Extremal Combinatorics and Universal Algorithms

Stefan David

Department of Pure Mathematics and Mathematical Statistics
University of Cambridge

This dissertation is submitted for the degree of
Doctor of Philosophy

Trinity College July 8, 2018

To the memory of my Grandfather

Declaration

I hereby declare that except where specific reference is made to the work of others, the contents of this dissertation are original and have not been submitted in whole or in part for consideration for any other degree or qualification in this, or any other university. This dissertation is my own work and includes nothing which is the outcome of work done in collaboration, except where indicated in the text.

Stefan David
July 8, 2018

Abstract

In this dissertation we solve several combinatorial problems in different areas of mathematics: automata theory, combinatorics of partially ordered sets and extremal combinatorics.

Firstly, we focus on some new automata that do not seem to have occurred much in the literature, that of solvability of mazes. For our model, a maze is a countable strongly connected digraph together with a proper colouring of its edges (without two edges leaving a vertex getting the same colour) and two special vertices: the origin and the destination. A pointer or robot starts in the origin of a maze and moves naturally between its vertices, according to a sequence of specific instructions from the set of all colours; if the robot is at a vertex for which there is no out-edge of the colour indicated by the instruction, it remains at that vertex and proceeds to execute the next instruction in the sequence. We call such a finite or infinite sequence of instructions an algorithm. In particular, one of the most interesting and very natural sets of mazes occurs when our maze is the square lattice $\mathbb{Z}^2$ as a graph with some of its edges removed. Obviously, we need to require that the origin and the destination vertices are in the same connected component and it is very natural to take the four instructions to be the cardinal directions. In this set-up, we make progress towards a beautiful problem posed by Leader and Spink in 2011 which asks whether there is an algorithm which solves the set of all such mazes.

Next, we address a problem regarding symmetric chain decompositions of posets. We ask if there exists a symmetric chain decomposition of a $2 \times 2 \times \ldots \times 2 \times n$ cuboid (k 2's) such that no chain has a subchain of the form $(a_1, \ldots, a_k, 0) \prec \ldots \prec (a_1, \ldots, a_k, n-1)$? We show this is true precisely when $k \geq 5$ and $n \geq 3$. This question arises naturally when considering products of symmetric chain decompositions which induce orthogonal chain decompositions — the existence of the decompositions provided in this chapter unexpectedly resolves the most difficult case of previous work by Spink on almost orthogonal symmetric chain decompositions (2017) which makes progress on a conjecture of Shearer and Kleitman. Moreover, we generalize our methods to other finite graded posets.

Finally, we address two different problems in extremal combinatorics related to mathematical physics. Firstly, we study metastable states in the Ising model. We propose a general model for 1-flip spin systems, and initiate the study of extremal properties of their stable states. By translating local stability conditions into Sperner-type conditions, we provide non-trivial upper bounds which are often tight for large classes of such systems. The last topic we consider is a deterministic bootstrap percolation type problem. More specifically, we prove several extremal results about fast 2-neighbour percolation on the two dimensional grid.

Acknowledgements

I am very thankful to my supervisor, Professor Béla Bollobás, for all the guidance and help he has provided over the years. His dedication and support towards me and all of his students helped us tremendously throughout our study and research. I am also extremely grateful for all the beautiful mathematics I have learned during our seminars and research trips all over the world. I would like to take this opportunity to thank him and his wonderful wife Gabriella for their hospitality wherever we were traveling together and I really hope that we also succeeded to make them feel at home when they visited me, Marius and our families in our native city.

I am grateful to my advisor, Professor Imre Leader for all the beautiful mathematics and wonderful problems he introduced me to, for his constant support, and for all the advice and guidance he has offered me during the last 5 years. I would also like to thank him for all the care he offered during our trips together with the whole research group to Memphis.

I would like to thank my closest collaborator and lifetime friend Marius Tiba for all the wonderful time we spent together during our excursions in the world of combinatorics and mathematics. I would also like to thank my other collaborators and my contemporary academic brothers Hunter Spink, Ivailo Hartarsky, Scott Binski, Andrew Carlotti, Peter van Hintum, Sebastian Koch, Shoham Letzter, Bhargav Narayanan, Julian Sahasrabudhe, Stephen Smith, Richard Snyder, István Tomon for all our mathematical discussions, mini-seminars and for everything I have learned from you over the years. I should also like to thank my older academic brothers Richard Johnson, Tomas Juškevičius, Teeradej Kittipassorn, Jonathan Lee for welcoming and introducing me to the group and Antonio Girao, Vytautas Gruslys for all our wonderful fun time together.

I am grateful to EPSRC, DPMMS and Trinity College for providing financial support for my work. Moreover, I would like to thank the University of Memphis and the IMT Institute for Advanced Studies in Lucca for inviting me over and for their hospitality.

Finally, I would like to thank my whole family for their continuous love and support throughout my studies. Without their help, this would not have been possible.

Contents

CHAPTER 1

Introduction

1. Structure

This dissertation is divided into five chapters. In the present chapter, we briefly present the problems and results that appear in the rest of the work. In the second chapter, which is the longest, we solve a problem in automata theory. In the third chapter we address a question regarding decomposition of posets into symmetric chains. In the last two chapters we consider two problems in extremal combinatorics, one regarding the metastable states in interaction spin glasses, and the last one in bootstrap percolation.

2. Solvability of Mazes by Blind Robots

Automata theory, the subject in discrete mathematics and theoretical computer science which is concerned with the study of a certain type of machines called automata, was introduced by von Neumann (see [42]) in 1966. Though studied for decades, recent important breakthroughs in automata theory, such as Trahtman's solution to the road coloring problem [41], have turned it into an important field in discrete mathematics and theoretical computer science. For a comprehensive introduction in the theory and other related subjects, see the book of Hopcroft, Motwani and Ullman [21].

Informally, an automaton is made of states, it receives inputs from a formal alphabet (which is a finite set of symbols), and for each state it determines to which state to switch when a new input is received. This model leads to the work done in Chapter 2, which is joint work with M. Tiba. There we introduce and study an automaton which so far has not received much serious attention, which is related to solvability of mazes. Though there were numerous maze solving questions asked over the years, giving raise to maze solving algorithms or shortest path algorithms for mazes, our model is quite different in nature and it is more difficult. It is motivated by a question of I. Leader and H. Spink which was passed to us by P. Balister and is discussed in detail in Section 1 of Chapter 2. After we introduce our model, we solve some particular instances of this problem. As our main result, we construct an algorithm for a robot to visit all accessible vertices in a set of mazes which arises as a collection of special subgraphs of $\mathbb{Z}^2$.

3. Products of posets with long chains

The consideration of symmetric chain decompositions of posets first started with Kleitman's proof of the Littlewood-Offord theorem on concentration of sums of Bernoulli random variables [24]. One of the key observations in that paper is that we can inductively create symmetric chain decompositions of the hypercube $Q_n = \{0,1\}^n$ (which can alternatively be viewed as the power set of $\{1,2,\ldots,n\}$), through a certain "duplication method". This observation is the special case with Q a two-element chain poset of a more general claim that given two posets P, Q with symmetric chain decompositions, we can decompose the product $P \times Q$ into symmetric chains by decomposing the rectangle posets formed by the product of a chain in P with a chain in Q. The literature is abundant with both necessary and sufficient conditions for the existence of symmetric chain decompositions on finite graded posets (see for example the works of Griggs [19] and Stanley [38]). However, the further study of commonalities between all symmetric chain decompositions is somewhat lacking, mostly due to the largely unstructured nature of a generic such decomposition of a typical poset.

In Chapter 3, which is joint work with H. Spink and M. Tiba, we address a problem which arises naturally when considering products of symmetric chain decompositions which induce orthogonal chain decompositions. The existence of the decompositions proved in this chapter unexpectedly resolves the most difficult case in a paper by H. Spink on almost orthogonal symmetric chain decompositions [37], which makes progress on a conjecture of Shearer and Kleitman [35]. We show that there exists a symmetric chain decomposition of a $2 \times 2 \times \ldots \times 2 \times n$ cuboid (k 2's) such that no chain has a subchain of the form $(a_1, \ldots, a_k, 0) \prec \ldots \prec (a_1, \ldots, a_k, n-1)$ precisely when $k \geq 5$ and $n \geq 3$. Moreover, we show how our methods generalize to other finite graded posets.

4. Metastable States in the Ising Model

The *Ising model* was introduced in 1920 by the physicist Wilhelm Lenz in [27], who passed it as a problem to his student, Ernst Ising (see [22]). Since then, this model has received serious attention in the literature, for example see Lee and Young [26], Glauber [16] and Kazakov [23]. For a comprehensive description of the model and closely related subjects, see the book of McCoy and Wu [29].

In the Ising model, a collection V of interacting particles are arranged in an underlying dependency graph G with vertex set V. Each particle $v \in V$ has a magnetic spin $\sigma_v \in \{\pm 1\}$ and it can interact only with its neighbours in G. The energy of a certain *spin configuration*, which is an assignment of ± 1 to each σ_v, is given by the Hamiltonian function

$$H = -\sum J_{ij}\sigma_i\sigma_j - \sum h_i\sigma_i,$$

where the J_{ij} are typically Gaussian random variables with $J_{ij} = 0$ if ij is not an edge of G, and h_i are constants corresponding to an external magnetic field. When the Hamiltonian is locally maximized, in the sense that for any v, negating σ_v strictly decreases the Hamiltonian, we say that the system is *metastable* or, equivalently, that we have a *metastable state*. An important question in mathematical physics is to understand the distribution of metastable states.

In Chapter 4, which is joint work with H. Spink and M. Tiba we initiate the study of the associated extremal problem, namely what is the maximal number of metastable states possible under various restrictions. To do so, we propose a generalization of these spin systems, capturing what we believe to be the combinatorial essence of the Ising model. Specifically, for each particle v, we show how the Ising model imposes constraints on the possible configurations of spins in $\Gamma(v)$ (the neighbourhood of v in G) for any metastable state. The conjunction of these constraints then imposes combinatorial conditions on the metastable states in $\{\pm 1\}^V$ which, as we will see, are very analogous to the Sperner antichain condition in extremal set theory.

5. Fast Bootstrap Percolation on the Grid

Cellular automata models are systems in which particles interact according to local and homogeneous rules and were introduced by von Neumann in 1966 (see [42]), but no general theory of such models has been developed until the recent work of Bollobás, Smith and Uzzell [9] (also see the paper of Bollobás, Duminil-Copin, Morris and Smith [8]). However, there are special cases that had been studied broadly, for example the bootstrap percolation model introduced in 1979 by Chalupa, Leath and Reich in [13], originating in the context of disordered magnetic systems.

In Chapter 5, which is joint work with S. Binski, we consider the deterministic 2-neighbour bootstrap percolation model on the grid and we address a question of B. Bollobás which asks about the minimal infection time in this set-up. We first present an easy but very nice argument establishing the exact minimal infection time for $n+1$ initially infected sites, and we then provide some general upper and lower bounds for the minimal infection time in terms of the number of initially infected sites.

CHAPTER 2

Solvability of Mazes by Blind Robots

1. Introduction

This chapter is joint work with Marius Tiba.

One of the long standing conjectures in automata theory is the road colouring problem introduced in 1970 by Adler, Goodwyn and Weiss in [2], [3]. This conjecture states that a strongly connected digraph $\overrightarrow{G}$ in which all vertices have the same out-degree has a synchronising colouring, provided $\overrightarrow{G}$ is aperiodic, i.e. the gcd of the lengths of all of its oriented cycles is one. A synchronising colouring of a strongly connected digraph $\overrightarrow{G}$ of uniform out-degree k is a labelling of the edges of $\overrightarrow{G}$ with colours $1, \ldots, k$ such that all the vertices have out-edges of all colours and for every vertex v of $\overrightarrow{G}$ there exists a word W_v in the alphabet of colours such that every path in $\overrightarrow{G}$ corresponding to W_v terminates at v. We note that the existence of a synchronising colouring makes it possible to reset the automaton to its original state after the detection of an error. In fact, this important property is the reason why the road coloring problem has received so much attention over the past few decades. There have been many positive partial results published over the years, such as Carbone [10], Friedman [15], and O'Brien [31]. In 2009, Trahtman made one of the most notable advances in the field by proving this conjecture in [41].

Another well-known related problem in the field is Černý's conjecture which appeared in [12] in 1964 and states that the length of the shortest synchronising word for any n-state deterministic finite automaton is bounded above by $(n-1)^2$ (for more details see Pin [33] and Trahtman [40]). There are many partial results concerning Černý's conjecture, see e.g. Grech and Kisielewicz [18] and Steinberg [39].

In this chapter we introduce and study a new model of automata which turns out to be deep and interesting, motivated by the following coffee time problem of Leader.

Problem 1.1. *Consider the classical* 8×8 *chessboard as a maze, where every small square is a room, such that between any two adjacent rooms there is either a wall that prevents the transit between them, or there is no wall and transit is possible. Additionally, the boundary of the board is formed only by walls. Say that a robot starts in one of the* 64 *squares and it receives a series of instructions from the cardinal directions:* north, south, east, west. *Each time the robot receives such an instruction, it executes it by moving to the corresponding adjacent room, provided there is no wall to prevent it from moving as instructed; if there is such a wall, the robot simply does not move and it continues with the following instruction. The robot does not give any feedback whether it moves or not when executing an instruction. Naturally, the maze can be regarded as a subgraph of the square lattice* 8×8 *where there is an edge between two vertices if and only if there is no wall between the corresponding squares. Without knowing the subgraph and the starting vertex of the robot, can one write a sequence of instructions such that at the end the robot is guaranteed to have visited all accessible vertices?*

To see the existence of such an algorithm, simply enumerate all the possible boards and solve them one by one, keeping track of the updated position of the robot when passing to a new board. Another related problem which can be solved similarly is the following.

Problem 1.2. *Consider a subgraph of some finite dimensional hypercubes* $Q_1, Q_2, \ldots$ *as a maze. Say that a robot starts in one of the vertices and it receives a sequence of instructions from the set of coordinate directions* $\pm e_1, \pm e_2, \ldots$. *Each time the robot receives such an instruction, it executes it by moving to* *the corresponding adjacent vertex, provided there is an edge between these two vertices; if there is no such edge, the robot simply does not move and it continues with the following instruction. Without knowing the subgraph and the starting vertex of the robot, can one write a sequence of instructions such that at the end the robot is guaranteed to have visited all accessible vertices?*

Problem 1.1 lead Spink and Leader to ask the following research question, which was later passed to us by P. Balister.

QUESTION 1.3. *What happens if in Problem 1.1 we replace the (finite)* 8×8 *square lattice with the infinite square lattice* $\mathbb{Z}^2$*?*

Having worked on it for a long time, we believe that Question 1.3 is extremely difficult. In this chapter we make progress towards answering this question, by constructing algorithms which solve certain subsets of mazes arising from the infinite square lattice $\mathbb{Z}^2$, thus establishing the following main result.

THEOREM 1.4. *There exists an infinite sequence of instructions for a robot to visit all accessible vertices in any maze for which the board is the graph* $\mathbb{Z}^2$ *with arbitrarily many horizontal edges removed but only finitely many vertical edges removed, and the columns with missing vertical edges are consecutive, i.e., they form an interval.*

We note that Theorem 1.4 follows immediately from two separate results, Theorem 3.1 and Theorem 3.2 which are of interest by themselves, and a rather technical result, Proposition 3.3.

This chapter is divided into nine sections. In Section 2 we start by developing a general set-up that encompasses a class of similar problems which we call "solvability of mazes by blind robots". We then return to the Leader-Spink problem and state all our main results in Section 3. In Section 4 we present a toy model that represents the foundation on which the general model is constructed. As part of this toy model, we prove Theorem 3.1; this allows us to introduce and investigate some generic algorithms that are used as building blocks in the proof of Theorem 3.2. In Section 5 we present a series of technical definitions that are used to construct a countable cover of the set of mazes in Theorem 3.2 with subsets of mazes that we can treat individually. In Section 6 and Section 7 we prove Theorem 3.2 by a suitable construction. We continue with the proof of the technical result and the proof of Theorem 1.4 in Section 8. Finally, in Section 9 we present several further directions of research and some of our conjectures.

2. Preliminaries

In this section we give a precise definition of our model and present some examples.

A *maze* is a quadruple (M, c, o, d), where M is a countable strongly connected digraph called the board and $c : E(M) \longrightarrow \mathbb{N}$ is a proper colouring of the directed edges of M, i.e. one in which the out-edges from any vertex have distinct colours. Further, o and d are two special vertices of M called the *origin* and the *destination*, respectively.

An *instruction* $I \in \mathbb{N}$ is an element from the set of colours $\mathbb{N}$. An *algorithm*

$$A = (I_i)_{i=1}^{n} \text{ or } A = (I_i)_{i=1}^{\infty}$$

is a finite or infinite sequence of instructions. Further, this is also how we define *finite* and *infinite* algorithms, respectively. A *subalgorithm* A' of an infinite algorithm A as above is a truncation of A of the form

$$A' = (I_i)_{i=k}^{j} \text{ or } A' = (I_i)_{i=k}^{\infty},$$

for some $k \leq j$. Similarly, a *subalgorithm* A' of a finite algorithm $A = (I_i)_{i=1}^{n}$ is a truncation of A of the form $A' = (I_i)_{i=k}^{j}$ for some $k \leq j \leq n$. Finally, a *robot* is an element form the set of vertices of M. In order to describe dynamically our process of visiting the graph we look at the following model.

Given an algorithm $A = (I_i)_{i=1}^{\infty}$ and a maze (M, c, o, d), the robot is initially o and then it changes its value or it *updates* to different vertices of M, as it follows the instructions $I_1, I_2, \ldots$ one by one in order: for $n \in \mathbb{N}$ the robot updates given the n-th instruction $I_n \in \mathbb{N}$ by changing its value from the current element v to the new vertex w if and only if there exists an oriented edge e of colour I_n from v to w; if there is no such oriented edge e, the robot's value remains v.

In order to use a more natural language, for the rest of the chapter we use the following notation. We view the robot as a pointer that indicates to different vertices of M and we say that it *starts* at o and then it *moves* between the vertices of M according to the instructions, as described above. In short, we say that the robot *follows the algorithm A in the maze* (M, c, o, d). Given a maze (M, c, o, d), a vertex v of M and an algorithm A, we say that the robot *visits* v as it follows A in (M, c, o, d) if the value of the robot is v at some point while it follows the algorithm A in the maze (M, c, o, d). We say that an algorithm A *solves* the maze (M, c, o, d) if the robot visits the destination d as it follows A in (M, c, o, d). Similarly, we say that an algorithm A *solves* a set $\mathcal{M}$ of mazes if it solves every maze in $\mathcal{M}$.

We remark that each connected graph can be regarded as a strongly connected digraph by doubling edges. This is done in order to allow us to give the desired colouring of the directed edges. Throughout the chapter all the boards of the mazes arise in this way and hence from now on we define the board of a maze to be a graph. Moreover, we omit the condition that the graph is connected and we require instead that the origin and the destination are in the same connected component of the graph. Finally, we call every vertex in the connected component of the origin *accessible*.

In this chapter we are interested in the fundamental question, whether there exist algorithms that solve certain natural sets of mazes. As we shall see from the arguments which appear in this chapter, and also from our conclusions and open questions in Section 9, this set up is quite rich in deep insights related to the phenomenon of state automata.

For example, we note that there is no algorithm that solves the set of all mazes. Indeed, let us assume for a contradiction that $A = (I_i)_{i=1}^{\infty}$ does the job. We construct M to be the path with vertices $v_0, v_1, \ldots$ and its only edges $v_i \to v_{i-1}$ and $v_{i-1} \to v_i$ for all $i \in \mathbb{N}$. We set $o = v_1$, $d = v_0$ and colour the edge $v_i \to v_{i+1}$ with colour I_i and the rest of the edges in any way that does not violate the proper colouring condition. A robot that starts in this maze and follows A will visit in order $v_1, v_2, v_3, \ldots$ as it follows $I_1, I_2, \ldots$, never reaching $v_0 = d$. As M was constructed to be strongly connected, we have reached a contradiction.

As another warm-up example let us note that for any countable set of mazes there exist algorithms that solve it. In particular, this solves Problem 1.1 and more importantly it shows that there exist algorithms that solve the set of *all* finite mazes. Indeed, let $(M_1, c_1, o_1, d_1), (M_2, c_2, o_2, d_2) \ldots$ be an enumeration of a countable set of mazes $\mathcal{M}$. Considering the strongly connectedness property, given any maze (M, c, o, d) one can write by inspection a finite algorithm that solves the maze. Then, let A_1 be any finite algorithm that solves (M_1, c_1, o_1, d_1); let o'_2 be the position of the robot after it follows the algorithm A_1 in (M_2, c_2, o_2, d_2); let A_2 be any finite algorithm that solves (M_2, c_2, o'_2, d_2) with origin o'_2; let o'_3 be the position of the robot after it follows the algorithm A_1A_2 in (M_3, c_3, o_3, d_3), etc. Continue in this way to create algorithms $A_1, A_2, \ldots$. We claim that the algorithm $A = A_1A_2\ldots$ obtained by concatenating $A_1, A_2, \ldots$ solves the set of mazes $\mathcal{M}$. Indeed, consider the maze $(M_i, c_i, o_i, d_i) \in \mathcal{M}$ for some $i \geq 2$. After the robot follows the initial subalgorithm $A_1A_2\ldots A_{i-1}$ of A it gets to the vertex o'_i of M_i and then after it follows A_i it gets to the destination point d_i. Trivially, for the maze (M_1, c_1, o_1, d_1), the robot gets to the destination point d_1 after it follows the initial subalgorithm A_1 of A. This shows that A solves $\mathcal{M}$.

We can see from these two examples that the most interesting cases of our model occur “in between”, when we consider natural uncountable sets of mazes for which we seek to construct algorithms to solve them. Let us present below an uncountable set of mazes, for which it is not hard to find such algorithms.

Let $Q = Q_1 \cup Q_2 \cup \ldots$ be the nested union of all finite dimensional hypercubes i.e. the graph with vertices all possible infinite $\{0, 1\}$ sequences with trailing zeros and edges between those pairs of vertices which differ in only one coordinate. Let $\mathcal{Q}$ be the set of mazes for which the board is a connected subgraph of Q and the colouring assigns to each directed edge the corresponding coordinate direction $\pm e_1, \pm e_2, \ldots$.

Our main aim in this chapter, though, is to solve a problem resembling Problem 1.1. One of the most fundamental sets of mazes is the set $\mathcal{M} = \mathcal{M}(\mathbb{Z}^2)$ for which the board is the square lattice $\mathbb{Z}^2$ considered as a graph with arbitrarily many edges deleted, the colouring assigns to each directed edge the corresponding cardinal direction from the set $\{N, S, E, W\} = \{S^{-1}, N^{-1}, W^{-1}, E^{-1}\}$, and the origin and destination are in the same connected component. From now on we define a maze to be a triple $(M, \mathbf{o}, \mathbf{d}) \in \mathcal{M}$, where M is the board, $\mathbf{o} = (x_o, y_o)$ is the origin, and $\mathbf{d} = (x_d, y_d)$ is the destination.

An *algorithm* A is a finite or infinite sequence of instructions. We say that a robot *follows* A in a given maze $(M, \mathbf{x})$ starting from $\mathbf{v} = (x, y)$ if it starts from $\mathbf{v}$ and then it executes in order one by one the instructions in A as described above.

We label the rows and columns of $\mathbb{Z}$ as $r_i = \{(x, i) \mid x \in \mathbb{Z}\}$ and $c_i = \{(i, y) \mid y \in \mathbb{Z}\}$, respectively. Finally, for a point $(x, y) \in \mathbb{Z}^2$ we refer to the x-coordinate as its *longitude* and to the y-coordinate as its *latitude*.

In Figure 1 we mark the destination point $(3, -2)$ with a cross and we note that in every maze there is a path from the origin to the destination point. When the robot follows the algorithm $SNWWN$ in M it gets to the point $(-1, 2)$ it follows the path $(0, 0)$, $(0, 0)$, $(0, 1)$, $(-1, 1)$, $(-1, 1)$, $(-1, 2)$; the robot does not move when it executes the first and fourth instructions, as there is no edge between $(0, 0)$ and $(0, -1)$ and between $(-1, 1)$ and $(-2, 1)$.

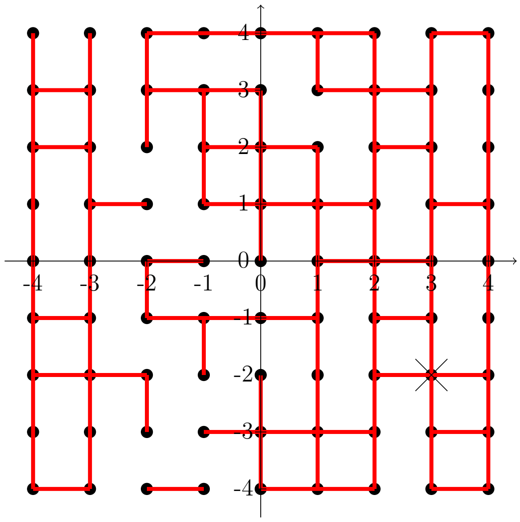

FIGURE 1. A representation of a piece of a general maze M, where edges are marked by red lines.

Regarding notation, we often create new algorithms by concatenations of instructions and other algorithms, and it is very convenient to use multiplication to denote concatenation. For example

$$SNSSNS = SNS^2NS = (SNS)^2$$

denotes the algorithm $A = (A_i)_{i=1}^6$ with $A_1 = S, A_2 = N, A_3 = S, A_4 = S, A_5 = N, A_6 = S$. For our convenience, let us further set $N^{-1} := S$, $S^{-1} := N$, $E^{-1} := W$, $W^{-1} := E$ by convention. For a finite algorithm A, we write $|A|$ for the number of instructions in A; similarly we write $|A|_I$ for the number of occurrences of I in A, for $I \in \{N, S, E, W\}$. For example, taking $A = (NSN)^2$ as above, $|A| = 6$ and $|A|_N = 4$.

3. Our Results

Our main result, Theorem 1.4, follows almost directly from Theorem 3.1, Theorem 3.2 and Proposition 3.3, all of which are interesting results on their own.

Theorem 3.1. *Let $\mathcal{C} \subseteq \mathcal{M}(\mathbb{Z}^2)$ be the set of all mazes for which the board has arbitrarily many horizontal edges removed but no vertical edges removed. Then there exists an algorithm that solves $\mathcal{C}$.*

Theorem 3.2. *Let $\mathcal{F} \subseteq \mathcal{M}(\mathbb{Z}^2)$ be the set of all mazes for which the board has arbitrarily many horizontal edges removed and nonzero finitely many vertical edges removed in consecutive columns. Then there exists an algorithm that solves $\mathcal{F}$.*

We should note that the proofs of Theorem 3.1 and Theorem 3.2 are both constructive. Moreover, and as one might expect, the proof of Theorem 3.2 turns out to be much more difficult than the proof of Theorem 3.1. One might see that Theorem 1.4 is stronger than both Theorem 3.1 and Theorem 3.2 as it shows that there exists algorithms that solves $\mathcal{C} \cup \mathcal{F}$. As a final note, from the way we prove Theorem 3.1 and Theorem 3.2, it is clear how Theorem 1.4 follows directly, obtaining a constructive proof of Theorem 1.4. However, Proposition 3.3 is a more general result which enables us to gain more insight and deduce further properties about our model, e.g. see Corollary 3.4.

In Section 4, in which we give the proof of Theorem 3.1, we also introduce some generic algorithms which constitute the main building blocks of the algorithm which solves $\mathcal{F}$. In Lemma 4.1, which is the key technical result of the chapter, we present the properties of these generic algorithms that will be used multiple times in the proof of Theorem 3.2.

Proposition 3.3. *Let $E(\mathbb{Z}^2)$ be the set of edges of $\mathbb{Z}^2$. We can regard any board of a maze as an indicator function $f : E(\mathbb{Z}^2) \longrightarrow \{0,1\}$. Hence, the set of boards of mazes equipped with the product topology is a compact metrizable space. Let $\{\mathcal{A}_i\}_{i=1}^{\infty}$, $\mathcal{A}_i \subseteq \mathcal{M}$ for all i, be a countable collection of mazes with the following properties:*

(1) for all $i \in \mathbb{N}$, all origins $o \in \mathbb{Z}^2$, all destination $d \in \mathbb{Z}^2$ and all paths P between o and d, the sets of boards $B_i = \{M \mid (M,o,d) \in \mathcal{A}_i, P \text{ is a subgraph of } M\}$ are compact;

(2) for all $i \in \mathbb{N}$, if $(M, o, d) \in \mathcal{A}_i$, then $(M, o', d') \in \mathcal{A}_i$ for all o', d' in the same connected component as o, d;

(3) for every i there exists an algorithm A_i that solves the set $\mathcal{A}_i$.

Then there exists an algorithm A that solves the set $\mathcal{A} = \bigcup_{i=1}^{\infty} \mathcal{A}_i$ and that furthermore guides the robot to visit the destination of every maze in the set infinitely often. Moreover, if we cut or add an initial segment to A, the algorithm obtained in this way has the same property.

Proposition 3.3, which is proved in Section 8, allows us to go back to the original problem regarding existence of algorithms that solve the set $\mathcal{M}$ of all mazes. The following immediate corollary which follows from Proposition 3.3 by taking $\mathcal{A} = \mathcal{A}_i = \mathcal{M}$ for all i ascertains the intuitive fact that there exists an algorithm such that if the robot follows it, the robot visits all accessible points in any given maze at least once if and only if there exists an algorithm such that if the robot follows it, the robot visits all accessible points in any given maze infinitely often.

Corollary 3.4. *The following statements are equivalent:*

(1) there exists an algorithm A_1 such that if the robot follows it in any given maze $(M, \boldsymbol{x}) \in \mathcal{M}$, the robot visits all accessible points of $(M, \boldsymbol{x})$ at least once;

(2) for any $d \in \mathbb{N}$ there exists a finite algorithm $A_2(d)$ such that if the robot follows it in any given maze $(M, \boldsymbol{x}) \in \mathcal{M}$, the robot visits all points at distance at most d from the origin in $(M, \boldsymbol{x})$;

(3) there exists an algorithm A_3 such that if the robot follows it in any given maze $(M, \boldsymbol{x}) \in \mathcal{M}$, the robot visits all accessible points of $(M, \boldsymbol{x})$ infinitely often.

Corollary 3.4 is an interesting result on its own, also because one could try to prove the existence of an algorithm that solves the set $\mathcal{M}$ of all mazes by constructing $A_2(1), A_2(2), \ldots$. Remarkably, even if $A_2(1)$ which we constructed exists, it is not trivial to find.

4. Our Toy Model

The aim of this section is to prove Theorem 3.1 and to introduce the general strategy and some generic algorithms that are used in the proof of Theorem 3.2 as well.

For our toy model we consider $\mathcal{C}$, the set of mazes with no vertical edges removed from Theorem 3.1. Without loss of generality, we assume that for any maze in $\mathcal{C}$ the origin is the point $(0,0)$. The main property of this set of mazes is that at each step of the algorithm we know the robot's latitude.

We start with a short proof of Theorem 3.1 and we then provide a more complex proof which introduces more profound ideas that are needed for our proof of Theorem 3.2.

Proof of Theorem 3.1. For any positive integer a, we define the oscillation $O(a)$ to be the algorithm $N^aS^{2a}N^a$. We begin by defining the class of algorithms *easy_move_east*:
$EME(a) := (O(a)\ E\ O(a)\ NES\ O(a)\ SEN\ O(a)\ N^2ES^2\ \ldots\ O(a)\ N^aES^a$
$O(a)\ S^aEN^a\ O(a))^a$.

The counterpart of this class of algorithms is *easy_move_west* defined as
$EMW(a) := (O(a)\ W\ O(a)\ NWS\ O(a)\ SWN\ O(a)\ N^2WS^2\ \ldots\ O(a)$
$N^aWS^a\ O(a)\ S^aWN^a\ O(a))^a$.

We note that for each a, $|EME(a)| = |EMW(a)|$ and we define the unbounded sequence of positive integers $(x_i)_{i=1}^n$ by the rules $x_1 = 1$, $x_2 = 2|EME(1)| = 46$, and in general $x_i = 2(|EME(x_1)| + |EME(x_2)| + \ldots + |EME(x_{i-1})|)$ for all $i > 1$.

We claim that the (infinite) algorithm $finish$ defined as

$$F = EMW(x_1)\ EME(x_2)\ EMW(x_3)\ldots$$

solves the class $\mathcal{C}$ of mazes. Indeed, let us assume without loss of generality that the destination point has longitude at least that of the origin. Let a be the smallest positive integer which is at least twice the difference in longitude between the destination point and the origin, at least the absolute value of the latitude of the destination and at least the smallest positive integer b such that every pair of consecutive columns at longitude between the origin and the destination point is connected by a horizontal edge at some latitude between $-b$ and b. Let x_k be any number in the sequence $(x_i)_{i=1}^n$ greater than a with k even, $k \geq 3$; such a k exists because the sequence $(x_i)_{i=1}^n$ is unbounded. We claim that if the robot starts in the origin and it follows the finite algorithm $EMW(x_1)EME(x_2)EMW(x_3)\ldots EMW(x_{k-1})EME(x_k)$, which is an initial segment of F, the robot visits the destination point.

Indeed, we note first of all that after the robot follows $EMW(x_1)EME(x_2)$ $EMW(x_3)\ldots EMW(x_{k-1})$ starting from the origin, it has a longitude at most that of the longitude of the origin (i.e. the robot is to the west of the origin). Indeed, after the robot follows $EMW(x_1)EME(x_2)EMW(x_3)\ldots EME(x_{k-2})$ in every maze in $\mathcal{C}$ it gets to a point $\mathbf{x} = (x, 0)$ on the x-axis, as $|EME(a)|_N = |EME(a)|_S$ and $|EMW(a)|_N = |EMW(a)|_S$ for all a and every maze in $\mathcal{C}$ has no vertical edges removed. Moreover, after the robot follows $EMW(x_1)EME(x_2)$ $EMW(x_3)\ldots EME(x_{k-2})$ in every maze in $\mathcal{C}$ it cannot be at longitude more than $\lambda := |EMW(x_1)EME(x_2)EMW(x_3)\ldots EME(x_{k-2})| = |EME(x_1)| + |EME(x_2)| + \ldots + |EME(x_{k-2})|$ and if it is at longitude $x > 0$, then every pair of consecutive columns at longitude between 0 and x are connected by a horizontal edge at some latitude between $-\lambda$ and λ. Therefore, if the robot starts at $(x', 0)$ which is any of the points $(0, 1), \ldots (0, x)$ and it follows the algorithm $O(x_{k-1})\ W\ O(x_{k-1})\ NWS\ O(x_{k-1})\ SWN \ldots O(x_{k-1})\ N^{x_{k-1}}WS^{x_{k-1}}$ $O(x_{k-1})\ S^{x_{k-1}}WN^{x_{k-1}}\ O(x_{k-1}) =: A$, then its longitude is at most $x' - 1$, as $x_{k-1} \geq \lambda$. As, again, $x_{k-1} \geq \lambda \geq x$, after the robot follows $A^{x_{k-1}}$ starting from $(x, 0)$, it has longitude at most 0. This shows that after the robot follows $EMW(x_1)EME(x_2)EMW(x_3)\ldots EMW(x_{k-1})$ starting from the origin, it gets to a point $(x_1, 0)$ with $x_1 \leq 0$. Finally, remark that the same argument still holds if we replace x_{k-1} by any value at least λ.

Our initial claim is that if the robot starts at $(x_1, 0)$ and it follows $EME(x_k)$, it visits the destination point. For this it is enough to show that the robot visits the column of the destination point, as the robot follows an oscillatory move $O(x_k)$ starting from latitude 0 in every column it visits. By applying the exactly same argument as above, after the robot follows $(O(x_k)\ E\ O(x_k)\ NES\ O(x_k)\ SEN \ldots O(x_k)\ N^{x_k}ES^{x_k}\ O(x_k)\ S^{x_k}EN^{x_k}\ O(x_k))^{x_k/2}$ starting from $(x_1, 0)$, it has longitude at least 0. From the fact that every pair of columns between the origin and the destination point is connected by a horizontal edge at some latitude between $-x_k$ and x_k and that the difference in longitude between the destination point and the origin is at most $x_k/2$, the robot is guaranteed to visit the column of the destination point by the same argument as above. This finishes the proof. $\square$

We are now ready to introduce our general strategy and present a more general proof of Theorem 3.1.

For a subset of mazes $\mathcal{C} \subseteq \mathcal{M}$, in order to construct an algorithm A that solves $\mathcal{C}$ we adopt the following natural strategy: we find a countable cover $\mathcal{C} = \bigcup_{i=1}^{\infty} \mathcal{C}_i$ such that for each $i \in \mathbb{N}$ and each finite algorithm X we are able to find a finite algorithm A^i_X such that the concatenated algorithm XA^i_X solves $\mathcal{C}_i$. Then we are able to find an algorithm A that solves $\mathcal{C}$. Indeed, we construct recursively the finite algorithms $(B_n)_{n\geq 0}$ with $B_0 = \varnothing$ and $B_n = A^n_{B_0B_1\ldots B_{n-1}}$, then we take $A = B_1B_2\ldots$.

Second proof of Theorem 3.1. We begin this proof by defining two classes of algorithms. The aim of the first one is to move the robot eastwards in a certain organised pattern and we call it *move_east*; it is defined as follows for all $a, e \geq 1$:
$ME(a, e) := (((((E)^e NES)^e SEN)^e N^2ES^2)^e \ldots S^aEN^a)^e$.

