Reconfiguration of Hamiltonian Cycles and Paths in Rectangular Grid Graphs

Albi Kazazi

A dissertation submitted to the Faculty of Graduate Studies
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

Graduate Program in Mathematics and Statistics
York University, Toronto, Ontario

September 2025

Abstract

An $m \times n$ *grid graph* is the induced subgraph of the square lattice whose vertex set consists of all integer grid points $\{(i, j) : 0 \leq i < m,\ 0 \leq j < n\}$. Let H and K be Hamiltonian cycles in an $m \times n$ grid graph G. We study the problem of reconfiguring H into K using a sequence of local transformations called *moves*. A *box* of G is a unit square face. A box with vertices a, b, c, d is *switchable* in H if exactly two of its edges belong to H, and these edges are parallel. Given such a box with edges ab and cd in H, a *switch move* removes ab and cd, and adds bc and ad. A *double-switch move* consists of performing two consecutive switch moves. If, after a double-switch move, we obtain a Hamiltonian cycle, we say that the double-switch move is *valid*. We prove that any Hamiltonian cycle H can be transformed into any other Hamiltonian cycle K via a sequence of valid double-switch moves, such that every intermediate graph remains a Hamiltonian cycle. This result extends to Hamiltonian paths. In that case, we also use single-switch moves and a third operation, the *backbite move*, which enables the relocation of the path endpoints.

Acknowledgements

I would like to thank my supervisor Neal Madras for his insightful and patient feedback. Through his mentorship, my time in the program has been one of growth, mathematical and beyond. I'm lucky and grateful to have had him as my advisor.

Thanks to Ada Chan who, as my master's supervisor, was very helpful, answering all my questions. It was that experience that made me decide to pursue a PhD.

Thanks to Paul Szeptycki and Youness Lamzouri, who taught early courses that made math feel approachable and worth the time.

Thanks to my family for all their support.

Table of Contents

Introduction

Reconfiguration problems consist of a set of configurations, or a state space, and a set of reconfiguration rules that can be used to change one configuration into another. We call an application of a rule a move or a step. Given this data, we may construct a reconfiguration graph, where configurations are vertices, and two vertices are adjacent if there is a move that can change the configurations they represent into one another. Some popular problems include the knight's tour on a chessboard, the 15-puzzle with sliding blocks, and the Rubik's cube.

Many well-known problems have been studied from a reconfiguration perspective, in particular with respect to the connectivity of the reconfiguration graph. One example is the reconfiguration of the k-colorings of a graph, where a k-coloring is an assignment of one of k colors to each vertex so that adjacent vertices receive different colors, and the allowed move is to recolor one vertex at a time [1, 2, 3, 4, 5]. Another example involves dominating sets of size at most k, where a dominating set is a set of vertices such that every vertex is either in the set or adjacent to a vertex in the set. Here, the allowed move is to add or remove a single vertex [6, 7]. A related example involves independent sets. An independent set is a set of vertices in a graph such that no two vertices in the set are adjacent. We place a token on each vertex of an independent set of size at most k and we allow one token to be slid at a time along an edge of the graph, provided that the resulting set is also an independent set of size k [8, 9]. For surveys on reconfiguration problems, see [10, 11].

Hamiltonian cycles: history and difficulty. The study of Hamiltonian paths and cycles in graphs dates back to Euler's 1736 resolution of the Königsberg bridge problem, which asked whether one could walk through the city crossing each of its seven bridges exactly once. This work established the foundation for graph theory. A century later, Hamilton introduced his icosian game in 1857, asking whether one could traverse all vertices of a dodecahedron exactly once and return to the starting point. Euler's work on the bridges problem led to a complete and easily verifiable characterization of graphs admitting Eulerian trails—graphs where every edge is traversed exactly once. Hamilton's problem, by contrast, remains notoriously difficult. Determining whether a graph contains a Hamiltonian cycle is NP-complete for general graphs, and few useful characterizations exist. Dirac's theorem provides some partial insight: it states that every graph with $n \geq 3$ vertices where each vertex has degree at least $n/2$ must contain a Hamiltonian cycle. Even for special graph classes, Hamiltonicity questions remain open; Barnette's conjecture [12] – that every 3-connected cubic planar bipartite graph is Hamiltonian – has stood unresolved since 1969. Hamiltonian cycle problems arise naturally in optimization through the traveling salesman problem, in recreational mathematics through the knight's tour on a chessboard, and as canonical examples in computational complexity theory, making their study both practically and theoretically significant.

Reconfiguration of Hamiltonian cycles. The reconfiguration problem for Hamiltonian cycles asks the following: Given any two Hamiltonian cycles H and K in a graph and an allowed type of move, is there a sequence $H_0, H_1, ..., H_r$ of Hamiltonian cycles, where $H = H_0$, $K = H_r$, such that for each $j \in \{1, \ldots r\}$ H_j is obtained by a single application of the allowed move on H_{j-1}? Often, the allowed move consists in adding and removing one or two pairs of edges from a given Hamiltonian cycle, in such a way that we obtain a new Hamiltonian cycle. For general graphs, this is a hard problem [13].

In this thesis, we restrict our attention to the Hamiltonian cycle and path reconfiguration problem on a class of graphs called grid graphs, which are subgraphs of the square lattice. An $m \times n$ *grid graph* G is the induced subgraph of the square lattice whose vertex set consists of all integer grid points $\{(i, j) : 0 \leq i < m,\ 0 \leq j < n\}$ with edges between vertices at distance 1. We call the unit-square faces of a grid graph *boxes*. We show that for an $m \times n$ grid graph, the reconfiguration graph of Hamiltonian cycles and the reconfiguration graph of Hamiltonian paths is connected under a small set of suitable moves.

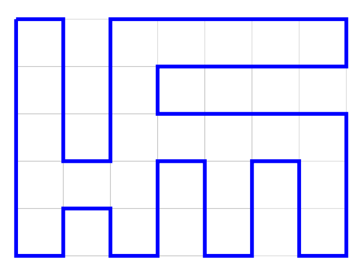

Fig. I.1. A 1-complex Hamiltonian cycle on an 8×6 grid graph.

The question of whether an $m \times n$ grid graph has a Hamiltonian path or cycle was first studied by Itai et al. in [14]. They showed that for an $m \times n$ grid graph to have a Hamiltonian cycle, it is necessary and sufficient that at least one of m and n is even. A *polyomino graph* is a superclass of $m \times n$ grid graphs, where every edge belongs to a box. We give a constructive definition at the start of Chapter 1. Umans and Lenhart [15] gave a polynomial-time algorithm to find a Hamiltonian cycle in polyomino graphs, if one exists. Chen et al. gave an efficient algorithm for constructing Hamiltonian paths in rectangular grid graphs [16].

Nishat and Whitesides [17] introduced the "flip" and "transpose" moves described below, and a complexity measure called "bend complexity" for Hamiltonian cycles in rectangular grid graphs. Roughly, a 1-complex Hamiltonian cycle is one in which every vertex of G is connected to the boundary via a straight line. They prove that using these two moves, it is possible to reconfigure any pair of 1-complex Hamiltonian cycles of G into one another.

We dispense with the need for bend complexity constraints, proving that any Hamiltonian cycle in a rectangular grid graph can be reconfigured into any other, using a more general move, which we call a double-switch move.

Let H be a Hamiltonian cycle of an $m \times n$ grid graph G. A box of G with vertices a, b, c, d is considered *switchable* in H if it has exactly two edges in H, and these edges are parallel. Let $abcd$ be a switchable box with edges ab and cd in H. We define a *switch move* on the box $abcd$ in H as follows: remove edges ab and cd and add edges bc and ad. If X is a switchable box in G in H, we denote a switch move as $\text{Sw}(X)$.

