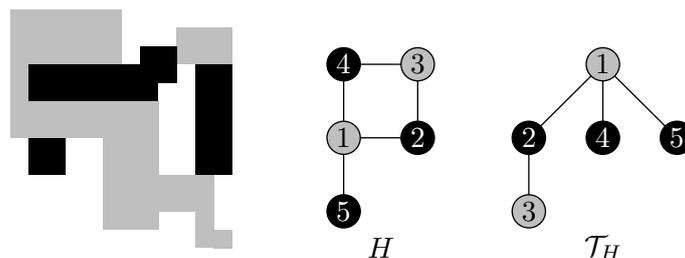


Figure 1: An example of tracks, the respective graph and a possible breadth-first-tree

This graph is clearly bipartite with a bicoloring given by the assignment component \mapsto topmost track. Consider the breadth-first tree \mathcal{T}_H (starting at node 1). Note that both H and \mathcal{T}_H can be calculated in time $O(mn)$.

Proposition 4. *The minimum number a of animals needed to create the given pattern equals the depth of \mathcal{T}_H , i.e. $a = 1 + \max_{v \in V(H)} \text{dist}(1, v)$.*

Proof. We show by induction that the cells visited by the previously described greedy algorithm after using n animals are exactly the cells represented by the nodes within distance $n - 1$ from the root. Again the case $n = 1$ is trivial. For $n \geq 2$ we can visit all the nodes “in the layer below”: since we have $E(H) \supseteq E(\mathcal{T}_H)$ the bicoloring of H gives a bicoloring of \mathcal{T}_H and thus all the nodes in the next layer have the same, namely the currently active, color, and are visited by the greedy algorithm.

Furthermore the greedy algorithm doesn’t reach any other nodes: the nodes visited so far are the previous layers. Assume component \mathcal{C} can be visited and take any path from the root to \mathcal{C} and throw away all the parts up to the last node that is already visited. The remaining nodes have to be of the currently active color, i.e. they belong to the component \mathcal{C} . Thus if \mathcal{C} hasn’t been visited before it is child of a node visited before and will be visited in this step. \square

This suffices for an $O(mn)$ solution (i.e. linear in the size of the input). However, instead of explicitly calculating components one can assign weights to all the edges in the original graph: any edge between two nodes of the same color gets weight 0 and any other edge weight 1. Then clearly the maximum distance of any node to the root (the upper left corner) equals a . The distances can be calculated by Dijkstra’s algorithm where we can use a **deque** instead of a **priority_queue** (0/1-Dijkstra) to achieve a linear running time.