We view $ME(a, e)$ as being composed from the multiple concatenation of $2a + 1$ different building blocks which we call *locomotory moves*: E, NES,

SEN, N^2ES^2, … N^aES^a, S^aEN^a. We constructed the class of algorithms *move_east* in such a way so that the following holds:

Let a, e be two natural numbers. Assume that the robot starts at the point $\mathbf{x} = (x, y)$ in any maze $M \in \mathcal{C}$ with no vertical edges removed. Take the maximal $k \leq e$ such that in M the columns c_i and c_{i+1} are joined at some latitude in $\{-a + y, \dots, a + y\}$ for all $x \leq i \leq x + k - 1$. Then, as the robot follows the algorithm $ME(a, e)$, it oscillates about the row r_y at latitudes between $y - a$ and $y + a$. After the algorithm is followed, the robot gets to a point $\mathbf{x'} = (x', y)$ with $x' \geq x + k$, in particular $x' = x + k$ if $k < e$. Moreover if we well order $\mathbb{Z}$ by $y < 1 + y < -1 + y < 2 + y < -2 + y < \dots$, then for all $x \leq i \leq x + k - 1$ the robot passes from the column c_i to the column c_{i+1} through the edge at the lowermost latitude with respect to this order.

This holds as a particular case of Lemma 4.1, which is a technical result used extensively, proved later in this section. One can also see how this statement follows from the construction of $ME(a, e)$, more specifically from the order in which the locomotory moves appear in the algorithm. The counterpart of *move_east* is called *move_west*, and we have:
$MW(a, e) := (((((W)^e NWS)^e SWN)^e N^2WS^2)^e \dots S^aWN^a)^e$.

The second class of algorithms that we define is called *oscillating_move_east*, which is a slight alteration of *move_east* formed by inserting the *oscillatory* algorithm $(N^bS^{2b}N^b)^e$ in between some locomotory moves; it is defined as follows for all $a, e \geq 1$ and $b \in \mathbb{Z}$:
$OME(a, e, b) := ((((((N^bS^{2b}N^b)^e E)^e NES)^e SEN)^e N^2ES^2)^e \dots S^aEN^a)^e$.

We note that in every maze with no vertical edge removed, after the robot follows the oscillatory algorithm $(N^bS^{2b}N^b)^e$, it gets back to the starting point. Therefore, for any parameters a, e, b, as the robot follows $OME(a, e, b)$ in any maze $M \in \mathcal{C}$, it has the same dynamics as it follows $ME(a, e)$ in M and in addition the robot visits some consecutive columns, beginning with the one which contains its starting point $\mathbf{x} = (x, y)$, at all latitudes between $y - b$ and $y + b$. Finally, we use the oscillatory algorithm $(N^bS^{2b}N^b)^e$ instead of $N^bS^{2b}N^b$

which works just as well for this purpose, only because we want $OME(a, e, b)$ to be a particular case of a much more general algorithm, $SME(a, e, K)$ that is defined later in this section.

The counterpart of *oscillating_move_east* is called *oscillating_move_west*, and we have:

$OMW(a, e, b) = ((((((N^b S^{2b} N^b)^e W)^e NWS)^e SWN)^e N^2 W S^2)^e \dots S^a W N^a)^e.$

We are now ready to prove the theorem using the general strategy described at the beginning of the section. In order to produce the desired countable cover, define $C_{n,\mathbf{x}}$ to be the set of all mazes with no vertical edges removed, with the destination point $\mathbf{x} = (x, y)$ and such that any two consecutive columns at longitude between 0 and x are joined at some latitude between $-n$ and n. Then $C = \bigcup_{n,\mathbf{x}} C_{n,\mathbf{x}}$ is a countable cover.

We let X be any finite algorithm and we fix the values $n, \mathbf{x}$. We now consider just the set of mazes $C_{n,\mathbf{x}}$ and we aim to construct an algorithm A such that XA solves $C_{n,\mathbf{x}}$, which by the discussion of our strategy at the beginning of the section is enough to conclude.

Say that the robot starts in any maze $M \in C_{n,\mathbf{x}}$ (as always, it starts in the origin) and it gets to the point $(a, 0)$ after it follows some finite algorithm Y. The following observation is crucial: for each pair $\{i, i+1\} \subset \{0, \dots, a\}$, the columns c_i and c_{i+1} are joined at some latitude in $\{-|Y|_S, \dots, |Y|_N\} \subseteq \{-|Y|, \dots, |Y|\}$. Therefore, after the robot follows the algorithm $Y\ ME(a, |Y|_W)$ in M it gets to some point $(a', 0)$ with $a' \geq 0$.

Now we build A as a concatenation of three algorithms $A := A_1 A_2 A_3$.

We construct $A_1 := S^{|X|_N - |X|_S}$; then after the robot follows the algorithm XA_1 in any maze $M \in C_{n,\mathbf{x}}$ it gets to the row r_0.

Define $a := \max\{|XA_1|_S, |XA_1|_N, n\}$; $e := |XA_1|_W + |x|$. Define $A_2 := ME(a, e)$. Then after the robot follows the algorithm XA_1A_2 in any maze $M \in C_{n,\mathbf{x}}$ it gets to some point $(x^+, 0)$ with $x^+ \geq x$.

Define $a := \max\{|XA_1A_2|_S, |XA_1A_2|_N, n\}$; $w := |XA_1A_2|_E + |x| + 1$; $b := |y|$. Define $A_3 := OMW(a, w, b)$. Then after the robot follows the algorithm $XA_1A_2A_3$ in any maze $M \in C_{n,\mathbf{x}}$, it gets to some point $(x^-, 0)$ with $x^- \leq x$ and it visits every intermediate column c_i with $x^- \leq i \leq x^+$ including c_x at every latitude in $\{-b, \ldots, b\}$ including y.

Therefore, after the robot follows $XA = XA_1A_2A_3$ in any maze $M \in C_{n,\mathbf{x}}$, it visits the destination point $\mathbf{x}$. Hence there exists an algorithm A such that XA solves $C_{n,xy}$. This finishes the proof. $\square$

We note that the missing vertical edges in the general model usually make the latitude of the robot unknown but it turns out that we can actually make use of the missing edges to regain the latitude of the robot. However, the unknown longitude and the missing edges require the robot to use a very subtle path to get to the destination point. As a result of these difficulties in the proof of Theorem 3.2 we need to make a much finer covering than in the proof of Theorem 3.1.

In the remainder of this section we introduce an algorithm which is a generalisation of $ME(a, e)$ and $OME(a, e, b)$ called *special_move_east* which is the main building block of the algorithms used in the general model. We then group all its properties in Lemma 4.1, which makes it one of the main results of the chapter. For $a, e \geq 1$ and a finite algorithm K we define:
$SME(a, e, K) := (((((K^eE)^eNES)^eSEN)^eN^2ES^2)^e...S^aEN^a)^e$. We view $SME(a, e, K)$ as being composed from the multiple concatenation of $2a+2$ different building blocks: the $2a+1$ locomotory moves E, NES, SEN, ... S^aEN^a and the special algorithm K.

Its counterpart, *special_move_west* is defined as:
$SMW(a, e, K) := (((((K^eW)^eNWS)^eSWN)^eN^2WS^2)^e...S^aWN^a)^e$.

Recall that $\mathcal{C} \subset \mathcal{M}$ is the set of mazes with no vertical edges removed. The following result encompasses the main properties of $SME(a, e, K)$ that are used countless times in the proof of Theorem 3.2.

LEMMA 4.1. *Let $a, e \geq 1$ and K be a finite algorithm such that for any maze $M \in \mathcal{C}$, if the robot follows K in M starting from the origin, it returns to the x-axis and it has a non-negative longitude. Let $\mathcal{C}' \subseteq \mathcal{C}$ be a subset of mazes for which there exists $0 \leq l \leq e-2$ with the following properties:*

(1) for every maze in $\mathcal{C}'$ and for any $0 \leq x \leq l$ the columns c_x and c_{x+1} are joined at some latitude between $-a$ and a;

(2) for any $\boldsymbol{v} = (x_v, 0)$ with $0 \leq x_v \leq l$, if the robot starts from $\boldsymbol{v}$ and follows K in any maze in $\mathcal{C}'$ it reaches some point $\boldsymbol{w} = (x_w, 0)$ with $x_v \leq x_w \leq l$ without visiting any vertex on the column c_{l+1}, i.e. without visiting any point of longitude at least $l+1$.

Then after the robot follows $SME(a, e, K)$ in any maze in $\mathcal{C}'$, (i) *it gets to some point $\boldsymbol{v} = (x_v, 0)$ on the x-axis with $x_v \geq l+1$;* (ii) *it does not pass from the column c_l to the column c_{l+1} for the first time while executing K;* (iii) *it passes from the column c_l to the column c_{l+1} for the first time while executing a locomotory move $N^m E S^m$, where $m \in \mathbb{Z}$ is the lowermost latitude with respect to the standard well order on $\mathbb{Z}: 0 < 1 < -1 < 2 < -2 < \ldots$ such that the columns c_l and c_{l+1} are joined at latitude m;* (iv) *immediately after this locomotory move $N^m E S^m$ is executed, the robot follows K.*

PROOF. Let M be any maze in $\mathcal{C}'$. As we prove the result for M, we make the convention that every time we say that the robot follows an algorithm, it follows that algorithm in M.

By the hypothesis on K, if the robot is on the x-axis and follows K (or $N^b E S^b$, $b \in \{-a, \ldots a\}$), it returns to the x-axis and its longitude does not strictly decrease. We fix x between 0 and l, so that the columns c_x and c_{x+1} are joined at some latitude b between $-a$ and a. Hence, if the robot starts from the point $(x, 0)$ and follows $N^b E S^b$ it gets to the point $(x+1, 0)$. Therefore, if the robot is on the x-axis at some longitude between 0 and l, then after each instance of the algorithm $(((((K^e E)^e NES)^e SEN)^e NNESS)^e \ldots S^a E N^a)^1$, the longitude of the robot increases by at least one. This proves *(i)*.

The conclusion *(ii)* follows directly from the hypothesis: indeed, for any $\mathbf{v} = (x_v, 0)$ with $0 \leq x_v \leq l$, if the robot starts from $\mathbf{v}$ and follows K, it gets at some point $\mathbf{w} = (x_w, 0)$ with $x_v \leq x_w \leq l$ without visiting any vertex on the column c_{l+1}, i.e. without visiting any point of longitude at least $l+1$.

From *(i)* and *(ii)* it follows that the robot passes for the first time from the column c_l to the column c_{l+1} while executing some instance of the move of the form N^bES^b, $-a \leq b \leq a$. Assume for a contradiction that $b \neq 0$ is not the lowermost latitude with respect to the well order on $\mathbb{Z}$ at which the columns c_l and c_{l+1} are joined. Let $b' \in \mathbb{Z}$ be the predecessor of b in the well order on $\mathbb{Z}$. Let Y be the largest initial segment of the algorithm $SME(a, e, K)$ strictly before this specific instance of this specific locomotory move, N^bES^b.

We define $A = (((((K^eE)^eNES)^eSEN)^eNNESS)^e...N^{b'}ES^{b'})^1$ and note that $A' := A^e = (((((K^eE)^eNES)^eSEN)^eNNESS)^e...N^{b'}ES^{b'})^e$ is a last segment of Y. Let B be the first segment of Y strictly before A', i.e. $Y = BA'$. For some $0 \leq x \leq l$ we denote by $(x, 0)$ the vertex where the robot gets if it starts from the origin and follows B. If the robot starts from $(x, 0)$ and follows A', it gets to the point $(l, 0)$. Also notice that if the robot starts from $(l, 0)$ and follows A, it gets to some point $(l', 0)$ with $l' \geq l+1$. Say the robot starts from the point $(x, 0)$ and follows the algorithm A^{e+1}. If the robot starts from the x-axis and follows A it advances eastwards at least 0 columns. When the robot starts from the x- axis and follows the $e+1$-th instance of A, it returns to the x-axis and advances eastwards at least one column. This means that if the robot starts from the x-axis and follows the w-th instance of A it returns to the x-axis and advances eastwards at least one column for each $1 \leq w \leq e+1$.

Therefore, if the robot starts from $(x, 0)$ and follows $A' = A^e$, it gets to the point $(l, 0)$ and advances eastwards at least e columns. This is a contradiction as $l + 1 \leq e$. This proves *(iii)*.

By *(iii)*, we know that the robot passes for the first time from the column c_l to the column c_{l+1} while executing the move N^mES^m. Assume K does not follow immediately that after this move is executed. Say Y is the first segment of the algorithm $SME(a, e, K)$ before and including this specific instance of this specific locomotory move, $N^m ES^m$.

We define $A = (((((K^e E)^e NES)^e SEN)^e NNESS)^e ... N^m ES^m)^1$ and note that $A' = A^e = (((((K^e E)^e NES)^e SEN)^e NNESS)^e ... N^m ES^m)^e$ is the last segment of Y. Let B be the first segment of Y strictly before A', i.e. $Y = BA'$. For some $0 \leq x \leq l$ we denote by $(x, 0)$ the vertex where the robot gets if it starts from the origin and follows B. If the robot starts from $(x, 0)$ and follows A', it gets to the point $(l + 1, 0)$. Say the robot starts from the point $(x, 0)$ and it follows the algorithm A^e. If the robot starts from the x- axis and follows A, it advances eastwards at least 0 columns. When the robot starts from the x-axis and follows the e-th instance of A, it returns to the x-axis and advances eastwards at least one column. This means that if the robot starts from the x-axis and follows the w-th instance of A it returns to the x-axis and advances eastwards at least one column for each $1 \leq w \leq e$.

Therefore, if the robot starts from $(x, 0)$ and it follows $A' = A^e$, it gets to the point $(l, 0)$ and advances eastwards at least e columns. This is a contradiction as $l + 2 \leq e$, proving *(iv)*. This finishes the proof. $\square$

We end this section with the following immediate corollary of Lemma 4.1.

Corollary 4.2. *Under the assumptions of Lemma 4.1, let us choose another order on $\mathbb{Z}$, say the n-*special order on $\mathbb{Z}$*: $0 < n < 1 < -1 < \dots$ instead of the well order on $\mathbb{Z}$ we considered in Lemma 4.1. Then if we construct*

$$SME^{(n)}(a, e, K) := ((((((K)^e N^n ES^n)^e E)^e NES)^e SEN)^e \dots S^a EN^a)^e,$$

the results in Lemma 4.1 still hold, with the amendment that after the robot follows $SME^{(n)}(a, e, K)$ in any maze in $\mathcal{C}$, it passes for the first time from the column c_l to the column c_{l+1} while executing $N^m ES^m$, where m is the lowermost latitude with respect to the n-special order on $\mathbb{Z}$.

5. The Cover

In the general model, let $\mathcal{F} \subset \mathcal{M}$ be the set of mazes with nonzero finitely many vertical edges removed in consecutive columns. Without loss of generality, we assume that for any maze in $\mathcal{F}$ the origin is the point $(0,0)$. In this section, we introduce a series of technical definitions that are used to classify the mazes in $\mathcal{F}$ in order to prove Theorem 3.2.

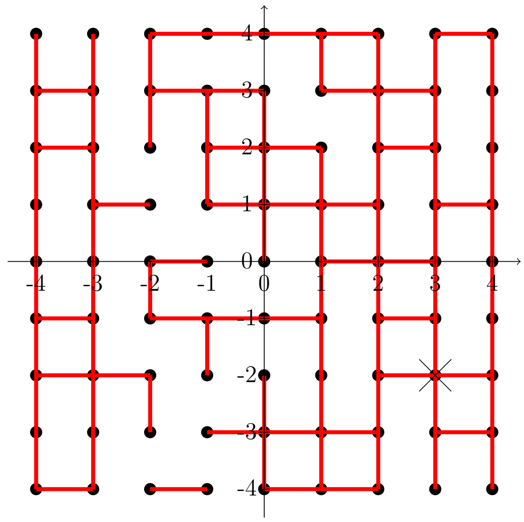

FIGURE 2. A local representation of a general maze $M \in \mathcal{F}$ that we use in order to illustrate our definitions. The destination point $(3,-2)$ is marked with an 'X'. We assume that there are no vertical edges removed from M other than the ones shown in the figure. For simplicity, we further assume that M is connected, though this may not be true for all mazes.

We recall that in order to construct an algorithm A that solves the set of mazes $\mathcal{F} \subset \mathcal{M}$ we adopt the following strategy: we find a countable cover $\mathcal{F} = \bigcup_{i=1}^{\infty} F_i$ with $(F_i)_{i\geq 1} \subseteq \mathcal{F}$ such that for each $i \in \mathbb{N}$ and each finite algorithm X we are able to find a finite algorithm A_X^i such that the concatenated algorithm XA_X^i solves F_i.

The aim of this section is to introduce the definitions that we need to use in order to construct the cover $(F_i)_{i\geq 1}$.

For any maze $M \in \mathcal{F}$ we denote by *HE*, *HNE*, *VE*, *VNE* a horizontal edge, horizontal non edge, vertical edge and vertical non edge, respectively. For M

as in Figure 2, between $(2,2)$ to $(3,2)$ there is a HE, between $(-1,-2)$ and $(0,-2)$ there is a HNE, between $(0,0)$ and $(0,1)$ there is a VE and between $(1,2)$ and $(1,3)$ there is a VNE.

From any maze $M \in \mathcal{F}$ we construct the maze $\overline{M} \in \mathcal{F}$ by adding all the possible VEs such that the connected component of the origin is unchanged as a graph in the process. In other words we add all the possible VEs with both endvertices not in the connected component of the origin. The new maze $\overline{M}$ has the nice property that the robot can get from the origin to one vertex of every VNE in $\overline{M}$. We note that an algorithm solves M if and only if it solves $\overline{M}$. Therefore, in order to prove Theorem 3.2 it is enough to construct an algorithm A which solves $\overline{\mathcal{F}} = \{\overline{M} \mid M \in \mathcal{F}\} \subseteq \mathcal{F}$.

The rest of the section will only address mazes in $\overline{\mathcal{F}}$, so for any maze $\overline{M} \in \overline{\mathcal{F}}$ we introduce the following definitions.

We define a *vertical strip* to be any subgraph of a maze in $\mathcal{F}$ obtained by taking the union between the restriction of the maze to a set C of consecutive columns and all the horizontal edges which have at least one vertex in C. For such a vertical strip S we call its restriction to C its *interior* and the complement of the interior in S the *boundary*. Let S be the smallest vertical strip that contains in its interior all the VNEs, the origin and the destination point. As there is only a finite number of VNEs, S contains a finite number of (consecutive) columns. For M as in Figure 2, S is the subgraph formed from the columns $c_{-2}, \ldots, c_3$ together with all the HEs between c_{-3} and c_{-2} and all the HEs between c_3 and c_4; in particular the vertex $(-3,-2)$ and the edge between $(3,1)$ and $(4,1)$ are in S, but the vertex $(-3,2)$ is not.

Considering the maze with all its HEs deleted, we can label the connected components obtained in this way by *upper infinite columns*, *lower infinite columns*, *infinite columns*, and *finite columns* accordingly. For M as in Figure 2, there are 4 upper infinite columns, e.g. the infinite path $(-2,2),(-2,3),\ldots$; there are also 4 lower infinite columns, e.g. the infinite path $(-2,-4),(-2,-5),\ldots$; the infinite columns are $c_{-3}, c_{-4}, \ldots$ and $c_2, c_3, \ldots$;

examples of finite columns are $(-2,1)$, the path $(-2,-1),(-2,0)$ or the path $(0,0),(0,1),(0,2),(0,3)$.

Considering only the HEs in S, we call a *pass* any of the following edges:

(1) the HE of smallest latitude with respect to the usual order on $\mathbb{Z}$ between two upper infinite columns, or between an upper infinite column and an infinite column, e.g. the edge between $(-2,4)$ and $(-1,4)$ or the edge between $(1,3)$ and $(2,3)$, respectively in Figure 2;

(2) the HE of largest latitude with respect to the usual order on $\mathbb{Z}$ between two lower infinite columns, or between an lower infinite column and an infinite column, e.g. the edge between $(0,-3)$ and $(1,-3)$ or the edge between $(1,1)$ and $(2,1)$, respectively in Figure 2;

(3) the HE of smallest latitude between two infinite columns with respect to the well order on $\mathbb{Z} : 0 < 1 < -1 < 2 < -2 < \ldots$, e.g. the edge between $(2,0)$ and $(3,0)$ in Figure 2.

Every maze has a finite number of VNEs, so every maze has a finite number of passes. We further note that between two consecutive columns in S there might not be a pass, if there is no HE between them. Finally, as a few more revealing examples, we note that in Figure 2 the edge between $(-3,1)$ and $(-2,1)$ is not a pass, and neither is the one between $(-4,-1)$ and $(-3,-1)$ which is not in S; however, the edge between $(3,1)$ and $(4,1)$ is in S and it is also a pass.

Furthermore, we define the following regions: the *obstacle strip* is the smallest vertical strip that contains all VNEs in its interior. For example, in Figure 2 the obstacle strip is formed from the columns $c_{-2}, \ldots c_1$ together with all the HEs incident with any vertex on c_{-2} or c_1. The *west strip* and *east strip* are the subgraphs situated at the left and right of the obstacle strip, respectively; they are both formed by a union of consecutive columns and all horizontal edges with both endpoints belonging to these columns. For example, in Figure 2 is formed from the columns $c_{-3}, c_{-4}, \ldots$ and the east strip is formed from the columns $c_2, c_3, \ldots$. We note that the obstacle strip and the east or west strip may intersect only in a certain set of vertices, i.e. the eastern or western endvertices of the edges that emerge on the right or left side of the obstacle strip, respectively; they have no edges in common.

We define the *primary rectangle* to be the induced subgraph contained in the smallest rectangle that contains the origin, the destination point, all the passes and all the VNEs. The primary rectangle is well defined, as there is a finite number of passes and VNEs. Let p be the smallest positive integer such that the primary rectangle is strictly contained in the interior of the square centred at the origin with the set of vertices $\{(\pm p, \pm p)\}$ (see Figure 3). We call p the *parameter of the primary rectangle*.

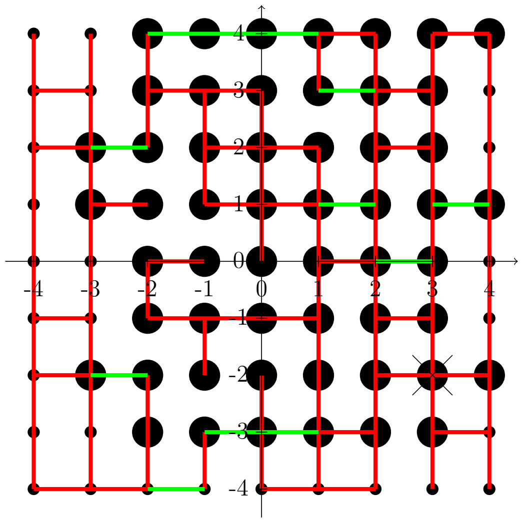

FIGURE 3. We assume that there are no vertical edges removed other than the ones shown in the figure. The destination point is $(3, -2)$. All the passes are marked with green edges. The primary rectangle has vertices $(-3, 4), (4, 4), (-4, 4), (-3, -4)$ and $p = 5$. The special vertices are drawn larger.

We define the *special vertices* to be all the vertices in S that are connected to the destination point and have the same latitude as an endpoint of a VNE (see Figure 3). Notice that there is a finite number of special vertices and label them $1, 2, \ldots, s$. We note that there must exist a path contained in the primary rectangle between each special vertex and the destination point. Indeed, the fact that all all VNEs are contained in the primary rectangle and the way we define passes allows us to find paths contained in the primary rectangle between the accessible infinite/upper and lower infinite and finite columns of the primary rectangle; this further allows us to find paths contained in the primary rectangle from the special points to the destination point.

Let l_i be the length of a shortest path contained in the primary rectangle from $i \in \{1, 2, \ldots, s\}$ to the destination point and set the following constant which depends only on the local configuration of the maze inside the primary rectangle:

$$l' = 1 + ((((l_1)2 + l_2)2 + l_3)2 + \ldots + l_{s-1})2 + l_s.$$

The *secondary rectangle* is obtained from the primary rectangle by augmenting it l' units in each of the four directions. Note that given the local configuration of the maze inside the secondary rectangle, we can construct a finite algorithm L' such that if the robot follows L' starting from any special point, it visits the destination point without leaving the secondary rectangle. Indeed, assume the robot starts at the special point labeled 1. We construct firstly a finite algorithm L_1 that takes the robot to the destination point with $|L_1| = l_1$. Then assume that the robot starts at the special point labeled 2 and that it first follows the algorithm L_1. We write the algorithm L_2 as a concatenation of two sub-algorithms. The first cancels the action of L_1 and brings the robot back to the special point 2 and the second sub-algorithm takes the robot further to the destination point. This can be done with at most $l_1 + l_2$ instructions, so without loss of generality $|L_2| \leq l_1 + l_2$. Moreover, if the robot starts at any of the special points labeled 1 or 2 and follows L_1L_2 it gets to the destination point. We continue in this way: given $L_1, L_2, \ldots, L_{i-1}$ and assuming that the robot starts at the special point i, we construct L_i as a concatenation of two sub-algorithms. The first brings the robot back to the special point i and the second takes the robot further to the destination point. This can be done with $|L_i| \leq (|L_1| + \ldots + |L_{i-1}|) + l_i$. Finally, take $L' = L_1L_2 \ldots L_s$ with $|L'| \leq l'$ which has the property that if the robot follows

L' starting from any special point, it visits the destination point. The role of adding 1 to the sum is that to ensure that the secondary rectangle augments non-trivially the primary rectangle.

In the rest of the section we define a series of very technical configurations. We group the mazes according to these configurations and obtain the desired countable cover at the end of the section. The importance of these configurations only becomes clear in Section 6 and Section 7, where we will recall them when appropriate.

For simplicity we use cardinal directions in our definitions. We say that *the row r_i is to the north of the row r_j* or *above row r_j*, provided $i > j$. By an *easternmost H(N)E e* with a certain property $\mathcal{P}$ we mean that e has $\mathcal{P}$ and no other H(N)E with $\mathcal{P}$ has longitude greater than e. These definitions easily extend to the other directions: *westernmost*, *uppermost*, *lowermost*. In pairings (e.g. "the lowermost easternmost HNE with $\mathcal{P}$") we always give priority to the first direction and then to the second one. For example, in order to find the uppermost easternmost HNE below all VNEs in the west strip, we first look for the row of highest latitude below all VNEs on which there is a HNE in the west strip and then on this row we pick the one HNE in the west strip with the largest longitude.

Define a *west bump* to be any of the easternmost HNE in the west strip or at the border between the west strip and the obstacle strip (i.e. with at least one vertex in the west strip) on a row that intersects some finite column. For example in Figure 2, the HNE between $(-4, 1)$ and $(-3, 1)$ and the HNE between $(-3, 2)$ and $(-2, 2)$ are both west bumps with the rows r_1 and r_2 intersecting the finite column $(-1, 1), (-1, 2), (-1, 3)$. Using symmetry, define similarly an *east bump*. We note that there are a finite number of west and east bumps.

If there exists a row which is a path when restricted to the west strip, but contains a HNE, then call the smallest such row with respect to the standard well order on $\mathbb{Z}$ a *magical west row*; define its *west cutoff* to be its westernmost HNE. Define similarly a *magical east row* and its *east cutoff*.

We define a *west pipe* to be any of the easternmost configurations in the west strip of three vertices $(x, y), (x+1, y), (x+2, y)$ where between (x, y) and $(x+1, y)$ there is a HE and between $(x+1, y)$ and $(x+2, y)$ there is a HNE, which can be at the border between the west strip and the obstacle strip. For example in Figure 2, $(-4, 2), (-3, 2), (-2, 2)$ is a west pipe. Note that a maze may have infinitely many west pipes. We define similarly an *east pipe* to be any of the westernmost configurations in the east strip of three vertices $(x, y), (x+1, y), (x+2, y)$ where between $(x+1, y)$ and $(x+2, y)$ there is a HE and between (x, y) and $(x+1, y)$ there is a HNE, which can be at the border between the east strip and the obstacle strip. For example in Figure 2, $(2, 1), (3, 1), (4, 1)$ is an east pipe.

Furthermore, we define the *special west pipe* to be the west pipe on the smallest row that has a west pipe, with respect to the standard well order on $\mathbb{Z}$, if such a row exists. Note that in Figure 2 the special west pipe may not be $(-4, -1), (-3, -1), (-2, -1)$ as we do not know from the picture whether there are west pipes on r_0 or r_1, but we do know that it is the west pipe on r_{-1}. We define similarly the *special east pipe*. Note that if a maze does not have any special west pipe, then in the west strip any row is either a path or it is the complement of an infinite path followed by a finite path.

We define an *almost empty west row* to be a row that in the west strip is the complement of an infinite path followed by a non-empty finite path. Thus, in Figure 2, both r_0 and r_1 cannot be almost empty west rows as the non-empty finite path in the west strip is missing for both of these columns; the edge between $(-3, 1)$ and $(-2, 1)$ does not belong to the west strip. We define similarly an *almost empty east row*. We define the *special almost empty west row* to be the smallest almost empty west row with respect to the standard well order on $\mathbb{Z}$, if such a row exists. We define the *west cutoff* of a special almost empty west row to be its easternmost HNE in the west strip. We define similarly the *special almost empty east row* and its *east cutoff*. For example, if in Figure 2 r_2 was the special almost empty east row, its east cutoff would be the edge between $(3, 2)$ and $(4, 2)$. Finally, we define an *empty west row* to be a row that in the west strip is empty; for the 'special' label in this context, we need in addition that the latitude of the row is large in absolute value. So we define the *special empty west row* to be the empty west row of smallest latitude, greater than $-3p$ (where the parameter of the primary rectangle, p, is defined above) with respect to the standard well order on $\mathbb{Z}$, if such a row exists. We define the *natural special empty west row* to be the empty west row of smallest latitude, without the additional constraint. We define similarly the *special empty east row* and the *natural special empty west row*.

We define the *upper west pass* to be the lowermost HE between the easternmost infinite column of the west strip and the westernmost upper infinite column with the property that its latitude k is greater than that of any pass in the obstacle strip, if such a HE exists. We define similarly the *upper east pass*, *lower west pass* and *lower east pass*. For example, in Figure 3 the edge between $(-3, -4)$ and $(-2, -4)$ is the lower west pass. Also, in Figure 4 the upper/lower west/east passes are the green edges.

Let us call the pair of columns at the border between the west strip and the obstacle strip (c_a, c_{a+1}), so c_a is in the west strip and c_{a+1} is in the obstacle strip. Let us call the pair of columns at the border between the obstacle strip and the east strip (c_b, c_{b+1}), so c_b is in the obstacle strip and c_{b+1} is in the east strip. We define the *west ascending chain* (if such a structure exists) to be the finite sequence of HEs: $HE_a, HE_{a+1}, \ldots, HE_b$ such that HE_a is the upper west pass and HE_m is the lowermost HE between the pair of columns (c_m, c_{m+1}) at latitude at least that of HE_{m-1} for $m = a + 1, \ldots, b$ (see Figure 4). Similarly, we define the *east ascending chain*, *west descending chain* and *east descending chain*. If a west ascending chain $HE_a, \ldots, HE_b$ exists with HE_b on some row r_t, we define the *upper west constant* $c_{uw} := t + p$, where p is the parameter of the primary rectangle. We define similarly the constants *lower west constant*, *upper east constant* and *lower east constant*.

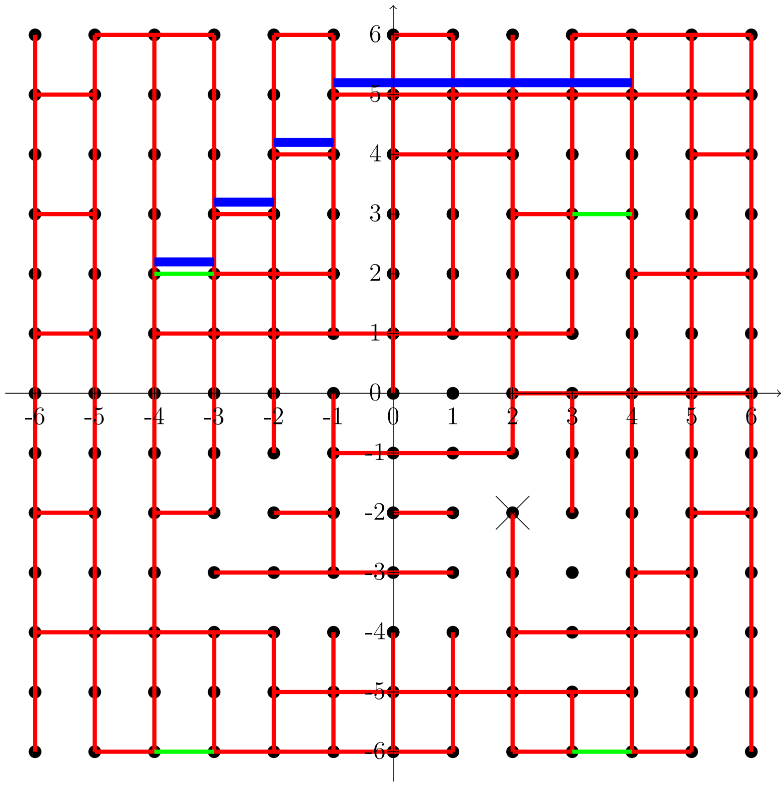

FIGURE 4. We assume that there are no vertical edges removed other than the ones shown in the figure. The green edges are the upper/lower west/east passes. Neither the HE $(-4,-2),(-3,-2)$ nor $(-4,1),(-3,1)$ is the upper west pass, as they are not above all the passes in the obstacle strip. The blue coloured edges in order from left to right form the west ascending chain.

Assume that the upper west pass is on some row r_k. We define an *upper west paired HNEs* to be any pair of HNEs with the same longitude in the west strip such that the upper HNE is at latitude k and the lower HNE is at latitude at most $k - c_{uw}$, where c_{uw} defined above is the upper-west constant. We define similarly the *upper east paired HNEs*, *lower west paired HNEs* and *lower east paired HNEs* with respect to the corresponding constants c_{ue}, c_{lw}, and c_{le}, respectively. We define the *special upper west paired HNEs* (if such a structure exists) to be the upper west paired HNEs with the uppermost and easternmost lower HNE. We recall that in all such instances we give priority to the first condition and then the second one. We define similarly the *special upper east paired HNEs*, *special lower west paired HNEs* and *special lower east paired HNEs*.

With k being as always the latitude of the upper west pass, we define the *upper west pipe* to be the west pipe on the row r_k, if one exists. We define similarly the *lower west pipe*, the *upper east pipe* and the *lower east pipe*. We define the *upper west cutoff* to be the easternmost HNE on the row r_k in the west strip, if one exists. We define similarly the *lower west cutoff*, the *upper east cutoff* and the *lower east cutoff*.

We define the *upper west HNE* (if such a structure exists) to be the lowermost westernmost HNE at the north-east of the uppermost westernmost VNE. We define similarly the *upper east HNE*, *lower west HNE* and the *lower east HNE*.

Being consistent with the constants a and b introduced in the definition of the west ascending chain, we define the parameters $h_{(m,m+1)}$ to be the latitude of the uppermost HE between two consecutive upper infinite columns or between an infinite column and an upper infinite column on c_m and c_{m+1} if such a HE exists and ∞ otherwise for $m = a, \dots, b$.

We define similarly the parameters $l_{(m,m+1)}$ to be the latitude of the lowermost HE between two consecutive lower infinite columns or between an infinite column and a lower infinite column on c_m and c_{m+1} if such a HE exists and infinity otherwise for $m = a, \dots, b$.

We define the parameters $w_1, e_1, w_2, e_2, w_3, e_3, w_4, e_4$ to be the latitude of the magical west/east row, the special almost empty west/east row, the special empty west/east row, and the natural special empty west/east row if such a configuration exists and infinity otherwise, respectively.

We finally define the *tertiary rectangle* to be the subgraph contained in the smallest rectangle that contains the secondary rectangle and all the west/east bumps, upper/lower west/east cutoffs, upper/lower west/east pipes, special west/east pipes, upper/lower west/east passes, west/east ascending/descending chains, upper/lower west/east paired HNEs and the upper/lower west/east HNEs.

As in the case of the primary rectangle, let q be the smallest positive integer such that the tertiary rectangle is strictly contained in the interior of the square centred at the origin with vertices $\{\pm\frac{q}{3}, \pm\frac{q}{3}\}$. We call q, together with the upper/lower west/east constants, $h_{(m,m+1)}, l_{(m,m+1)}$ for $m = a, \ldots, b$, $w_1, e_1, w_2, e_2, w_3, e_3, w_4, e_4$ the *parameters of the tertiary rectangle*. Therefore, when we construct an algorithm by inspecting the tertiary rectangle, we have access to the subgraph contained in the tertiary rectangle and the values of all its parameters.

We group the mazes in $\overline{\mathcal{F}}$ according to agreeing on the destination point, the subgraph contained in the square $\{\pm q, \pm q\}$ and the set of parameters of the tertiary rectangle. We thus obtain a countable cover $\overline{\mathcal{F}} = \bigcup_{i=1}^{\infty} F_i$. It is obvious directly from the definitions that we set above that such a construction is achievable.

All of these definitions are used in the following section to prove Theorem 3.2 and the relevant ones will be recalled where appropriate.

6. The General Model, Preliminaries

In this section and Section 7 we prove Theorem 3.2. Considering that the proof is very complex, we split it into two parts: in this part, we present the set-up and we show that we are able to assume without loss of generality that the robot is in the east/west strip on the x-axis or it has already visited the destination point; in Section 7 we show that we are able to write an algorithm that further guides the robot to visit the destination point.