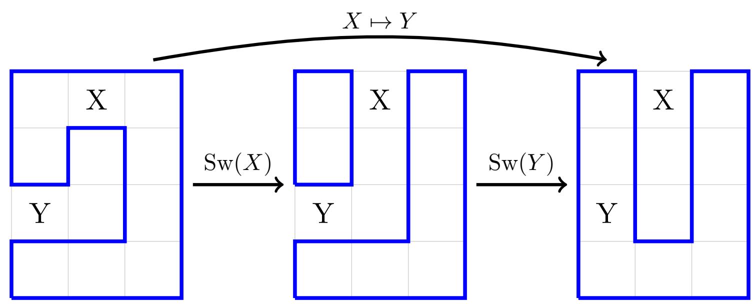

Fig. I.2. Illustration of each switch of a general double-switch move in a 4×5 grid graph.

A *double-switch move* is a pair of switch operations where we first switch X and then find another switchable box Y and switch it, and denote it by $X \mapsto Y$. If, after a double-switch move, we obtain a new Hamiltonian cycle, we call the move a *valid move.* See Figure I.2

Let $X = abcd$ and $Y = dcef$ be boxes sharing the edge cd of G. Assume that the edges ab, fd, dc and ce belong to H, and that the edges fe, ad and bc do not. A *flip* move consists in removing the edges fd, ce and ab, and adding the edges ad, bc and fe. Effectively, this is the same as the move $X \mapsto Y$, obtained by first switching X and then switching Y. See Figure I.3.

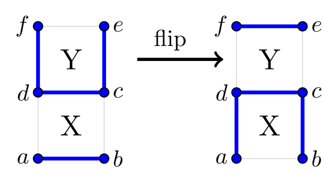

Fig. I.3. A flip move.

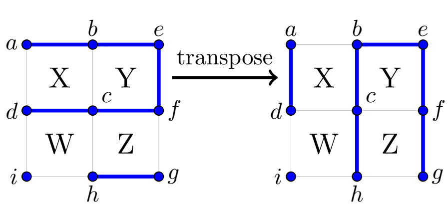

Fig. I.4. A transpose move.

Consider the four boxes $X = abcd$, $Y = cbef$, $Z = cfgh$ and $W = dchi$ that are incident on the vertex c. Note that X and Y share the edge cb, Y and Z share cf, Z and W share ch, and W and X share cd. Assume that the edges ab, be, ef, fc, cd and hg belong to H and that the edges ad, bc and fg do not. A *transpose* move consists in switching X and then switching Z. See Figure I.4.

Nishat in [18] showed that flip and transpose moves are always valid. The more general double-switch moves are sufficient for constructing algorithms that reconfigure arbitrary Hamiltonian cycles in grid graphs. This comes at the added cost of verifying the validity of each move. We provide such reconfiguration algorithms and prove the existence of all required moves.

Theorem I.1. *Let H and K be any two Hamiltonian cycles in an $m \times n$ grid graph G with $n \geq m$. Then there exists a sequence of at most n^2m valid double-switch moves that reconfigures H into K.*

See [19] for an illustration.

We also give an algorithm by which we are able to reconfigure any two Hamiltonian paths in an $m \times n$ grid graph. This requires a move that can relocate the end-vertices of a path, but switch moves do not relocate end-vertices. To address this, we use a move called a *backbite* move, first introduced by Mansfield in 1982 [20].

Let H be a Hamiltonian path $v_1, \ldots, v_r$ of an $m \times n$ grid graph G, and let v_s be adjacent to v_1, $s \neq 2$. If we add the edge $\{v_1, v_s\}$, we obtain a cycle $v_1, \ldots, v_s, v_1$, and a path $v_s, \ldots, v_r$. Now, if we remove the edge $\{v_{s-1}, v_s\}$, we obtain a new Hamiltonian path $H' = (H \setminus \{v_{s-1}, v_s\}) \cup \{v_1, v_s\}$. This operation is called a backbite move. See Figure I.5.

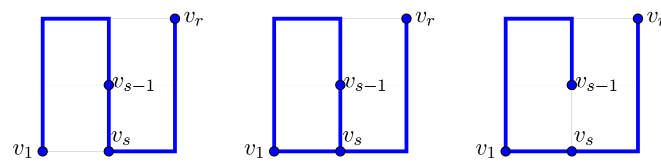

Fig. I.5. An illustration of a backbite move.

Theorem I.2. *Let H and K be any two Hamiltonian paths in an $m \times n$ grid graph G with $n \geq m$. Then there exists a sequence of at most $n^2m + O(n^2)$ valid switch, double-switch, and backbite moves that reconfigures H into K.*

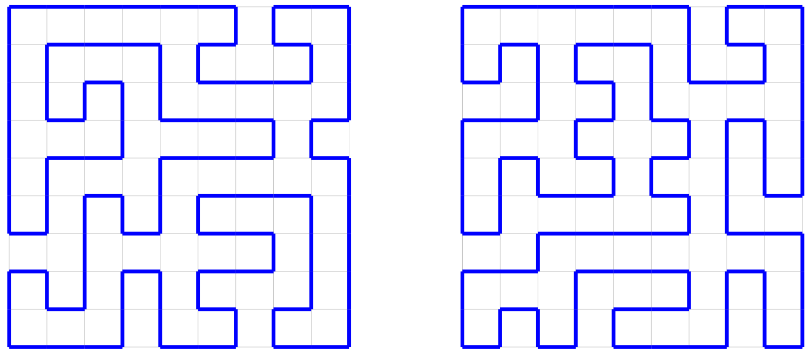

Figure I.6. Two distinct Hamiltonian cycles on 10×10 grid graphs.

Motivation. A self-avoiding walk is a walk in a lattice where every vertex is unique. A Hamiltonian path in a grid graph is an example of a self-avoiding walk. Madras and Slade in [21] present a comprehensive and rigorous study of self-avoiding walks. One application of Theorems I.1 and I.2 is in chemical physics, drawing from the theory of self-avoiding walks. Researchers in [22], [23], [24], and [20] use Monte Carlo methods to study statistical properties of polymer chains, which they abstracted as cycles and paths in the three-dimensional square lattice. They use self-avoiding walks to model how a flexible polymer chain is arranged in a liquid solution. A polymer chain's concentration is the fraction of vertices of the lattice that are occupied by the vertices (monomers) of the polymer. The authors consider maximally concentrated polymers (high-density polymers), where all the space is occupied by the polymer. These can be naturally represented as Hamiltonian paths or cycles. They view the set of Hamiltonian cycles of a rectangular grid graph as the state space of a Markov chain, with the double-switch move being the transition mechanism. Given a Hamiltonian cycle (a state in the state space), we choose two switchable boxes at random and perform a double-switch move. If the move is valid, then the new state is the resulting Hamiltonian cycle. Otherwise, we remain at the initial state and choose another pair of switchable boxes. The idea is that after a sufficiently large number of transitions, we obtain a sequence of many different states, which represents a reasonable random sample of the entire state space.

The goal of these methods is to obtain a uniform random sample of all Hamiltonian cycles of the grid. It follows from the theory of Markov chains that for the Monte Carlo simulation to generate such a distribution of Hamiltonian cycles, the chain must be *reversible* and *irreducible*. Reversibility is satisfied when every move has an inverse, which is true here since the inverse of a double-switch move is obtained by switching the same pair of boxes again. Irreducibility is satisfied when we can reconfigure between any two Hamiltonian cycles using valid moves. Irreducibility ensures that the limit

$$\lim_{n\to\infty} P^n(i,j) = \pi(j)$$

exists, and that it satisfies

$$(1) \quad \sum_i \pi(i) = 1, \qquad (2) \quad \pi(j) = \sum_i \pi(i) P(i,j),$$

where $\pi(j)$ is the long-term probability of Hamiltonian cycle j occurring. Two other conditions, *aperiodicity* (the chain can "stay put" with positive probability, which happens here when moves are rejected) and finiteness of the state space (obviously true for a finite grid graph) are necessary and clearly satisfied but warrant mention for completeness. Furthermore, π is the only nonnegative solution of (1) and (2), and we call π the *equilibrium* distribution.