Following our strategy, we assume that we are given F_i and a finite algorithm X and we aim to construct a finite algorithm A such that XA visits the destination point of F_i. We construct the algorithm A from several sub-algorithms treated in separate subsections, each with a specific task: in **Part I** we position the robot in the east strip, in **Part II** we position the robot in the west strip at latitude 0 and in **Part III** we guide the robot through the destination point. In each part we consider a finite number of cases for the subsets F_i so that although a sub-algorithm depends quantitatively on F_i and X, it does depend qualitatively only on the case. We treat each case separately. According to their degree of generality, we label the broader cases as "Propositions", and the more specific cases as "Claims".

In fact, at the end of most parts, we prove something more. We show how our methods can be generalised in order to produce algorithms that achieve the desired goal in the more general case of a finite number the VEs removed (in other words, we do not need the condition that the VNEs are in consecutive columns). The only case where we do not provide such a generalisation is Case 4 of Part III.

PROOF OF THEOREM 3.2. Let F_i be any of the classes of mazes defined above and assume we are given a finite algorithm X. Let $\lambda := |X|$.

6.1. Reset latitude.

Part 0. We recall the finite algorithm L' defined in Section 5 for a particular maze M, which had the property that if the robot starts at any special point of M and follows L', it visits the destination point. Take $M \in F_i$ and construct its L' as described in Section 5. We claim that for any $M' \in F_i$, the algorithm L' has the same property in M', i.e. if the robot starts at any special point of M' and follows L', it visits the destination point of M'. This follows from the fact that all the mazes in F_i share the destination point, the secondary rectangle and in particular the set of special points. Therefore, we pick this L' as a representative for the set of mazes F_i.

We now define the algorithm

$$L = L_E = L' \ N^{\varepsilon} \ ME(|L'N^{\varepsilon}|, |L'N^{\varepsilon}|),$$

where the correcting constant $\varepsilon \in \mathbb{Z}$ is picked such that $|L'N^{\varepsilon}|_N = |L'N^{\varepsilon}|_S$ and therefore $|L|_N = |L|_S$. We recall that the algorithm *move_east*, ME used in

constructing L was defined in Section 4:
$ME(a,e) := (((((E)^e NES)^e SEN)^e N^2 ES^2)^e \dots S^a EN^a)^e$. Let $l := |L|$. The counterpart of $L = L_E$ is

$$L_W = L' \ N^\varepsilon \ MW(|L'N^\varepsilon|, |L'N^\varepsilon|),$$

and as before we have $|L_W|_N = |L_W|_S$ and also $|L_W| = l$.

The algorithm L is a generic algorithm used in several other algorithms below. We remark that if the robot starts from a special point and it follows L in any maze in F_i, it gets to the destination point; this property is inherited from L'. We further note that if the robot is at the origin on a maze with no VNEs and it follows L, then it returns to the x-axis and its longitude does not decrease. These properties are crucial in order to apply Lemma 4.1.

We finally note that all mazes in F_i also share the same parameter of the primary rectangle p and parameter of the tertiary rectangle q and we keep this notation consistent for the rest of the proof.

Part I. The algorithm *rough_positioning_east* defined in this part aims to either position the robot in the east strip or to make the robot visit the destination point. We define

$$RPE := ME(\lambda + p, \lambda + p) \ N^{l+\lambda+4p} \ L \ S^{2(l+\lambda+4p)} \ L.$$

PROPOSITION 6.1. *For any maze in F_i, after the robot follows the algorithm X RPE, it is either in the east strip or it has visited the destination point.*

PROOF. Pick any maze in F_i. We claim that by our choice of parameters of ME, after the robot follows $X \ ME(\lambda + p, \lambda + p)$, it is either in the east strip or in the obstacle strip, but not in the west strip. Indeed, assume for a contradiction that after the robot follows $X \ ME(\lambda + p, \lambda + p)$, it is in the west strip. Denote by $\mathbf{x} = (x, y)$ the position of the robot after it follows X starting from the origin. By assumptionm $\mathbf{x}$ must be in the west strip as the algorithm ME has no instruction W. Therefore, as the robot follows $ME(\lambda + p, \lambda + p)$, it does not visit any endvertex of a VNE. We recall that all mazes in $F_i \subset \overline{\mathcal{F}}$ have the property that for every VNE at least one of its vertices is accessible, hence the westernmost column of the obstacle strip c_{a+1} is accessible from $\mathbf{x}$. The robot starts in the origin which is at most p units in longitude away from the obstacle strip as the primary rectangle contains the origin and all the VNEs. Hence the column $c_{y+\lambda+p}$ is not in the west strip. Moreover, every pair of consecutive columns at longitude between y and $a+1$ are connected by a HE at some latitude between $x+\lambda+p$ and $x-\lambda-p$, as the primary rectangle contains all the passes and the VNEs. Therefore, if the robot starts from $\mathbf{x}$ and follows $ME(\lambda+p, \lambda+p)$, it gets to a longitude at least $y+\lambda+p$, which is not in the west strip. This contradiction proves the claim.

Hence, after the robot follows $X\ ME(\lambda+p, \lambda+p)$, its longitude is at least $a+1$ and so it is either in the east strip or in the obstacle strip. In the former case, after the robot follows $X\ RPE$, it remains in the east strip. Indeed, while the robot follows $N^{l+\lambda+4p}\ L$ starting in the east strip, its latitude is too high to meet any VNE and on a maze with no VNE if the robot follows L its longitude does not decrease so the robot remains in the east strip. Therefore, after the robot follows also $S^{2(l+\lambda+4p)}\ L$ its latitude is too low to meet any VNE, and it remains in the east strip by a similar argument. To conclude, if the robot gets to the east strip after it follows the initial segment $X\ ME(\lambda+p, \lambda+p)$, then it remains in the east strip after it follows $X\ RPE$.

In the latter case, after the robot follows $X\ ME(\lambda+p, \lambda+p)$, it is in the obstacle strip either in (1) a lower infinite column or a finite column or (2) an upper infinite column. In case (1), after the robot follows $X\ ME(\lambda+p, \lambda+p)\ N^{l+\lambda+4p}$ it gets to a special point. Therefore after the robot follows $X\ ME(\lambda+p, \lambda+p)\ N^{l+\lambda+4p}\ L$, it gets to the destination point. In case (2), while the robot follows $N^{l+\lambda+4p}\ L$ it does not meet any VNE and its longitude does not decrease, so after it follows the initial segment $X\ ME(\lambda+p, \lambda+p)\ N^{l+\lambda+4p}\ L$, it is either in the east strip or in the obstacle strip in an upper infinite column.

In both cases, it is clear that after the robot follows $X\ RPE$ it is either in the east strip or it has visited the destination point. $\square$

Remark. In the first part of the proof of Proposition 6.1 we argue that the parameters $(\lambda + p, \lambda + p)$ of ME are large enough for the robot to have longitude at least $a + 1$. The key of this argument is two-fold: firstly, all the passes are in the primary rectangle which has parameter p; secondly, if the robot starts from the origin and follows the algorithm X with $|X| = \lambda$, it can not advance more than λ columns east or west and any two consecutive columns between its initial and final position are connected at latitude no more than λ in absolute value. We do not expand this argument every time we use it, but instead we use the phrase "by our choice of parameters" to mark that the same reasoning is used in similar instances to prove that the robot advances westwards/eastwards to the desired longitude.

To finish **Part I**, we note that although we used in the proof of Proposition 6.1 the fact there are no infinite columns in the obstacle strip, a variation of RPE can be used to position the robot in the east strip, even if we drop this assumption. This note is important, because it shows that **Part I** can be generalised to improve Theorem 3.2 by dropping the consecutive column condition for the finite number of VNEs. To present this variation, assume that infinite columns are allowed in the obstacle strip, i.e. the (finitely many) VNEs need not be in consecutive columns.

We define now the algorithm RPE' that generalises RPE as described above. It is formed by $\lambda + p$ subalgorithms $S_1, \dots S_{\lambda+p}$ concatenated in order. We define

$$S_i = N^{\lambda_i+p+2l}\ L\ S^{\mu_i+p+2l}\ ME(\gamma_i, 1),$$

for $i = 1, \dots, \lambda + p$. The parameters $\lambda_i, \mu_i, \gamma_i \in \mathbb{N}$ are chosen to be at least the number of instructions written in the whole algorithm until they occur, for example we can take $\lambda_1 = |X| = \lambda$, $\mu_1 = |X\ N^{\lambda_1+p+2l}\ L|$, $\gamma_1 =$

$|X\ N^{\lambda_1+p+2l}\ L\ S^{\mu_1+p+2l}|$, etc. Finally, let

$$RPE' = S_1\ S_2\ \dots\ S_{\lambda+p}.$$

Note that for the mazes we consider, we first replace all the VNEs with VEs which do not change the connected component of the origin, so every pair of consecutive columns in the obstacle strip must be connected by an accessible HE. The reason why RPE' indeed generalises RPE is similar to the argument in the proof of Proposition 6.1: here, after the robot follows every S_i it either moves at least one column to the east or it has visited the destination point.

Moving on from this digression, by Proposition 6.1 we may assume that we are given F_i and a finite algorithm X with $\lambda = |X|$ such that after the robot follows X in any maze in F_i, it is either in the east strip or it has visited the destination point. Without loss of generality, we assume that the robot is in the east strip and our aim is to build a finite algorithm A such that XA solves F_i.

Part II. The algorithm *reset_latitude_west* defined in this part aims to either position the robot in the west strip on the x-axis (i.e. at latitude 0) or to make the robot visit the destination point.

Case (1). We assume that the mazes in F_i do not contain a pass between the obstacle strip and the east strip. Then from the assumptions on the mazes in F_i, the east strip is connected to a finite column in the easternmost column of the obstacle strip c_b. This follows from the fact that for every VNE of every maze in F_i, at least one of its vertices is accessible from the origin. Let R be the lowermost finite column in c_b such that there exists a HE between R and the east strip. Let $\mathbf{v} = (b, i)$ be the lowermost vertex of the finite column R. In this case, we define the algorithm

$$RLW := S^{\lambda+p+l}\ SMW(2\lambda + 2p + l, \lambda + 2p, L).$$

Claim 6.2. *For any maze in F_i, after the robot follows the algorithm $X\ RLW$, it visits the destination point.*

Proof. After the robot follows $X\ S^{\lambda+p+l}$, it is in the east strip at a certain point $\mathbf{x} = (x, j)$, with $j \leq i - l$. By the choice of parameters and by Lemma 4.1, while the robot follows $SMW(2\lambda + 2p + l, \lambda + 2p, L)$ it advances westwards in the east strip oscillating about the row r_j. It passes for the first time from the column c_{b+1} to the column c_b not while executing L, but while executing a locomotory move. Moreover, if we well order $\mathbb{Z}$ by $j < 1 + j < -1 + j < 2 + j < -2 + j < \ldots$, then the robot passes for the first time from the column c_{b+1} to the column c_b through the smallest HE with respect to this order and so it gets to the point $\mathbf{v}$, which is a special point. Immediately afterwards, it follows L and it reaches the destination point. The conclusion follows. □

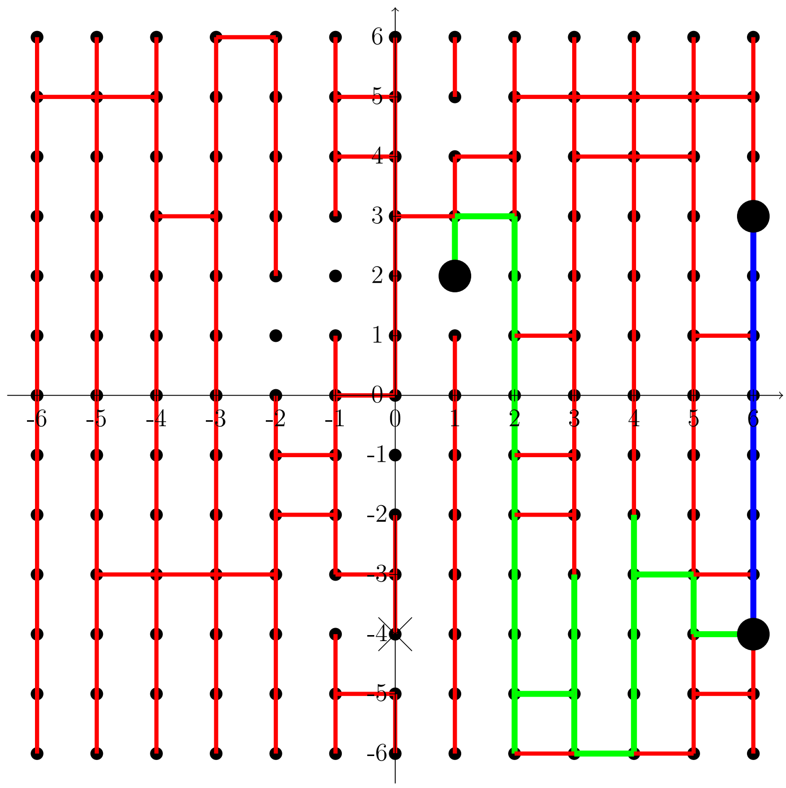

Figure 5. **Part II, Case (1).** There is no pass between the obstacle strip and the east strip. We assume that there are no VEs removed other than the ones shown in the figure.

Let us examine the example shown in Figure 5. Let us suppose that the position of the robot after following X is $(6, 3)$. The first segment $S^{\lambda+p+l}$ of RLW takes the robot to some very small latitude j such that if it follows L starting from any point of r_j in the east strip, it will always remain in the east strip. This can be done, as there is no pass between the obstacle strip and the east strip. For our example, we may assume that after the robot follows $XS^{\lambda+p+l}$ it gets to the point $(6, -4)$, though this latitude should be much smaller. The green route to the special point $(1, 2)$ is the route of the robot if it would follow the algorithm MW. The algorithm SMW used in RLW generalises MW by inserting the algorithm L between locomotory moves. However, L is constructed in such a way that if the robot follows it while it is in the east strip, its longitude does not increase. Moreover, the latitude of the robot is so small that it will never pass from the column c_{b+1} to c_b while following L. By Lemma 4.1 and the choice of parameters of SMW, the robot reaches the special point $(1, 2)$ while executing a locomotory move. Immediately afterwards, it executes L and it gets to the destination point. Finally, we remark that when the robot reaches the obstacle strip from the east strip for the first time, it does not enter the finite column $(1, 2), (1, 3), (1, 4)$ via the HE $(1, 4), (2, 4)$ or indeed it does not enter any other finite column which is above R. This follows from the order of locomotory moves in SMW, i.e. priority is given to smaller latitudes.

Case (2). We assume without loss of generality that the mazes in F_i contain a pass π between the easternmost lower infinite column and the east strip. In this case, we define the algorithm

$$RLW := S^{\lambda+p}\ SMW(2\lambda + 2p, \lambda + 2p, K)\ N^{2\lambda+6p+l}\ L_W\ S^{2p+l-k},$$

where $K = N^{2\lambda+4p}S^{2\lambda+4p}$ and k is the latitude of the lowermost special vertex.

PROPOSITION 6.3. *For any maze in F_i, after the robot follows the algorithm $X\ RLW$, it is either in the west strip on the x-axis or it has visited the destination point.*

Proof. After the robot follows $X\ S^{\lambda+p}$ it is in the east strip at a certain latitude say j, smaller than the latitude of the pass π. By the choice of parameters and by Lemma 4.1, while the robot follows $SMW(2\lambda+2p, \lambda+2p, K)$ it advances westwards in the east strip oscillating about the row r_j. It passes for the first time from the east strip to the obstacle strip not while executing K. Moreover, if we well order $\mathbb{Z}$ by $j < 1+j < -1+j < 2+j < \dots$, then the robot passes from the east strip to the easternmost lower infinite column in the obstacle strip through the smallest HE with respect to this order. Immediately afterwards, it follows K and gets at latitude $2\lambda+4p$ below the easternmost lowermost special vertex. By the choice of parameters the robot advances westwards only through lower infinite columns while in the obstacle strip. Therefore, after the robot follows $X\ S^{\lambda+p}\ SMW(2\lambda+2p, \lambda+2p, K)$, it is either (1) in the west strip at latitude $2\lambda+4p$ below the lowermost special vertex, i.e. at latitude $k-2\lambda-4p$ or (2) in the obstacle strip in a lower infinite column c_m at latitude $2\lambda+4p$ below some special vertex (see Figure 6).

In case (1), while the robot follows $N^{2\lambda+6p+l}\ L_W$ its latitude is too large for it to hit any VNE and after it follows $N^{2\lambda+6p+l}\ L_W$ its longitude does not increase, so it remains in the west strip. Hence, after it follows $X\ RLW$, the robot is in the west strip on the x-axis.

In case (2), after the robot follows $N^{2\lambda+6p+l}$ it gets to a special point, more specifically to the uppermost vertex of the lower infinite column c_m. Immediately afterwards, it follows L_W and it reaches the destination point. The conclusion follows. $\square$

In **Part II** we see an example on how we divide all the sets of mazes F_i in two classes in such a way that our algorithm RLW depends qualitatively only on the class. This is why we treat each class in a separate case. In **Part III** the principle is the same, but we need to consider many more cases and write a different algorithm for each one of them.

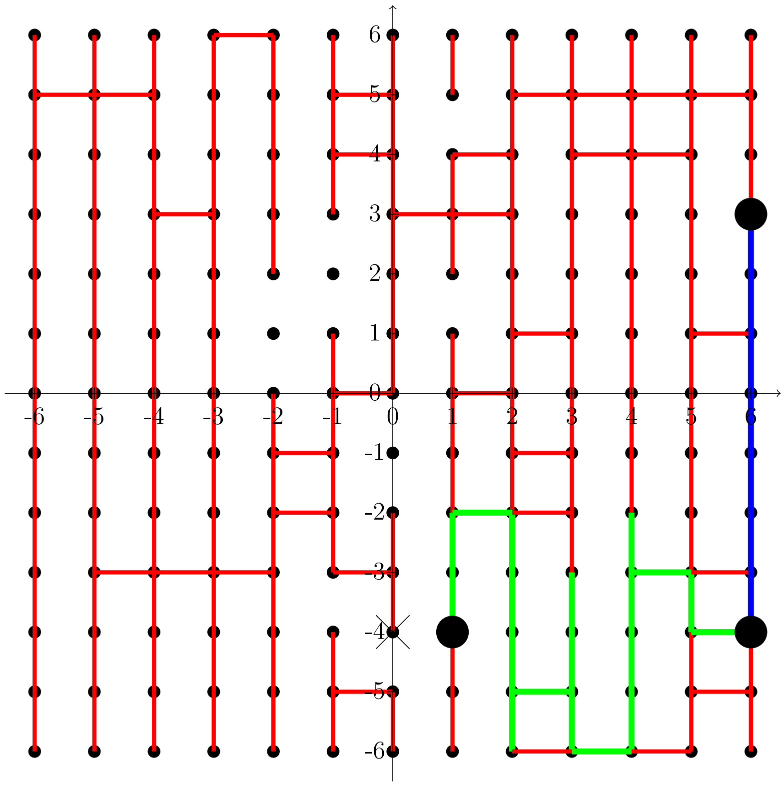

FIGURE 6. **Part II, Case (2).** There is a pass π between a lower infinite column and the east strip. We assume that there are no VEs removed other than the ones shown in the figure.

Let us examine the example shown in Figure 6. Let us suppose that the position of the robot after following X is $(6, 3)$. The pass π is the HE $\{(1, 0), (2, 0)\}$. The first segment $S^{\lambda+p}$ of RLW takes the robot at a latitude lower than that of the pass π. For our example, we may assume that after the robot follows $XS^{\lambda+p}$ it gets to the point $(6, -4)$, though this latitude should be much smaller. While the robot is in the east strip, after it executes $K = N^{2\lambda+4p}S^{2\lambda+4p}$, it returns to the starting point. By the choice of parameters, the robot enters the easternmost lower infinite column at longitude b for the first time via a locomotory move (in our case, $b = 1$). Ignoring, as we may, the action of K in the east strip, the path of the robot to the column c_b is coloured in green. Immediately after the robot enters the column c_b, it executes K which sets it latitude so small that the parameters of SMW are not large enough to make the robot visit any other configurations in the obstacle strip other than the lower infinite columns.

To finish **Part II**, we note that although we used in this part the fact there are no infinite columns in the obstacle strip, a variation of RLW can be used to position the robot in the west strip on the x-axis, even if we drop this assumption. This note is important, because it shows that **Part II** can also be generalised to improve Theorem 3.2 by dropping the consecutive column condition for the finite number of VNEs. To present this variation, assume that infinite columns are allowed in the obstacle strip, i.e. the (finitely many) VNEs need not be in consecutive columns.

We begin with the remark that **Case (1)** considered above can be treated in the exact same way with or without infinite columns in the obstacle strip, so we may assume without loss of generality that **Case (2)** holds, i.e. that every maze in the class of mazes we consider contain a pass π between a lower infinite column and the east strip. We recall that by the $\overline{()}$ transformation we apply on mazes, there are always passes between any two consecutive infinite columns in the obstacle strip. We now need to consider 2 cases: (i) there exist passes between all consecutive lower infinite columns and between consecutive lower infinite columns and infinite columns; this case can be treated similarly with **Case (2)** above; (ii) there exist two entities, one of which is a lower infinite column and the other is either a lower infinite column or an infinite column with no pass between them. In this case we define the algorithm RLW' which generalises RLW as described above,

$$RLW' := S^{\lambda+p}\ OMW(2\lambda + 2p, \lambda + p, 2\lambda + 4p)\ N^{2\lambda+4p}\ L.$$

The reason why RLW' indeed generalises RLW in this case is that after the robot follows $X\ S^{\lambda+p}\ OMW(2\lambda + 2p, \lambda + p, 2\lambda + 4p)$, it remains trapped in the lower infinite column or infinite column in the obstacle strip with largest longitude m which is not connected with the lower infinite column or infinite column at longitude $m-1$. The robot's latitude is $2\lambda + 4p$ below the lowermost

VNE at longitude between m and b. Hence from this starting position, after the robot follows $N^{2\lambda+4p}$ it gets to a special vertex (by definition) and therefore, $X\ RLW'$ takes the robot to the destination point.

Moving on from this digression, by **Case (1)** and **Case (2)**, we may assume that we are given F_i and a finite algorithm X with $\lambda = |X|$ such that after the robot follows X it is either in the west strip on the x-axis or it has visited the destination point. Without loss of generality, we assume that the robot is in the west strip on the x-axis and our aim is to build a finite algorithm F such that XF solves F_i. □

7. The General Model, Finish

Proof of Theorem 3.2, continued. In this section we define the algorithm $finish$ which makes the robot visit the destination point.

Part III. Case (1). We assume that the destination point is in an infinite column in the west strip. We define the algorithm:

$$F = MW(p, 2p)\ OME(\lambda + \mu, \lambda + \mu, p),$$

where $\mu = |MW(p, 2p)|$.

Claim 7.1. *For any such maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. After the robot follows $X\ MW(p, 2p)$ it is in the west strip, to the west of the origin or it has already visited the destination point. By the choice of parameters and by the consequence of Lemma 4.1 when applied to the particular algorithm OME, after it follows $X\ F$, the robot visits the destination point. □

To finish **Case (1)**, we note that we note that although we used the fact there are no infinite columns in the obstacle strip, a variation of F can be used in order to make the robot visit the destination point, even if we drop this assumption. To present this variation, we assume that infinite columns are allowed in the obstacle strip, i.e. the (finitely many) VNEs need not be in consecutive columns. Let us assume for now that the destination point is in an infinite column in the east strip or obstacle strip.

In this case, given any finite algorithm A we will construct a finite algorithm $U(A)$ with the following 2 properties: if the robot starts in the origin of any maze in F_i, it follows A and it gets to the west of the destination point then (1) after the robot follows $U(A)$, either its latitude strictly increases or the robot remains stuck in a finite column, or upper/lower infinite column at some longitude i with no HE connecting that column to points at longitude $i+1$; (2) as the robot follows $U(A)$, if the robot visits the infinite column which contains the destination point, then the robot visits the destination point. We will construct our algorithm U from bricks of the form

$$B(k, A) = N^{|A|+2p} S^{2|A|+4p} N^{|A|+2p} N^k E S^k$$

where k is an integer and A is a finite algorithm. Every time we insert a brick $B(k, A)$ as a subalgorithm of $X\ F$, we take A to be the entire algorithm written until that instance of $B(k, A)$. Hence, every brick depends on the length of the algorithm written up to it in $X\ F$. With this convention, from now on we shall drop the second argument from the definition of a brick and let $B(k) = B(k, A)$. We note in advance that the aim of the first segment $N^{|A|+2p} S^{2|A|+4p} N^{|A|+2p}$ of a brick $B(k, A)$ is the following: for any maze in F_i, if the robot is in the same column as the destination point after it follows a finite algorithm A, if the robot then follows $N^{|A|+2p} S^{2|A|+4p} N^{|A|+2p}$, it visits the destination point. Hence we regard the first segment $N^{|A|+2p} S^{2|A|+4p} N^{|A|+2p}$ of a brick just as an oscillation large enough to make the robot visit the destination point after it reaches the right longitude.

We note that for any k, if the robot follows $B(k)$ starting in any point of any maze, its longitude does not decrease. We define four types of steps by concatenating bricks, so each of the steps also have this property.

The first step U_1 is designed to have the following property: if the robot starts in the origin of any maze in F_i, it follows a finite algorithm A and it gets to an infinite column strictly at the west of the destination point, if the robot then follows U_1, its longitude strictly increases. For instance we can take

$$U_1(A) := B(-|A| - p)\ B(-|A| - p + 1) \ldots B(|A| + p),$$

where A is always taken to be the entire algorithm written before the occurrence of this step. With this convention, we drop the argument A and U_1 has the desired property (cf. steps 3 and 4 below). We also note that formally, the first brick in $U_1(A)$ is $B(-|A|-p, A)$, the second brick is $B(-|A|-p+1, A\ B(-|A|-p, A))$, etc.

The second step U_2 is designed to have the following property: if the robot starts in the origin of any maze in F_i, it follows a finite algorithm A and it gets to a finite column at longitude i which is connected to any point at longitude $i+1$ by a HE (i.e. no HEs emerging in the east part of the finite column), if the robot then follows U_2, its longitude strictly increases. For instance we can take

$$U_2(A) := B(0)\ B(1) \ldots B(2p),$$

where the definition of $U_2(A)$ does not depend on A, as every finite column has at most $2p$ vertices. Therefore, let $U_2 = U_2(A)$ have the desired property (cf. steps 3 and 4 below).

The third and forth step U_3 and U_4 are designed to have the following property: if the robot starts in the origin of any maze in F_i, it follows a finite algorithm A and it gets to an upper infinite or lower infinite column at longitude i which is connected to any point at longitude $i+1$ by a HE, if the robot then follows U_3 or U_4, respectively, its longitude strictly increases. For instance we can take both U_3 and U_4 to be a concatenation of $2p$ bricks in the following way

$$U_3(A) = B(-|A|-2p)\ B(-|A|-|B(-|A|-2p)|-2p+1)$$
$$B(-|A|-|B(-|A|-2p)|-|B(-|A|-|B(-|A|-2p)|-2p+1)|-2p+2)\ldots,$$
$$U_4(A) = B(|A|+2p)\ B(|A|+|B(|A|+2p)|+2p-1)$$
$$B(|A|+|B(|A|+2p)|+|B(|A|+|B(|A|+2p)|+2p-1)|+2p-2)\ldots,$$

where A is always taken to be the entire algorithm written before the occurrence of this step. With this convention, we drop the argument A and U_3, U_4 have the desired property. Indeed, let's assume that the robot is in an upper infinite column $c = (i,y),(i,y+1),\ldots$ at longitude i which is connected to any point at longitude $i+1$ by a HE, and it follows U_3. Let j be the smallest non-negative integer such that the vertices $(i,y+j)$ and $(i+1,y+j)$ are connected by a HE. From the definition of passes and the primary rectangle, we first note that $-p \leq y+j \leq p$ and $j \leq 2p$. As the robot follows the first brick $B(-|A|-2p) = N^{|A|+2p}S^{2|A|+4p}N^{|A|+2p}S^{|A|+2p}EN^{|A|+2p}$ in U_3, it oscillates in c, executing an E instruction at the vertex (i,y) in c. If $j=0$ we are done; otherwise, after the robot follows $B(-|A|-2p)$, it gets at the vertex $(i, y+|A|+2p)$. Therefore, we can track the position of the robot as it follows the second brick $B(-|A|-|B(-|A|-2p)|-2p+1)$ in U_3, and we observe that it oscillates in c, executing an E instruction at the vertex $(i,y+1)$ in c. We continue in the same way; as $-p \leq y+j \leq p$, $j \leq 2p$, we are done.

Let us make one more remark regarding these steps. If the robot follows a concatenation of bricks and it reaches a finite column, an upper infinite column or a lower infinite column at longitude i with no HE connecting it to points at longitude $i+1$, the robot remains stuck in that structure while it follows the rest of the algorithm. Let us define

$$U(A) = U_1U_2U_3U_4,$$

or formally $U(A) = U_1(A)\ U_2(A\ U_1(A))\ U_3(A\ U_1(A)\ U_2(A\ U_1(A)))\ U_4(\ldots)$. As usual, every time we use the algorithm $U(A)$ as a subalgorithm, we take A to be the entire algorithm written before the occurrence of $U(A)$, so with this convention we drop the argument of U. Therefore, it is clear that the algorithm U has the two promised properties at the beginning of the case: if the robot starts in the origin of any maze in F_i, it follows a finite algorithm A and it gets to the west of the destination point then (1) after the robot follows U, either its latitude strictly increases or the robot remains stuck in a finite column, or upper/lower infinite column at some longitude i with no HE connecting that column to points at longitude $i+1$; (2) as the robot follows U, if the robot visits the infinite column which contains the destination point, then the robot visits the destination point. Furthermore, let $V(A) = \underbrace{UU\ldots U}_{\lambda+p}$, or formally

$$V(A) = \underbrace{U(A)\ U(A\ U(A))\ U(A\ U(A)\ U(A\ U(A)))\ldots}_{\lambda+p \text{ terms}}.$$

We finally define the algorithm

$$F = V(X)\ N^{|V(X)|+l+p}\ L\ S^{2|V(X)|+l+2p}L.$$

Let us see that indeed, after the robot follows $X\ F$, it visits the destination point. We may assume without loss of generality that after the robot follows X it is in the west strip on the X axis, clearly to the west of the destination point (which is assumed to be in the east strip or in the obstacle strip). From property (1) of U, after the robot follows $X\ V(X)$, it either visits the destination point or it remains stuck in a finite, lower or upper infinite column at longitude i, at the west of the destination point, with no HE connecting it to points at longitude $i+1$. In the first two cases, it is clear that after the robot follows $X\mathrm{V}(X)\ N^{|V(X)|+l+p}\ L$ it visits the destination point. In the third case, the robot is stuck in an upper infinite column c after it follows $X\mathrm{V}(X)\ N^{|V(X)|+l+p}$. We claim that the robot returns to the same vertex in c after it follows $X\mathrm{V}(X)\ N^{|V(X)|+l+p}\ L$. Indeed, as the robot follows L its latitude is too large to meet any VNE, so by the construction of L the latitude of the robot does not change after it follows L. Moreover, the fact that c has no HE connecting it to points at longitude $i+1$ makes the robot return at longitude i after it follows L. After that, the robot follows $S^{2|V(X)|+l+2p}$ and it gets to a special point, and then it follows L which further takes it to the destination point.

Moving on from this digression, we have solved **Case (1)** in which the destination point is in an infinite column in the west strip. By the symmetry of this case and **Part II**, we similarly solve the case when the destination point is in an infinite column in the east strip. Likewise, the generalisation of **Case (1)** proved at the end of the section generalises to the case when the destination point is in an infinite column in the west strip. In fact, in the generalisation of **Case (1)** we could have only considered the case when the destination point is in an infinite column in the obstacle strip, as **Case (1)** itself works just as well even in the generalised set up, when the destination point is either in the west strip or east strip. That would not have simplified the argument, though.

Case (2). We assume that the destination point is in the obstacle strip in a finite column, upper infinite column or lower infinite column and it is connected to the west strip via a path through a (finite) sequence of finite columns. Let $R_1, R_2, \dots, R_k$ be a sequence of finite columns and R be a finite column, upper infinite column or lower infinite column such that R contains the destination point and there exists a HE between the west strip and R_1, between R_m and R_{m+1} for $1 \le m \le k$, where by convention $R_{k+1} = R$. Let $\mathbf{w} = (a+1, u)$ be the uppermost point of the finite column R_1.

We consider the following sub-cases:

2(i) We assume that there exists a row r_i that intersects the finite column R_1 and that it has a west bump. We recall that the west bumps are the easternmost HNEs with at least one vertex in the west strip on a row that intersects some finite column. Assume first that the eastern vertex $\mathbf{v}$ of that west bump is in the west strip. By inspecting the longitude of the west bump and the primary rectangle we can construct an algorithm of the form $H' := \prod_{m=1}^{h} N^{k_m} N^{-k_m} E^{\varepsilon_m}$, where $\varepsilon_m \in \{-1, 1\}$ and k_m is an integer for all $1 \leq m \leq h$, such that if the robot starts at v and follows H', it visits the destination point. Indeed if the robot is at some specified latitude in the finite column R_m and follows $N^{k_m} N^{-k_m} E$ for suitable k_m, it gets to some specified latitude in the finite column R_{m+1}. Let $H = H' E^{|H'|}$. We define the algorithm

$$F = N^i \; W^q \; SME(\lambda + q, \lambda + q, H).$$

Proposition 7.2. *For any maze in F_i, after the robot follows X F, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows X N^i, it is on the row r_i. Hence after the robot follows X N^i W^q it is on the row r_i at a longitude at most that of $\mathbf{v}$. By the choice of parameters and by Lemma 4.1, while the robot follows $SME(\lambda + q, \lambda + q, H)$ it advances eastwards in the west strip oscillating about row r_i and passing through the smallest HE with respect to the well order on $\mathbb{Z}$: $i < 1+i < -1+i < 2+i < \ldots$. Considering that $|H|_N = |H|_S$, while the robot is in the west strip, after it follows H its latitude does not change and its longitude does not decrease. It eventually arrives at the point $\mathbf{v}$ on r_i not while executing H (from the form of H and the shape of the maze which has a HNE with its eastern vertex at v). Immediately after the robot reaches $\mathbf{v}$, it follows H and it gets to the destination point (see Figure 7). $\square$

Let us examine the example shown in Figure 7. In this example we take $R_1 = (1,1), (1,2), (1,3)$, $R_2 = (2,0), (2,1), (2,2)$, $R_3 = (3,1), (3,2)$, $R_4 = (4,-3), ...(4,1)$, $R_5 = R = (3,-1), (3,-2), (3,-3)$ and $r_i = r_2$ is the row that intersects R_1 with its west bump $(-3,2), (-2,2)$ and $\mathbf{v} = (-2,2)$. In general, we do not require R to be a finite column. Let $H' = EEEENN^{-1}EEN^3N^{-3}E^{-1}N^2N^{-2}E$ and note that if the robot starts at $\mathbf{v}$

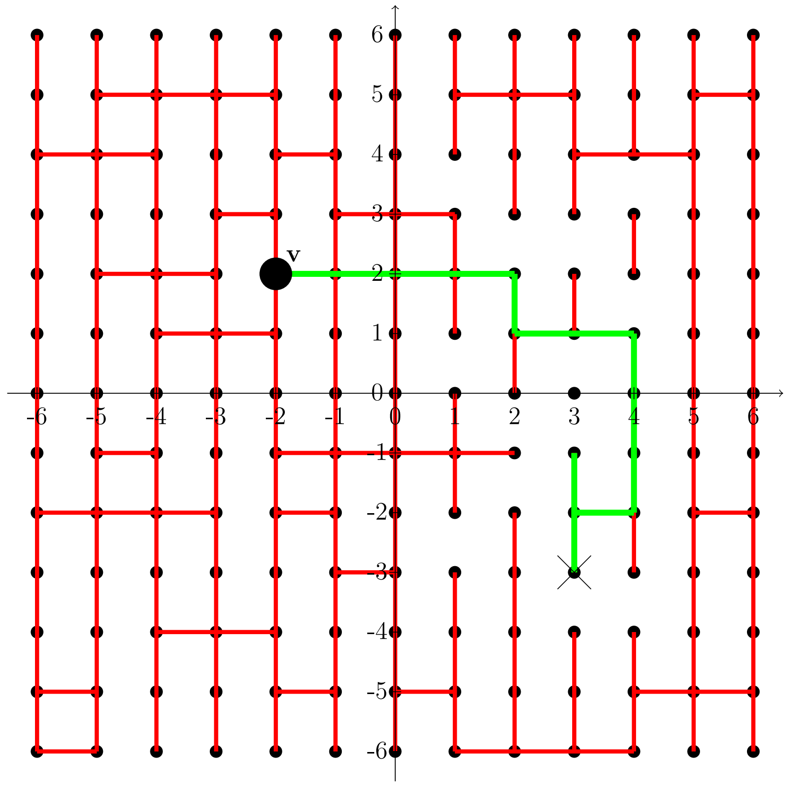

FIGURE 7. **Part III, Case (2)(i).** We assume that there exists a row r_i that intersects the finite column R_1 and that r_i has a west bump. We assume that there are no VEs removed other than the ones shown in the figure.

and follows H' it follows the green path and gets to the destination point. However, if the robot starts at the west of $\mathbf{v}$ on r_i and it follows H', its longitude is always strictly smaller than that of $\mathbf{v}$. Therefore, it does not hit any VNE and as $|H'|_N = |H'|_S$ its latitude does not change. In our case, $H = H'E^{20}$ which has the extra property that after the robot follows H in a maze with no VNEs, its longitude does not decrease - in fact, it can be proven that there always exits a certain H' that has this property itself, but we do not wish to complicate the argument.

In the case that there exists a west bump positioned at the border between the obstacle strip and the west strip on a row r_i that intersects R_1, we consider r_j to be a row on which there exists a HE between the west strip and R_1. As before, let $\mathbf{v}$ be the eastern vertex of the west bump, $\mathbf{v} \in \mathbb{R}_1$. We recall the algorithm

$$SME^{(j-i)}(a,e,H) := ((((((H)^e N^{j-i} E S^{j-i})^e E)^e NES)^e SEN)^e \ldots S^a EN^a)^e$$

introduced in Corollary 4.2. We further define the algorithm

$$F = N^i\ W^q\ SME^{(j-i)}(\lambda+q, \lambda+q, H).$$

Claim 7.3. *For any maze in F_i after the robot follows $X\ F$, it visits the destination point.*

Proof. The conclusion follows by the same reasoning as in the proof of Proposition 7.2 and by Corollary 4.2. □

2(ii) We assume that the previous case does not hold, so every row that intersects the column R_1 does not have a west bump, i.e. each such row is a path in the west strip. In addition, we assume there exists a special west pipe on some row r_j. We recall that the west pipes are easternmost configurations in the west strip formed by a HE followed by a HNE. Denote by $\mathbf{v}$ the easternmost vertex of the HE of the special west pipe. Assume without loss of generality that $j > u$, where $\mathbf{w} = (a+1, u)$ is the uppermost point of the finite column R_1.

We start by defining a new algorithm called *west_pipe_finder*:

$$WPF(a,e) := (E^e W S^a E N^a)^e,$$

with its counterpart *east_pipe_finder*. This is used directly in the final algorithm F and it will be analysed later (see Figure 8).

We then define the algorithm

$$K = S^{j-u} E^d N^{j-u} W S^{j-u} W^d N^{j-u} E,$$

where d is the difference in longitude between c_{a+1} and $\mathbf{v}$.

Claim 7.4. *For any maze in F_i, if the robot starts at $\boldsymbol{v}$ and follows K it reaches a certain known point $\boldsymbol{z}$ (given the tertiary rectangle) on the row r_u.*

PROOF. Starting at $\mathbf{v}$, after the robot follows S^{j-u} it gets on the row r_u; after it follows $S^{j-u}E^d$ it gets to $\mathbf{w}$; after it follows $S^{j-u}E^dN^{j-u}$ it remains fixed at $\mathbf{w}$; after it follows $S^{j-u}E^dN^{j-u}W$ it gets to (a, u), on the row r_u to the west of $\mathbf{w}$; finally, while it executes $S^{j-u}W^dN^{j-u}E$ starting at (a, u) it does not leave the square $\{(\pm q, \ \pm q)\}$; while it executes both the subalgorithms S^{j-u} and N^{j-u} of $S^{j-u}W^dN^{j-u}E$ it does not hit any VNE (see Figure 8). The conclusion follows. □

Remark. It is easy to check that if the robot starts from the easternmost vertex **v'** of a HE followed by a HNE on r_j with **v'** strictly at the west of $\mathbf{v}$ and it follows K, then the robot remains in the west strip while following K and after it follows K, it returns back to the starting point **v'** (see Figure 8). The algorithm K was constructed specifically to have this property, together with the one proved in the Claim above.

By inspecting the tertiary rectangle, we construct an algorithm H' of the form $H' = \prod_{i=1}^{h} N^{k_i}N^{-k_i}E^{\epsilon_i}$, where $\epsilon_i \in \{-1, 1\}$ and k_i is an integer for all $1 \leq i \leq h$, such that if the robot starts at $\mathbf{z}$ and it follows H' it visits the destination point. Let $H = H'E^{|H'|}$. We observe that if the robot starts from the easternmost vertex **v'** of a HE followed by a HNE on the row r_j in the west strip and follows H, it remains in the west strip and it returns to the same point **v'**. We finally define the algorithm:

$$F = N^u \ W^q \ N^{j-u} \ (WPF(j-u, \lambda+q) \ K \ H \ S^{j-u} \ E \ N^{j-u})^{\lambda+q}.$$

PROPOSITION 7.5. *For any maze in F_i, after the robot follows X F, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X \ N^u \ W^q \ N^{j-u}$ it gets on the row r_j, to the west of the point $\mathbf{v}$. While the robot is at the west of the point $\mathbf{v}$ on the row r_j, after each instance of $WPF(j-u, \lambda+q)$ it advances eastwards to the easternmost vertex **v'** of a HE followed by a HNE on the row r_j. While **v'** is strictly at the west of $\mathbf{v}$, the robot follows the algorithm $K\ H$ and returns back to **v'**; while the robot follows the algorithm $S^{j-u}EN^{j-u}$ it advances one unit to the east of **v'** on the row r_j. By the choice of parameters, the robot eventually arrives at **v'** $=$ **v**. Immediately afterwards, it follows $K\ H$ and it visits the destination point (see Figure 8). ☐

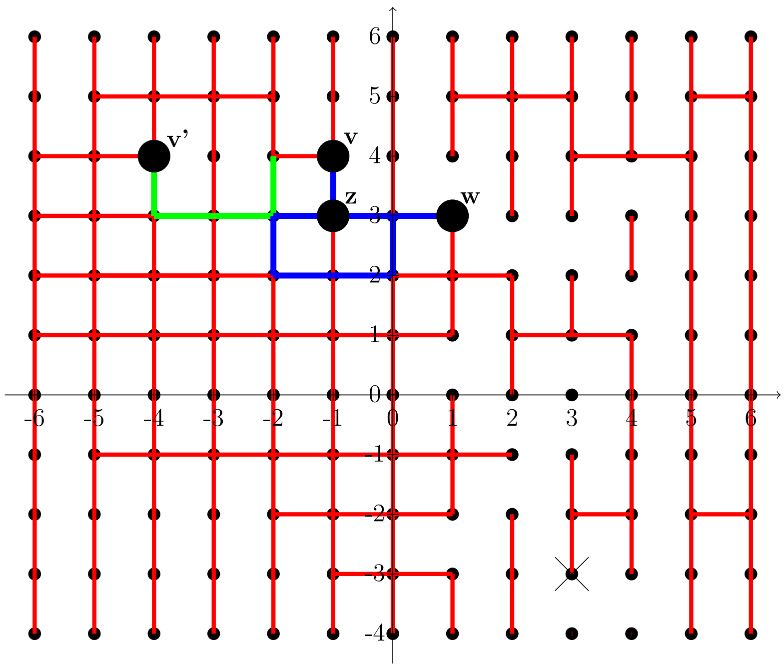

FIGURE 8. **Part III, Case (2)(ii).** Every row that intersects R_1 does not have a west bump and there exists a special west pipe on some row r_j. We assume that there are no VEs removed other than the ones shown in the figure.

Let us examine the example shown in Figure 8. In this example we have $r_j = r_4$ and so the special west pipe is $(-2, 4), (-1, 4), (0, 4)$ with $\mathbf{v} = (-1, 4)$ and $\mathbf{w} = (1, 3)$. Let us observe that after the robot follows $WPF(1, 1000)$ starting at $(-6, 4)$ it gets to **v'** $= (-4, 4)$ which is the middle vertex of the "fake" west pipe $(-5, 4), (-4, 4), (-3, 4)$. Further note that after the robot follows $WPF(1, 1000)$ starting at $(-3, 4)$ it gets to **v**. For this example we have $K = SE^2NWSW^2NE$ and after the robot follows K starting from **v** it gets to $\mathbf{z} = (-1, 3)$, which is indeed on $r_u = r_3$ (see the blue walk). In addition, note that if the robot starts from **v'** and follows K it gets back to **v'** (see the green circuit). We can take $H' = EENSENSEEN^3S^3E^{-1}N^2S^2E$ which has the required form and the property that after the robot starts from $\mathbf{z} = (-1, 3)$ and follows H' it visits the destination point. The reader may assume that the robot starts at $(-4, 0)$ and it follows $F = N^3W^2N(WPF(1, 1000)\ K\ H\ SEN)^{10}$ to see how the algorithm F solves the maze: after it follows N^3W^2N, the robot gets to $(-6, 4)$; further, after the first iteration of $WPF(1, 1000)\ K\ H\ SEN$, it gets to $(-3, 4)$ as K and H do not change the position of the robot while it is strictly at the west of $\mathbf{v}$; after the second iteration of $WPF(1, 1000)\ K\ H\ SEN$, the robot visits the destination point.

2(iii) We assume there exists a magical west row r_j. We recall that a magical west row is a row which is a path when restricted to the west strip, and it contains a HNE; its west cutoff is its westernmost HNE. Denote by $\mathbf{v}$ the westernmost vertex of the west cutoff of r_j. Then, by inspecting the tertiary rectangle, we can construct an algorithm K such that if the robot starts from $\mathbf{v}$ and follows K it gets to the destination point. We define the algorithm

$$F = N^j E^{\lambda+q} K.$$

CLAIM 7.6. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^j$, it gets on the row r_j. Therefore, after it follows $X\ N^j\ E^{\lambda+q}$ the robot gets to the point $\mathbf{v}$. Hence, after the robot follows $X\ F$ it gets to the destination point. $\square$

2(iv) We assume that every row that intersects the finite column R_1 does not have a west bump and there exists a special almost empty west row r_j. We recall that a special almost empty west row is a row that in the west strip is the complement of an infinite path followed by a non-empty finite path; its west cutoff is its easternmost HNE in the west strip. We recall that $\mathbf{w} = (a + 1, u)$ is the uppermost point of R_1 and let $\mathbf{v}$ be the easternmost vertex of the west cutoff of r_j. Then, by inspecting the tertiary rectangle, we can construct an algorithm K such that if the robot starts from $\mathbf{v}$ and follows K it gets to the destination point. We define the algorithm

$$F = N^j\ W^{\lambda+q}\ (S^{j-u}EN^{j-u}W)^{\lambda+q}\ K.$$

CLAIM 7.7. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^j\ W^{\lambda+q}$ it gets on the row r_j to the west of the point $\mathbf{v}$. While the robot follows one instance of $S^{j-u}EN^{j-u}W$ it returns on the row r_j and advances one unit eastwards if it is at the westernmost vertex of a HNE; it returns to the same point if it is at the westernmost vertex of a HE. By the choice of exponent, after the robot follows $N^j\ W^{\lambda+q}\ (S^{j-u}EN^{j-u}W)^{\lambda+q}$ it remains stuck at the point $\mathbf{v}$. Immediately afterwards, it follows K and it gets to the destination point (see Figure 9). $\square$

Let us examine the example shown in Figure 9. In this example we have $R_1 = (1,1),(1,2),(1,3)$, so r_1, r_2, r_3 are paths in the west strip, moreover $j = -1$, so r_{-1} is the special almost empty west row. Its west cutoff is the HNE $\{(-4,-1), \mathbf{v} = (-3,-1)\}$. We construct an algorithm K by inspecting the tertiary rectangle such that if the robot starts from $\mathbf{v}$ and follows K, it gets to the destination point. For example we may take $K = N^3E^5SE^2S^3WS$. We may assume that the robot starts at $(-5,0)$ and it follows $F = N^{-1}\ W^{100}$ $(S^{-4}EN^{-4}W)^{100}\ K$. After the robot follows $N^{-1}\ W^{100}$, it gets to $(-5,-1)$ on the row $r_j = r_{-1}$ at a longitude not greater than that of $\mathbf{v}$. Let us see what is the position of the robot after it follows one instance of $(S^{-4}EN^{-4}W)$, starting from r_{-1}: while it starts strictly at the west of $\mathbf{v}$, its longitude increases by 1 (see the blue path); if it starts at $\mathbf{v}$, it comes back to $\mathbf{v}$ (see the green path). The exponent of $(S^{-4}EN^{-4}W)$ is large enough for the robot to reach $\mathbf{v}$ after

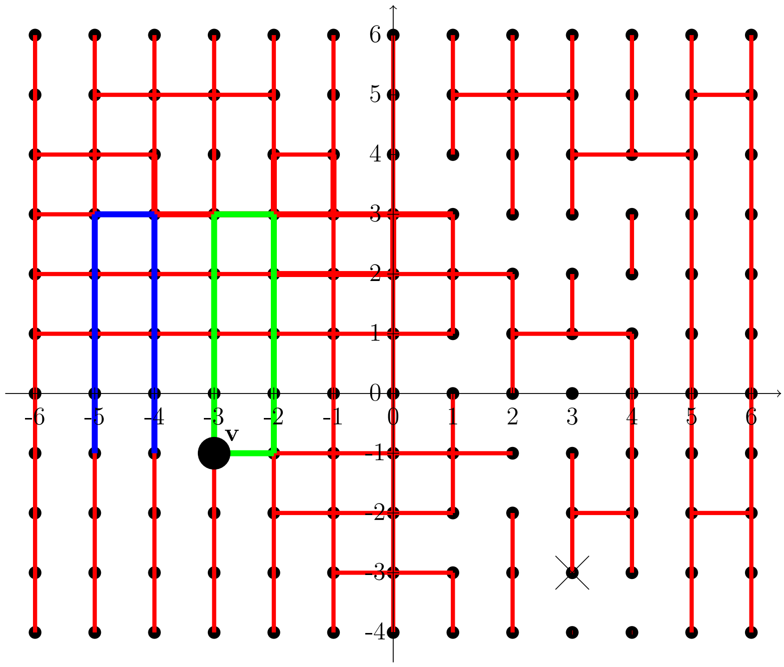

FIGURE 9. **Part III, Case (2)(iv).** Every row that intersects R_1 does not have a west bump, i.e. all such rows are paths in the west strip and there exists a special almost empty west row r_j. We assume that there are no VEs removed other than the ones shown in the figure.

it follows $(S^{-4}EN^{-4}W)^{100}$. After that, the robot follows K and it visits the destination point.

2(v) We assume that every row that intersects the column R_1 does not have a west bump. In addition we assume that there exists a special empty west row r_{w_3}. We recall that an empty west row is a row that in the west strip is empty and the special empty west row is the empty west row of smallest latitude greater than $-3p$ with respect to the standard well order on $\mathbb{Z}$. We recall that $\mathbf{w} = (a+1, u)$ is the uppermost point of the finite column R_1 and let $\mathbf{v}$ be the easternmost vertex in the west strip on the row r_{w_3}. We may assume without loss of generality that $w_3 > u$.

By inspecting the primary rectangle, we construct an algorithm H' of the form $H' = \prod_{m=1}^{h} N^{k_m} N^{-k_m} E^{\epsilon_m}$, where $\epsilon_m \in \{-1, 1\}$ and k_m is an integer with $|k_m| \leq 2p$ for all $1 \leq m \leq h$, such that if the robot starts at $\mathbf{w}$ and it follows H', it visits the destination point (see H' in Figure 8). Let $H = H' W^{|H'|}$. We note that if the robot is in the origin in a maze with no VNEs and it follows H it returns to the x-axis and its latitude does not increase. We further note that if the robot starts from $\mathbf{v}$ and it follows H, it oscillates about latitude w_3 without hitting any VNE and at the end it returns back to the starting point $\mathbf{v}$.

We define the algorithm

$$F = N^{w_3}(S^{w_3-u}EN^{w_3-u}H)^{\lambda+q}.$$

CLAIM 7.8. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^{w_3}$ it gets on the row r_{w_3} at the west of the point $\mathbf{v}$. While the robot follows each instance of $S^{w_3-u}EN^{w_3-u}H$ in the west strip, it advances eastwards one unit making an oscillation about the row r_{w_3}. By the choice of exponent, after a certain instance of $S^{w_3-u}EN^{w_3-u}H$, the robot eventually gets to the point $\mathbf{v}$. Immediately afterwards, it follows another instance of $S^{w_3-u}EN^{w_3-u}H$ and it gets to the destination point. Indeed, if the robot starts at the point $\mathbf{v}$ and it follows $S^{w_3-u}EN^{w_3-u}$, it gets to the point $\mathbf{w}$. If the robot starts at $\mathbf{w}$ and it follows H, it gets to the destination point. The conclusion follows. $\square$

2(vi) This is the final case, where we may assume all of the following: every row that intersects the column R_1 does not have a west bump; there does not exist a west pipe; there does not exist a magical west row; there does not exist a special almost empty west row; there does not exist a special empty west row. Then every row at latitude greater than $-3p$ with respect to the well order on $\mathbb{Z}$ is a path in the west strip and indeed a path in the maze; every row that intersects the finite column R_1 is a path in the west strip and indeed a path in the maze; each row at latitude at most $3p$ with respect to the standard well order on $\mathbb{Z}$ is known to be either a path or the complement of a path in the west strip. We recall that $\mathbf{w} = (a+1, u)$ is the uppermost point of the finite column R_1.

By inspecting the primary rectangle we can construct an algorithm H' of the form $H' = \prod_{m=1}^{h} N^{k_m} N^{-k_m} E^{\epsilon_m}$, where $\epsilon_m \in \{-1, 1\}$ and k_m is an integer with $|k_m| \leq 2p$ for all $1 \leq m \leq h$, such that if the robot starts at $\mathbf{w}$ and follows H' it visits the destination point (see H' in Figure 8). Let $H = H'E^r$, where r is an integer such that if the robot follows H on a maze without meeting any VNE and HNE then it returns back to its starting point. We construct the algorithm

$$F = N^u (EN^{6p} H S^{6p})^{\lambda + p}.$$

CLAIM 7.9. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^u$ it gets in the west strip on the row r_u. While the robot executes one instance of $EN^{6p}HS^{6p}$ it advances one unit eastwards in the west strip on the row r_u without meeting any VNE or HNE. Indeed, every row at latitude greater than $3p$ is a path in the maze. The robot eventually eventually gets at $\mathbf{w}$. Immediately afterwards, it follows N^{6p}, remaining at $\mathbf{w}$ and then H, hence it gets to the destination point. □

This finally solves **Case (2)** in which the destination point was connected with the west strip by a finite number of finite columns. It is immediate to see that the presence of infinite columns in the obstacle strip does not affect any of the arguments made in this case.

Case (3). We assume that the destination point is in the obstacle strip and there exists some parameter $h_{(i,i+1)} < \infty$. We recall that this is equivalent to the existence of a pair of consecutive upper infinite columns (or a consecutive upper infinite column and an infinite column at the border of the obstacle strip and either the east or west strip) which are not connected by HEs at arbitrarily high latitudes. By symmetry, treating this case also solves the case in which there exists some parameter $l_{(i,i+1)} < \infty$.

3(i) We assume $h_{(a,a+1)} < \infty$. We recall that the pair of columns (c_a, c_{a+1}) is at the border between the west strip and the obstacle strip and we also recall that the pair of columns (c_b, c_{b+1}) is at the border between the obstacle strip and the east strip. We assume without loss of generality that there exists a HE between the west strip and a finite column or a lower infinite column (otherwise we are done by **Part I**). Let R be a finite column or a lower infinite column on the column c_{a+1} such that there exists a HE between the west strip and R on some row r_c. Let $\mathbf{w}$ be the uppermost vertex of R. Let $j = h_{(a,a+1)} + l$ and $\mathbf{v} = (a, j)$ be the easternmost point on the row r_j in the west strip. We recall the generic algorithm

$$SME^{(j-c)}(a, e, L) = ((((((L)^e S^{j-c} EN^{j-c})^e E)^e NES)^e SEN)^e \ldots S^a EN^a)^e.$$

We define the algorithm

$$F = N^j SME^{(j-c)}(\lambda + j + q, \lambda + j + q, L).$$

Claim 7.10. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows $X\ N^j$ it is in the west strip on the row r_j. By the choice of parameters and by Corollary 4.2, while the robot follows $SME^{(j-c)}(\lambda + j + q, \lambda + j + q, L)$ it advances eastwards in the west strip oscillating about row r_j. After the robot starts from some point on the row r_j in the west strip and follows L its longitude does not decrease and it remains in the west strip. It eventually gets to the point $\mathbf{v}$. After the robot starts from $\mathbf{v}$ and follows $S^{j-c}EN^{j-c}$ it gets to the point $\mathbf{w}$. Immediately afterwards, it follows L and gets to the destination point. $\square$

3(ii) Consider the pair of consecutive columns (c_i, c_{i+1}) which is not at the border between the west strip and the obstacle strip. Assume there are not arbitrarily high HEs between the columns c_i and c_{i+1}, i.e. $h_{(i,i+1)} < \infty$. Assume further that there exists a pass on some row r_c between the west strip and an upper infinite column R (see the case **3(i)**). We define $K = S^{\lambda+2q+|h_{(i,i+1)}|+1}N^{\lambda+2q+|h_{(i,i+1)}|+1}$. We define the algorithm

$$F = N^c\ SME(\lambda+q, \lambda+q, K)\ S^{\lambda+2q+|h_{(i,i+1)}|+1}\ L.$$

CLAIM 7.11. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^c$ it gets in the west strip on the row r_c. While the robot follows $SME(\lambda+q, \lambda+q, K)$ it advances eastwards in the west strip oscillating about the row r_c. It eventually enters the upper infinite columns R. Immediately afterwards it executes K and gets to latitude at least $\lambda + q + |h_{(i,i+1)}| + 1$ in R. While the robot is in the obstacle strip and follows SME it advances eastwards through upper infinite columns at latitudes greater than $h_{(i,i+1)} + 1$. Hence the robot remains stuck in some column c_j with $j \leq i$ at latitude $\lambda+2q+|h_{(i,i+1)}|+1$ above the highest VNE in the columns c_r with $a \leq r \leq j$. After that the robot follows $S^{\lambda+2q+|h_{(i,i+1)}|+1}$ and gets to a special point. Therefore, after the robot follows $X\ F$ it gets to the destination point. □

This finally solves **Case (3)** in which the destination point is in the obstacle strip in a finite or infinite column and there exists some parameter $h_{(i,i+1)} < \infty$. Moreover, the case in which there exists some parameter $l_{(i,i+1)} < \infty$ is tackled similarly by symmetry. Finally, it is immediate to see that the presence of infinite columns in the obstacle strip does not affect any of the arguments made in this case.

Case (4). This is the final case, in which we may assume that **Case (3)** does not hold and the destination point is in the obstacle strip in a finite column, an upper infinite column or lower infinite column and it is connected to the west strip by a (finite, possibly empty) sequence of finite columns followed by a (finite, non-empty) sequence of upper infinite columns, in this order starting from the destination point and advancing towards the west strip. Indeed, we may assume that the west strip is accessible by **Part II**. The case in which is the destination point is not in the obstacle strip is tackled in **Case (1)**. Furthermore, if we assume that the destination point is in the obstacle strip, it may either be reachable from the west strip through a finite sequence of finite columns tackled in **Case (2)** or otherwise it must be reachable from the west strip through a finite sequence of upper/lower infinite and finite columns which contains at least one upper or lower infinite column. Choose any such finite sequence of columns which leads to the destination point starting from the west strip and call the last upper or lower infinite column in the sequence c; this may either be the last element of the sequence or it might be followed by a finite sequence of finite columns. By **Case (3)** we may assume that there are horizontal edges between consecutive upper infinite columns and between consecutive lower infinite columns at latitudes arbitrarily high and low, respectively. Hence, assuming without loss of generality as we may that c is an upper infinite column, c can be reached from the west strip through a finite sequence of upper infinite columns. Therefore, the last case that we tackle is the one in which we assume that the destination point is connected to the west strip by a (finite, possibly empty) sequence of finite columns followed by a (finite, non-empty) sequence of upper infinite columns, in this order starting from the destination point and advancing towards the west strip.

The condition that **Case (3)** does not hold means that in this case we assume that all the parameters $h_{(i,i+1)}$ and $l_{(i,i+1)}$ are all infinity for $a \leq i \leq b$; in particular, this implies that there exists a west ascending chain. We recall that the pair of columns (c_a, c_{a+1}) are at the border between the west strip and the obstacle strip; we also recall that the pair of columns (c_b, c_{b+1}) are at the border between the obstacle strip and the east strip. We further recall that a west ascending chain is a finite sequence of HEs: $HE_a, HE_{a+1}, \ldots, HE_b$ such that HE_a is the upper west pass (i.e. the lowermost HE between the west strip and the upper infinite column on c_{a+1} above all passes in the obstacle strip) and HE_m is the lowermost HE between the pair of columns (c_m, c_{m+1}) at latitude at least that of HE_{m-1} for $m = a+1, \dots, b$. In this case, we take $R_{a+1}, R_{a+2}, \dots, R_n$ to be a finite non-empty sequence of upper infinite columns and $R_{n+1}, \dots, R_k$ to be a finite possibly empty sequence of finite columns and finally we take R to be a finite, upper infinite or lower infinite column such that R contains the destination point and there exists a HE between the west strip and R_{a+1}, between R_m and R_{m+1} for $a+1 \leq m \leq k-1$ and between R_k and R. By the discussion at the beginning of the case, we may assume that such a sequence has the extra property that R_m is on the column c_m for $a+1 \leq m \leq n$. Moreover, if R_{n+1} exists we may assume that $R_{n+1} \in c_{n+1}$; indeed, $R_{n+1} \in c_{n+1}$ or $R_{n+1} \in c_{n-1}$ and if $R_{n+1} \in c_{n-1}$ then we can use the symmetry of the argument in **Part II** to assume that the robot is in the east strip on the x-axis. From that perspective, we can use the arguments from the case that we are treating with $R_{n+1} \in c_{n+1}$. Obviously, if R_{n+1} does not exist, by the same argument we may assume that R is in c_{n+1}. Finally, say that the row r_i contains the upper west pass and note that the upper west pass is above all passes in the obstacle strip and therefore, as **Case (3)** does not hold, it is above all special vertices.

4(i) We assume there exists a magical west row r_j. We recall that a magical west row is a row which is a path when restricted to the west strip, and it contains a HNE; its west cutoff is its westernmost HNE. We see in the end that our argument also solves the case when there exists a magical east row. Denote by $\mathbf{v}$ the westernmost vertex of the west cutoff of r_j. Then, by inspecting the tertiary rectangle, we can construct an algorithm K such that if the robot starts from $\mathbf{v}$ and follows K it gets to the destination point. We construct the algorithm

$$F = N^j E^{\lambda+q} K.$$

Claim 7.12. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows $X\ N^j$, it gets on the row r_j. Therefore, after it follows $X\ N^j\ E^{\lambda+q}$ the robot gets to the point $\mathbf{v}$. Hence, after the robot follows $X\ F$ it gets to the destination point. $\square$

Clearly, in this case we may easily drop the general assumption that $R_{n+1} \in c_{n+1}$. Therefore, this argument also solves the case when there exists a magical east row.

4(ii) We assume there exists a special almost empty west row r_j and we call $\mathbf{v}$ the easternmost vertex of the west cutoff of r_j. We recall that an almost empty west row is a row that in the west strip is the complement of an infinite path followed by a non-empty finite path; its west cutoff is its easternmost HNE in the west strip. Then, by inspecting the tertiary rectangle, we can construct an algorithm K such that if the robot starts from $\mathbf{v}$ and follows K it gets to the destination point (see Figure 10).

We then define the algorithm *auxiliary_move_east*,

$$AME(a,e) = ((NESW)(SENW)(N^2ES^2W)(S^2EN^2W)\dots(S^aEN^aW))^e.$$

We finally define the algorithm

$$F := N^j\ W^q\ AME(\lambda+q, \lambda+q)\ K.$$

Claim 7.13. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows $X\ N^j\ W^q$ it gets on the row r_j at a longitude at most that of $\mathbf{v}$. By the choice of parameters, while the robot follows $AME(\lambda+q, \lambda+q)$ it advances eastwards in the west strip oscillating about the row r_j and it remains stuck at the point $\mathbf{v}$. Hence, after the robot follows $X\ F$, it reaches the destination point (see Figure 10). $\square$

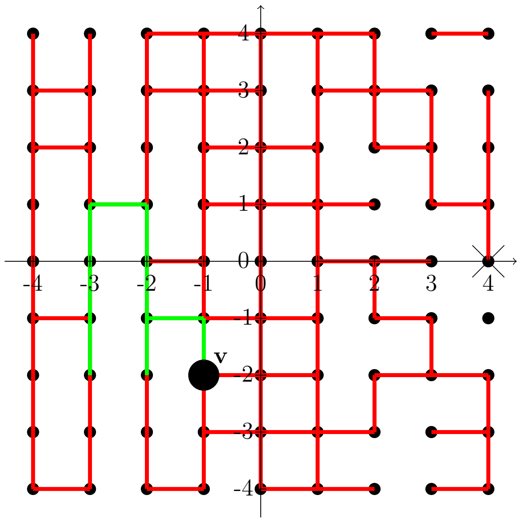

FIGURE 10. **Part III, Case (4)(ii).** There exists a special almost empty west row r_j and let $\mathbf{v}$ be the easternmost vertex of the west cutoff of r_j. We assume that there are no VEs removed other than the ones shown in the figure.

Let us examine the example shown in Figure 10. In this example, let $r_j = r_{-2}$ and so $\mathbf{v} = (-1, -2)$. Let us see how the robot gets to $\mathbf{v}$ after it follows $AME(5, 5)$ starting from $(-3, -2)$. As long as the robot is at the west of $\mathbf{v}$, the W instructions do not decrease the longitude of the robot, as there are no HEs on r_j at the west of $\mathbf{v}$. Therefore, the robot takes the green path to $\mathbf{v}$. Once the robot gets at $\mathbf{v}$, every subsequent subalgorithm of the form N^iES^iW takes it back to $\mathbf{v}$: after N^iES^i, the robot is either at $\mathbf{v} = (-1, -2)$ or $(0, -2)$; immediately afterwards, the robot follows W and the presence of a HE between $(-1, -2)$ and $(0, -2)$ guarantees that the robot returns back to $\mathbf{v}$. Immediately after the robot follows AME and it gets to $\mathbf{v}$, it follows K and it visits the destination point. For our example we can take $K = N^3E^2N^2E^2S^2ES$.

4(iii) We assume there exist the special upper west paired HNEs. We recall the following definitions: let $HE_a, \ldots, HE_b$ be the west ascending chain with HE_a being the upper west pass say on some row r_i and also say that HE_b is on some row r_t. Then $c_{uw} = t + p$ is the upper west constant, where p is the parameter of the primary rectangle. The upper west paired HNEs are any pair of HNEs with the same longitude, in the west strip, such that the upper HNE is at latitude i, on the same row as the upper west pass, and the lower HNE is at latitude at most $i - c_{uw}$. For the special upper west paired HNEs, we choose the upper west paired HNEs with the uppermost easternmost lower HNE. In this subcase, we assume that there exist the special upper west paired HNEs, with the upper HNE on the row r_i and the lower HNE on the row r_j, $j \leq i - c_{uw}$.

Let the point $\mathbf{v}$ be the easternmost vertex of the upper HNE of the pair and let the point $\mathbf{t}$ be the easternmost vertex of the upper west pass. We pick any HE between the upper infinite column R_n and the finite column R_{n+1} at latitude say ν. In the case that R_{n+1} does not exist, we pick the lowermost HE between the upper infinite column R_n and R at latitude say ν. Let the point $\mathbf{w}$ be the vertex in the infinite column R_n at latitude $\nu + i - j$. Then the eastern vertex of HE_{n-1} which has a latitude of at most t by definition is in the column c_n below $\mathbf{w}$; indeed, $\nu + i - j \geq \nu + t + p$ and $\nu + p \geq 0$. Finally, let the point $\mathbf{z}$ be the uppermost vertex of the finite column R_{n+1} if R_{n+1} exists. In the following argument, we assume that R_{n+1} exists and it will be clear how this also naturally treats the case when R_n is connected to R which contains the destination point. For an illustration of all these definitions in a concrete example, see Figure 11.

In what follows, we will construct 5 algorithms $K_1, \ldots, K_5$, by inspecting the tertiary rectangle.

We start by constructing a finite algorithm K_1 of the form
$K_1 = \prod_{m=1}^{h_1} S^{\epsilon_m} E N^{\epsilon_m}$, where $\epsilon_m \in \{0, i-j\}$ for all $1 \leq m \leq h_1$, such that after the robot follows K_1 starting from the point $\mathbf{v}$ it gets to the point $\mathbf{t}$. We make use of the fact that in the west strip at the east of the special upper west paired HNEs at each given longitude at least one of the rows r_i and r_j contains a HE. Clearly, $\epsilon_{h_1} = 0$.

We construct a finite algorithm K_2 of the form
$K_2 = (\prod_{m=a+1}^{n-1} S^{k_m} N^{k_m} E) S^{k_n} N^{k_n}$, where k_m is a positive integer for all $a+1 \leq m \leq n$, such that if the robot starts from the point $\mathbf{t}$ and it follows K_2, it

gets to the point $\mathbf{w}$. More specifically, if the robot is in the upper infinite column R_m in the column c_m at the easternmost end of HE_{m-1} and it follows $S^{k_m}N^{k_m}E$ it gets in the upper infinite column R_{m+1} in the column c_{m+1} at the easternmost end of HE_m, for $a+1 \leq m \leq n-1$; if the robot is in the upper infinite column R_n in the column c_n at the easternmost end of HE_{n-1} and it follows $S^{k_n}N^{k_n}$ it gets to the point $\mathbf{w}$.

We construct an algorithm $K_3 = S^{i-j}EN^{i-j}$, such that if the robot starts from $\mathbf{w}$ and it follows K_3, it gets to the point $\mathbf{z}$.

We construct an algorithm K_4 of the form $K_4 = (\prod_{m=n+1}^{k} N^{k_m}S^{k_m}E^{\epsilon_m})$ $N^{k_{k+1}}\ N^{-k_{k+1}}$, where $\epsilon_m \in \{-1,1\}$ and k_m is an integer for all $n+1 \leq m \leq k+1$, such that if the robot starts from the point $\mathbf{z}$ and it follows K_4, it visits the destination point. More specifically, if the robot is at some specified latitude in the finite column R_m and it follows $N^{k_m}N^{-k_m}E^{\epsilon_m}$, it gets to some specified latitude in the finite column R_{m+1} for $n+1 \leq m \leq k$, where by convention we write R_{k+1} for R. If the robot is at some specified latitude inside R and follows $N^{k_{k+1}}N^{-k_{k+1}}$ it visits the destination point.

We define the algorithm $K_5 = E^{|K_4|}$.

We define the algorithm $K = K_1K_2K_3K_4K_5$. Note that if the robot is on the row r_i strictly at the west of the point $\mathbf{v}$ and it follows K it returns on the row r_i strictly at the west of $\mathbf{v}$. Indeed, by examining $K_1, \ldots, K_5$ one by one, we conclude that if the robot starts strictly at the west of $\mathbf{v}$, while executing K it can only change its longitude at latitudes i or j. Thus, the existence of the special west paired HNEs prevents the robot from reaching a longitude at least that of $\mathbf{v}$. In particular, if the robot is on the row r_i strictly at the west of the point $\mathbf{v}$ and it follows K it does not meet any VNE, so it is easy to see that it returns back to the row r_i. Finally, the only W instructions in K could appear as part of K_4, which is followed by $K_5 = E^{|K_4|}$ in K; therefore if the robot is on the row r_i strictly at the west of the point $\mathbf{v}$ and it follows K its longitude does not decrease. If the robot starts at the point $\mathbf{v}$ and it follows

K, then it visits the destination point; this follows directly from the definitions of $K_1, \ldots, K_5$ (see Figure 11).

Finally, we construct the algorithm

$$F = N^i \; MW(i-j,q) \; SME(\mu+\lambda+2q, \mu+\lambda+2q, K),$$

where $\mu = |MW(i-j,q)|$.

PROPOSITION 7.14. *For any maze in F_i, after the robot follows X F, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows the algorithm X N^i $MW(i-j,q)$ it gets on the row r_i at a longitude at most that of the point $\mathbf{v}$. By the choice of parameters and by Lemma 4.1, while the robot follows $SME(\mu+\lambda+2q, \mu+\lambda+2q, K)$, it advances eastwards in the west strip oscillating about the row r_i. After defining K, we checked that it satisfies the conditions required in order to apply Lemma 4.1. Therefore, by Lemma 4.1, the robot gets for the first time to the point $\mathbf{v}$ not while executing K, but while executing a locomotory move. Immediately afterwards, it follows K and it gets to the destination point. The conclusion follows. □

Let us examine the example shown in Figure 11. In this example we have $a = 1$ and $b = 5$. The upper west pass between c_1 and c_2 is $HE_1 = (1,1),(2,1)$, above all the passes in the obstacle strip; the west ascending chain is coloured green. The chosen path from the west strip to the destination point goes through $R_2 = \{(2,-3),(2,-2),\ldots\}$, then R_3, $R_n = R_4$, $R_5 = \{(5,-1),(5,-2),(5,-3),(5,-4)\}$, $R_6 = \{(4,-4),(4,-5)\}$, $R = R_7 = \{(3,-4),(3,-5)\}$. The point $\mathbf{z} = (5,-1)$ is the uppermost vertex of R_5. For the purpose of this example, let us assume $c_{uw} = 6$, although this should be larger. To find the upper west paired HNEs, the set of HNEs on r_1 is the set of all possible candidates for the upper HNE in the pair. To find the second HNE in the pair, we look on $r_{1-6} = r_{-5}$ (i.e. at latitude $i - c_{uw}$) to find a matching HNE at the same longitude with one on r_1 and we choose the easternmost

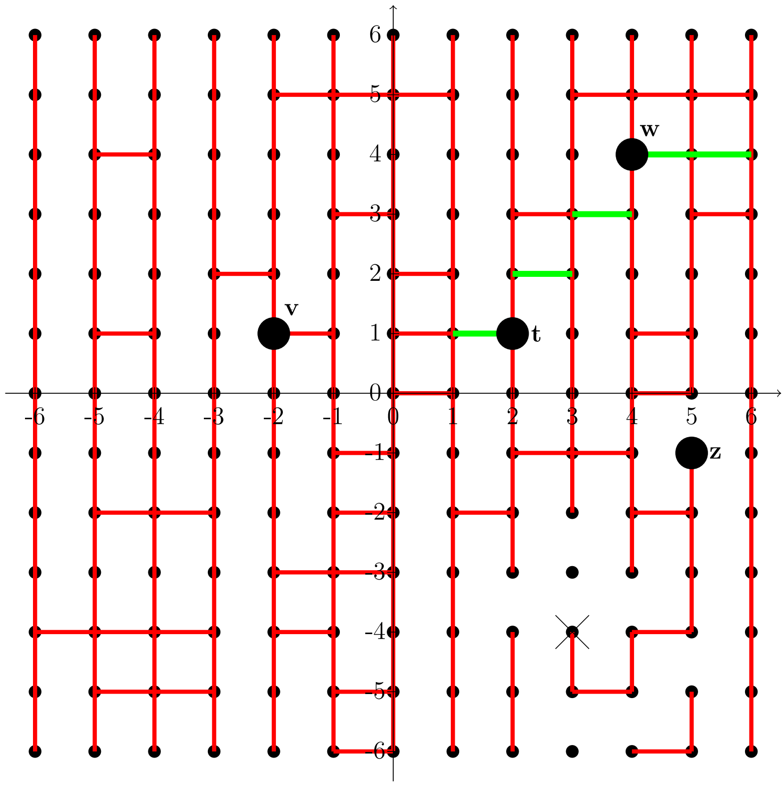

FIGURE 11. **Part III, Case (4)(iii).** We assume there exist the special upper west paired HNEs. We assume that there are no VEs removed other than the ones shown in the figure.

one. If none such HNE exists, we repeat the same process on r_{-6}, then on r_{-7} and so on. In this example, we find the upper west paired HNEs to be $(-3,1),(-2,1)$ and $(-3,-5),(-2,-5)$. Therefore $\mathbf{v}=(-2,1)$ and $\mathbf{t}=(2,1)$. The only HE between R_4 and R_5 is $(4,-2),(5,-2)$ at latitude $\nu=-2$, so $\mathbf{w}=(4,4)$. Then, if the robot follows $K_1=ES^6EN^6EE$ starting from $\mathbf{v}$ it gets to $\mathbf{t}$; if the robot follows $K_2=S^5N^5ES^5N^5ES^7N^7$ starting from $\mathbf{t}$ it gets to $\mathbf{w}$, passing through the green edges from R_2 to R_3 and from R_3 to R_4; if the robot follows $K_3=S^6EN^6$ starting from $\mathbf{w}$ it gets to $\mathbf{z}$; if the robot follows $K_4=N^3S^3WNSWNSE$ starting from $\mathbf{z}$ it gets to the destination point; $K_5=E^{13}$. Finally, we remark that if the robot follows $K=K_1K_2K_3K_4K_5$ starting on any point of r_1 strictly at the west of $\mathbf{v}$, then it returns on r_1 strictly at the west of $\mathbf{v}$.