Since each move has an inverse, the probability $P(i,j)$ of going from a Hamiltonian cycle i to another one j via a single move is the same as the probability of going from j to i. This means that the transition matrix P is symmetric: $P(i,j) = P(j,i)$. We can verify that the uniform distribution satisfies equation (2). If π is uniform (that is, $\pi(i) = c$ for all i), then it follows that

$$\pi(j) = \sum_i c\, P(i,j) = c \sum_i P(i,j) = c.$$

Since the uniform distribution π satisfies (1) and (2), it is *an* equilibrium distribution, and because the equilibrium distribution of an irreducible, aperiodic chain with finite state space is unique, it is *the* equilibrium distribution. The harder condition to prove is irreducibility.

The authors in [22, 23, 24, 20] assume irreducibility but do not prove it. This dissertation proves that irreducibility holds for all rectangular grid graphs, thereby confirming that the corresponding Markov Chain Monte Carlo method indeed produces a uniform equilibrium distribution. We further extend this result to Hamiltonian paths by introducing the single-switch and backbite moves in addition to the double-switch move. For a more detailed discussion on Monte Carlo methods and reconfiguration of self-avoiding walks, see Chapter 9 in [21].

Related work Nishat, Whitesides, and Srinivasan extended the result of [17] to 1-complex Hamiltonian paths in rectangular grid graphs [25, 26, 27], and to 1-complex Hamiltonian cycles in L-shaped grid graphs [28]. The authors define a 1-complex s,t Hamiltonian path to be a 1-complex Hamiltonian path that begins and ends at diagonally opposite corners s and t of a rectangular grid graph. We note that Corollary 5.4.1 extends the results in [25], [26], and [27] to arbitrary s,t Hamiltonian paths.

We prove the existence of move sequences that reconfigure Hamiltonian cycles and paths in rectangular grid graphs (Theorems I.1 and I.2). The proofs rely on structural properties that Hamiltonian paths impose on grid graphs, which Chapter 1 analyzes in detail. Chapter 2 presents the reconfiguration algorithm for cycles (Theorem I.1[1],), whose correctness depends on two algorithms: MLC and 1LC. Chapter 3 proves these algorithms exist and shows that the 1LC proof requires a technical lemma (Lemma 3.13) handling a difficult family of configurations. Chapter 4 analyzes this family and proves Lemma 3.13, while also establishing results needed for Chapter 5. Chapter 5 extends the reconfiguration result to Hamiltonian paths (Theorem I.2, restated there as Theorem 5.9), building on the cycle algorithms (Chapter 2), the structural analysis (Chapter 1), and the tools from Chapter 4.

[1] We prove Theorem 2.1 which is a slight generalization of Theorem 1.I.

1 Structural properties of Hamiltonian paths and cycles on grid graphs

A *grid graph* is a subgraph of the integer grid $\mathbb{Z}^2$. A *lattice animal* is a finite connected subgraph of $\mathbb{Z}^2$. A *Hamiltonian path (cycle)* of a graph G is a path (cycle) that visits each vertex of the graph exactly once. Assume that G has a cut vertex v. Then G cannot have a Hamiltonian cycle. Let G_1 and G_2 be the components of $G \setminus v$. Let H_1 be a Hamiltonian path of $G \setminus G_1$ and let H_2 be a Hamiltonian path of $G \setminus G_2$ such that H_1 and H_2 have v as an end-vertex. Then a Hamiltonian path H of G can be obtained by concatenating H_1 and H_2. Since H_1 and H_2 are smaller than H, they are easier to find and reconfigure. It follows that a graph that cannot be decomposed in this manner must be 2-connected. Thus, from here on, we will restrict our attention to 2-connected grid graphs.

We will show that, for a subclass of 2-connected grid graphs called rectangular grid graphs, we can reconfigure a Hamiltonian path into any other Hamiltonian path, by making a sequence of moves. We define rectangular grid graphs further down and we define moves in Section 1.4. We remark that Section 1.4 requires only the definitions from the opening of this chapter and the definition of a polyomino graph from Section 1.1. Readers interested in understanding the moves before proceeding through the more dense technical Sections 1.1–1.3 may prefer to read Section 1.4 immediately.

Suppose H and K are two distinct Hamiltonian paths of G. We will show that there exists a sequence of moves $\mu_1, \ldots, \mu_r$ and a sequence of Hamiltonian paths $H_0, H_1, \ldots, H_r$, where $H = H_0$ and $K = H_r$, such that for each $j \in \{1, \ldots, r\}$ μ_j reconfigures H_{j-1} into H_j.

Notation. Let G be a grid and let H be a Hamiltonian path of G. We will call the unit square faces of the integer grid *boxes* (see Figure 1.2). We denote by Boxes(G) the set of boxes of a grid graph G that have all four of their edges in G. From here on we will assume that G is finite. We will need some definitions to navigate G and H. We will position G in the first quadrant of the integer grid in such a way that the westernmost vertices of G have x-coordinate zero and the southernmost vertices of G have y-coordinate zero. We use the x and y coordinates to describe a rectangle in the graph and denote it $R(k_1, k_2; l_1, l_2)$. This rectangle corresponds to the rectangle on the plane which is the Cartesian product of closed intervals $[k_1, k_2] \times [l_1, l_2]$. We will denote a box of G by $R(k, l)$, where k and l are the coordinates of the corner of the box that is closest to the origin.

We denote a vertex v by $v(k, l)$, where k and l are the vertex coordinates. We denote a horizontal edge e by $e(k_1, k_2; l_1)$, where k_1, k_2 are the x-coordinates of the vertices of e, and l_1 is the y-coordinate of the vertices of e. Similarly, we write $e(k_1; l_1, l_2)$ for vertical edges. For a horizontal line consisting of r consecutive edges, we write $e(k, k+r; l)$; we denote vertical lines analogously. It will be convenient to use the notation $\{u, v\}$ to denote edges of G, and the notation (u, v) to denote directed edges of G. For a directed edge $e = (u, v)$, u is said to be the *tail* of e and v is said to be the *head* of e.

Let $P(u, v)$ and $P(v, w)$ be paths. We denote the concatenation of $P(u, v)$ with $P(v, w)$ as $P(u, v), P(v, w)$ (a comma indicates concatenation).

The rest of Chapter 1 contains definitions and technical results used in Chapters 2-5. A flowchart at the end of the chapter illustrates dependencies within Chapter 1 and connections to later chapters.

Our reconfiguration strategy relies on controlling which edges belong to the Hamiltonian cycle by applying moves. To analyze when moves can add or remove specific edges while preserving the Hamiltonian property, we need to understand how H decomposes the grid graph into components. Chapter 1 explores the structure that arises from this decomposition.

Consider a Hamiltonian path H on the $m \times n$ *grid graph* G in Figure 1.4. Section 1.2 describes how H decomposes G into H-components. Since these components need not be rectangular, Section 1.1 first defines polyomino graphs as a generalization of rectangular grids.

Section 1.3 proves technical results about H-component structure needed in Chapters 3-5. Section 1.4 gives conditions for valid moves and describes their effects on edge orientations and path structure. See the Flowchart 1 at the end of the chapter for detailed dependencies.

1.1 Polyomino graphs

In this section we prove preliminary results about polyomino graphs. Corollary 1.1.5, a consequence of Jordan's Curve Theorem, is used throughout later proofs. Results 1.1.7-1.1.11 characterize structural properties of polyominoes needed in Section 1.3 and Chapters 3-5.

Definitions. Let G be a planar graph with face set F, and let $F' \subseteq F$. Analogous to a vertex-induced subgraph, we define the *face-induced subgraph* $G[F']$ to be the subgraph of G consisting of all vertices and edges incident to the faces in F'.

We restrict our attention to a subclass of 2-connected, finite, box-induced subgraphs of the square lattice that we call *polyomino graphs* (see Figure 1.2). We define polyomino graphs recursively as follows:

1. The polyomino graph of order 1 is the box-induced subgraph on a single box. This is the cycle C_4.

2. Let $\mathcal{P}_k$ denote the set of polyomino graphs of order k. Let $\{X_1, \ldots, X_k\}$ be a set of boxes in $\mathbb{Z}^2$, and let the polyomino graph $G \in \mathcal{P}_k$ be the box-induced subgraph on this set of boxes. We may obtain a polyomino graph $G' \in \mathcal{P}_{k+1}$ by taking the box-induced subgraph on the set $\{X_1, \ldots, X_k, X\}$, where $X \in \mathbb{Z}^2 \setminus \{X_1, \ldots, X_k\}$ shares an edge with some box in $\{X_1, \ldots, X_k\}$.

Remark. It is possible for two polyomino graphs $G \in \mathcal{P}_k$ and $G' \in \mathcal{P}_{k+r}$, for some $r > 0$, to be equal as subgraphs of the square lattice, since different sets of boxes can induce the same subgraph, as we shall now show. Consider a polyomino $G \in \mathcal{P}_k$, constructed recursively using the procedure above. Then G is the box-induced subgraph on a set of boxes $\{X_1, \ldots, X_k\}$ of $\mathbb{Z}^2$.

We define a *one-hole* of G to be a box $X \in \mathbb{Z}^2$ such that:

1. $X \notin \{X_1, \ldots, X_k\}$, and
2. X shares an edge with four distinct boxes in $\{X_1, \ldots, X_k\}$. See Figure 1.1.

In this case, all four edges and all four vertices of X are already contained in G, so $X \in \text{Boxes}(G) \setminus \{X_1, \ldots, X_k\}$. We may now define $G' \in \mathcal{P}_{k+1}$ as the box-induced subgraph on $\{X_1, \ldots, X_k, X\}$. Since the addition of X introduces no new edges or vertices to the graph, we have $G = G'$.

It will be convenient to define here a generalization of a polyomino graph called a *polyking* graph. A polyking graph of order k is defined to be a box-induced subgraph on k boxes, constructed recursively in the same way as a polyomino graph, except that each new box is only required to share a vertex (rather than an edge) with an existing box.

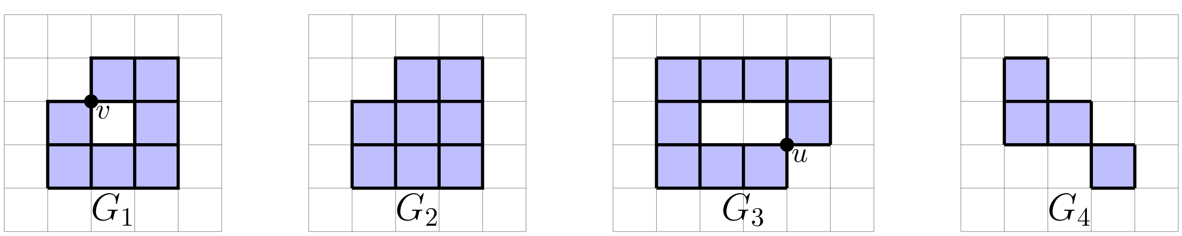

Fig. 1.1. $G_1 \in \mathcal{P}_7$ is a polyomino with a one-hole. $G_2 \in \mathcal{P}_8$ is a polyomino equal to G_1. G_3 is a polyomino with a polyking junction at the vertex u. Note that v in G_1 is not a polyking junction because it is incident on three boxes of G. G_4 is a polyking.

Let G be a polyomino graph. We say that G has a *polyking junction* at a vertex $v(k, l)$ if there are exactly two boxes Z and Z' of G incident on $v(k, l)$ and $V(Z) \cap V(Z') = \{v(k, l)\}$.

From here on, we will use the term *polyomino* to mean polyomino graph, and *polyking* to mean polyking graph.

Let G be a polyomino, and let X_1, X_2 be two distinct boxes of G. If X_1 and X_2 share an edge of G, we say that X_1 and X_2 are adjacent. Define a *walk of boxes in* G to be a sequence $X_1, \dots, X_r$ of boxes in G, not necessarily distinct, such that for all $j \in \{1, 2, \dots, r-1\}$, either X_j is adjacent to X_{j+1} or $X_j = X_{j+1}$. We denote such a walk by $W(X_1, X_r)$. For each $j \in \{1, \dots, r-1\}$, we call the edge of G shared by X_j and X_{j+1} a *gluing edge of* $W(X_1, X_r)$, whenever X_j and X_{j+1} are distinct boxes. If for all $i, j \in \{1, 2, \dots, r\}$ with $i \neq j$, X_i is distinct from X_j, we call the sequence a *path of boxes in* G and denote it by $P(X_1, X_r)$. The length of $P(X_1, X_r)$ is $r-1$. A *cycle of boxes in* G is a walk $X_1, \dots, X_r$ such that $X_1 = X_r$ and for all $i, j \in \{1, \dots, r-1\}$, $X_i \neq X_j$. Note that, by construction, for any two boxes X and Y in G, there is at least one path of boxes $P(X, Y)$ in G. We say that G is *box-path connected*.

A graph G is said to be k-*connected* if it remains connected after the removal of any $k-1$ vertices.

Lemma 1.1.1. Every face of a 2-connected planar graph is bounded by a cycle. (Proposition 2.1.5 in Mohar and Thomassen, Graphs on Surfaces, p.21.[29]) □

Lemma 1.1.2. Polyominoes are 2-connected.

Proof. Let G be a polyomino. Suppose we remove the vertex w from G. We want to show that G is still connected. Let $u, v \in G \setminus w$ be incident on the boxes X_u and X_v of G, respectively. Since polyominoes are box-path-connected, there is a path of boxes $P(X_u, X_v)$ in G. We show that we can find a path of vertices $P(u, v)$ in $G \setminus w$ by induction on the length of the shortest path of boxes $P(X_u, X_v)$.

Base Case. Suppose $P(X_u, X_v)$ in G has length zero. Then $P(X_u, X_v) = X_u = X_v$ is the cycle C_4, which is 2-connected.

Inductive Case. Assume that if $P(X_u, X_v)$ in G has length k then we can find a path of vertices $P(u, v)$ in $G \setminus w$. Suppose $P(X_u, X_v)$ has length $k+1$, let the box Y be the $(k+1)^{\text{st}}$ box of $P(X_u, X_v)$, and let $\{u_1, u_2\}$ be the gluing edge between Y and X_v. Remove w from G. Either w is one of u_1 and u_2, or it is not.

If the latter, by inductive hypothesis, there is a path $P(u, u_1)$ contained in $P(X_u, Y)$ and avoiding w. By base case, there is a path $P(u_1, v)$ contained in X_v and avoiding w. Then, $P(u, u_1)$, $P(u_1, v)$ is a (u, v)-path in $G \setminus w$ contained in $P(X_u, X_v)$.

Assume the former. By symmetry, we may assume without loss of generality that $w = u_1$. By inductive hypothesis, there is a path $P(u, u_2)$ contained in $P(X_u, Y)$ and avoiding w. By base case, there is a path $P(u_2, v)$ contained in X_v and avoiding w. Then, $P(u, u_2)$, $P(u_2, v)$ is a (u, v)-path in $G \setminus w$ contained in $P(X_u, X_v)$. □.

Observation 1.1.3. Let G be a polyomino. Then:

(a) Every edge of G is incident on a box of G.

(b) A box-path-connected set of boxes is a polyomino.