4(iv) We assume that there do not exist some special upper west paired HNEs and there exists an upper west pipe on the row r_i which contains the upper west pass. The upper west pipe is the west pipe (the easternmost configuration of a HE followed by a HNE) on the row r_i. Let the point $\mathbf{v}$ in the west strip be the easternmost vertex of the HNE of the upper west pipe. Let the point $\mathbf{t}$ be the easternmost vertex of the upper west pass. Consider the finite sequence of HEs in the west ascending chain $HE_a, HE_{a+1}, \ldots, HE_b$. Let the point $\mathbf{w}$ in the upper infinite column R_n in c_n be the westernmost vertex of HE_n. Let $HE_{special}$ be a HE between the upper infinite column R_n and the finite column R_{n+1}. As in case **4(iii)**, if R_{n+1} does not exist, let $HE_{special}$ be the lowermost HE between R_n and R. Let the constant d be the difference in latitude between HE_n and $HE_{special}$, with $d \geq 0$ from the definition of the upper west pass which is above all passes and special vertices in the obstacle strip. Let $\mathbf{z}$ be the uppermost point in the finite column R_{n+1} if R_{n+1} exists. In the following argument, we assume that R_{n+1} exists and it will be clear how this also naturally treats the case when R_n is connected to R which contains the destination point.

In what follows, we will construct 5 algorithms $K_1, \ldots, K_5$, by inspecting the tertiary rectangle.

We start by constructing the algorithm $K_1 = (WS^{c_{uw}}EN^{c_{uw}})^{h_1}E^{h_2}$, where h_1 and h_2 are positive integers, such that if the robot starts from the point $\mathbf{v}$ and follows K_1 it gets to the point $\mathbf{t}$. We make use of the fact that in the west strip at each given longitude at least one of the rows r_i and r_j, $j = i - c_{uw}$ contains a HE. We also make use of the fact that in the west strip the section of the row r_i at the east of the upper west pipe is the complement of a path, followed by a path (which is nonempty from the existence of the upper west pass). However, we remark that if the robot starts on r_i strictly at the west of $\mathbf{v}$ and it follows K_1, it always remains strictly at the west of $\mathbf{v}$, due to the HNE of the west pipe and the fact that there are no VNEs at the west of $\mathbf{v}$.

We construct the algorithm $K_2 = (\prod_{m=a+1}^{n} S^{k_m}N^{k_m}E)W$, where k_m is a positive integer for all $a+1 \leq m \leq n$, such that if the robot starts from the point $\mathbf{t}$ and follows K_2 it gets to the point $\mathbf{w}$. More specifically, if the robot is in the upper infinite column R_m in the column c_m at the easternmost point of HE_{m-1} and follows $S^{k_m}N^{k_m}E$, it gets in the upper infinite column R_{m+1} in the column c_{m+1} at the easternmost end of the HE_m, for $a+1 \leq m \leq n$. After the robot follows the last instruction in the product, $S^{k_m}N^{k_m}E$, it gets to c_{n+1} at the easternmost point of HE_n and so after it follows the last instruction in K_2, that is W, the robot gets to the point $\mathbf{w}$.

We define the algorithm $K_3 = S^d E N^d$, such that if the robot starts from $\mathbf{w}$ and it follows K_3, it gets to the point $\mathbf{z}$. However, we remark that if the robot starts on r_i strictly at the west of $\mathbf{v}$ and it follows K_2 K_3, it always remains strictly at the west of $\mathbf{v}$. Indeed, while the robot follows K_2 starting strictly at the west of $\mathbf{v}$, the HNE of the west pipe prevents it from visiting longitudes greater than that of $\mathbf{v}$. Hence, the robot could only potentially get to a large longitude by reaching $\mathbf{v}$ after it follows K_3; however, this is impossible as the last instruction in K_2 is W.

We construct the algorithm $K_4 = (\prod_{m=n+1}^{k} N^{k_m}N^{-k_m}E^{\epsilon_m})N^{k_{k+1}}N^{-k_{k+1}}$, where $\epsilon_m \in \{-1,1\}$ and k_m is an integer for all $n+1 \leq m \leq k+1$, such that if the robot starts from the point $\mathbf{z}$ and follows K_4 it passes through the destination point. More specifically if the robot is at some specified latitude in the finite column R_i and it follows $N^{k_i}N^{-k_i}E^{\epsilon_i}$, it gets to some specified latitude in the finite column R_{i+1} for $n+1 \leq i \leq k$, where by convention we write R_{k+1} for R. If the robot is at some specified latitude in R and it follows $N^{k_{k+1}}N^{-k_{k+1}}$ it passes through the destination point. However, we remark that if the robot starts on r_i strictly at the west of $\mathbf{v}$ and it follows K_4 it always remains strictly at the west of $\mathbf{v}$, as the robot follows the E instructions at latitude i and the HNE of the west pipe prevents it from visiting longitudes greater than that of $\mathbf{v}$.

We finally construct the algorithm $K_5 = E^{|K_4|+1}$ and note that if the robot starts on r_i strictly at the west of $\mathbf{v}$ and it follows K_5 it always remains strictly at the west of $\mathbf{v}$.

We define the algorithm $K = K_1K_2K_3K_4K_5$. Note that if the robot starts on the row r_i strictly at the west of the point $\mathbf{v}$ and it follows K then it returns on the row r_i strictly at the west of $\mathbf{v}$. Indeed, the last part follows by the remarks we made on $K_1, \ldots, K_5$ individually and the first part follows from the fact that the robot does not meet any VNEs if it starts on the row r_i strictly at the west of $\mathbf{v}$ and it follows K. If the robot starts at the point $\mathbf{v}$ and it follows K, then it visits the destination point; this follows directly from the definitions of $K_1, \ldots, K_5$. Finally, we claim that if the robot starts on the row r_i strictly at the west of the point $\mathbf{v}$ and it follows K, its longitude does not decrease. Indeed, the only W instructions in K occur either in K_4, which is followed by K_5 specifically designed to negate them or as the last instruction in K_2, which is preceded by an E instruction. Therefore, the claim holds (see Figure 12).

Finally, we define the algorithm

$$F = N^i \; MW(c_{uw}, q) \; SME(\mu + \lambda + 2q, \mu + \lambda + 2q, K),$$

where $\mu = |MW(c_{uw}, q)|$.

Proposition 7.15. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows XN^i it gets in the west strip on the row r_i. While the robot follows the algorithm $MW(c_{uw}, q)$ it gets on the row r_i at $\mathbf{v}$ or to the west of $\mathbf{v}$. By the choice of parameters and by Lemma 4.1, if the robot is in the west strip on the row r_i at the west of the point $\mathbf{v}$ and it follows $SME(\mu+\lambda+2q, \mu+\lambda+2q, K)$, it advances eastwards oscillating about the row r_i. While the robot is on the row r_i strictly at the west of $\mathbf{v}$ and it follows K, it remains on the row r_i strictly at the west of $\mathbf{v}$. After defining K, we checked that it satisfies the conditions required in order to apply Lemma 4.1. Finally, the robot reaches the point $\mathbf{v}$ not while executing K, but while executing a locomotory move in SME.

Immediately afterwards, the robot follows K and it gets to the destination point. The conclusion follows. $\square$

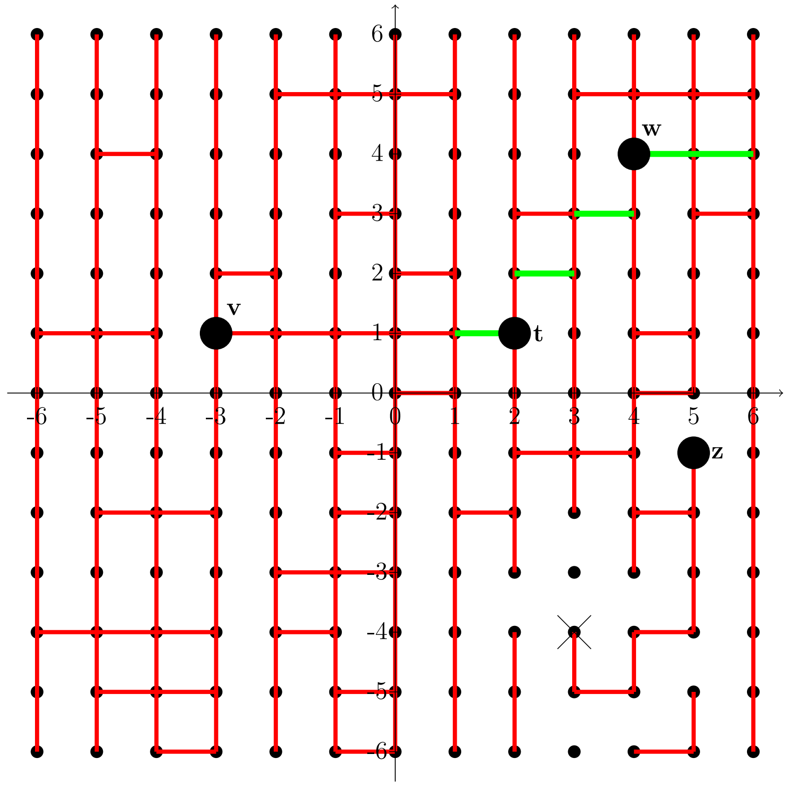

FIGURE 12. **Part III, Case (4)(iv).** We assume that there do not exist some special upper west paired HNEs and there exists an upper west pipe on the row r_i. We assume that there are no VEs removed other than the ones shown in the figure.

Let us examine the example shown in Figure 12. For this example we have $a = 1$ and $b = 5$. We have $r_i = r_1$ with the upper west pipe $\{(-5,1),(-4,1),(-3,1)\}$. The points $\mathbf{v}$, $\mathbf{t}$, $\mathbf{w}$ and $\mathbf{z}$ are marked on the figure and the west ascending chain is coloured green. We also have $HE_{special} = \{(4,-2),(5,-2)\}$, $d = 6$ and let us assume for this example that $c_{uw} = 6$, though this value should be larger. Then, if the robot follows $K_1 = E^5$ starting from $\mathbf{v}$ it gets to $\mathbf{t}$; if the robot follows $K_2 = (S^5N^5E)^2\ (S^7N^7E)\ W$ starting from $\mathbf{t}$ it gets to $\mathbf{w}$; if the robot follows $K_3 = S^6EN^6$ starting from $\mathbf{w}$ it gets to $\mathbf{z}$; if the robot follows $K_4 = (N^4S^4W)(NSW)NS$ starting from $\mathbf{z}$ it gets to the destination point; $K_5 = E^{15}$. We define $K = K_1K_2K_3K_4K_5$ and note that if the robot follows K starting from $\mathbf{v}$ it visits the destination point, but if the robot follows K starting on $r_i = r_1$ strictly at the west of $\mathbf{v}$, it returns on r_i strictly at the west of $\mathbf{v}$.

4(v) We assume that there does not exist a magical west row, there does not exist a special almost empty west row, there does not exist an upper west pipe, there do not exist the special upper west paired HNEs, but there exists an upper west cutoff. We recall that the upper west cutoff is the easternmost HNE in the west strip on the row r_i which contains the upper west pass. Then the row r_i is the complement of a path in the west strip and all the rows r_k with $k \leq j = i - c_{uw} \leq -p$ are paths in the west strip and indeed paths in the entire maze (from the non existence of the special upper west paired HNEs). Let $\mathbf{v} = (a, i)$ be the easternmost vertex of the row r_i in the west strip. Let $\mathbf{z} = (a+1, z)$ be the uppermost vertex of the westernmost lower infinite column in the column c_{a+1}. Let $\mathbf{w} = (a, z - c_{uw})$. By inspecting the tertiary rectangle, we can construct an algorithm K that takes the robot from $\mathbf{v}$ to the destination point.

We define the algorithm

$$F := N^i \ (S^{i-j} E N^{2i-2j} S^{i-j} W)^{\lambda+q} \ N^{i+c_{uw}-z} \ K.$$

CLAIM 7.16. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^i$ it gets in the west strip on the row r_i. By the choice of exponents, while the robot follows $(S^{i-j} E N^{2i-2j} S^{i-j} W)^{\lambda+q}$ it gets to the point $\mathbf{w}$ and remains stuck there. Indeed, while the robot follows each instance of $S^{i-j} E N^{2i-2j} S^{i-j} W$, it advances one unit to the east, oscillating about the row r_i until it gets to $\mathbf{v}$. Immediately afterwards, it follows $S^{i-j} E N^{2i-2j} S^{i-j} W$ and gets to $\mathbf{w}$. After the robot gets to $\mathbf{w}$, after each other instance of $S^{i-j} E N^{2i-2j} S^{i-j} W$, the robot gets back to $\mathbf{w}$. If the robot starts at $\mathbf{w}$ and it

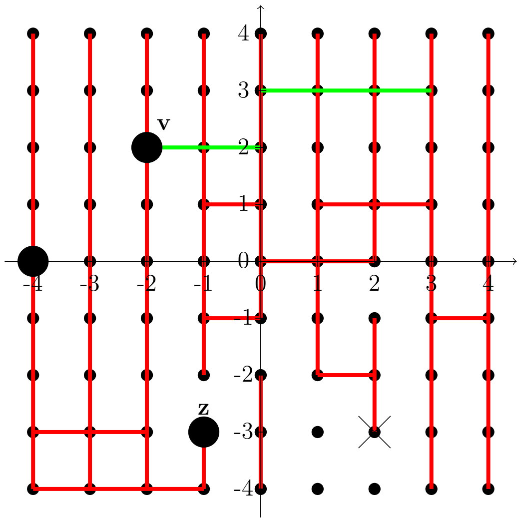

FIGURE 13. **Part III, Case (4)(v).** There does not exist a magical west row, there does not exist a special almost empty west row, there does not exist an upper west pipe, there do not exist the special upper west paired HNEs, but there exists an upper west cutoff. We assume that there are no VEs removed other than the ones shown in the figure.

follows $N^{i+c_{uw}-z}$, it gets to $\mathbf{v}$. Therefore, after the robot follows $X\ F$, it gets to the destination point. The conclusion follows. $\square$

Let us examine the example shown in Figure 13. In this example, the row $r_i = r_2$ is the complement of a path in the west strip and all the rows r_k with $k \leq j = i - c_{uw} \leq -p$ are paths in the west strip and indeed paths in the entire maze. For the purpose of this example, we can take c_{uw} to be any large constant, say $c_{uw} = 100$. The points $\mathbf{v}$ and $\mathbf{z}$ are marked on the figure, $z = -3$, $j = i - c_{uw} = -98$ and $\mathbf{w} = (-2, -103)$. Let us see what is the path of the robot as it follows $F = N^2(S^{100}EN^{200}S^{100}W)^4N^{105}K$ starting from $(-4, 0)$, where K is any algorithm that takes the robot from $\mathbf{v}$ to the destination point. When the robot follows $S^{100}EN^{200}S^{100}W$ starting from $(-4, 2)$, it first reaches a row which is a path after it executes S^{100}, so its longitude increases by 1 after it executes $S^{100}E$; so after the robot executes $S^{100}EN^{200}S^{100}$ it is back on $r_2 = r_i$ with its latitude increased by one, at $(-3, 2)$; the W instruction at the end does not change the longitude of the robot, as r_2 is the complement of a path in the west strip. Similarly, after the robot follows $S^{100}EN^{200}S^{100}W$ starting from $(-3, 2)$ it gets to $\mathbf{v} = (-2, 2)$. After the robot follows $S^{100}EN^{200}S^{100}W$ starting from $\mathbf{v}$, it enters the lower infinite column on c_{a+1}: after $S^{100}E$ it is at $(a + 1, j) = (-1, -98)$; after the robot follows $S^{100}EN^{200}S^{100}$, it is at $(-1, -103)$; finally, after the robot follows $S^{100}EN^{200}S^{100}W$, it is at $\mathbf{w} = (-2, -103)$. Similarly, we can see that after the robot follows each subsequent instance of $S^{100}EN^{200}S^{100}W$ starting at $\mathbf{w}$, it returns to $\mathbf{w}$. After the robot follows enough instances of $S^{100}EN^{200}S^{100}W$ to reach $\mathbf{w}$, it follows $N^{i+c_{uw}-z} = N^{105}$ and it reaches $\mathbf{v}$; immediately afterwards, the robot follows K and it reaches the destination point.

4(vi) We assume there exists an upper west HNE on some row r_j. We recall that the upper west HNE is the lowermost westernmost HNE at the north-east of the uppermost westernmost VNE. We further assume there does not exist a magical west row, there does not exist a magical east row and there does not exist an upper west cutoff. Then all the rows r_m with $i \leq m < j$ are paths in the maze (from the minimality of j and the non-existence of a magical east row). Let $\mathbf{v}$ be the western vertex of the upper west HNE. Let $\mathbf{w} = (x_w, y_w)$ be the upper vertex of the uppermost westernmost VNE. Then $\mathbf{v}$ is at the east of $\mathbf{w}$. By inspecting the tertiary rectangle, we construct an algorithm K which takes the robot from $\mathbf{v}$ to the destination point (see Figure 14).

We define the algorithm

$$F = N^i(ES^{j-y_w}N^{j-y_w})^{\lambda+q}K.$$

Claim 7.17. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows XN^i it gets in the west strip on the row r_i. In the west strip, while the robot follows $ES^{j-y_w}N^{j-y_w}$ it advances eastwards oscillating about the row r_i. In the obstacle strip, while the robot follows $ES^{j-y_w}N^{j-y_w}$ it advances eastwards, potentially increasing its latitude as it meets VNEs. It eventually

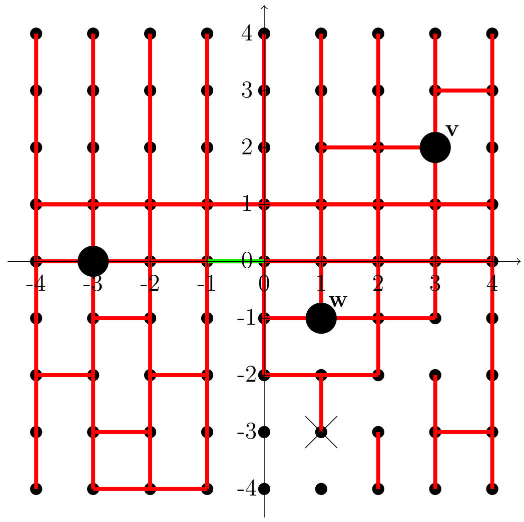

FIGURE 14. **Part III, Case (4)(vi).** We assume there exists an upper west HNE on some row r_j. We further assume there does not exist a magical west row, there does not exist a magical east row and there does not exist an upper west cutoff and that there are no VEs removed other than the ones shown in the figure.

gets on the row r_j and remains stuck at the point $\mathbf{v}$. Therefore, after the robot follows $X\ F$ it gets to the destination point. The conclusion follows (see Figure 14). $\square$

Let us examine the example shown in Figure 14. In this example, all the rows r_m with $i \leq m < j$ are paths in the maze. In this example, the upper west pass is coloured green, the uppermost westernmost VNE is $\{(1, -2), (1, -1)\}$, the upper west HNE is $\{(3, 2), (4, 2)\}$ and so $j = 2$. The vertices $\mathbf{v}$ and $\mathbf{w}$ are marked on the figure. We can take $K = S^3(WS)^2$, so if the robot follows K starting from $\mathbf{v}$ it visits the destination point. Let us observe how the robot follows $F = (ES^3N^3)^{10}K$ starting from $(-3, 0)$. As long as the robot is in the west strip, each instance of ES^3N^3 increases its longitude by one. Eventually, the robot gets to $(0, 0)$. After that, the robot follows S^3N^3 and it gets to $(0, 1)$. Considering that every row at latitude between $i = 0$ and $j = 2$ is a path in the maze, every further instance of ES^3N^3 increases the longitude of the robot by one, until it arrives at $\mathbf{v} = (3, 2)$, as its latitude is determined by the uppermost

VNEs at the west of $\mathbf{v}$. Once the robot reaches $\mathbf{v}$, we can see that after each instance of ES^3N^3, the robot returns to $\mathbf{v}$. Finally, the robot follows K and it visits the destination point.

4(vii) We assume there does not exist a magical west row, there does not exist a magical east row, there does not exist an upper west cutoff, there does not exist an upper west HNE, but there does exist a special west pipe on some row r_j. We recall that the special west pipe is the west pipe (the easternmost configuration in the west strip of a HE followed by a HNE) on the smallest row that has a west pipe with respect to the standard well order on $\mathbb{Z}$. Then all the rows r_m with $m \geq i$ are paths in the maze (from the non existence of an upper west HNE and the non existence of a magical east row).

Let $\mathbf{v} = (x_v, j)$ be the eastern vertex of the HE of the special west pipe. Let $\mathbf{w} = (a+1, y_w)$ be the lowermost vertex of the westernmost upper infinite column R_{a+1}. Let $\mathbf{t} = (x_v, i)$ be the vertex at the intersection between the column c_{x_v} and the row r_i. Let $\mathbf{z} = (n+1, y_z)$ be the uppermost vertex of the finite column R_{n+1} or the uppermost vertex of the lower infinite column $R_{n+1} = R$ that contains the destination point. The special case that the destination point is in the upper infinite column $R_{n+1} = R$ is much more easy and we will make a note on how to solve it before defining the finish algorithm F. Let $HE_{special}$ be a HE on some row r_γ between the upper infinite column R_n and the finite column R_{n+1}. Let **v'** be the eastern vertex of the HE of any "fake west pipe", i.e. a configuration in the west strip on r_j that is formed by a HE followed by a HNE, strictly at the west of the special west pipe (see Figure 15).

We define the algorithm $K_1 = N^{i-j}E^{a+1-x_v}S^{2i-j-y_w}N^{2i-j-y_w}W^{a+1-x_v}S^{i-j}$ with the property that if the robot starts from $\mathbf{v}$ and follows K_1 it passes through the point $\mathbf{w}$ and gets to the point $\mathbf{t}$. However, if the robot starts at **v'** and it follows K_1 then it returns at **v'**. The second statement follows from the fact that the robot moves at every instruction in K_1: indeed, while the robot executes N^{i-j} starting from **v'**, it is in the west strip which contains no VNEs, so it changes its latitude to i; considering that r_i is a path in the maze, when the robot continues to follow E^{a+1-x_v}, its longitude increases by exactly $a+1-x_v$ which is the exact difference in longitude between **v** and the westernmost column in the obstacle strip, c_{a+1}; as **v'** is strictly at the west of **v**, we conclude that after the robot follows $N^{i-j}E^{a+1-x_v}$ starting from **v'**, it is still in the west strip on the row r_i which is a path in the maze; hence, if the robot follows K_1 starting from **v'**, it gets back to **v'**. Similarly, we can show the first statement about K_1, that if the robot starts from **v** and follows K_1 it gets to the point **t**; in this case, we note that the only instructions in K_1 that do not change the position of the robot are instructions of type S from the group S^{2i-j-y_w} that occur immediately after the robot reaches **w** (see Figure 15).

We define the algorithm $K_2 = E^{n+1-x_v}WS^{i-\gamma}EN^{i-\gamma}$ such that if the robot starts from **t** and follows K_2 it gets to the point **z**. This is clear as the robot starts on r_i which is a path, so after it follows $E^{n+1-x_v}W$ it gets at the point (n, i) and so after it follows K_2 it is in R_{n+1}; moreover, as the upper west pass at latitude i is above all the passes in the obstacle strip and so, in this case, also above all the VNEs, the robot actually gets to **z** in R_{n+1} after it follows K_2 starting from **t**. However, if the robot follows K_2 starting from **v'**, it does not move after it follows E^{n+1-x_v} and its longitude decreases by 1 after it follows $E^{n+1-x_v}W$. Hence, if the robot follows K_2 starting from **v'**, it either gets back to **v'** or it gets to the western neighbour of **v'** (see Figure 15).

By inspecting the tertiary rectangle, we construct the algorithm K_3 of the form $K_3 = (\prod_{m=n+1}^{k} N^{k_m}N^{-k_m}E^{\epsilon_m})N^{k_{k+1}}N^{-k_{k+1}}$, where $\epsilon_m \in \{-1, 1\}$ and k_m is an integer for all $n+1 \le m \le k+1$, such that if the robot starts from the point **z** and follows K_3 it passes through the destination point. More specifically, if the robot is at some specified latitude in the finite column R_m and follows $N^{k_m}N^{-k_m}E^{\epsilon_m}$ it gets to some specified latitude in the finite column R_{m+1} for $n+1 \le m \le k$, where by convention we write R_{k+1} for R. If the robot is at some specified latitude inside R and it follows $N^{k_{k+1}}N^{-k_{k+1}}$, it visits the destination point.

We construct the algorithm $K_4 = E^{|K_3|+1}$. We note that from the structure of a fake west pipe and its position in the west strip, if the robot starts either at **v'** or at the western neighbour of **v'** and it follows K_3K_4, it gets to **v'**.

We define the algorithm $K = K_1K_2K_3K_4$ with the property that if the robot starts at $\mathbf{v}$ and it follows K, it passes through the destination point. However, if the robot starts at **v'** and it follows K, it gets back to **v'**. In the special case when $\mathbf{z}$ does not exist and so the destination point $(n+1, \delta)$ is in the upper infinite column $R_{n+1} = R$ we define $K_2' = E^{n+1-x_v}N^{\delta-i}S^{\delta-i}$. In this case we define $K = K_1K_2'$ instead and we note that, as before, if the robot starts at $\mathbf{v}$ and it follows K, it passes through the destination point; moreover, if the robot starts at **v'** and it follows K, it gets back to **v'**.

We recall the algorithm $WPF(a, e) := (E^eWS^aEN^a)^e$, defined in the case **2(ii)**. Finally, we define the algorithm

$$F = N^i\ W^{\lambda-x_v}\ S^{i-j}\ (WPF(j-i, 2\lambda+q)KN^{i-j}ES^{i-j})^{2\lambda+q}.$$

CLAIM 7.18. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

PROOF. We may assume without loss of generality that after the robot follows $X\ N^i\ W^{\lambda-x_v}\ S^{i-j}$ it gets in the west strip on the row r_j at the west of the point $\mathbf{v}$. While the robot follows each instance of $WPF(j-i, 2\lambda+q)$ it advances eastwards to the easternmost vertex **v'** of a HE of a fake west pipe on the row r_j. If **v'** is strictly at the west of $\mathbf{v}$, after the robot follows the algorithm K it returns to the point **v'**; after the robot follows the algorithm $N^{i-j}ES^{i-j}$ starting from **v'**, it advances to the east of **v'** on the row r_j. By the choice of parameters, the robot eventually gets to the point **v'** $= \mathbf{v}$. Immediately afterwards, it follows K and it gets to the destination point. The conclusion follows. □

Let us examine the example shown in Figure 15. The upper west pass is coloured green and it is on the row $r_i = r_0$. From the assumptions, it

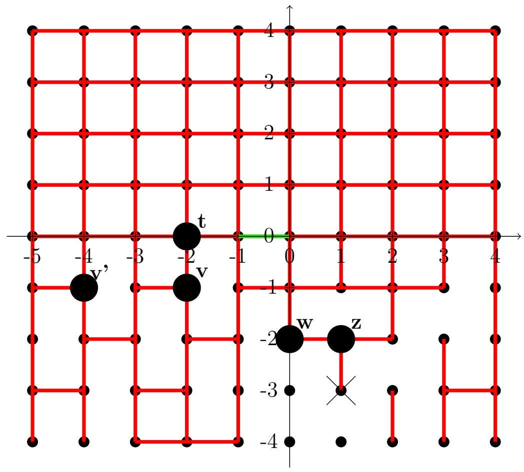

FIGURE 15. **Part III, Case (4)(vii).** We assume there does not exist a magical west row, there does not exist a magical east row, there does not exist an upper west cutoff, there does not exist an upper west HNE, but there does exist a special west pipe on some row r_j. We assume that there are no VEs removed other than the ones shown in the figure.

follows that for every $m \geq i$, the row r_m is a path in the maze. The special west pipe is $\{(-3,-1),(-2,-1),(-1,-1)\}$ on $r_j = r_{-1}$. We take R_{n+1} to be $\{(1,-2),(1,-3)\}$, accessible from $R_n = \{(0,-2),(0,-1),\ldots\}$ via $HE_{special} = \{(0,-2),(1,-2)\}$ on $r_\gamma = r_{-2}$. Then, if the robot follows $K_1 = NE^2S^3N^3W^2S$ starting from **v**, it gets to **t** passing from **w**; however, note that if the robot follows K_1 starting from **v'** (which is the eastern vertex of the HE of the "fake west pipe" $\{(-5,-1),(-4,-1),(-3,-1)\}$ on r_j strictly at the west of **v**), it returns to **v'**. If the robot follows $K_2 = E^3WS^2EN^2$ starting from **t**, it gets to **z**; however, if the robot follows K_2 starting from **v'** it gets back to **v'**; in general, we are certain that if the robot follows K_2 starting from **v'** it either gets back to **v'** or to the western neighbour of **v'**. If the robot follows $K_3 = NSW$ starting from **z** it visits the destination point. In this case, $K_4 = E^4$. Therefore, if the robot follows K_3K_4 starting either from **v'** or from the western neighbour of **v'**, it gets to **v'**.

4(viii) We assume that there does not exist a magical west row, there does not exist a magical east row, there does not exist an upper west cutoff, there does not exist an upper west HNE, but there exists a natural special empty west row on r_j. Then, as in **4(vii)**, all the rows r_m for $m \geq i$ are paths in the maze. Let $\mathbf{v} = (a, j)$ be the easternmost vertex of the row r_j in the west strip. Let $\mathbf{z} = (a+1, \gamma)$ be the lowermost vertex of the westernmost upper infinite column R_{a+1}. Let $\mathbf{w} = (a, 2i - \gamma)$. By inspecting the tertiary rectangle, we construct an algorithm K that takes the robot from $\mathbf{v}$ to the destination point.

We define the algorithm

$$F = N^j \ (N^{i-j} E S^{3i-2\gamma-j} N^{2i-2\gamma} W)^{\lambda+q} \ S^{2i-\gamma-j} \ K.$$

Claim 7.19. *For any maze in F_i, after the robot follows $X\ F$, it visits the destination point.*

Proof. We may assume without loss of generality that after the robot follows $X\ N^j$, it gets in the west strip on the row r_j. While the robot follows each instance of $N^{i-j} E S^{3i-2\gamma-j} N^{2i-2\gamma} W$, it advances eastwards one unit making an oscillation about the row r_j. By the choice of exponent, the robot eventually gets to the point $\mathbf{v}$. Immediately afterwards, it follows $N^{i-j} E S^{3i-2\gamma-j} N^{2i-2\gamma} W$ and gets to the point $\mathbf{w}$. The robot remains stuck at $\mathbf{w}$, i.e. while it follows each instance of $N^{i-j} E S^{3i-2\gamma-j} N^{2i-2\gamma} W$, it gets back to $\mathbf{w}$ (see Figure 16). Hence after the robot follows $X\ N^i\ W^{\lambda-a}\ S^{i-j}\ (N^{i-j} E S^{3i-2\gamma-j} N^{2i-2\gamma} W)^{\lambda+q}\ S^{2i-\gamma-j}$, it gets to $\mathbf{v}$. Hence, after the robot follows $X\ F$, it gets to the destination point. The conclusion follows. $\square$

Let us examine the example shown in Figure 16. Let us suppose that the robot starts at $(-3, 0)$ and it follows $F = N^{-2}(N^2 E S^6 N^4 W)^{10} S^4 K$, where $K = N(ES)^2$ is an algorithm with the property that if the robot follows it starting from $\mathbf{v}$ it reaches the destination point. While the robot is on $r_j = r_{-2}$ strictly at the west of $\mathbf{v}$, its longitude increases by one after each instance of $N^2 E S^6 N^4 W$. After the robot reaches $\mathbf{v}$ and it follows $N^2 E S^6 N^4 W$, it gets to $\mathbf{w}$. If the robot follows $N^2 E S^6 N^4 W$ starting from $\mathbf{w}$ it gets back to $\mathbf{w}$.

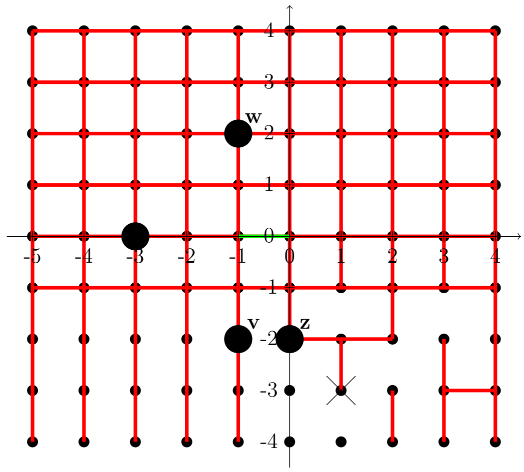

FIGURE 16. **Part III, Case (4)(viii).** We assume that there does not exist a magical west row, there does not exist a magical east row, there does not exist an upper west cutoff, there does not exist an upper west HNE, but there exists a natural special empty west row on $r_{-2} = r_j$. We assume that there are no VEs removed other than the ones shown in the figure.

4(ix) As a final case, we may assume that there does not exist a magical west/east row, there does not exist a special west pipe, there does not exist a natural special empty west row, there does not exist a special almost empty west row. Then all the rows are paths in the maze and hence the maze does not contain any HNE. Therefore, both the latitude and the longitude of the robot are known and, by inspecting the primary rectangle, we can write an algorithm F that takes the robot from its known position to the destination point. The conclusion follows.

This finally solves **Case (4)** in which the destination point is connected to the west strip by a (finite, possibly empty) sequence of finite columns followed by a (finite, non-empty) sequence of upper infinite columns.

We have therefore treated all possible cases, as detailed in the arguments above. This completes the proof of Theorem 3.2. $\square$

8. Proof of Proposition 3.3 and Theorem 1.4

In this short section we present a proof of the slightly technical but easy Proposition 3.3 and then we finally establish Theorem 1.4.

The following observation represents the main idea of the proof.

Observation 8.1. *Let o, d be fixed vertices in $\mathbb{Z}^2$ and let $\mathcal{B}$ be a set of subgraphs of $\mathbb{Z}^2$ which is compact in the product topology. Let A be a possibly infinite algorithm that solves the set of mazes $\mathcal{A} = \{(B, o, d) \mid B \in \mathcal{B}\}$. Then there exists a finite initial segment A_0 of A that solves $\mathcal{A}$.*

Proof. Assume for a contradiction that there does not exists such an initial segment A_0. For each $i \geq 1$, let A_i be the initial segment of A with the first i instructions. By assumption, for each $i \geq 1$ there exists a board $B_i \in \mathcal{B}$ such that A_i does not solve B_i. By compactness there exists a subsequence $(B_{i_j})_{j\geq 1}$ such that $\lim_{j\to\infty} B_{i_j} = B_0 \in \mathcal{B}$ in the product topology. As A solves B_0, there exists an initial segment A_0 of A which solves (B_0, o, d). As $\lim_{j\to\infty} B_{i_j} = B_0 \in \mathcal{B}$, A_0 solves $(B_{i_j})_{j\geq 1}$ for all $j \geq |A_0|$ sufficiently large. This gives the desired contradiction. □

We are now ready to prove Proposition 3.3 and Theorem 1.4.