Theorem 1.1.4 (Jordan's Curve Theorem for polygons[29]**).** Let Q be a simple polygon in the plane. Then the set $\mathbb{R}^2 \setminus Q$ consists of two disjoint subsets, called the "interior" (Int) and "exterior" (Ext), each of which has Q as its boundary. Moreover, any two points in Int (or in Ext) can be joined by a polygonal path that does not intersect Q, and any polygonal path joining a point of Int to a point of Ext must intersect Q □.

We record here a useful consequence of Jordan's Curve Theorem for polygons (JCT).

Corollary 1.1.5 Let Q be a simple polygon and let p_1 and p_2 be points in the plane not on Q. If the segment $[p_1, p_2]$ intersects Q exactly once at a point q that is not a vertex of Q, then one of p_1 and p_2 is in Ext and the other is in Int. □

Definitions. Let G be a polyomino. We will denote the cycle that bounds the outer face Outer(G) of a polyomino G by $B_0(G) = B_0$. Define the *enclosure of* G to be the set of boxes contained in the region of the plane bounded by B_0, and denote it by Encl(G). By JCT, B_0 divides the plane into two disjoint sets of boxes: Outer(G) and Encl(G).

We define the *boundary of* G to be the edge-induced subgraph on the set of edges that are incident on a box of G and a box of $\mathbb{Z}^2$ that is not a box of G, and denote it by B(G). We define a *boundary box* of G to be a box that is incident on B(G) but that is not a box of G. Denote the set of all boundary boxes of G by BBoxes(G). See Figure 1.2.

Consider the set of boxes Encl(G) $\setminus$ Boxes(G). Let $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$ be the set of maximal box-path connected sets of boxes of Encl(G)$\setminus$Boxes(G). We call each member of $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$ a *hole* of G and we denote the Encl(G) $\setminus$ Boxes(G) by Holes(G).

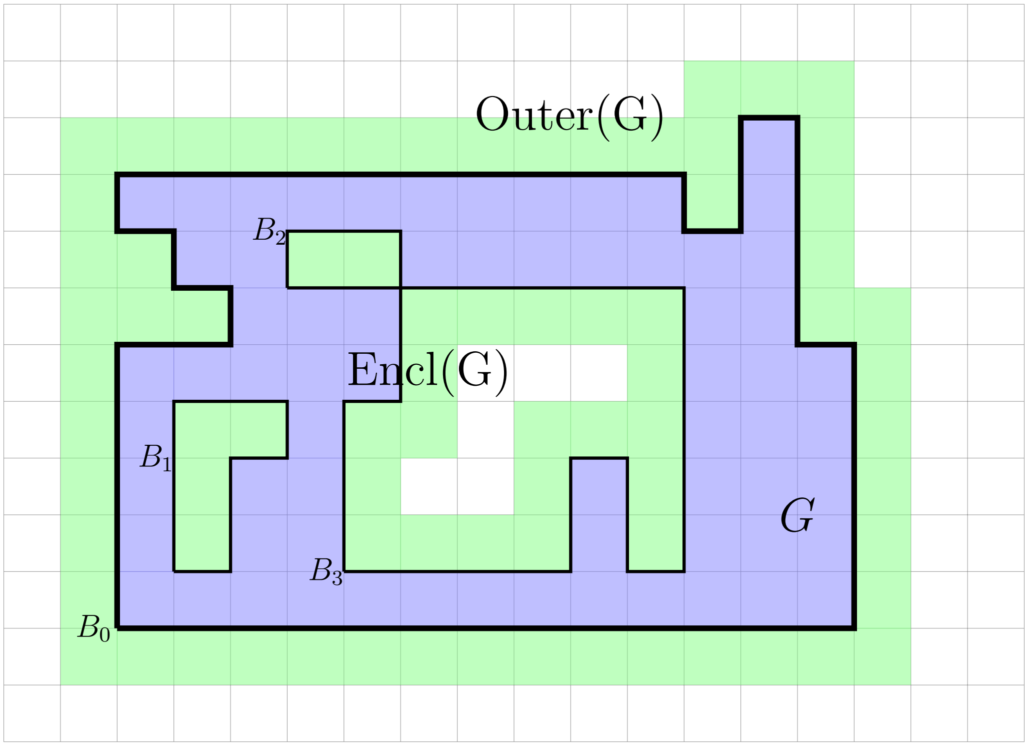

Figure 1.2. A polyomino G. Boxes of G, Boxes(G), are in blue. Boundary boxes of G, BBoxes(G), are in green. Boundary of G, $B(G)$, is in black. B_0 is in thick black. For $i \in \{1, 2, 3\}$, $B_i = B_0(\mathcal{O}_i)$ in thin black bounds the hole $\mathcal{O}_i$. See also Observation 1.1.9.

We define a *simply connected polyomino* G to be a polyomino G such that Boxes(G) = Encl(G). Let G be a simply connected polyomino such that $B(G)$ is an $m \times n$ vertices rectangle. That is, the horizontal side of $B(G)$ has length $m - 1$ and the vertical side of $B(G)$ has length $n - 1$. Then we call G an *$m \times n$ rectangular grid graph without holes ($m \times n$ grid graph)*. Note that this is the opposite of the usual convention for matrices. We have the following hierarchy of lattice animals:

$m \times n$ grid graph $m \times n$ grid graphs $\subset$ simply connected polyominoes $\subset$ polyominoes $\subset$ 2-connected grid graphs $\subset$ lattice animals. See Figure 1.3.

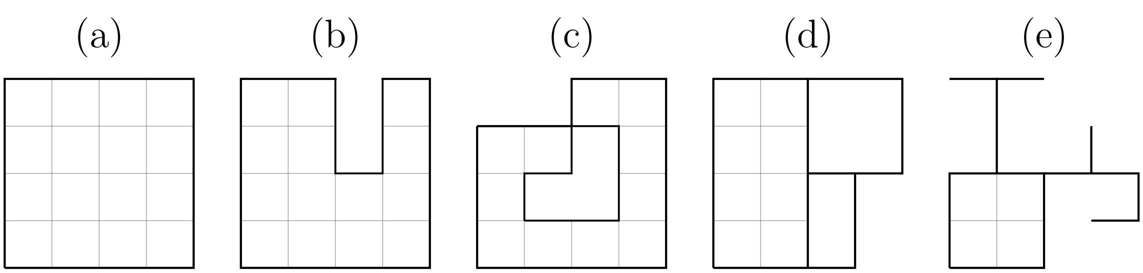

Figure 1.3. (a) is an $m \times n$ grid graph; (b) is a simply connected polyomino; (c) is a polyomino; (d) is a 2-connected grid graph; (e) is a lattice animal.

Observation 1.1.6. Let G be a polyomino with holes $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$. Then:

(a_1) $\mathrm{Boxes}(G) \cap \mathrm{Outer}(G) = \emptyset$.

(a_2) $\mathrm{Boxes}(G) \subseteq \mathrm{Encl}(G)$.

(b_1) For each $i \neq j$, $E(B_0(\mathcal{O}_i)) \cap E(B_0(\mathcal{O}_j)) = \emptyset$.

(b_2) For each $i \in \{1, .., q\}$, $E(B_0(G)) \cap E(B_0(\mathcal{O}_i)) = \emptyset$.

(c) For each $i \in \{1, .., q\}$, $B_0(\mathcal{O}_i) \subset \mathrm{Encl}(G)$.

(d) For each $i \in \{1, .., q\}$, $\mathcal{O}_i$ is a polyomino. □

Lemma 1.1.7. Let G be a polyomino with holes $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$. Then:

(a) An edge e is in $B_0(G)$ iff e is incident on a box of $\mathrm{Outer}(G)$ and on a box of $\mathrm{Boxes}(G)$.

(b) For each $i \in \{1, .., q\}$, $E(B_0(\mathcal{O}_i)) \subset E(B(G))$. Furthermore, every edge $e \in E(B_0(\mathcal{O}_i))$ is incident on a box of $\mathcal{O}_i$ and on a box of G.