Proof of Proposition 3.3. By hypotheses (1) and (3) and by Observation 8.1, for all i, all origins $o \in \mathbb{Z}^2$, all destination $d \in \mathbb{Z}^2$ and all paths P between o and d, there exists a finite initial segment $A_{i,P}$ of A_i that solves the set of mazes $\{(M, o, d) \mid (M, o, d) \in \mathcal{A}_i, P \leq M\}$ that contain the path P (this set of mazes might be empty). By hypothesis (2), for all i, all origins $o \in \mathbb{Z}^2$ and all $j \in \mathbb{N}$, there exists a finite initial segment $A_{i,o,j}$ of A_i that guides the robot to visit all accessible points at distance at most j from the origin o in the set of mazes $\{(M, o, d) \mid (M, o, d) \in \mathcal{A}_i\}$ that have origin o (notice that here the destination d plays no role so we might as well drop it). But then for all $i, j, k \in \mathbb{N}$, there exists a finite initial segment $A_{i,j,k}$ of A_i such that for any origin o at distance at most k from $\mathbf{0}$ in the graph $\mathbb{Z}^2$, the algorithm guides the robot to visit all accessible points at distance at most j from the origin o in the set of mazes $\{(M,o) \mid (M,o) \in \mathcal{A}_i\}$ that have origin o.

In order to construct the algorithm A, we define the algorithms B_i recursively to be $B_i = A_{f(i),2|B_1 \ldots B_{i-1}|+1,2|B_1 \ldots B_{i-1}|+1}$, $f = (1,1,2,1,2,3,\ldots)$ and take $A := B_1 B_2 \ldots$. Clearly, the algorithm A has the desired properties. $\square$

Proof of Theorem 1.4. In Proposition 3.3, let $\mathcal{A}_1 = \mathcal{C}$ be the set of all mazes with no VNEs and for all $i \geq 2$ let $\mathcal{A}_i \subseteq \mathcal{F}$ be the set of all mazes with finitely many VNEs in consecutive columns, all of which are in the finite box $[-i,i]^2$. Then hypothesis (1) of Proposition 3.3 easily holds, hypothesis (2) is trivial, and hypothesis (3) follows from Theorem 3.1 for $\mathcal{A}_1$ and Theorem 3.2 for $\mathcal{A}_i$, $i \geq 2$. $\square$

9. Open Problems

As we emphasised in the proof of Theorem 1.4, we strongly believe that there exists an algorithm which solves the set of all mazes with arbitrarily many HNEs and finitely many VNEs. The only case in our proof where an argument for this result breaks down is **Case 4** of **Part III**. We believe that this problem, together with Conjecture 9.1 below could be solved using similar techniques with those developed in this chapter.

Conjecture 9.1. *There exists an algorithm that solves the set of all mazes with arbitrarily many HNEs and arbitrarily many VNEs in one column.*

Furthermore, we believe the following positive result to hold.

Conjecture 9.2. *Consider the subset $\mathcal{N} \subseteq \mathcal{M}(\mathbb{Z}^2)$ of mazes in which the connected component of the origin is a simple (possibly infinite) path. Then there exists an algorithm that solves $\mathcal{N}$.*

In the opposite direction, we believe the following to be true.

Conjecture 9.3. *There is no algorithm that solves the class $\mathcal{M}(\mathbb{Z}^2)$ of all mazes.*

From another perspective, let us call $\mathcal{M}_k \subseteq \mathcal{M}(\mathbb{Z}^2)$ the set of mazes for which the destination is at distance k from the origin. From Corollary 3.4, the following conjecture is equivalent to Conjecture 9.3.

Conjecture 9.4. *There exists a k for which $\mathcal{M}_k$ is not solvable.*

Perhaps the following stronger results also hold.

Conjecture 9.5. *Let $\mathcal{N}_3 \subset \mathcal{M}(\mathbb{Z}^2)$ be the set of all mazes for which there are only HNEs between the pairs of columns (c_{-4}, c_{-3}) and (c_3, c_4). Then there is no algorithm that solves $\mathcal{N}_3$.*

Conjecture 9.6. *Conjecture 9.4 holds for $k = 10$.*

The intuition behind Conjecture 9.5 is that for us it does not look plausible to navigate the robot in a coordinated way between infinitely many finite columns, even if we make additional assisting assumptions. Conjecture 9.5 is one of the main reasons why we think Conjecture 9.3 holds.

Finally, we strongly believe that the classes of mazes in higher dimensions arising from the lattice $\mathbb{Z}^k$ with suitable mild restrictions should represent a captivating further study.

CHAPTER 3

Products of posets with long chains

1. Introduction

This chapter is joint work with Hunter Spink and Marius Tiba.

A finite graded poset $(P, \leq)$ is a finite poset equipped with a *rank function* $\mathrm{rk} : P \to \mathbb{N} \cup \{0\}$ such that the rank of every minimal element is 0, and if y covers x, i.e. $x < y$ and if $x \leq z \leq y$ then $z = x$ or $z = y$, denoted by $x \prec y$, then $\mathrm{rk}(y) = \mathrm{rk}(x) + 1$. The rank of P, denoted by $\mathrm{rk}(P)$, is the maximal value of rk on P. A graded poset P is said to be *rank-symmetric* if the number of elements of rank r is the same as the number of elements of rank $\mathrm{rk}(P) - r$. If P has a unique maximal/minimal element, then we will denote them by $\min_P$ and $\max_P$.

A *symmetric chain* in P is a chain which for some r consists of exactly one element of ranks $r, r+1, \ldots, \mathrm{rk}(P) - r$. A *symmetric chain decomposition* of P is a partition of P into symmetric chains. Let $\mathbf{m}$ be the m-element chain poset $0 \to 1 \to \ldots \to m - 1$.

The first attempt to study multiple symmetric chain decompositions simultaneously on a given poset occurred perhaps in 1979 when Shearer and Kleitman [35] found the minimum probability that two randomly chosen elements contain each other in Q_n for an arbitrary probability distribution. To make their proof work, they needed two of what they called "orthogonal chain decompositions" of Q_n, which are simply two decompositions of the n-dimensional hypercube Q_n into $\binom{n}{\lfloor n/2 \rfloor}$ chains so that any chain in one decomposition intersects any other chain in the other decomposition in at most one element. Their construction proceeds by slightly modifying on Q_n two "almost orthogonal symmetric chain decompositions" — two symmetric chain decompositions which satisfy the orthogonal intersection condition except for the maximal chain in both decompositions, which must intersect in precisely their top and bottom elements. Since Shearer and Kleitman's paper in 1979 where they further conjectured that there are $\lfloor n/2 \rfloor + 1$ orthogonal decompositions of Q_n, no progress has been made on the conjecture until Spink [37]. In [37], it was shown that three orthogonal decompositions can be constructed for all sufficiently high dimensional hypercubes, and additionally they arise from three almost orthogonal symmetric chain decompositions.

The strategy pursued in [37] was as follows. Suppose that for $1 \leq j \leq l$, we have almost orthogonal symmetric chain decompositions $\mathcal{F}_i^j$ of Q_{n_i} for $i = 1, 2, \ldots, r$. Then to create l almost orthogonal symmetric chain decompositions in $Q_{n_1+\ldots+n_r}$, we aim to give symmetric chain decompositions of the cuboids in $\prod_i \mathcal{F}_i^j$ in such a way that the chains from cuboids in $\prod_i \mathcal{F}_i^j$ and chains from cuboids in $\prod_i \mathcal{F}_i^{j'}$ intersect in at most one element when $j \neq j'$ (except of course for the two maximal chains, which we require to intersect in just their top and bottom elements).

To put the questions addressed in this chapter in the proper context, we consider the most difficult case from [37]. Suppose $l = k+1$ and take the product of a 2-element chain from each $\mathcal{F}_i$ for $1 \leq i \leq k$ with the maximal chain in $\mathcal{F}_{k+1}$. Let n be the size of a maximal chain in the last hypercube. We then have two cuboids of the form $P(k,n) = \underbrace{\mathbf{2} \times \mathbf{2} \times \ldots \times \mathbf{2}}_{k} \times \mathbf{n}$, with the property that their intersection is either empty, or is $\{x\} \times \{\min(Q_{n_{k+1}}), \max(Q_{n_{k+1}})\}$, where x is some element of $Q_{n_1} \times \ldots \times Q_{n_k}$. To avoid the situation of having two chains intersect in at least two elements, it suffices to decompose $P(k,n)$ such that no subchain of a chain has the form $(a_1, \ldots, a_k, 0) \prec \ldots \prec (a_1, \ldots, a_k, n-1)$. In $P(k,n)$, we call a symmetric chain containing such a subchain *taut.* More generally, given a finite graded poset P, we say a symmetric chain in $P \times \mathbf{n}$ is *taut* if it contains for some $p \in P$ a subchain of the form $p \times 0 \prec p \times 1 \prec \ldots \prec p \times (n-1)$.

From this, the most natural question that arises is whether there is a symmetric chain decomposition of $P(k,n)$ without a taut chain. One of the main results of this chapter, Theorem 2.1, completely answers this question.

The answer is very surprising. For the family of posets $P(k,n)$ with $k \leq 4$, i.e. for the posets $\mathbf{2} \times \mathbf{n}$, $\mathbf{2} \times \mathbf{2} \times \mathbf{n}$, $\mathbf{2} \times \mathbf{2} \times \mathbf{2} \times \mathbf{n}$, and $\mathbf{2} \times \mathbf{2} \times \mathbf{2} \times \mathbf{2} \times \mathbf{n}$, every symmetric chain decomposition has a taut chain. For $k \geq 5$ and $n \geq 3$ however, we will explicitly construct in Section 4 decompositions with no taut chains by boot-strapping decompositions of $P(5,3)$, $P(5,4)$ and $P(5,5)$ using more general results about finite graded posets we prove in the remaining sections. These decompositions turn out to be very hard to find, as they are completely ad hoc, and finding them was the biggest challenge in proving the above main result.

One of the general bootstrapping results we prove in Theorem 2.2 is that if P is a finite graded poset with rank function rk, then for $m, n \geq \mathrm{rk}(P)+1$, the symmetric chain decompositions of the posets $P \times \mathbf{m}$ and $P \times \mathbf{n}$ are in natural bijection, and furthermore, this bijection preserves tautness of chains.

Also, if we additionally stipulate that P has a unique maximal/minimal element, then there is a canonical $\mathrm{rk}(P)+1$ to 1 surjection from symmetric chain decompositions of $P \times \mathbf{rk(P)+1}$ to symmetric chain decompositions of $P \times \mathbf{rk(P)}$ which send taut chains to taut chains. Under a mild additional hypothesis, if a symmetric chain decomposition of $P \times (\mathbf{rk(P)+1})$ with no taut chains exists, then there exists a symmetric chain decomposition of $P \times \mathbf{rk(P)}$ with no taut chains.

All posets in this chapter are finite graded posets; P will always refer to a finite graded poset.

This chapter is divided into four sections. In Section 2, we state our main results. In Section 3, we prove the main results pertaining to general finite graded posets. In Section 4, we explicitly construct symmetric chain decompositions with no taut chains for $P(3,5)$, $P(4,5)$, and $P(5,5)$. By previous results, we complete the proof of Theorem 2.1 on which $P(k,n)$ have symmetric chain decompositions without taut chains.

2. Main Results

The central result of this chapter is Theorem 2.1, proved in Section 4.

THEOREM 2.1. *There exists a symmetric chain decomposition of $P(k,n)$ with no taut chain if and only if $k \geq 5$ and $n \geq 3$.*

Importantly, for a fixed number of $\mathbf{2}$'s, making n very large does not aid us in constructing decompositions with no taut chains.

Most of our considerations generalize under mild conditions to arbitrary posets P in place of Q_k, which we consider in Section 3. In particular, we prove the following two theorems which we later apply to Q_k in the proof of Theorem 2.1. These theorems would allow one to answer the analogous question for $P \times \mathbf{n}$ in a similar way, reducing the problem to a finite computation.

THEOREM 2.2. *Let P be a rank-symmetric poset P, and let $m, n \geq rk(P) + 1$. Then there is a canonical bijection between the set of symmetric chain decompositions of $P \times \mathbf{m}$ and of $P \times \mathbf{n}$ which bijects decompositions with taut chains.*

THEOREM 2.3. *Let P be a rank-symmetric poset with a unique maximum and minimum element. Then there is a $(rk(P) + 1)$ to 1 surjection from the set of symmetric chain decompositions of $P \times$ ***(rk(P)+1)*** to the set of symmetric chain decompositions of $P \times$ ***rk(P)*** such that the pre-image of a decomposition without a taut chain contains only decompositions without taut chains. Furthermore, if P additionally has at least two elements of rank $rk(P)-1$ connected to the maximal element of P, then $P \times$ ***rk(P)*** has a decomposition without taut chains if and only if this is true for $P \times$ ***(rk(P)+1)***.*

REMARK. The hypothesis on the elements of rank $\mathrm{rk}(P) - 1$ in Theorem 2.3 is needed for example when $P = \mathbf{3}$.

$$
\begin{array}{ccccc}
 & & & & 1 \\
 & & & 4 & 1 \\
 & & 6 & 4 & 1 \\
 & 4 & 6 & 4 & 1 \\
1 & 4 & 6 & 4 & 1 \\
1 & 4 & 6 & 4 & 1 \\
1 & 4 & 6 & 4 & \\
1 & 4 & 6 & & \\
1 & 4 & & & \\
1 & & & &
\end{array}
$$

FIGURE 1. Pictorial Representation of $P \times \mathbf{n}$ in the case of $Q_4 \times \mathbf{6}$

3. Proofs of general results

In this section we prove the main results, postponing the completion of the proof of Theorem 2.1 until Section 4.

DEFINITION 3.1. *In the poset $P \times \mathbf{n}$, we define a packet to be the collection of elements of a given rank and $\mathbf{n}$-coordinate. The rank of a packet Λ, denoted $rk(\Lambda)$, is the common rank of elements of Λ. We also define $[p, r] := (p, r - rk(p)) \in P \times \mathbf{n}$ whenever $rk(p) \leq r \leq rk(p) + n$ (so $[p, r]$ is the unique element of $P \times \mathbf{n}$ with P-coordinate p and rank r).*

Consider the map from $P \times \mathbf{n} \to \mathbb{Z}^2$ given by $(p, r) \mapsto (\mathrm{rk}(p), \mathrm{rk}(p) + r)$. Note that $\mathrm{rk}(p) + r$ is just the rank of (p, r) in $P \times \mathbf{n}$, so $[p, r]$ gets sent to $(\mathrm{rk}(p), r)$. Under this map, the elements of $P \times \mathbf{n}$ are identified into their packets. If we label each point in the image of this map with the number of points in the corresponding packet, we call this the *pictorial representation* of $P \times \mathbf{n}$. Figure 1 depicts the pictorial representation of $Q_4 \times \mathbf{6}$ — we use the pictorial representation extensively for many of our proofs.

In the pictorial representation, the row $y = k$ contains all elements of rank k in $P \times \mathbf{n}$, the column $x = l$ contains all elements with P-coordinate of rank l, and the diagonal $y = x + r$ contains all elements of $\mathbf{n}$-coordinate r. A chain which skips no ranks connects a sequence of packets, with each packet following the previous one either vertically up one packet (which is uniquely determined in $P \times \mathbf{n}$ as increasing the $\mathbf{n}$-coordinate by 1), or diagonally up-right one packet (which corresponds to moving up in the P-coordinate).

LEMMA 3.2. *If $P \times \mathbf{n}$ has a decomposition with no taut chain, and Q is a poset with a symmetric chain decomposition, then $(P \times Q) \times \mathbf{n}$ has a decomposition with no taut chains.*

PROOF. We take the product of each non-taut chain in $P \times \mathbf{n}$ with each chain in the symmetric chain decomposition of Q and decompose each resulting rectangle into symmetric chains arbitrarily. Then the resulting chains are symmetric, non-taut, and give a symmetric chain decomposition of $(P \times Q) \times \mathbf{n}$ as desired. $\square$

COROLLARY 3.3. *If $P(k, n)$ has a symmetric chain decomposition with no taut chain, then so does $P(k', n)$ for any $k' \geq k$.*

LEMMA 3.4. *If $P \times \mathbf{n}$ has a symmetric chain decomposition, then P must be rank-symmetric.*

If furthermore $P \times \mathbf{n}$ has a symmetric chain decomposition into non-taut chains, then

- *if rk(P) is even, the size of the middle rank of P does not exceed the sum of all the sizes of lower ranks, and*
- *if rk(P) is odd, the common size of the middle ranks of P does not exceed twice the sum of the sizes of all ranks strictly before the middle ranks.*

PROOF. If $P \times \mathbf{n}$ has a symmetric chain decomposition, then by the rank-symmetry of $P \times \mathbf{n}$, we find that P is rank-symmetric (by arguing inductively from the smallest rank up).

Suppose $\mathrm{rk}(P)$ is even and we have a decomposition of $P \times \mathbf{n}$ into non-taut chains. Note that $\mathrm{rk}(P \times \mathbf{n}) = \mathrm{rk}(P) + \mathrm{rk}(\mathbf{n}) = \mathrm{rk}(P) + n - 1$. Let Λ be the packet of elements $(p, n-1) \in P \times \mathbf{n}$ with $\mathrm{rk}(p) = \mathrm{rk}(P)/2$. As $\mathrm{rk}(\Lambda) = \mathrm{rk}(P)/2 + n - 1$, a symmetric chain which contains an element

$(p, n-1)$ of Λ must also contain an element of the form $[q_p, \mathrm{rk}(P)/2]$. We have $\mathrm{rk}(q_p) < \mathrm{rk}(P)/2$, as if $\mathrm{rk}(q_p) = \mathrm{rk}(P)/2$, then the $\mathbf{n}$-coordinate of this element is 0 and $p = q_p$, so the symmetric chain is taut. Hence, as the q_p's are distinct, if the number of elements p of middle rank in P exceeds the number of elements of lower rank, then there are not enough elements q_p to accommodate the chains passing through elements of Λ.

Finally, if $\mathrm{rk}(P)$ is odd, we apply the above argument to the even-ranked poset $P \times \mathbf{2}$ (using Lemma 3.2 with $Q = \mathbf{2}$). $\square$

Corollary 3.5. *If $k \leq 4$ or $n \leq 2$, then every symmetric chain decomposition of $P(k, n)$ contains a taut chain.*

Proof. For $n = 1$ the result is trivial, and for $n = 2$ the maximal chain is always taut. For $k = 1, 2$, the result is trivial by inspection. For $k = 3, 4$, Lemma 3.4 applies. $\square$

Note that for $P = Q_k$ with $k \geq 5$, Lemma 3.4 does not apply. Now we are ready to prove Theorem 2.2 and Theorem 2.3.

Proof of Theorem 2.2. In the pictorial representation, when $n \geq \mathrm{rk}(P) + 1$, we have $n - \mathrm{rk}(P)$ consecutive rows in the middle at $y = \mathrm{rk}(P), \mathrm{rk}(P) + 1, \dots, n-1$, each consisting of $\mathrm{rk}(P) + 1$ packets at points (x, y) with $x = 0, 1, \dots, \mathrm{rk}(P)$. Furthermore, in these consecutive rows, for a fixed x, the number of elements in the packets at $(x, \mathrm{rk}(P)), \dots, (x, n-1)$ are the same. Each of these rows corresponds to a rank in $P \times \mathbf{n}$, and hence these ranks have the same number of elements, so a symmetric chain decomposition when restricted to any pair of adjacent rows must biject the elements between them.

As the chains can only move vertically up and diagonally up-right, and any two of these rows have identical packet sizes, this bijection is clearly only possible by having all of the chains move vertically up across this block of rows.

Now if $m \geq \mathrm{rk}(P) + 1$ we can modify a symmetric chain decomposition for $P \times \mathbf{n}$ to create one for $P \times \mathbf{m}$ as follows. Write each chain C in the decomposition of $P \times \mathbf{n}$ as the disjoint union of chains $C_1 \cup C_2 \cup C_3$, where C_2 is the subchain of elements in this middle block of rows, of the form $[p, \mathrm{rk}(P)] \prec [p, \mathrm{rk}(P)+1] \prec \ldots \prec [p, n-1]$, C_1 is the subchain of elements of lower rank than those of C_2, and C_3 is the subchain of elements of higher rank than those of C_2. We modify C to become a chain in $P \times \mathbf{m}$ by replacing C_2 with $[p, \mathrm{rk}(P)] \prec [p, \mathrm{rk}(P)+1] \prec \ldots \prec [p, m-1]$, and shifting C_3 by adding $m-n$ to the last coordinate of each element in C_3.

Finally, it is easy to see that this process preserves tautness of chains between $P \times \mathbf{n}$ and $P \times \mathbf{m}$. $\square$

Proof of Theorem 2.3. In the pictorial representation of $P \times (\mathbf{rk(P)+1})$, call M the middle row with packets at $(0, \mathrm{rk}(P)), \ldots (\mathrm{rk}(P), \mathrm{rk}(P))$, M^- the row right below the middle row with packets at $(0, \mathrm{rk}(P)-1), \ldots (\mathrm{rk}(P)-1, \mathrm{rk}(P)-1)$, and M^+ the row right above with packets at $(1, \mathrm{rk}(P)+1), \ldots, (\mathrm{rk}(P), \mathrm{rk}(P)+1)$. From the locations of the packets, the number of elements in the packets in M is 1 more than that in M^- and that in M^+, as P has a unique maximum and minimum element. Hence, there is a unique chain of length 1 in M in some packet Λ, and the remaining elements in M biject with those in M^- and in M^+. By working from left to right in M, we get the numbers of chains connecting pairs of packets from M^- to M are all completely determined by Λ: all packets in M which are to the right of Λ receive precisely one chain diagonally from M^-, and all other chains between M^- and M are vertical. Similarly, we get all packets in M which are to the left of Λ send one chain diagonally to M^+, and all other chains between M and M^+ are vertical.

Hence regardless of where Λ is, every element in M^- is connected to an element in M^+ whose P-coordinate has rank at most 1 higher. We can thus modify a symmetric chain decomposition of $P \times (\mathbf{rk(P)+1})$ to one for $P \times \mathbf{rk(P)}$ as follows. Ignore the chain of length 1, and for every other chain, decompose it as $D^- \cup D \cup D^+$ with D containing the element in M, D^- containing all elements of lower rank in the chain than those of D, and D^+ containing all elements of higher rank in the chain than those of D. To construct the chain in

$P \times \mathbf{rk(P)}$, we remove D, and decrease the second coordinate of all elements of D^+ by 1.

It is easy to check if a chain was taut, then it remains taut, and all newly constructed chains are still symmetric chains in $P \times \mathbf{rk(P)}$.

We now verify that the map above from the set of symmetric chain decompositions of $P \times (\mathbf{rk(P)+1})$ to the set of symmetric chain decompositions of $P \times (\mathbf{rk(P)})$ is a $\mathrm{rk}(P) + 1$ to 1 surjection. Suppose we have a symmetric chain decomposition $\mathcal{S}$ of $P \times \mathbf{rk(P)}$, viewed as a directed graph via $\prec$. Denote by N^- and N^+ the two middle rows in the pictorial representation of $P \times \mathbf{rk(P)}$. Consider the directed graph G (with loops) on P defined by taking the restriction of $\mathcal{S}$ to $N^- \cup N^+$, and projecting this induced directed subgraph onto the P-coordinate. As $\mathcal{S}$ induces a bijection between N^- and N^+, all vertices in G except min and max have in-degree and out-degree 1. Also, min has out-degree 1 and in-degree 0, while max has in-degree 1 and out-degree 0. Every directed edge in G is either a loop, or increases rank by 1 in P. From this observation, we can trivially deduce that G consists of one directed maximal chain (from $\min_P$ to $\max_P$) and loops on the remaining vertices.

We show now that there exists a canonical equivalence between symmetric chain decompositions $\mathcal{S}'$ of $P \times (\mathbf{rk(P)+1})$ that are mapped to $\mathcal{S}$, and matchings f between the edges of G and their endpoints.

Set m^-, m, m^+, n^-, and n^+ to be the ranks of M^-, M, M^+, N^-, and N^+ respectively ($n^+ - 1 = n^- = m^- = m - 1 = m^+ - 2$).

Suppose first that we have such a matching f and construct $\mathcal{S}'$ as follows. Identify the restriction of $\mathcal{S}'$ up to rank m^- with the restriction of $\mathcal{S}$ up to rank n^-. Similarly, identify the restriction of $\mathcal{S}'$ from rank m^+ onwards with the restriction of $\mathcal{S}$ from rank n^+ onwards. All that remains now is to identify the 1-element chain, and correctly join up the ends of the chains in M^- with the starts of the chains in M^+. Consider a directed edge e from p to q in G, corresponding to $[p, n^-] \prec [q, n^+]$ from N^- to N^+ in $P \times \mathbf{rk(P)}$. Then we create the chain $[p, m^-] \prec [f(e), m] \prec [q, m^+]$ from M^- to M^+ in $P \times (\mathbf{rk(P)+1})$.

Finally, there is a unique vertex v in G which no edge matches to. We create the 1-element chain $[v, m]$ in $\mathcal{S}'$. All of these chains do not intersect, as f is an injection, and f misses v. Also, all chains in $\mathcal{S}'$ are symmetric.

Conversely, suppose we have a symmetric chain decomposition $\mathcal{S}'$ of $P \times$ **(rk(P)+1)** which maps to $\mathcal{S}$, and create the matching between the edges of G and their endpoints as follows. Given a directed edge e from p to q in G corresponding to $[p, n^-] \prec [q, n^+]$ in $P \times \mathbf{rk(P)}$, consider the chain in $\mathcal{S}'$ which connects $[p, m^-]$ to $[q, m^+]$ in $P \times (\mathbf{rk(P)+1})$. We define $f(e)$ so that $[f(e), m]$ is the intermediate point on this chain. Clearly this is a matching, as the chains between M^- and M^+ are disjoint, and $\mathrm{rk}(q)$ is at most 1 higher than $\mathrm{rk}(p)$ so $f(e) = p$ or q.

These two maps are inverses of each other, proving the equivalence. As there are $\mathrm{rk}(P) + 1$ matchings on G (coming from the $\mathrm{rk}(P) + 1$ possible matchings on the edges of the long chain in G), the conclusion follows.

Finally, suppose we have a decomposition of $P \times$ **(rk(P)+1)** with no taut chain, and P has at least 2 elements of rank $\mathrm{rk}(P) - 1$ connected to $\max_P$. A taut chain in $P \times \mathbf{rk(P)}$ is created in exactly the following cases. Either the maximal chain in $P \times$ **(rk(P)+1)** has a subchain of the form $(\min_P, 0) \prec (\min_P, 1) \prec \ldots \prec (\min_P, \mathrm{rk}(P) - 1)$, or a subchain of the form $(\max_P, 1) \prec (\max_P, 2) \prec \ldots \prec (\max_P, \mathrm{rk}(P))$. Disconnect $(\min_P, 0)$ and $(\max_P, \mathrm{rk}(P))$ from the maximal chain. Connect $(\max_P, \mathrm{rk}(P))$ to an adjacent element with second coordinate also $\mathrm{rk}(P)$ which does not belong to the chain containing $(\min_P, \mathrm{rk}(P) - 1)$ (this is possible as there are at least 2 choices by the hypothesis on P), and add a connection from $(\min_P, 0)$ to the chain which $(\max_P, \mathrm{rk}(P))$ now belongs to. This new configuration of symmetric chains now avoids the two cases which would cause taut chains to appear in $P \times \mathbf{rk(P)}$, without creating any taut chains in $P \times \mathbf{rk(P)+1}$. This finishes the proof.

$\square$

1			110000	111000	111100	111110		
2			011000	011100	011110	011111		
3			001100	001110	101110	101111		
4			000110	100110	110110	110111		
5			100010	110010	111010	111011		
6	000000	100000	101000	101001	111001	111101	111111	111112
7		010000	010100	010101	011101	011102	011112	
8		001000	001010	001011	001111	001112	101112	
9		000100	100100	100101	100111	100112	110111	
10		000010	010010	010011	110011	110012	111012	
11		000001	100001	110001	110002	110102	111102	
12			010001	011001	011002	011012		
13			001001	001101	001102	101102		
14			000101	000111	000112	010112		
15			000011	100011	100012	101012		
16			000002	100002	101002	111002		
17				010002	010102			
18				001002	001012			
19				000102	100102			
20				000012	010012			
21				110100	110101			
22				011010	011011			
23				101100	101101			
24				010110	010111			
25				101010	101011			

TABLE 1. Symmetric chain decomposition of $P(5,3)$ with no taut chains

4. Proof of Theorem 2.1

By Corollary 3.5, we only have left to construct symmetric chain decompositions of $P(k,n)$ for $k \geq 5$, and $n \geq 3$ with no taut chains. In the tables below, we give decompositions with no taut chains for $k = 5$, $n = 3, 4, 5$. Theorem 2.3 then yields such a decomposition for $k = 5$, $n = 6$, and Theorem 2.2 then yields such a decomposition for $k = 5$ and all $n \geq 3$. Finally from this, Corollary 3.3 can then be used to get such decompositions for all $k \geq 5$ and $n \geq 3$.

In the tables below, the rows give the symmetric chains in $Q_5 \times \mathbf{n}$, written in coordinates. Aiding in the finding of the decompositions below were the packet descriptions, and the natural $\mathbb{Z}/5\mathbb{Z}$ action on the points of $Q_5 \times \mathbf{n}$.

1	000000	100000	101000	101100	101101	101102	111102	111103	111113
2		010000	010100	010110	010111	010112	011112	011113	
3		001000	001010	101010	101011	101012	101112	101113	
4		000100	100100	110100	110101	110102	110112	110113	
5		000010	010010	011010	011011	011012	111012	111013	
6			011000	011001	011002	011102	011103		
7			001100	001101	001102	001112	001113		
8			000110	000111	000112	100112	100113		
9			100010	100011	100012	110012	110013		
10				111000	111100	111110			
11				011100	011110	011111			
12				001110	101110	101111			
13				100110	110110	110111			
14				110010	111010	111011			
15		000001	100001	101001	111001	111101	111111	111112	
16			010001	010101	010102	010103	110103		
17			001001	001011	001012	001013	011013		
18			000101	100101	100102	100103	101103		
19			000011	010011	010012	010013	010113		
20			000002	100002	101002	101003	101013		
21				010002	010003	011003			
22				001002	001003	001103			
23				000102	000103	000113			
24				000012	000013	100013			
25				000003	100003	110003			
26					011101				
27					001111				
28					100111				
29					110011				

TABLE 2. Symmetric chain decomposition of $P(5,4)$ with no taut chains

From Table 1, Table 2, and Table 3, the proof of Theorem 2.1 is complete.

As Theorem 2.1 completely solves the question for $Q_k \times \mathbf{n}$, one direction of further study would be to investigate other natural families of posets in a similar way using Theorem 2.2 and Theorem 2.3.

1		000001	000002	000003	000004	000014	100014	110014	110114	
2			010001	010002	010003	010004	010014	010114		
3			100001	100002	100003	100004	110004	111004		
4		010000	110000	110001	110002	110003	110103	110104	111104	
5			001001	001002	001003	001004	001104	101104		
6				011001	011002	011003	011004			
7				101001	101002	101003	101004			
8			101000	111000	111001	111002	111003	111013		
9			000101	000102	000103	000104	100104	100114		
10				010101	010102	010103	010104			
11			100100	100101	100102	100103	100113			
12		000100	001100	001101	011101	011102	011103	011104	011114	
13				101100	101101	101102	101103			
14				110100	111100	111101	111102			
15			000011	000012	000013	001013	001014	101014		
16				100011	100012	100013	101013			
17				110010	110011	110012	110013			
18	000000	001000	001010	001011	001012	001013	001113	001114	101114	111114
19		000010	010010	010011	011011	011012	011013	011014	111014	
20			011000	011010	111010	111011	111012	111013		
21				000111	000112	000113	000114			
22			010100	010110	010111	010112	010113	011113		
23			000110	100110	110110	110111	110112	110113		
24				011100	011110	011111	011112			
25				001110	001111	101111	101112			
26		100000	100010	101010	101110	111110	111111	111112	111113	
27					110101	110102				
28					100111	100112				
29					101011	101012				
30					010012	010013				
31					001102	001103				

TABLE 3. Symmetric chain decomposition of $P(5,5)$ with no taut chains

CHAPTER 4

Metastable States in the Ising Model

1. Introduction

This chapter is joint work with Hunter Spink and Marius Tiba.

The Ising model has received serious attention in the literature of statistical mechanics, for example see Lee and Young [26], Glauber [16] and Kazakov [23]. Also see the recent work of Addario-Berry [1]. For a comprehensive description of the model and closely related subjects, see the book of McCoy and Wu [29].

In this chapter, following the Ising model, we consider a collection V of interacting particles which are arranged in an underlying dependency graph G with vertex set V. Each particle $v \in V$ has a magnetic spin $\sigma_v \in \{\pm 1\}$ and it can interact only with its neighbours in G according to certain rules.

In [36], Spink established a conjecture attributed to Holzman (see [32]), that any quadratic function on the cube Q_n has at most $\binom{n}{\lfloor n/2 \rfloor}$ local maxima. This problem classically corresponds to the Sherrington-Kirkpartick model from mathematical physics, the particular case of the Ising model for which the dependency graph G is complete, often used as a toy example in the theory of spin glasses due to its simpler properties. The aim of this chapter is to generalise the work done in [36], thus capturing what we believe to be the combinatorial essence of the Ising model.

We start by describing our *general model*, while pointing out how it relates to the Ising model. In the general model, we also consider a collection V of interacting particles which are arranged in an underlying dependency graph G with vertex set V and we assume that each particle v has a spin $\sigma_v \in \{\pm 1\}$. A particle can be either stable or unstable and the system is said to be *stable* if each of the individual particles is stable. To accurately capture the behaviour of the Ising model, we further assume that for each neighbour w of v in G we have set either a *ferromagnetic* (positive) or *antiferromagnetic* (negative) correlation between their spins, given by $c_{vw} \in \{\pm 1\}$. Physical considerations force $c_{vw} = c_{wv}$, so we also impose this restriction in our model. The stability of v is governed by the following axioms.

(1) Given all c_{vw}'s, the stability of v only depends on σ_v and σ_w for $w \in \Gamma(v)$.
(2) For any given state of $\Gamma(v)$, there is at most one choice of σ_v which makes v stable.
(3) If v is stable and $w \in \Gamma(v)$, then v remains stable if flipping σ_w increases $c_{vw}\sigma_v\sigma_w$.
(4) If v is unstable and $w \in \Gamma(v)$, then v remains unstable if flipping σ_w decreases $c_{vw}\sigma_v\sigma_w$.

We now point out how our general model is designed to address an extremal problem concerning the maximal number of metastable states in the Ising model, under various restrictions. To do so, we recall from Chapter 1 that the energy of a given a spin configuration (or *state*) $(\sigma_v)_{v\in V} \in \{\pm 1\}^V$ in the Ising model is given by the Hamiltonian $H = -\sum J_{ij}\sigma_i\sigma_j - \sum h_i\sigma_i$, where the J_{ij} are typically Gaussian random variables with $J_{ij} = 0$ if ij is not an edge of G, and h_i are constants corresponding to an external magnetic field. When the Hamiltonian is locally maximized, in the sense that for any v, negating σ_v strictly decreases the Hamiltonian, the system is called metastable or that the state is *metastable*. We note that if we take the correlations c_{vw} in our general model to be the opposite sign of J_{vw} in the Ising model with magnetic field, then the metastable states in the Ising model satisfy the above stability axioms for each $v \in V$. When we refer to the Ising model, we do so from the perspective of our extremal problem: given a dependency graph G, the interactions J_{ij} and the external magnetic field h_i, we seek the maximal number of metastable states. Being consistent with the standard nomenclature from physics, we also call the instance of the Ising model for which G is a complete graph the *Sherrington-Kirkpatrick model*.

We also note a slight asymmetry between stability and non-stability, arising from degenerate situations where for example the Hamiltonian of the Ising model at a given state is not influenced by flipping σ_v — this subtlety never arises in our maximisation problems, and whenever convenient, we can safely assume there is always exactly one choice of σ_v in the second axiom, in which case the third and fourth axioms are identical. In other words, we say that if the Hamiltonian is not changed by flipping σ_v then v is unstable.

To translate our model to the language of extremal set theory, we identify our vertex set V with $[n] = \{1, 2, \ldots, n\}$, and our states with subsets of $[n]$ via $(\sigma_v)_{v=1}^n \mapsto \{v \in [n] \mid \sigma_v = 1\}$. The dependency graph G is then a graph on $[n]$, and the correlations c_{vw} can be thought of as a two-colouring of the edges of G by $\{\pm 1\}$. For any v we let $\Gamma'(v) \subseteq \Gamma(v)$ be the set of neighbours of v with $c_{vw} = 1$. We further let $\mathcal{F}_v \subseteq \mathcal{P}(n)$ be all the states which stabilize v, for every v so then the set of all stable states is given by $\mathcal{F} = \cap_{v \in [n]} \mathcal{F}_v$.

Our stability axiomatization for v implies that for $A, B \in \mathcal{F}_v$ we cannot have $v \in A \setminus B$ and that

$$(A \cap \Gamma(v)) \Delta \Gamma'(v) \supseteq (B \cap \Gamma(v)) \Delta \Gamma'(v).$$

As we note below, we will refer to this as the *stability condition at* v. Indeed, if $v \in A$, then if the spins in $\Gamma(v)$ makes v stable, then v cannot destabilize if we flip some of σ_w with $w \in \Gamma(v)$ in places where $A \cap \Gamma(v)$ disagrees with $\Gamma'(v)$ (which we could do to destabilize A to $B \cup \{v\}$ restricted at $v \cup \Gamma(v)$, obtaining a contradiction). Conversely, if we only impose this condition then all but the first stability axiom for v necessarily hold. We can make the first condition hold by adding into $\mathcal{F}_v$ all states with the same spins in $v \cup \Gamma(v)$ as a state in $\mathcal{F}_v$, and since we are only concerned with the maximal size of $\mathcal{F}$, we can (and do) ignore this issue, and just consider the combinatorial condition above on each $\mathcal{F}_v$.