(c) For each $i \in \{1, .., q\}$, $\mathrm{Encl}(\mathcal{O}_i) \subset \mathrm{Encl}(G)$.

(d) For every edge $e_X \in B_0(G)$ there is a box $X \in \mathrm{Boxes(G)} \setminus \bigcup_{i=1}^{q} \mathrm{Encl}(\mathcal{O}_i)$ such that X is incident on e_X.

(e) For each $i \in \{1, .., q\}$, $\mathrm{Boxes(G)} \cap \mathrm{Encl}(\mathcal{O}_i) = \emptyset$.

(f) The holes of G are simply connected polyominoes.

Proof of (a). Let e be an edge of B_0. Let Y be the box contained in $\mathrm{Outer}(G)$ that is incident on e, and let X be the box contained in $\mathrm{Encl}(G)$ that is incident on e. By Observation 1.1.3(a), X must be a box in $\mathrm{Boxes}(G)$. Conversely, assume that e is incident on a box of $\mathrm{Outer}(G)$ and a box of $\mathrm{Encl}(G)$. Then e is contained in the cycle that bounds $\mathrm{Outer}(G)$. That is, e is an edge of B_0. End of proof for (a).

Proof of (b). Let $e \in E(B_0(\mathcal{O}_i))$. By Observation 1.1.6(d), $\mathcal{O}_i$ is a polyomino. By part (a), e is incident on a box of $\mathcal{O}_i$ and a box of $\mathrm{Outer}(\mathcal{O}_i)$. Let X be the box of $\mathrm{Outer}(\mathcal{O}_i)$ incident on e, and let Y be the box of $\mathcal{O}_i$ incident on e. By Observation 1.1.6(b_2), X cannot be in $\mathrm{Outer}(G)$. Otherwise, since $Y \in \mathcal{O}_i \subset \mathrm{Encl}(G)$ and $X \in \mathrm{Outer}(G)$, it would follow that $e \in B_0(G)$. This would imply $e \in E(B_0(G)) \cap E(B_0(\mathcal{O}_i)) \neq \emptyset$, contradicting ($b_2$). Therefore, $X \in \mathrm{Encl}(G)$. Note that if $X \in \mathrm{Encl}(G) \setminus \mathrm{Boxes}(G)$, this would contradict the maximality of $\mathcal{O}_i$ (since X is adjacent to $\mathcal{O}_i$, we could simply include it in $\mathcal{O}_i$). Thus, X must be a box of G. Therefore, by definition, $e \in E(B(G))$. End of proof for (b).

Proof of (c). By Observation 1.1.3(b), $\mathrm{Encl}(\mathcal{O}_i)$ is a polyomino. Let $X \in \mathrm{Encl}(\mathcal{O}_i)$. Then $X \in \mathrm{Encl}(G)$ or $X \in \mathrm{Outer}(G)$. Assume, for contradiction, that $X \in \mathrm{Outer}(G)$. Let

$Y \in \mathrm{Encl}(\mathcal{O}_i)$ be adjacent to a boundary box Y' of $\mathcal{O}_i$. Consider the path of boxes $P(X, Y)$. This path is either contained in $\mathrm{Outer}(G)$ or not.

CASE 1: $P(X, Y) \subset \mathrm{Outer}(G)$. Let e be the edge that Y and Y' share. Then $e \in E(B_0(\mathcal{O}_i))$. By Lemma 1.1.7(b), $e \in E(B(G))$. Since $Y \in \mathrm{Outer}(G)$, by Lemma 1.1.7(a), $e \in E(B_0(G))$. But this contradicts Observation 1.1.6(b_2). End of Case 1.

CASE 2: $P(X, Y) \not\subset \mathrm{Outer}(G)$. Let $P(X, Y) = Z_1, \ldots, Z_s$. Then there must be a box Z_j in $P(X, Y)$ such that $P(Z_1, Z_j)$ is entirely contained in $\mathrm{Outer}(G)$, while $P(Z_1, Z_{j+1})$ is not. Let f be the edge that Z_j and Z_{j+1} share. Then, by definition, $f \in E(B_0(G))$, and by Lemma 1.1.7(a), Z_{j+1} is a box of G. Since $Z_j \in \mathrm{Encl}(\mathcal{O}_i)$ and $Z_{j+1} \in \mathrm{Boxes}(G)$, it follows that $f \in E(B_0(\mathcal{O}_i))$. But this contradicts Observation 1.1.6(b_2). End of proof for (c).

Proof of (d). Let $Z \in \mathrm{Boxes}(G)$ be a box incident on an edge $e \in B_0(G)$, and let Z' be the other box incident on e. If Z is in $\mathrm{Boxes}(G) \setminus \bigcup_{i=1}^q \mathrm{Encl}(\mathcal{O}_i)$, we are done. We will now show that the alternative is impossible.

For contradiction, assume that there is some $i \in \{1, \ldots, q\}$ such that $Z \in \mathrm{Boxes}(G) \cap \mathrm{Boxes}(\mathcal{O}_i)$. By definition, $Z' \in \mathrm{Outer}(G)$. If there is some $i \in \{1, \ldots, q\}$ such that $e \in B_0(\mathcal{O}_i)$, this contradicts Observation 1.1.6(b_2). Otherwise, if no such i exists, then Z' must also belong to $\mathrm{Boxes}(\mathcal{O}_i)$. But then, by part (c), $Z' \in \mathrm{Encl}(G)$, which contradicts the assumption that $Z' \in \mathrm{Outer}(G)$. End of proof for (d).

Proof of (e). Assume, for contradiction, that there is some $i \in \{1, \ldots, q\}$ such that $\mathrm{Boxes}(G) \cap \mathrm{Encl}(\mathcal{O}_i) \neq \emptyset$. Let X be in $\mathrm{Boxes}(G) \cap \mathrm{Encl}(\mathcal{O}_i)$. By part (d), $\mathrm{Boxes}(G) \setminus \bigcup_{i=1}^q \mathrm{Encl}(\mathcal{O}_i) \neq \emptyset$. Let $Y \in \mathrm{Boxes}(G) \setminus \bigcup_{i=1}^q \mathrm{Encl}(\mathcal{O}_i)$.

Since $Y \in \mathrm{Boxes}(G) \setminus \bigcup_{i=1}^q \mathrm{Encl}(\mathcal{O}_i)$, and $X \in \mathrm{Boxes}(G) \cap \mathrm{Encl}(\mathcal{O}_i)$, X and Y are on different sides of $B_0(\mathcal{O}_i)$, and they both belong to $\mathrm{Boxes}(G)$. Then there is a path of boxes $P(X, Y)$ in G with $Z_1 = X$, and so that it contains t boxes, with $t \geq 2$. By the Jordan Curve Theorem, there is $j \in \{1, \ldots, t-1\}$ such that $P(Z_1, Z_j) \subset \mathrm{Encl}(\mathcal{O}_i)$, but $P(Z_1, Z_{j+1}) \not\subset \mathrm{Encl}(\mathcal{O}_i)$.

It follows that the gluing edge e between Z_j and Z_{j+1} is in $B_0(\mathcal{O}_i)$. Then, by part (b), one of Z_j and Z_{j+1} is in G, and the other is in $\mathcal{O}_i$. But this contradicts the assumption that $P(X, Y)$ is entirely contained in G. End of proof for (e).