To that end, we call the above condition on $\mathcal{F}_v$ the *stability condition at* v. We say that a collection $\mathcal{E} \subseteq \mathcal{P}(n)$ is *admissible* for a given G and c_{vw} coloring if $\mathcal{E} \subseteq \bigcap_{v=1}^{n} \mathcal{F}_v$. We are investigating the maximal cardinality of $\mathcal{E}$ over all admissible collections. It is in this form that we study the extremal properties of the general model and we bound from above the number of stable states. As the general model is potentially less restrictive than the models that it generalizes, including the Ising model, any upper bound on the size of an admissible collection in the general model implies upper bounds for these other models. We emphasise that in the general model we only consider the combinatorial stability condition, as it turns out that we do not lose much with this generalisation, when considering the extremal problem for the Ising model.

Finally, of particular combinatorial interest, we will also prove a broad generalization of the LYMB inequality of Yamanoto [43], Meshalkin [30], Bollobás [7] and Lubell [28] for set systems, which arises naturally when studying the general model. The LYMB inequality states that if $\mathcal{A} \subseteq \mathcal{P}(n)$ is an antichain, then $\sum_r \frac{|\mathcal{A} \cap [n]^{(r)}|}{\binom{n}{r}} \leq 1$ and it is the particular case of our Theorem 3.2 for G being the graph with one vertex.

This chapter is divided into six sections. In Section 2 we make some preliminary observations and give a few definitions; in Section 3 we state our main results; in Section 4 we prove our results concerning either purely ferromagnetic or purely antiferromagnetic interactions; in Section 5 we prove our remaining results on the Sherrington-Kirkpatrick model (i.e. complete dependency graphs) and on bipartite dependency graphs; in Section 6 we discuss directions of further research.

2. Preliminary Observations and Definitions

In this section we continue the discussion started in the introduction by presenting a list of preliminary observations that will be used throughout the chapter, which also build on our intuition about the general model. Finally, we recall a few classical definitions needed to state our main results.

1. If for each $v \in [n]$, $\mathcal{F}_v$ satisfies the stability condition, then the stability conditions for each vertex are still satisfied if we replace $\mathcal{F}_v$ with $\mathcal{F} = \cap_{v\in[n]}\mathcal{F}_v$ for each v.

This allows us to consider, instead of n families of sets $\mathcal{F}_v$ each with a stability condition, the family $\mathcal{F}$ with a stability condition for each $v \in [n]$.

2. If G is the complete graph on $[n]$, and all $c_{vw} = -1$, then a collection of states $\mathcal{E}$ is admissible if and only if it is an antichain in $\mathcal{P}(n)$.

Indeed, let us carefully consider the stability condition associated to $v \in [n]$. As all $c_{vw} = -1$, the stability condition is that it is not possible that for $A, B \in \mathcal{F}_v$, we have $v \in A \setminus B$, and $A \supseteq B$. For distinct $A, B \in \mathcal{E}$, the conditions together imply that we cannot have $A \supseteq B$. Indeed, the stability condition fails for any $v \in A \setminus B$. Conversely, for a stability condition to fail for distinct $A, B \in \mathcal{E}$, we must have A and B comparable.

3. For a given graph G on $[n]$ and a subset $A \subseteq [n]$, there is a canonical size preserving bijection between collections of stable states for a given $\{\pm 1\}$-colouring c_{vw} and stable states for the $\{\pm 1\}$-colouring $\tilde{c}_{vw}$, where $\tilde{c}_{vw} = c_{vw}$ if either both or none of v, w lie in A, and $\tilde{c}_{vw} = -c_{vw}$ otherwise.

This bijection is given by taking each state X, and mapping it to the symmetric set difference $X\Delta A$, corresponding to the hypercube automorphism of flipping each of the states in A.

4. Replacing G with a (coloured) subgraph cannot increase the maximum number of stable states.

This is clear, as the stability conditions become stronger as some forbidden inclusions of sets in $\mathcal{P}(n)$ become forbidden inclusions of sets restricted to a proper subset of $[n]$. Hence, for example, a bound for $G = K_n$ is an upper bound for any G.

At the end of this section, we set the following notation for the rest of the chapter. We let $A^{(r)}$ denote the set of all r-element subsets of A. Given a graph G, $\alpha(G)$ is the independence number of G. Given a digraph G, $\tilde{\alpha}(G)$ is the largest cardinality of a set of vertices $V \subseteq G$ such that for all $v, w \in V$ we do not have both $v \to w$ and $w \to v$.

Given a graph G and an integer n, a *Sperner* (G, n)*-family* $\{\mathcal{F}_x\}$ is a collection of antichains in $\mathcal{P}(n)$ indexed by vertices $x \in G$ such that if xy is an edge of G, then the elements of $\mathcal{F}_x$ and $\mathcal{F}_y$ are incomparable.

3. Our Results

We start with the following extremal result.

THEOREM 3.1. *Given a Sperner* (G, n)*-family* $\{\mathcal{F}_x\}$, $\sum_x |\mathcal{F}_x| \leq \alpha(G)\binom{n}{n/2}$, *and the upper bound is tight.*

Theorem 3.1 comes as a consequence of the following stronger result, as $\binom{n}{n/2}$ is the largest binomial coefficient amongst $\binom{n}{0}, \binom{n}{1}, \ldots, \binom{n}{n}$.

THEOREM 3.2. *Given a Sperner* (G, n)*-family* $\{\mathcal{F}_x\}$, $\sum_x \sum_r \frac{|\mathcal{F}_x \cap [n]^{(r)}|}{\binom{n}{r}} \leq \alpha(G)$.

This is a generalisation of the well-known LYMB inequality. We use these theorems to obtain the following result on solely antiferromagnetic interactions.

COROLLARY 3.3. *Given a graph G on $[n]$ and a collection of vertices $A \subset [n]$, suppose that $c_{vw} = -1$ whenever vw is an edge of G. Consider the digraph $\Lambda_{G,A}$ with vertex set $\mathcal{P}(V(G) - A)$ given by joining B to C if there exists $i \in B \setminus C$ such that $\Gamma(i) \cap (C \setminus B) = \varnothing$. Then the size of an admissible collection of states for G is bounded above by $\tilde{\alpha}(\Lambda_{G,A})\binom{|A|}{\lfloor |A|/2 \rfloor}$.*

When G is the complement of a k-clique, we show that the lower bound given by the above theorem when A is the set of vertices not in the k-clique is in fact extremal. Our proof is reminiscent of Sak's proof [34] of the maximal number of elements in a set system with no k-chain.

THEOREM 3.4. *Let G be the complement of a k-clique inside K_n and let us assume that all $c_{vw} = -1$. Then the number of stable states is at most the sum* *of the middle $k+1$ binomial coefficients $\binom{n-k}{r}$, if all exist, or 2^{n-k} otherwise. Furthermore, when n is even, there exists an instance of the Ising model with external magnetic field which attains this bound with all $c_{vw} = -1$, and an instance without an external magnetic field if we do not require all $c_{vw} = -1$.*

When all of the interactions are ferromagnetic, we derive the following as a consequence of of results of Leader and Long [25] on set systems with forbidden differences of size 1.

THEOREM 3.5. *Let G be a complete graph, and let us assume that all $c_{vw} = 1$. Then the maximal size of an admissible collection is at most $(2+o(1))\frac{1}{n}\binom{n}{\lfloor n/2 \rfloor}$, and there is an admissible collection of size within a factor of 2 of this upper bound.*

Next, we study the Sherrington-Kirkpatrick model, which is the particular case of the Ising model for which G is a complete graph, with or without external magnetic fields. The first part of Theorem 3.6 was shown by Spink in [36], and our contribution is to solve the analogous problem in the absence of a magnetic field, which imposes an extra combinatorial condition (symmetry) on our general model.

THEOREM 3.6. *In the Sherrington-Kirkpatrick model, the maximal number of metastable states is $\binom{n}{\lfloor n/2 \rfloor}$ in the presence of an external magnetic field. In the absence of such a field, the maximum is still $\binom{n}{n/2}$ for n even, but decreases to $2\binom{n-1}{(n-3)/2}$ when n is odd.*

Next, we derive the following via an application of König's Theorem.

THEOREM 3.7. *Let G be a bipartite graph. Then we have an upper bound of $2^{n-\alpha(G)}$ on the number of metastable states, and we can attain equality in the Ising model with no external magnetic field with any choice of signs for the J_{ij}'s.*

The following immediate corollary of Theorem 3.7 pertains to the case of most interest to lattice models. It has been noted previously for subsets of lattices where there is a clear perfect matching on the edges, though the general result seems to have escaped notice.

COROLLARY 3.8. *Given a graph G formed by an induced subgraph of $\mathbb{Z}^d$ of size n, we have a sharp upper bound of $2^{n-\alpha(G)}$ on the number of metastable states in the general model and the Ising model with no external magnetic field with any choice of signs for the J_{ij}'s.*

4. Ferromagnetic and Antiferromagnetic Spin Models

We begin with our proof of Theorem 3.2.

PROOF OF THEOREM 3.2. We apply a compression argument, reducing the problem to the case when all $\mathcal{F}_x$ contain only subsets of size $\lfloor n/2 \rfloor$. The result is then clear, as each element of size $\lfloor n/2 \rfloor$ can lie in at most $\alpha(G)$ of the $\mathcal{F}_x$.

Preserving the (G, n)-Sperner condition forces us to modify a naive compression however. To that end, we show how to do upper compression first for the layers below $\lfloor n/2 \rfloor$, and lower compression for the layers above works identically. Let us suppose that no $\mathcal{F}_x$ contains a subset of size strictly less than r for some $r < \lfloor n/2 \rfloor$.

Given an element S in the upper shadow of $(\bigcup \mathcal{F}_x) \cap [n]^{(r)}$, let R_S be the set of vertices $x \in G$ such that the upper shadow of $\mathcal{F}_x$ contains S, and let I_S be an independent set of G lying inside R_S of largest size. Our compression procedure will be to add S to all $\mathcal{F}_x$ with $x \in I_S$ for each such S, and then remove all subsets of size r from all $\mathcal{F}_x$.

Firstly, it is clear that this preserves the antichain condition within each $\mathcal{F}_x$. It is also clear that this preserves the incomparability condition in edges corresponding to G, as in the union of the two $\mathcal{F}_x$ corresponding to an edge we have performed a partial upper shadow.

It remains to be seen that we have not decreased $\sum_x \sum_r \frac{|\mathcal{F}_x \cap [n]^{(r)}|}{\binom{n}{r}}$. Indeed, we first note that $(n-r)\sum_x |\mathcal{F}_x \cap [n]^{(r)}|$ counts the number of elements in the upper shadows of all the $\mathcal{F}_x \cap [n]^{(r)}$ with multiplicities. For each S in the upper shadow, we claim $(r+1)|I_S|$ is larger than the multiplicity of S in this sum. Indeed, there are $r+1$ elements in the lower shadow of S, and a given $T \in [n]^{(r)}$ in the lower shadow of S appears in at most $|I_S|$ of the $\mathcal{F}_x$ since the x with $T \in \mathcal{F}_x$ forms an independent set in R_S.

Thus, $\sum_S \frac{|I_S|}{\binom{n}{r+1}} \geq \sum_x \frac{|\mathcal{F}_x \cap [n]^{(r)}|}{\binom{n}{r}}$ as desired. □

We now address Corollary 3.3.

Proof of Corollary 3.3. We identify $\mathcal{P}(n)$ with $\mathcal{P}(A) \times \mathcal{P}([n]-A)$.

Consider the digraph $\Lambda_{G,A}$ as defined in Corollary 3.3 and associate with each vertex $v \in \mathcal{P}([n]-A)$ a family $\mathcal{F}_v \subseteq \mathcal{P}(A)$ with the convention that the admissible collections are given by $\bigcup_v \{B \cup v \mid B \in \mathcal{F}_v\}$. Firstly, we must have that each $\mathcal{F}_v$ is an antichain by Observation **2**. Moreover, by the definition of $\Lambda_{G,A}$, whenever there is an oriented arrow from $v \in \mathcal{P}([n]-A)$ to $w \in \mathcal{P}([n]-A)$, the stability condition forbids to have an element of $\mathcal{F}_w$ as a subset of an element of $\mathcal{F}_v$. Now the upper bound follows from a direct application of Theorem 3.1. □

The proof of the upper bound in Theorem 3.4 which we give below is a nice application of classical finite set system results. It would be interesting for us to know how close to the upper bound we can actually get in the absence of an external magnetic field (as the underlying dependency graph of our construction does not have all of its interactions antiferromagnetic).

Proof of Theorem 3.4. Recall that we seek an upper bound when all $c_{vw} = -1$, and the dependency graph G is the complement of a k-clique. Let B denote the set of vertices in the k-clique. Consider the poset map $\mathcal{P}(n) \cong 2^{[n]-B} \times 2^B \to 2^{[n]-B} \times [k+1]$, where the final map is simply the rank mapping. We will show that the images of two distinct stable states

$X, Y \in \mathcal{P}(n)$ are mapped to incomparable elements of $2^{[n]-B} \times [k+1]$. This yields the desired upper bound, as it is equal to the size of the largest antichain in this poset.

Indeed, suppose that for two stable states X, Y, the image of X contains the image of Y. Then $|X \cap B| \geq |Y \cap B|$ and $X \cap ([n] - B) \supseteq Y \cap ([n] - B)$.

Suppose first that $X \cap B \neq Y \cap B$. Then there exists $v \in B$ with $v \in X \setminus Y$ by the inequality. But then $X \cap ([n] - B) \supseteq Y \cap ([n] - B)$, directly contradicts the stability condition at v.

Suppose instead that $X \cap B = Y \cap B$. Then X strictly contains Y. Taking $v \in X \setminus Y \subseteq [n] - B$, we get that $X \supseteq Y$ directly contradicts the stability condition at v.

Hence, we get the desired upper bound.

For n even, our antiferromagnetic Ising model construction proceeds as follows. We take $J_{ij} = -1$ for all i, j with ij an edge of G, and on B (which we can assume to be $[k]$), we let $h_i = 0$ for $i \notin B$, and $h_i = k + 1 - 2i$ for $i \in B$.

When there is no magnetic field and n is even, the following construction yields the correct number of states, but fails to have all interactions antiferromagnetic. We let $B = [k]$, and take $J_{ij} = -1$ for all ij an edge of G except those of the form ik, to which we assign $J_{ik} = k + 2 - 2i$. $\square$

We now derive Theorem 3.5 using the results of Leader and Long [25].

Proof of Theorem 3.5. The conditions on admissibility are precisely the same as saying that we do not have $|A \setminus B| = 1$ for any $A, B \in \mathcal{F}$, and maximizing a family $\mathcal{F}$ with this condition is precisely the problem addressed in [25] with the conclusions as stated in Theorem 3.5.

Indeed, the admissibility condition says that we never have $v \in A \setminus B$ with $A \setminus \{v\} \subseteq B$. $\square$

5. Sherrington-Kirkpatrick and Bipartite Dependencies

We start by proving Theorem 3.6.

Proof of Theorem 3.6. The upper bounds have been shown in [36] in the cases that there is an external magnetic field (giving $\binom{n}{\lfloor n/2 \rfloor}$), and if there is at least one ferromagnetic (or independent by a continuity argument) edge (giving $2\binom{n-1}{(n-3)/2}$). The only remaining case is when all $c_{vw} = -1$, and there are no external magnetic fields. Since all $c_{vw} = -1$, the stability conditions are equivalent to saying that $\mathcal{F}$ is an antichain in $\mathcal{P}(n)$. The additional information we get when there is no external magnetic field is that if a state A is stable or unstable, then the same is true of the complement A^c. Indeed, this follows from the fact that the Hamiltonian is invariant under negating all σ_v. Hence, we have an antichain $\mathcal{F}$ with all elements distinct from $\mathcal{F}^c$. This is classically known to bound from above the size of $\mathcal{F}$ in the case n is odd by $2\binom{n-1}{(n-3)/2}$ by a local LYMB compression argument extremely similar to the proof of Theorem 4.2 (the additional input needed is that once one has compressed to the two middle layers, the Erdős-Ko-Rado Theorem [14] that the maximal size of an intersecting family in $[n]^{(r)}$ is $\binom{n-1}{r-1}$ for $r \leq n/2$ implies the result).

To attain all of these bounds in the Sherrington-Kirkpatrick model, when n is even we take $J_{ij} = -1$ and $h_i = 0$. When n is odd and we are allowed an external magnetic field, we take $J_{ij} = -1$ and $h_i = \epsilon$. When n is odd and we have no external magnetic field, we take $J_{ij} = -1$ if $2 \leq i < j$, and $J_{1j} = 1 + \epsilon$. $\square$

We now show the sharp upper bound Theorem 3.7 for bipartite dependency graphs. The key observation is that there are trivial upper and lower bounds, which by König's Theorem are in fact equal.

Proof of Theorem 3.7. If we specify the spins on the complement of an independent set, then by the first stability axiom, there is at most one choice of spins on the independent set to make the resulting system stable. Hence, we have an upper bound on the number of stable states as $2^{n-\alpha(G)}$.

For the lower bound, we note that if we have a matching of size m in G, then in the Ising model we can make all J_{ij} non-zero on this matching and 0 outside the matching, then each edge of the matching will contribute a factor of 2 to the number of stable states, giving 2^m stable states. Letting $m(G)$ be the maximal size of a matching in G, we get a lower bound of $2^{m(G)}$. As $\alpha(G) + m(G) = n$ by König's Theorem, the result follows. $\square$

6. Further Directions of Research

It would be extremely interesting to extend the antiferromagnetic results to the general model. We make the following bold conjecture.

Conjecture 6.1. *The upper bound in Corollary 3.3 continues to hold even without the assumption that all interactions are antiferromagnetic.*

Of course, the most general open question about our model would be as follows, though perhaps it is too ambitious as phrased.

Question 6.2. *For each of the $\{\pm 1\}$-coloured graphs on n vertices, what is the maximal number of stable states?*

It appears to us that triples of vertices which up to the automorphism mentioned in Section 2 are pairwise ferromagnetic can only hinder the number of strict local maxima. This was concretely noted in [36], that to get the number of stable states to be within a $(1 - \frac{1}{n})$ fraction of $\binom{n}{n/2}$, we require all interactions to be strictly anti-ferromagnetic up to the automorphism mentioned in Section 2. An interesting refinement of the previous question is as follows.

Question 6.3. *What can we say about the maximal number of stable states in the Sherrington-Kirkpatrick model as the number of triples with Ferromagnetic interactions (up to automorphisms) increases?*

In the extreme case, we saw that when all interactions were ferromagnetic, we had an upper bound on the size of an admissible family that was $\frac{2}{n}$ of what it was when we had all interactions anti-ferromagnetic.

QUESTION 6.4. *How many stable states can there be in the ferromagnetic Ising model? Clearly there is a lower bound of* $2^{\lfloor n/2 \rfloor}$ *given by the construction in Theorem 3.5. How close is this to optimal?*

Finally, we can consider lattices other than $\mathbb{Z}^d$, in particular, ones whose dependency graphs are not bipartite.

QUESTION 6.5. *What is the largest admissible collection when* G *corresponds to the triangular lattice? What about other lattices?*

We have the following conjecture regarding the triangular lattice.

CONJECTURE 6.6. *For the triangular lattice, the maximal number of stable states is* $3^{n/3}$*, given by partitioning the lattice into triangles.*

CHAPTER 5

Fast Bootstrap Percolation on the Grid

1. Introduction

This chapter is joint work with Scott Binski.

Although its origins are in physics, in the context of particle systems, bootstrap percolation is one of the most studied instances of cellular automata in combinatorics. Let G be a graph, whose vertices are referred to as *sites*. Translated to the language of combinatorics, r-neighbour bootstrap percolation on G can be regarded as the following infection process: we start with a set $S_1 \subseteq V(G)$ whose elements we call *initially infected sites (i.i.s.)*, leaving the remaining sites initially healthy. Then, the healthy sites of G get infected in rounds, provided they have at least r infected neighbours. Formally, for $t = 2, 3, \ldots$, as long as $S'_{t-1} = \bigcup_{i=1}^{t-1} S_i \neq V(G)$, define

$$S_t := \{v \in V(G) \mid \ |\Gamma(v) \cap S'_{t-1}| \geq r\} \setminus S'_{t-1}$$

to be the set of sites that get infected at time t. Further define the *percolation time* k to be infinity if $\bigcup_{i=1}^{\infty} S_i \neq V(G)$ and otherwise define k to be the largest positive integer such that $S_k \neq \varnothing$. In the later case, we say that S_1 *percolates*.

Almost all results in the literature concern probabilistic bootstrap percolation, for example see Aizenman and Lebowitz [4], Balogh, Bollobás, Duminil-Copin and Morris [5], Cerf and Manzo [11], Gravner, Holroyd and Morris [17], and Holroyd [20]. In the deterministic world, Benevides and Przykucki addressed a question of Bollobás and studied the maximum percolation time of the 2 neighbour bootstrap process on the $n \times n$ square grid, showing in [6] that it is $\frac{13}{18}n^2 + O(n)$. In this chapter we address another related question

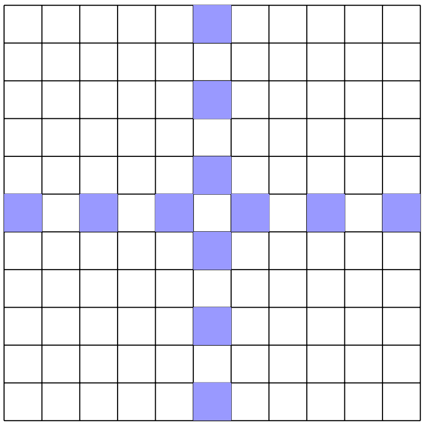

FIGURE 1. The initial configuration that percolates in time $k = n - 1$ for $n \equiv 3 \pmod 4$, $s_1 = n + 1$

of Bollobás about the minimal percolation time of the 2 neighbour bootstrap process on the $n \times n$ square grid, as a function of the number of i.i.s., $|S_1|$.

For the rest of the chapter, unless stated otherwise, we only consider the 2 neighbour bootstrap process on the $n \times n$ grid model and sets of i.i.s. S_1 that percolate. For $t \geq 1$, we define $s_t := |S_t|$ to be the number of sites infected precisely in round t and we emphasise that, according to our definition, S_t's are disjoint sets.

Our first result addresses the case when $|S_1|$ is small. It is well known that the minimal size of a percolating set S_1 is n, and in Theorem 1.1 we give the exact minimum percolation time when $s_1 = n$ and $s_1 = n + 1$.

THEOREM 1.1. *In the* 2*-neighbour bootstrap percolation model on the square grid $n \times n$, the minimal percolation times k for $s_1 = n, n + 1$ are as follows:*

(1) for $s_1 = n$, we have $k = n$;

(2) for $s_1 = n + 1$, we have $k = n$, except for the case $n \equiv 3 \pmod 4$, when $k = n - 1$ and there exists only one possible S_1 that percolates in time $n - 1$ (see Figure 1):

$$S_1 = \{(\frac{n+1}{2}, 2i - 1) \mid 1 \leq i \leq \frac{n+1}{2}\} \cup \{(2i - 1, \frac{n+1}{2}) \mid 1 \leq i \leq \frac{n+1}{2}\}.$$

We then move to the general case and provide a construction which establishes a general upper bound for the minimum number of i.i.s. needed to percolate in time k. Thus, we improve the naive bound $n + \frac{n(n-k)}{k}$ obtained by the natural construction of taking S_1 to be the union of evenly spaced diagonals.

THEOREM 1.2. *Let n, k be given, $k < n$. Then there exists a set $S_1 \subseteq n \times n$ of i.i.s. which percolates in time k of size:*

- $s_1 = 3n - 2k$, *for* $n \leq 2k$;
- $s_1 \leq n + \frac{(n+k-2)(n+k+1)}{2k-1}$, *for* $n > 2k$.

Finally, we present a short argument which establishes a general lower bound.

THEOREM 1.3. *The size of every set S_1 of i.i.s. that percolates on the $n \times n$ grid in time k satisfies:*

- $s_1 \geq \frac{n^2}{3} + \frac{2n}{9} - \frac{2}{9}$, *for* $k = 2$;
- $s_1 \geq \frac{n^2 + (\sqrt{4k-1}-2)(n-1)}{2k-1}$, *for* $k \geq 3$.

Although our bounds are close for k fixed and n large, we believe (especially for $k = n - o(1)$) that the lower bound can be improved and that the actual value is very close to our upper bound in Theorem 1.2.

This chapter is divided into five sections. In Section 2, Section 3, and Section 4 we prove Theorem 1.1, Theorem 1.2, and Theorem 1.3, respectively. In Section 5 we give our conclusions and discuss further directions of research.

2. Small initial configurations

In this section we give a combinatorial argument, establishing Theorem 1.1. To this end, we introduce the "semiperimeter" function, which arises naturally when establishing the fact that there are no percolating sets of sizes smaller than n. For each $t = 1, \ldots, k$ we define f_t to be twice the total number of

C	N	N	N	N	N	C
W	I	I	I	I	I	E
W	I	I	I	I	I	E
W	I	I	I	I	I	E
W	I	I	I	I	I	E
W	I	I	I	I	I	E
C	S	S	S	S	S	C

FIGURE 2. Regions of the grid. The corners are marked with 'C', the interior sites are marked 'I', the sites on the north boundary are marked 'N', etc. The centre is coloured blue.

infected sites at time t minus the number of sets of adjacent infected sites:

$$f_t := 2|\bigcup_{i=1}^{t} S_i| - |\{\{v_1, v_2\} \mid v_1, v_2 \in \bigcup_{i=1}^{t} S_i \text{ and } v_1 \in \Gamma_{[n]^2}(v_2)\}|.$$

Considering that for $t < k$, every site in S_{t+1} has 2 or more neighbours in $\bigcup_{i=1}^{t} S_i$, we immediately conclude that $(f_t)_{t=1}^{k}$ is non-increasing. A trivial calculation gives $f_k = 2n$, hence $f_1 \geq 2n$ which shows that $s_1 \geq n$. Hence, a percolating set S_1 has size at least n.

Before proving Theorem 1.1 we need to border some regions of the $n \times n$ grid, $n \geq 7$ (see Figure 2): a *corner* is any of the four points $(1,1), (1,n), (n,1), (n,n)$; the *north boundary* is the set of $n-2$ points $\{(i,n) \mid 2 \leq i \leq n-1\}$ and similarly we define the *east, south* and *west boundary*; the *boundary* is the union of the north, east, south and west boundary; the *interior* contains $(n-2)^2$ sites and is defined to be the entire grid, except the boundary and the corners; for odd n the *centre* is the single site $(\frac{n+1}{2}, \frac{n+1}{2})$ and for even n the *central sites* are the four sites $(\frac{n}{2}, \frac{n}{2}), (\frac{n}{2}+1, \frac{n}{2}), (\frac{n}{2}, \frac{n}{2}+1), (\frac{n}{2}+1, \frac{n}{2}+1)$. Note that the centre or central sites also belong to the interior. Moreover, for a point (i,j) let its *northern neighbour* be the point $(i, j+1)$ if it belongs to the square grid and similarly define its eastern, southern and western neighbour.

Let S_1 be a percolating set and recall that the sets S_i's are disjoint, so label each site v with the index $l(v) = i$ of the set S_i in which v belongs. A (maximal) *increasing path* of length L is a path $v_1, \ldots, v_L$ in the grid such that

$l(v_1) < l(v_2) < \ldots < l(v_L)$ and v_L has no neighbours of larger label. A *truncated increasing path* allows v_L to have neighbours of larger label. The reason for introducing increasing paths is that the existence of an increasing path of length L shows that $k \geq L$. We construct long increasing paths dynamically, by specifying certain rules that they have to follow in order to guarantee the desired length. It is therefore natural in this context to define for a site v a *valid* neighbour to be a neighbour w of v such that $l(v) < l(w)$.

PROOF OF THEOREM 1.1. We assume $n \geq 7$, as for smaller values of n we can check the result by hand.

(1). The set of sites on the main diagonal $\{(i,i) \mid 1 \leq i \leq n\}$ which has cardinality n and percolating time $k = n$ gives an upper bound for the minimal percolation time. Now take any set S_1 with $s_1 = n$ and percolating time k; we need to show that $k \geq n$.

Consider the non-increasing sequence $(f_t)_{t=1}^{k}$ defined above which has the property, as always, that $f_k = 2n$. As $s_1 = n$, the maximum value of f_1 is $2n$ by definition. Therefore, we must have that $f_1 = 2n$, which is attained iff no two sites in S_1 are neighbours; moreover $2n = f_1 = f_2 = \ldots = f_k = 2n$, hence every site in S_i has exactly two neighbours in $S_1 \cup S_2 \cup \ldots \cup S_{i-1}$ and no two sites in S_i are neighbours for $i = 2, 3, \ldots, k$. Indeed, otherwise $f_i < f_{i-1}$.

This implies that every site v has at most two neighbours with labels strictly smaller than $l(v)$ and that there are no two neighbours v, w such that $l(v) = l(w)$. In particular, this means that every site in the interior has at least 2 neighbours of larger label and every site on the boundary has at least one neighbour of larger label; the corners may have both neighbours of smaller label.

If n is odd we can guarantee the existence of an increasing path of length at least n - and therefore prove $k \geq n$ - by taking v_1 to be the centre and then keep adding valid vertices to the increasing path while this is possible. By the above remarks, such a path can only end in a corner; the shortest path from the centre to a corner has length n, so we are done.

There is a small technicality for n even that the distance from any of the 4 central sites to the closest corner is $n-1$ instead of n. For brevity, in part (2) of the theorem, we assume in general that n is odd and after proving the result in this case we mention that for n even, when there are 4 central sites, the result follows from a "*max* 4 argument". A *max* 4 argument can be applied in the following set-up:

- we can find a suitable increasing path of desired length starting in the central site for n odd; and
- for n even we can provide the same argument starting in some central site whose label is at least 2,

which yields the desired increasing path for n even. To illustrate how this simple argument works, we will write it in below for the case which arises in part (1).

If n is even, the 4 central squares cannot be all in S_1 as S_1 does not contain adjacent sites. So let v_1 be a central site with $l(v_1) \geq 2$. We construct an increasing path as before starting at v_1 which must have length at least $n-1$ by the same argument. Considering that $l(v_1) \geq 2$, this again shows $k \geq n$ and we are done.

(2). We have the upper bound $k \leq n$ from part (1). Except for the special case

$$S_1 = \{(\frac{n+1}{2}, 2i-1) \mid 1 \leq i \leq \frac{n+1}{2}\} \cup \{(2i-1, \frac{n+1}{2}) \mid 1 \leq i \leq \frac{n+1}{2}\},$$

which is treated at the end, we need to show that if we take S_1 not as above with $s_1 = n+1$ we have that $k \geq n$.

We have $f_1 \leq 2n+2$ and $f_k = 2n$, but considering that $f_1 - f_k \leq 2$, we can generalise the argument above to conclude that every site has at most two neighbours of smaller label and that there are no two neighbours v, w such that $l(v) = l(w)$ with at most one of the following exceptions occurring:

(1) there exist one or two pairs of sites (ν_1, ν_2) and (ν_3, ν_4) such that $l(\nu_1) = l(\nu_2)$, $l(\nu_3) = l(\nu_4)$ with $\nu_1, \dots, \nu_4$ distinct;

(2) there exists a site ν with 3 neighbours of smaller label and there's a pair (ν_1, ν_2) of neighbours of equal label with ν, ν_1, ν_2 distinct;

(3) there exists one or two different sites ν_1, ν_2 which have 3 neighbours of smaller label;

(4) there exists one site ν with two neighbours of equal label, or one neighbour of equal label and three of smaller label or four neighbours of smaller label, except for the case when n odd, ν is the centre and has all 4 neighbours of smaller label;

(5) the special case when n is odd and the central site ν has all 4 neighbours of smaller label.

This follows from the definition of f_i by analysing all the possible ways in which we can have $f_k \geq f_1 - 2$. As in part (1), we are looking for long increasing paths in each case.

Case (1). Assume that n is odd and that there are two pairs of sites (ν_1, ν_2) and (ν_3, ν_4) such that $l(\nu_1) = l(\nu_2)$, $l(\nu_3) = l(\nu_4)$ and $\nu_1, \dots, \nu_4$ are all distinct. It is clear how the argument presented below also deals with the much easier case when only one such pair exists.

If all of ν_i's are in the interior, it is true that every vertex in the interior and on the boundary has at least one neighbour of larger label, so we are done by the same argument as in (1).

Say that a pair of sites (ν, ν') with at least one site on the boundary *belongs* to the north, east, south or west boundary if at least one of ν, ν' is on the north, east, south or west boundary respectively. Clearly, any such pair belongs to only one of the four boundary regions.

If at least one of the sites in the pair (ν_1, ν_2) is on the boundary and both ν_3 and ν_4 are in the interior, the increasing path as defined in (1) may end either in the corners or in any of ν_1, ν_2 which lie on the boundary. Hence, an increasing path constructed as before may have length $L < n$. However we can construct an increasing path of length n starting in the central site as follows: firstly, assume without loss of generality that (ν_1, ν_2) belongs to the north boundary. Take v_1 to be the centre and note that all vertices in the interior except maybe ν_3 and ν_4 have at least 2 neighbours of larger label. Keep adding valid vertices to the increasing path always ignoring northern valid neighbours, except when the path passes through ν_3 or ν_4 and we may be forced to pick a northern valid neighbour. As $n \geq 7$, this increasing path hits the boundary for the first time in the east, south or west boundary region, so when it ends in the corners or in ν_1, ν_2 it has length $L \geq n$ and we are done.

If both pairs (ν_1, ν_2) and (ν_3, ν_4) have at least one vertex on a boundary, we consider the following subcases:

i). The pairs (ν_1, ν_2) and (ν_3, ν_4) belong to the same boundary, say the north boundary. In this case construct as before the increasing path by always ignoring north neighbours while in the interior, which will force the path to hit the boundary for the first time in the east, south or west boundary region and we are done as before.

ii). The pair (ν_1, ν_2) belongs to the north boundary and the pair (ν_3, ν_4) belongs to the south boundary. We claim that we can construct an increasing path which hits the boundary for the first time either in the east or the west boundary. Take the starting vertex of the path to be the centre and until the increasing path hits the boundary for the first time, always pick an eastern or western valid neighbour whenever available. If neither is valid, we must still have 2 choices of continuing the path which have to be the northern and southern neighbour. If this situation occurs at least once pick the northern neighbour first and then in all subsequent occurrences alternate between the southern and the northern neighbour to keep the path in the strip

$$\{(i, \frac{n+1}{2}) \mid 1 \leq i \leq n\} \cup \{(i, \frac{n+3}{2}) \mid 1 \leq i \leq n\}.$$

This guarantees that the path hits the boundary for the first time either in the east or the west boundary and we are done as before.

Obviously, the same argument works when the pair (ν_1, ν_2) belongs to the east boundary and the pair (ν_3, ν_4) belongs to the west boundary.

iii). The cases i) and ii) don't occur, so without loss of generality the pair (ν_1, ν_2) belongs to the north boundary and the pair (ν_3, ν_4) belongs to the east boundary. In this case the idea is to avoid both northern and eastern neighbours and if this is not possible it will give us enough information about the board to construct an increasing path of the desired length.

Claim. Assume that a truncated increasing path P starts in the centre, never hits the boundary and ends at any interior site with a valid southern or western neighbour. Then there exists an increasing path containing P that reaches the boundary for the first time in the southern boundary or the western boundary and we are done as before.

Proof of Claim. Call south, west good directions and north, east bad directions and recall that for every interior site v different from $\nu_1, \dots \nu_4$ we have two choices of valid neighbours.

If a truncated increasing path that never hits the boundary reaches any interior vertex v (including $\nu_1, \dots \nu_4$) which has a valid neighbour v' in a good direction, we can continue the path with v' and we claim that starting with v' we can always choose a valid neighbour in a good direction until the increasing path hits the boundary. Indeed, by induction, assume that the path reaches an interior site v_i from v_{i-1} via a good direction, i.e. say without loss of generality that v_i is the southern neighbour of v_{i-1}. If v_i is not one of $\nu_1, \dots \nu_4$, there are two choices of valid neighbours of v_i, none of which is north, as v_{i-1} is the northern neighbour of v_i and $l(v_{i-1}) < l(v_i)$; therefore, one of the choices is a good direction.

If on the other hand v_i is one of $\nu_1, \dots \nu_4$, it cannot be ν_1 or ν_2 because the pair (ν_1, ν_2) belongs to the north boundary and v_{i-1}, the northern neighbour of v_i is an interior point. So assume $v_i = \nu_3$, then our only choice of a valid neighbour must be in a good direction, as the eastern neighbour of $v_i = \nu_3$ is ν_4 and we reached v_i from its northern neighbour v_{i-1}, so we cannot have valid neighbours in any of the bad directions.

Hence such an increasing path hits the boundary for the first time in the southern or eastern boundary. This finishes the proof the claim. $\square$

By the claim, it follows that we may assume that the valid neighbours of the centre are its northern and eastern neighbours then their valid neighbours are their northern and eastern neighbours and then inductively for every site in the whole north-eastern interior of the board $\{(i,j) \mid \frac{n+1}{2} \leq i,j \leq n-1\}$ its valid neighbours are its northern and eastern neighbours. In particular we remark that the label of the centre is at least 2 and so we aim for an increasing path of length $n-1$. Moreover, this strong property also implies that the following path of length $n-2$, $(\frac{n+1}{2}, \frac{n+1}{2}), (\frac{n+1}{2}+1, \frac{n+1}{2}), \ldots, (n-1, \frac{n+1}{2}), (n-1, \frac{n+1}{2}+1), \ldots, (n-1, n-1)$, is a truncated increasing path. However, $(n-1, n-1)$ is in the interior so even if it is one of the $\nu_i's$ it cannot be a terminal vertex for an increasing path, so we are done.