Proof of (f). We will check that for each $i \in \{1, \ldots, q\}$, $\mathrm{Encl}(\mathcal{O}_i) \subset \mathrm{Boxes}(\mathcal{O}_i)$. If $\mathrm{Encl}(\mathcal{O}_i)$ has no holes, we are done, so assume that $\mathrm{Encl}(\mathcal{O}_i)$ has a hole U. By part (b), there is an edge $e \in B_0(U)$ that is incident on the boxes $Z \in U$ and $Z' \in \mathcal{O}_i$. By maximality of holes (of G, in this case), Z does not belong to any hole of G. Then $Z \in \text{Boxes}(G)$ or $Z \in \text{Outer}(G)$. If the former, then, since Z is also in $\text{Encl}(\mathcal{O}_i)$, this contradicts part (e). And if the latter, then, since $Z' \in \text{Boxes}(\mathcal{O}_i) \subset \text{Encl}(\mathcal{O}_i)$, e belongs to $B_0(G)$. But then, since neither Z, nor Z' belong to G, this contradicts part (a). End of proof for (f). □

Lemma 1.1.8. Let G be a simply connected polyomino. Then the boundary of G is B_0.
Proof. Let $e \in B(G)$. Then e is incident on a box X of G and on a boundary box Y of G. By Observation 1.1.6(a_2), $X \in \text{Boxes}(G)$ implies that $X \in \text{Encl}(G)$. Either Y belongs to $\text{Encl}(G)$ or Y belongs to $\text{Outer}(G)$. If $Y \in \text{Encl}(G)$, since G is a simply connected polyomino, then Y is a box of G. But then Y is both a box of G and a boundary box of G, which is a contradiction. It must be the case that $Y \in \text{Outer}(G)$. By Lemma 1.1.7(a), e belongs to B_0.
□

Observation 1.1.9. Let G be a polyomino with holes $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$, bounded by the cycles $B_1, \ldots, B_q$. It follows from Lemma 1.1.7(f) that $B(G)$ is the disjoint union of $B_0, B_1, \ldots, B_q$. For each $i \in \{0, \ldots, q\}$, we call B_i a *boundary component of* G.

Corollary 1.1.10. G is a simply connected polyomino if and only if $B(G) = B_0(G)$.

Proof. By Lemma 1.1.8 we only need to show that if $B(G) = B_0(G)$ then G is a simply connected polyomino. We prove the contrapositive. Suppose that G is not simply connected. Then the set $D = \text{Encl}(G) \setminus \text{Boxes}(G)$ is nonempty. Let $\{\mathcal{O}_1, \ldots, \mathcal{O}_q\}$ and $B_0, B_1, \ldots, B_q$ be as in Observation 1.1.9. Then $q \geq 1$. By Observation 1.1.9, B_1 and B_0 are disjoint. It follows that $B(G) \neq B_0(G)$. □

Lemma 1.1.11. Simply connected polyominos have no polyking junctions.

Proof. Let G be a simply connected polyomino. By Corollary 1.1.10 $B(G) = B_0(G)$. By Lemma 1.1.1, $B_0(G)$ is a cycle. For a contradiction, assume that G has a polyking junction at a vertex v. Observe that this implies that all edges incident on v are boundary edges of G. But then $B_0(G)$ has a vertex of degree 4, contradicting that it is a cycle. □

1.2 The H-decomposition of G and the follow-the-wall construction

This section describes the H-decomposition of G in detail and proves that H-components have a tree-like structure. We introduce the follow-the-wall construction, a procedure for navigating H-components that appears extensively in later proofs (Chapters 1, 3, 4, and 5).

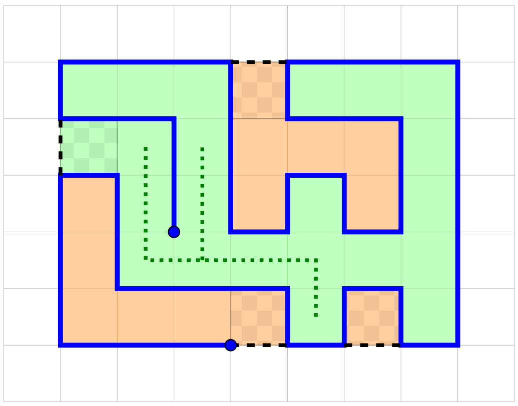

Fig. 1.4. An 8×6 grid graph with a Hamiltonian path in blue; a self-adjacent H-component of G shaded green; the non-self-adjacent H-components of G shaded orange; a non-maximal H-tree dotted dark green; the necks of H-components are checkered; the neck edges are dashed black; BBoxes$(G) = G_{-1} \setminus G$ in white.

Definitions. A *walk* (of length r) in a graph G is an alternating sequence $v_0e_1v_1e_2 \dots e_rv_r$ of vertices and edges. Define a *lazy walk* to be a sequence of edges and vertices where every edge is in between two vertices and in between every two edges there is a vertex or multiple copies of a vertex. That is, a lazy walk is roughly a walk in which consecutive vertices can be the same, allowing the walk to remain at a vertex for one or more steps without traversing any edges.

Let G be any polyomino and let H be any subgraph in G. Let X_1, X_2 be two adjacent boxes of G. If $E(X_1) \cap E(X_2) \cap E(H) = \emptyset$, we say that X_1 and X_2 are *H-neighbours* or X_1 is *H-adjacent* to X_2. Define an *H-walk of boxes in G* (H-walk) to be a sequence $X_1, \dots, X_r$ of boxes in G, not necessarily distinct, such that for all $j \in \{1, 2, \dots, r-1\}$, X_j is an H-neighbour of X_{j+1} or $X_j = X_{j+1}$. If for all $i, j \in \{1, 2, \dots, r\}$ with $i \neq j$, X_i is distinct from X_j, we call the sequence an *H-path of boxes in G* and denote it by $P(X_1, X_r)$. If in addition, for all $i, j \in \{1, 2, \dots, r-1\}$ with $i \neq j$, $E(X_i) \cap E(X_j) \cap E(H) = \emptyset$, we say that $P(X_1, X_r)$ is a *non-self-adjacent H-path of boxes in G*. Otherwise, if $E(X_i) \cap E(X_j) \cap E(H) \neq \emptyset$, we say that $P(X_1, X_r)$ is a *self-adjacent H-path of boxes in G*. Let $r \geq 4$. Define an *H-cycle C of boxes in G* (H-cycle) to be a set $X_1, X_2, \dots, X_r = X_1$ of boxes in G such that for each $j \in \{1, 2, \dots, r-1\}$, X_j is an H-neighbour of X_{j+1} and the boxes $X_1, \dots, X_{r-1}$ are distinct. We note that every box of C has exactly two gluing edges. Proposition 1.2.1. will show that if H is a Hamiltonian cycle of G, then there are no H-cycles of boxes in G.

By definition, BBoxes$(G) \cap$ Boxes$(G) = \emptyset$. We use the notation G_{-1} to denote the union of the graph on BBoxes(G) and G. We extend the definitions of H-walks, H-cycles, H-paths, and self-adjacency to G_{-1}. See Figure 1.2, where G_{-1} consists of the green boxes and blue boxes

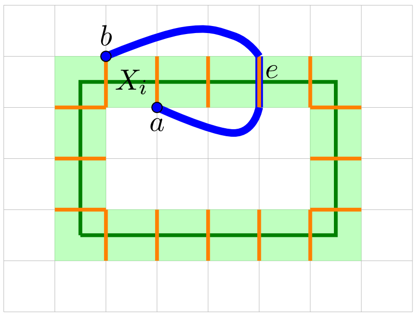

Fig. 1.5. $P(a,b)$ in blue, gluing edges of C in orange, Q in dark green, C shaded light green.

Proposition 1.2.1. Let G be a polyomino and let H be a Hamiltonian path of G. Then every H-cycle in G_{-1} is contained in BBoxes(G).

Proof. For a contradiction, assume that there is an H-cycle C in G_{-1} with boxes $X_1, X_2, \ldots, X_r = X_1$ that has a box X_i contained in G. Let $c_1, c_2, \ldots, c_r = c_1$ be the centers of the boxes of C. That is, if $X_j = R(k,l)$ then $c_j = (k + \frac{1}{2}, l + \frac{1}{2})$. Observe that for each $j \in \{1, 2, \ldots, r-1\}$ $\{1, 2, \ldots, r-1\}$, $[c_j, c_{j+1}]$ intersects the gluing edge of X_j and X_{j+1} and $[c_j, c_{j+1}]$ intersects no other edge of G_{-1}.