A max 4 argument is now used to establish the result when n is even. This finishes **Case (1)**.

Cases (2), (3). These two are both easier versions of **Case (1)**; when one or two pairs (ν, ν') of neighbours of equal label are replaced by one or two sites ν_1, ν_2 with three neighbours of smaller label we can extend the previous definition of "a pair belonging to (the north) boundary" also to "a vertex ν_1 with three smaller neighbours belonging to (the north) boundary" and use the exact path-avoiding techniques in **Case (1)**. The simplification comes from the fact that we replace one/two pairs which had two"special" vertices each by only one/two special vertex which we aim to avoid.

Case (4). We assume that n is odd and there exists one site ν with two neighbours of equal label, or one neighbour of equal label and three of smaller label or four neighbours of smaller label, except for the case when ν is the centre and has all 4 neighbours of smaller label.

If ν is in the centre and has at least one neighbour of equal label, say ν' then a (classical) increasing path starting at ν' has length at least $n-1$. Moreover, ν has at least two neighbours of smaller label, so $2 \leq l(\nu) = l(\nu')$ and we are done.

Assume now that ν has two neighbours of equal label, say ν_1, ν_2 - which by our approach it turns out to be the most interesting case.

If ν has all neighbours in the interior, then we are done by constructing an increasing path which never goes to ν whenever reaches one of its neighbours. Indeed, as all neighbours are interior, this restriction still leaves one valid neighbour except ν for all neighbours of ν.

If ν has one or two neighbours on the boundary, then without loss of generality it belongs to the region $\{(i, n-1) \mid 2 \leq i \leq n-1\}$ which we call the *north second boundary.* An increasing path can only end in a corner, in ν or in any of ν_1, ν_2 that lie on the boundary. Then we aim to construct the same increasing path as in **Case(1)**(i) by always avoiding the direction north while in the interior.

If the southern neighbour of the centre is valid, then we make it the second vertex of our increasing path and while in the interior we never pick a valid northern neighbour. We hit the boundary for the first time in the east, south or west boundary region and this guarantees that our increasing path has length at least n.

If the southern neighbour of the centre is not valid, then the label of the centre is at least 2 and as before while in the interior we never pick a valid northern neighbour. We hit the boundary for the first time in the east, south or west boundary region and this guarantees that our increasing path has length at least $n-1$, so the final label is at least n.

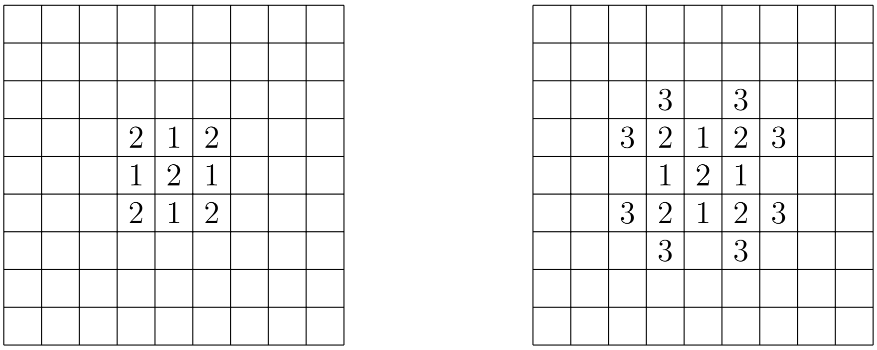

FIGURE 3. Determined labels 1

Moreover, we can now see that the argument above works just as well if we replace neighbours of ν of equal label by neighbours of ν of smaller label. By the max 4 argument, we can also deal with the case when n is even. The case when n is even and one of the four central sites has all neighbours of smaller label is also straightforward. This finishes the case.

Case (5). The only case left to consider is when n is odd and there exists a site ν in the centre with all four neighbours smaller than it. By starting an increasing path in any neighbour of the centre we can prove that the minimal infection time is at least $n-1$ in this case. To finish part (2) of the proof, we show that if $n \equiv 1 \pmod 4$ then the minimal infection time is n and if $n \equiv 3 \pmod 4$ then the minimal infection time is $n-1$ and the only set S_1 of i.i.s. which percolates in time $n-1$ is the set

$$S_1 = \{(\frac{n+1}{2}, 2i-1) \mid 1 \le i \le \frac{n+1}{2}\} \cup \{(2i-1, \frac{n+1}{2}) \mid 1 \le i \le \frac{n+1}{2}\}.$$

Remark. Assume that we have percolation in time $n-1$ for some set of i.i.s. S_1, $s_1 = n+1$, n odd. Then every site ν except the centre must satisfy property $\mathcal{P}$: n is at least the sum between the distance from ν to the closest corner added to the label of ν. Indeed, if any ν doesn't satisfy $\mathcal{P}$ we can start an increasing path at ν which contradicts that $k = n-1$; we only need to specify that for the four neighbours of the centre, from the two choices of valid neighbours we never choose the centre. Also recall that except the centre which has four neighbours of smaller label we have that every site ν must satisfy property $\mathcal{Q}$,

			3	2	3			
		3	2	1	2	3		
		2	1	2	1	2		
		3	2	1	2	3		
			3	2	3			

		4	3	2	3	4		
		3	2	1	2	3		
		2	1	2	1	2		
		3	2	1	2	3		
		4	3	2	3	4		

FIGURE 4. Determined labels 2

i.e. ν has at most two neighbours of smaller label and no neighbours of equal label. It turns out that this together with property $\mathcal{P}$ completely determines the label of each site. For example consider the 5×5 central region

$$\{(i,j) \mid \frac{n-3}{2} \leq i, j, \leq \frac{n+5}{2}\}.$$

We start arguing from the centre. By $\mathcal{P}$, all of the four neighbours of the centre must have label 1 hence the centre must have label 2 - see Figure 3, left. Next look at the site $(\frac{n-1}{2}, \frac{n-3}{2})$: by $\mathcal{Q}$ it cannot have label 1 or 2 and by $\mathcal{P}$ it cannot have label greater than 4, so the site $(\frac{n-1}{2}, \frac{n-3}{2})$ must have label 3, exactly as all the other 7 symmetric sites labeled 3 - see Figure 3, right. The sites $(\frac{n+1}{2}, \frac{n-3}{2})$, $(\frac{n+1}{2}, \frac{n+5}{2})$, $(\frac{n-3}{2}, \frac{n+1}{2})$, $(\frac{n+5}{2}, \frac{n+1}{2})$ have label 1 or 2 by $\mathcal{P}$, which must be 2 by $\mathcal{Q}$ - see Figure 4, left. Finally, the sites $(\frac{n-3}{2}, \frac{n-3}{2})$, $(\frac{n+5}{2}, \frac{n-3}{2})$, $(\frac{n-3}{2}, \frac{n+5}{2})$, $(\frac{n+5}{2}, \frac{n+5}{2})$ must have label at least 4 by $\mathcal{Q}$ and at most 4 by $\mathcal{P}$ - see Figure 4, right. This shows that $n \neq 5$ as the site $(\frac{n+1}{2}, \frac{n+5}{2})$ has label 2 and only one neighbour of label 1. The general induction step is presented below:

Assume that for odd i, $2i+1 < n$ we have determined that the labels the sites

$$\{(i,j) \mid \frac{n+1}{2} - i \leq i, j \leq \frac{n+1}{2} + i\}$$

are as follows (see Figure 5):

- the sites $(\frac{n+1}{2}, \frac{n+1}{2} - i), (\frac{n+1}{2}, \frac{n+1}{2} - i + 2), \ldots, (\frac{n+1}{2}, \frac{n+1}{2} + i)$ and $(\frac{n+1}{2} - i, \frac{n+1}{2}), (\frac{n+1}{2} - i + 2, \frac{n+1}{2}), \ldots, (\frac{n+1}{2} + i, \frac{n+1}{2})$ all have label 1;

	…	$\frac{n-2i+1}{2}$	…		$\frac{n-3}{2}$	$\frac{n-1}{2}$	$\frac{n+1}{2}$	$\frac{n+3}{2}$	$\frac{n+5}{2}$	…		$\frac{n+2i+1}{2}$	…
$\vdots$													
$\frac{n+2i+1}{2}$		$2i$	$2i-1$	$\cdots$	$i+2$	$i+1$	1	$i+1$	$i+2$	$\cdots$	$2i-1$	$2i$	
$\vdots$		$2i-1$	$\ddots$	$\cdots$	$i+1$	i	$\vdots$	i	$i+1$	$\cdots$	$\iddots$	$2i-1$	
		$\vdots$	$\vdots$	6	5	4	1	4	5	6	$\vdots$	$\vdots$	
$\frac{n+5}{2}$		$i+2$	$i+1$	5	4	3	2	3	4	5	$i+1$	$i+2$	
$\frac{n+3}{2}$		$i+1$	i	4	3	2	1	2	3	4	i	$i+1$	
$\frac{n+1}{2}$		1	$\cdots$	1	2	1	2	1	2	1	$\cdots$	1	
$\frac{n-1}{2}$		$i+1$	i	4	3	2	1	2	3	4	i	$i+1$	
$\frac{n-3}{2}$		$i+2$	$i+1$	5	4	3	2	3	4	5	$i+1$	$i+2$	
		$\vdots$	$\vdots$	6	5	4	1	4	5	6	$\vdots$	$\vdots$	
$\vdots$		$2i-1$	$\iddots$	$\cdots$	$i+1$	i	$\vdots$	i	$i+1$	$\cdots$	$\ddots$	$2i-1$	
$\frac{n-2i+1}{2}$		$2i$	$2i-1$	$\cdots$	$i+2$	$i+1$	1	$i+1$	$i+2$	$\cdots$	$2i-1$	$2i$	
$\vdots$													

FIGURE 5. **Case 5**, the first phase of the induction.

- the sites $(\frac{n+1}{2}, \frac{n+1}{2} - i + 1), (\frac{n+1}{2}, \frac{n+1}{2} - i + 3), \ldots, (\frac{n+1}{2}, \frac{n+1}{2} + i - 1)$ and $(\frac{n+1}{2} - i + 1, \frac{n+1}{2}), (\frac{n+1}{2} - i + 3, \frac{n+1}{2}), \ldots, (\frac{n+1}{2} + i - 1, \frac{n+1}{2})$ all have label 2;
- all other sites $\nu \in \{(i,j) \mid \frac{n+1}{2} - i \leq i, j \leq \frac{n+1}{2} + i\}$ have label the distance from ν to the centre.

Then, given that $2i + 1 < n$, we can deduce more labels as follows: all the sites $\{(j, \frac{n+1}{2} + i) \mid \frac{n+1}{2} + 1 \leq j \leq \frac{n+1}{2} + i\}$ have two neighbours of smaller label: their southern neighbour and their western neighbour; therefore, by $\mathcal{Q}$ each site in $\{(j, \frac{n+1}{2} + i + 1) \mid \frac{n+1}{2} + 1 \leq j \leq \frac{n+1}{2} + i\}$ has label

	$\cdots$	$\frac{n-2i-1}{2}$	$\frac{n-2i+1}{2}$	$\cdots$		$\frac{n-3}{2}$	$\frac{n-1}{2}$	$\frac{n+1}{2}$	$\frac{n+3}{2}$	$\frac{n+5}{2}$	$\cdots$		$\frac{n+2i+1}{2}$	$\frac{n+2i+3}{2}$	$\cdots$
$\vdots$															
$\frac{n+2i+3}{2}$		$2i+2$	$2i+1$	$2i$	$\cdots$	$i+3$	$i+2$		$i+2$	$i+3$	$\cdots$	$2i$	$2i+1$	$2i+2$	
$\frac{n+2i+1}{2}$		$2i+1$	$2i$	$2i-1$	$\cdots$	$i+2$	$i+1$	1	$i+1$	$i+2$	$\cdots$	$2i-1$	$2i$	$2i+1$	
		$2i$	$2i-1$	$\ddots$	$\cdots$	$i+1$	i	$\vdots$	i	$i+1$	$\cdots$	$\iddots$	$2i-1$	$2i$	
$\vdots$		$\vdots$	$\vdots$	$\vdots$	6	5	4	1	4	5	6	$\vdots$	$\vdots$	$\vdots$	
$\frac{n+5}{2}$		$i+3$	$i+2$	$i+1$	5	4	3	2	3	4	5	$i+1$	$i+2$	$i+3$	
$\frac{n+3}{2}$		$i+2$	$i+1$	i	4	3	2	1	2	3	4	i	$i+1$	$i+2$	
$\frac{n+1}{2}$			1	$\cdots$	1	2	1	2	1	2	1	$\cdots$	1		
$\frac{n-1}{2}$		$i+2$	$i+1$	i	4	3	2	1	2	3	4	i	$i+1$	$i+2$	
$\frac{n-3}{2}$		$i+3$	$i+2$	$i+1$	5	4	3	2	3	4	5	$i+1$	$i+2$	$i+3$	
		$\vdots$	$\vdots$	$\vdots$	6	5	4	1	4	5	6	$\vdots$	$\vdots$	$\vdots$	
$\vdots$		$2i$	$2i-1$	$\iddots$	$\cdots$	$i+1$	i	$\vdots$	i	$i+1$	$\cdots$	$\ddots$	$2i-1$	$2i$	
$\frac{n-2i+1}{2}$		$2i+1$	$2i$	$2i-1$	$\cdots$	$i+2$	$i+1$	1	$i+1$	$i+2$	$\cdots$	$2i-1$	$2i$	$2i+1$	
$\frac{n-2i-1}{2}$		$2i+2$	$2i+1$	$2i$	$\cdots$	$i+3$	$i+2$		$i+2$	$i+3$	$\cdots$	$2i$	$2i+1$	$2i+2$	
$\vdots$															

FIGURE 6. **Case 5**, the second phase of the induction.

greater than the label of its southern neighbour. Moreover, by $\mathcal{P}$ each site in $\{(j, \frac{n+1}{2}+i+1) \mid \frac{n+1}{2}+1 \leq j \leq \frac{n+1}{2}+i\}$ must have label exactly one greater than the label of its southern neighbour. By symmetry, and then applying a similar reasoning to the four 'corner' sites $(\frac{n+1}{2}-i-1, \frac{n+1}{2}-i-1), (\frac{n+1}{2}-i-1, \frac{n+1}{2}+i+1), (\frac{n+1}{2}+i+1, \frac{n+1}{2}-i-1), (\frac{n+1}{2}+i+1, \frac{n+1}{2}+i+1)$ we have determined the labels of all sites $\{(i,j) \mid \frac{n+1}{2}-i-1 \leq i, j \leq \frac{n+1}{2}+i+1\}$ except for the four sites $(\frac{n+1}{2}, \frac{n+1}{2}-i-1), (\frac{n+1}{2}, \frac{n+1}{2}+i+1), (\frac{n+1}{2}-i-1, \frac{n+1}{2}), (\frac{n+1}{2}+i+1, \frac{n+1}{2})$ which are only known to have label at most $i+1$ (see Figure 6).

Note that n cannot be equal to $2i+3$. Indeed, the site $(\frac{n+1}{2}, \frac{n+1}{2}-i-1)$ has label at most $i+1$ by $\mathcal{P}$ and the fact that it has a neighbour of label $i+2$, but then the site $(\frac{n+1}{2}, \frac{n+1}{2} - i - 1)$ itself has only one neighbour of smaller or equal label which has label 1 and that gives the desired contradiction. We emphasise that this is the part of the argument which shows that if $n \equiv 1 \pmod 4$ then the minimal infection time cannot be $n-1$ and so it must be n.

Therefore we have that $n > 2i+3$ and by the exact same reasoning, this time using that the four sites $(\frac{n+1}{2}, \frac{n+1}{2} - i - 1), (\frac{n+1}{2}, \frac{n+1}{2} + i + 1), (\frac{n+1}{2} - i - 1, \frac{n+1}{2}), (\frac{n+1}{2} + i + 1, \frac{n+1}{2})$ have label at most $i+1$ (instead of knowing their exact label as before) we determine the labels of all the sites $\{(i,j) \mid \frac{n+1}{2} - i - 2 \leq i, j \leq \frac{n+1}{2} + i + 2\}$ except for the eight sites:

- the four sites $(\frac{n+1}{2}, \frac{n+1}{2} - i - 1), (\frac{n+1}{2}, \frac{n+1}{2} + i + 1), (\frac{n+1}{2} - i - 1, \frac{n+1}{2}), (\frac{n+1}{2} + i + 1, \frac{n+1}{2})$ of label at most $i+1$;
- the four sites $(\frac{n+1}{2}, \frac{n+1}{2} - i - 2), (\frac{n+1}{2}, \frac{n+1}{2} + i + 2), (\frac{n+1}{2} - i - 2, \frac{n+1}{2}), (\frac{n+1}{2} + i + 2, \frac{n+1}{2})$ of label at most $i+2$.

We focus on $(\frac{n+1}{2}, \frac{n+1}{2} - i - 2)$ and note that actually its label is smaller than the label of its northern neighbour, in order for its northern neighbour (which cannot be initially infected) to be infected. But then $(\frac{n+1}{2}, \frac{n+1}{2} - i - 2)$ has at most one potential neighbour of smaller label, i.e. its southern neighbour, which shows that $(\frac{n+1}{2}, \frac{n+1}{2} - i - 2)$ is initially infected, i.e. it has label 1. This also implies that $(\frac{n+1}{2}, \frac{n+1}{2} - i - 1)$ has label 2. By symmetry we have determined the labels of all the sites

$$\{(i,j) \mid \frac{n+1}{2} - i - 2 \leq i, j \leq \frac{n+1}{2} + i + 2\},$$

which agree with our induction statement. The base case $i = 1$ is in Figure 3, left. So either $n = 2(i+2)+1$ and we are done, as this model is obviously valid, or $n > 2(i+2)+1$ and we can continue, increasing the value of i by 2 at each step. This finishes the proof. $\square$

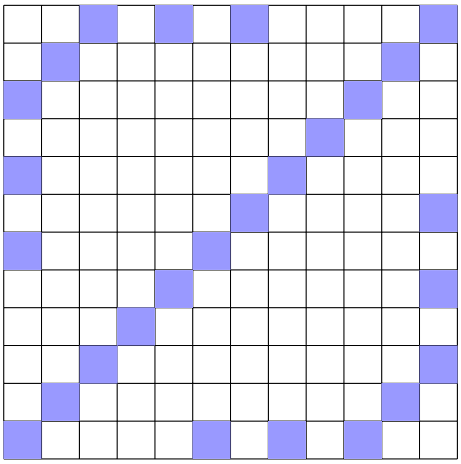

FIGURE 7. Our construction of S_1 for the upper bound, $n = 12, k = 5$.

3. The Upper Bound

In this section we aim to find sets S_1 of i.i.s. of small size that percolate in time k and prove Theorem 1.2. In Figure 7 we show graphically our construction for $n = 12, k = 5$.

PROOF OF THEOREM 1.2. For the proof of the first part of the theorem, we assume $n \geq 2k$ and then we define S_1 as a union of three diagonals:

$$S_1 = \{(i,i) \mid 1 \leq i \leq n\} \cup \{(k+i,i) \mid 1 \leq i \leq n-k\} \cup \{(i,k+i) \mid 1 \leq i \leq n-k\}.$$

For the proof of the second part of the theorem, we assume $n > 2k$ and we formally define S_1 as a union of $1 + 2\lceil \frac{n-k}{2k-1} \rceil$ diagonals :

$$D_1 = \{(i,i) \mid 1 \leq i \leq n\}$$

$$D_2 = \{(2k-1+i,i) \mid 1 \leq i \leq n-(2k-1)\} \ , \ D_2' = \{(i,2k-1+i) \mid 1 \leq i \leq n-(2k-1)\}$$

$$D_3 = \{(2(2k-1)+i,i) \mid 1 \leq i \leq n-2(2k-1)\} \ , \ D_3' = \{(i,2(2k-1)+i) \mid 1 \leq i \leq n-2(2k-1)\}$$

$$\vdots$$

$$D_j = \{((j-1)(2k-1)+i,i) \mid 1 \leq i \leq n-(j-1)(2k-1)\} \ ,$$

$$D_j' = \{(i,(j-1)(2k-1)+i) \mid 1 \leq i \leq n-(j-1)(2k-1)\}$$

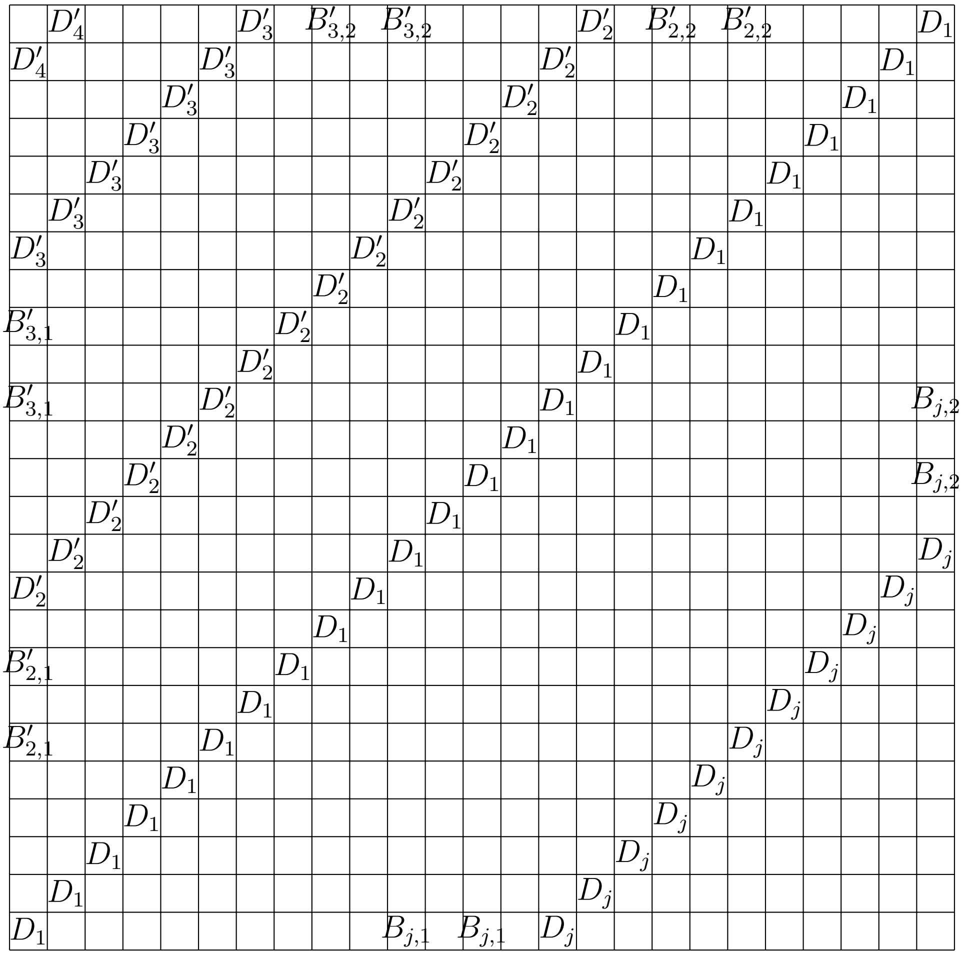

FIGURE 8. Formal construction of S_1.

$$\vdots$$

together with 2 groups of $\lceil \frac{k-1}{2} \rceil$ boundary sites for each of the diagonals $D_2, D'_2, D_3, D'_3, \ldots$ - labeled for clarity $B_{i,1}$ and $B_{i,2}$ for the diagonal D_i (see Figure 8). For k odd they are defined as follows:

$$B_{2,1} = \{(1+(2k-1)-2i, 1) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B_{2,2} = \{(n-(2k-1)+2i, n) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$B'_{2,1} = \{(1, 1+(2k-1)-2i) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B'_{2,2} = \{(n, n-(2k-1)+2i) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$B_{3,1} = \{(1+2(2k-1)-2i, 1) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B_{3,2} = \{(n-2(2k-1)+2i, n) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$B'_{3,1} = \{(1, 1+2(2k-1)-2i) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B'_{3,2} = \{(n, n-2(2k-1)+2i) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$\vdots$$

$$B_{j,1} = \{(1{+}j(2k{-}1){-}2i, 1) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B_{j,2} = \{(n{-}j(2k{-}1){+}2i, n) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$B'_{j,1} = \{(1, 1{+}j(2k{-}1){-}2i) \mid 1 \le i \le \frac{k-1}{2}\}\ ,\ B'_{j,2} = \{(n, n{-}j(2k{-}1){+}2i) \mid 1 \le i \le \frac{k-1}{2}\}$$

$$\vdots$$

For k even, the boundary sites B's are defined similarly with the range of i changing from $i \in \{1, 2, \dots, \frac{k-1}{2}\}$ to $i \in \{1, 2, \dots, \frac{k-2}{2}\} \cup \{\frac{k-1}{2}\}$. It is not hard to see that such an S_1 percolates in time k and has size

$$s_1 \le n + \frac{(n-2k+1+k)(n+k-2)}{2k-1} + 2 \cdot 2 \cdot \frac{k}{2} \cdot \frac{n+k-2}{2k-1},$$

i.e. $s_1 \le n + \frac{(n+k-2)(n+k+1)}{2k-1}$, as claimed. This finishes the proof. □

4. The Lower Bound

In this section, we give a short argument which establishes the rough lower bound on s_1 given in Theorem 1.3.

Proof of Theorem 1.3. For $1 \le i < j \le n$, denote by $s_{i,j}$ the total number of pairs (v, w) where $v \in S_i$ and $w \in S_j$ are neighbours. Moreover, let s'_i denote the number of sites infected at time i that lie on the boundary of the grid, including the four corners. Then, considering that every site on the boundary has at most three neighbours in the square grid, the following set of inequalities is straightforward from the definition of the model:

$$4s_1 - s'_1 \ge s_{1,2} + s_{1,3} + \dots + s_{1,k}$$

$$\frac{s_{1,2} - s'_2}{2} \ge s_2 - \frac{s'_2}{2} \ge \frac{s_{2,3} + s_{2,4} + \dots + s_{2,k}}{2}$$

$$\frac{s_{1,3} + s_{2,3} - s'_3}{2} \ge s_3 - \frac{s'_3}{2} \ge \frac{s_{3,4} + \dots + s_{3,k}}{2}$$

$$\vdots$$

$$\frac{s_{1,k-2} + \dots + s_{k-3,k-2} - s'_{k-2}}{2} \ge s_{k-2} - \frac{s'_{k-2}}{2} \ge \frac{s_{k-2,k-1} + s_{k-2,k}}{2}$$

$$\frac{s_{1,k-1} + \ldots + s_{k-2,k-1} - s'_{k-1}}{2} \geq s_{k-1} - \frac{s'_{k-1}}{2} \geq \frac{s_{k-1,k}}{2}$$
$$\frac{s_{1,k} + \ldots + s_{k-1,k}}{2} \geq s_k$$

Using these inequalities, we establish the following.

Claim. For $2 \leq j \leq k$ we have

$$s_2 + s_3 + \ldots + s_j \leq 2s_1 + s_2 + \ldots + s_{j-1} - \frac{1}{2}(s'_1 + \ldots + s'_{j-1}).$$

Proof. Let $2 \leq j \leq k$, then

$$s_2 + s_3 + \ldots + s_j \leq \frac{s_{1,2}}{2} + \frac{s_{1,3} + s_{2,3}}{2} + \ldots + \frac{s_{1,j} + s_{2,j} + \ldots + s_{j-1,j}}{2} =$$
$$= \frac{s_{1,2} + s_{1,3} + \ldots + s_{1,j}}{2} + \frac{s_{2,3} + s_{2,4} + \ldots + s_{2,j}}{2} + \ldots + \frac{s_{j-1,j}}{2} \leq$$
$$\leq \frac{s_{1,2} + \ldots + s_{1,k}}{2} + \frac{s_{2,3} + \ldots + s_{2,k}}{2} + \ldots + \frac{s_{j-1,j} + \ldots + s_{j-1,k}}{2} \leq$$
$$\leq 2s_1 + s_2 + \ldots + s_{j-1} - \frac{1}{2}(s'_1 + \ldots + s'_{j-1}).$$

This completes the proof of the claim. □

Therefore, by applying the claim $k-1$ times for $j = k, k-1, \ldots, 2$ we have

$$n^2 = s_1 + s_2 + \ldots + s_k \leq (2k-1)s_1 - \frac{1}{2}((k-1)s'_1 + (k-2)s'_2 + \ldots + s'_{k-1}).$$

Hence, in order to conclude the proof we need to show that for any $k \geq 3$ and any configuration of i.i.s. that percolate on the grid we have that

$$Q = (k-1)s'_1 + (k-2)s'_2 + \ldots + s'_{k-1} \geq (\sqrt{4k-1} - 2)(2n-2),$$

and that $Q = s'_1 \geq \frac{1}{3}(4n-4)$ for $k = 2$.

Being consistent with the notation in Section 2, we label each vertex with index of the set S_i in which it lies. We note that, as any vertex on the boundary has at most three neighbours, we have that for any $i \geq 2$, a necessary condition for a vertex of label i to lie on the boundary is that it has a neighbour on the boundary of strictly lower label. Using just this necessary condition, we get that in between two consecutive vertices of label 1 on the boundary, the sequence of labels of vertices must be unimodal. Therefore, by a further optimisation which takes into account the fact that vertices of larger label have a smaller weight in Q, we get that a configuration of labels of vertices on the boundary that minimises Q is obtained when we pave the boundary with blocks of labels $k, k-1, \dots, m, 1, m, m+1, \dots k$, for some $2 \leq m \leq k$ which minimises

$$f(m) = \frac{(k-1) + 2[(k-m) + (k-m-1) + \ldots + 1]}{1 + 2(k-m+1)}.$$

For example, when $k = 2$, it is clear that $m = k = 2$ minimises $f(m)$, so a lower bound for Q is given by $\frac{1}{3}(4n-4)$, obtained when we pave the boundary with blocks of labels $2, 1, 2$. Hence, in this case we get $s_1 \geq \frac{n^2}{3} + \frac{2n-2}{9}$ as promised.

For $k \geq 3$, we have by the same reasoning that $Q \geq (4n-4)f(m)$ for any $2 \leq m \leq k$, $m \in \mathbb{R}$ that minimises $f(m)$. By considering the first and second derivative of f, we get that for $k \geq 3$, $f(m)$ is minimised for $m = k - \frac{1}{2}\sqrt{4k-1} + \frac{3}{2}$ which gives $f(m) = \frac{1}{2}\sqrt{4k-1} - 1$. This gives the desired lower bound

$$Q \geq (\sqrt{4k-1} - 2)(2n-2),$$

whenever $k \geq 3$, finishing the proof. $\square$

5. Concluding Remarks

Clearly, the proof of Theorem 1.3 can be generalised to higher dimensions and/or different values of r for the r-neighbour percolation. However, the lower bound obtained in this way is pretty rough. Therefore, it would be very interesting to see an improvement of our bounds, especially the lower bound, even in some special cases. In particular, we think that one could obtain a better lower bound on the minimum infection time even when $s_1 = n + O(1)$ and $s_1 = O(n)$.

Moreover, one could investigate (maybe using the techniques developed in Section 2) the precise minimum infection time when $s_1 = n + 2$.

Bibliography

1. L. Addario-Berry, L. Devroye, G. Lugosi, and R. I. Oliveira, *Local optima of the Sherrington-Kirkpatrick Hamiltonian*, arXiv, 2017.
2. R. L. Adler, L.W. Goodwyn, and B. Weiss, *Equivalence of topological Markov shifts*, Israel Journal of Mathematics **27** (1977), 49–63.
3. R.L. Adler and B. Weiss, *Similarity of automorphisms of the torus*, Memoirs of the American Mathematical Society **98** (1970).
4. M. Aizenman and J.L. Lebowitz, *Metastability effects in bootstrap percolation*, Journal of Physics A **21** (1988), 3801–3813.
5. J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris, *The sharp threshold for bootstrap percolation in all dimensions*, SIAM Journal on Discrete Mathematics **29** (2012), 224–251.
6. F. Benevides and M. Przykucki, *Maximum percolation time in two-dimensional bootstrap percolation*, Transactions of the American Mathematical Society **364** (2014), 2667–2701.
7. B. Bollobás, *On generalized graphs*, Acta Mathematica Academiae Scientiarum Hungaricae **16 (3–4)** (1965), 447–452.
8. B. Bollobás, H. Duminil-Copin, R. Morris, and P.J. Smith, *Universality of two-dimensional critical cellular automata*, preprint (2017).
9. B. Bollobás, P.J. Smith, and A.J. Uzzell, *Monotone cellular automata in a random environment*, Combinatorics, Probability and Computing **24(4)** (2015), 687–722.
10. A. Carbone, *Cycles of relatively prime length and the road coloring problem*, Israel Journal of Mathematics **123** (2001), 303–316.
11. R. Cerf and F. Manzo, *The threshold regime of finite volume bootstrap percolation*, Stochastic Processes and their Applications **101** (2002), 69–82.

12. J. Černý, *Poznámka k homogénnym experimentom s konečnými automatmi*, Matematicko-fyzikálny časopis Slovenskej Akadémie Vied **14** (1964), 208–216.
13. J. Chalupa, P.L. Leath, and G.R. Reich, *Bootstrap percolation on a Bethe lattice*, Journal of Physics C **12** (1979), L31–L35.
14. P. Erdős, C. Ko, and R. Rado, *Forbidding a set difference of size* 1, Quarterly Journal of Mathematics. Oxford. Second Series **12** (1961), 313–320.
15. J. Friedman, *On the road coloring problem*, Proceedings of the American Mathematical Society **110** (1990), 1133–1135.
16. R.J. Glauber, *Time dependent statistics of the ising model*, Journal of Mathematical Physics **4** (1963), 294.
17. J. Gravner, A.E. Holroyd, and R. Morris, *A sharper threshold for bootstrap percolation in two dimensions*, Probability Theory and Related Fields **18** (2010), 1–23.
18. M. Grech and A. Kisielewicz, *The Černý conjecture for automata respecting intervals of a directed graph*, Discrete Mathematics and Theoretical Computer Science **15(3)** (2013), 61–72.
19. J. R. Griggs, *Sufficient conditions for a symmetric chain order*, SIAM Journal on Applied Mathematics **32** (1977), 807–809.
20. A. Holroyd, *Sharp metastability threshold for two-dimensional bootstrap percolation*, Probability Theory and Related Fields **125** (2003), 195–224.
21. J.E. Hopcroft, R. Motwani, and J.D. Ullman, *Introduction to automata theory, languages, and computation (3rd edition)*, Pearson Education, 2007.
22. E. Ising, *Beitrag zur Theorie des Ferromagnetismus*, Zeitschrift für Physik **31** (1925), 253–258.
23. V.A. Kazakov, *Ising model on a dynamical planar random lattice: Exact solution*, Physics Letters A **119 (3)** (1986), 140–144.
24. D.J. Kleitman, *On a lemma of Littlewood and Offord on the distribution of linear combinations of vectors*, Advances in Mathematics **5** (1970), 251–259.

25. I. Leader and E. Long, *Intersection theorems for systems of finite sets*, Discrete Applied Mathematics **169** (2014), 219–224.
26. T. D. Lee and C. N. Yang, *Statistical theory of equations of state and phase transitions. ii. lattice gas and Ising model*, Physical Review **87** (1952), 410.
27. W. Lenz, *Beiträge zum Verständnis der magnetischen Eigenschaften in festen Körpern*, Physikalische Zeitschrift **21** (1920), 613–615.
28. D. Lubell, *A short proof of Sperner's lemma*, Journal of Combinatorial Theory **1 (2)** (1966), 299.
29. B.M. McCoy and T.T. Wu, *The two-dimensional Ising Model*, second ed., Harvard University Press, 2004.
30. L.D. Meshalkin, *Generalization of Sperner's theorem on the number of subsets of a finite set*, Theory of Probability and its Applications **8 (2)** (1963), 203–204.
31. G.L. O'Brien, *The road coloring problem*, Israel Journal of Mathematics **39** (1981), 145–154.
32. R. O'Donnell, *Open problems in analysis of boolean functions*, https://arxiv.org/pdf/1204.6447.pdf (2012).
33. J.E. Pin, *On two combinatorial problems arising from automata theory*, Annals of Discrete Mathematics **17** (1983), 535–548.
34. M. Saks, *A short proof of the existence of k-saturated partitions of partially ordered sets*, Advances in Mathematics **33(3)** (1979), 207–211.
35. J. Shearer and D.J. Kleitman, *Probabilities of independent choices being ordered*, Studies in Applied Mathematics **60(3)** (1979), 271–275.
36. H. Spink, *Local maxima of quadratic boolean functions*, Combinatorics Probability and Computing **25(4)** (2016), 633–640.
37. ———, *Orthogonal symmetric chain decompositions*, preprint (2017+).
38. R. Stanley, *Quotients of Peck posets*, Order **1** (1984), 29–34.
39. G. Steinberg, *The Černý conjecture for one-cluster automata with prime length cycle*, Theoretical Computer Science **412(39)** (2011), 5487–5491.

40. A.N. Trahtman, *Notable trends concerning the synchronization of graphs and automata, in Cologne-Twente workshop on graphs and combinatorial optimization*, Electronic Notes in Discrete Mathematics **25** (2006), 173–175.

41. ———, *The road coloring problem*, Israel Journal of Mathematics **172** (2009), 51–60.

42. J. von Neumann, *Theory of self-reproducing automata*, University of Illinois Press, Champaign (1966).

43. K. Yamamoto, *Logarithmic order of free distributive lattice*, Journal of the Mathematical Society of Japan **6** (1954), 343–353.