We will first show that the set of segments $[c_j, c_{j+1}]$ is a non-self-intersecting polygon Q. Since $c_1 = c_r$, Q is a polygon. For a contradiction, assume that Q is self intersecting, so there are points c_{j-1}, c_j, c_{j+1} and c_i such that the segments $[c_{j-1}, c_j]$, $[c_j, c_{j+1}]$ and $[c_j, c_i]$ are edges of Q. But then X_j has three gluing edges, a contradiction.

By JCT, Q divides the plane into two subsets Int and Ext such that any two points within a subset can be joined by a path that does not intersect Q while a path joining a point of Int to a point of Ext must intersect Q. Let $V(\text{Int})$ be the vertices of G_{-1} contained in Int and let $V(\text{Ext})$ be the vertices of G_{-1} contained in Ext. Note that any box of C contains at least one vertex in $V(\text{Int})$ and one vertex in $V(\text{Ext})$, so both sets are nonempty. In particular, this is true for X_i. For definiteness, assume that $a \in V(X_i)$ is contained in $V(\text{Int})$ and $b \in V(X_i)$ is contained in $V(\text{Ext})$. By JCT, $V(\text{Int}) \cup V(\text{Ext}) = V(G)$ and $V(\text{Int}) \cap V(\text{Ext}) = \emptyset$. Consider the subpath $P(a,b)$ of H. By JCT again, there is an edge $e \in P(a,b) \subset H$ intersecting Q at some segment $[c_j, c_{j+1}]$. But then e is a gluing edge of C, so e cannot belong to H. See Figure 1.5. □

Let G be a polyomino and let H be a Hamiltonian path or cycle of G. Let $X_1, \ldots, X_r$ be a set of boxes in G such that for any $i, j \in \{1, 2, \ldots, r\}$, there is an H-path $P(X_i, X_j)$ between X_i and X_j contained in G. We say that X_i and X_j are *H-path-connected* in G and the set of boxes $\{X_1, \ldots, X_r\}$ is an *H-path connected set of boxes in G*. If an H-path-connected set of boxes in G contains no cycles of boxes we call it an *H-tree*. An *H-component* of G is a maximal H-path connected set of boxes.

Let J be an H-subtree of an H-component of G. We say that J is non-self-adjacent if it contains no self-adjacent paths of boxes. Otherwise we say that J is self-adjacent. We define a *neck of* J to be a box N_J of J that is incident on a boundary edge e_J of G such that $e_J \notin H$. We call e_J a *neck-edge* of J. Note that the other box incident on e_J must be in $G_{-1} \setminus G$. See Figure 1.4.

Corollary 1.2.2. Let G be a polyomino and let H be a Hamiltonian path of G. Then the H-path $P(X,Y)$ in G is unique, if it exists. Furthermore, H partitions the boxes of G into H-path-connected H-components which are maximal *H-trees.* □

The uniqueness of H-paths contained in G is a key structural property of the decomposition that a Hamiltonian path H imposes on the boxes of a polyomino G. We note that such an H-path exists if and only if X and Y belong to the same H-component. Most H-paths we consider in the results that follow lie entirely within a single H-component. We consider H-paths that contain boxes in G_{-1} only a handful of times. This uniqueness property is used implicitly throughout the dissertation—appearing tacitly to avoid cluttering the text with excessive references.

Corollary 1.2.3[2] Let G be a polyomino and let H be a Hamiltonian path of G. Then every H-cycle in $\mathbb{Z}^2$ is contained in $\text{Boxes}(\mathbb{Z}^2) \setminus \text{Boxes}(G)$. □

Definitions. Let G be a graph. A *trail* is a walk in G where all edges are distinct. A trail where the first and last vertices coincide is called a *closed trail* or a *circuit.* A *directed walk* (of length s) is an alternating sequence $v_0e_1v_1e_2 \dots e_sv_s$ of vertices and directed edges such that for $j \in \{1, \dots, s\}$, the directed edge e_j has tail v_{j-1} and head v_j. A *directed trail* is a directed walk where all directed edges are distinct. Note that the edges (u,v) and (v,u) in a directed trail are distinct edges. Similarly, a directed trail $v_0e_1v_1e_2 \dots e_sv_s$ and its reversal $v_se_s \dots e_2v_2e_1v_1$ are distinct directed trails. We will use the notation $\overrightarrow{K}$ to denote directed trails.

Let the box X be incident on a directed edge (u,v) of the integer grid. Then X and its vertices not incident on the edge e_j are either on the *right* or on the *left* side of (u,v). See Figure 1.6 for an illustration and Note 1.2.5 at the end of this section for a more precise definition of a box being on the right or left side of a directed edge. Note that if X is on the right side of (u,v), then the other box incident on (u,v), say X', is on the right side of (v,u).

[2] This Corollary is only used once by Lemma 1.3.1.

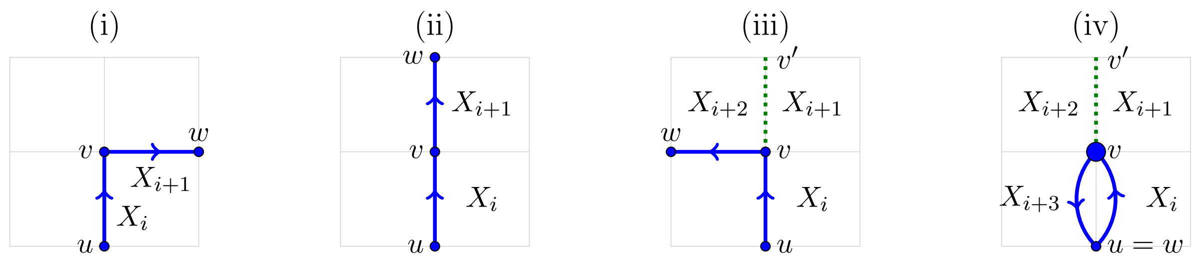

Fig. 1.6. (i) w is right of e_j. (ii) $w \neq u$ is collinear with e_j. (iii) w is left of e_j. (iv) $w = u$ is collinear with e_j.

Let H be a path in a polyomino G, that is not necessarily Hamiltonian. Consider the directed multigraph H^* where each edge $\{u, v\}$ of H is replaced by the directed edges (u, v) and (v, u). Define an H^**-trail* to be a directed trail $\overrightarrow{K} = e_1, \ldots, e_s$ in H^*, such that if the edge $(u, v) = e_j$ and $(v, u) = e_{j+1}$, then v is an end-vertex of H. We call a closed H^*-trail an H^**-circuit*. So, H^*-trails are trails that can only "turn around" at an end-vertex.

Now we will use an H^*-trail $\overrightarrow{K}$ to construct an H-walk of boxes in G_{-1} that we will call *the right H-walk induced by* $\overrightarrow{K}$ and denote it by $W_{\text{right}}(X_1, X_r)$, where X_1 and X_r are the end-boxes of $W_{\text{right}}(X_1, X_r)$. Roughly, $W_{\text{right}}(X_1, X_r)$ will be the walk of boxes that a "walker" would encounter as they followed along the side of $\overrightarrow{K}$ when starting on the right side of the first edge e_1 of $\overrightarrow{K}$. We will call this construction the *follow-the-wall* construction (FTW). This is very similar to the well-known hand-on-the-wall maze-solving algorithm.

Let X be the box of G_{-1} on the right of the edge e_1 of $\overrightarrow{K}$ and let e_j be the j^{th} edge of $\overrightarrow{K}$. Then $X = X_1$ is the first box of the H-walk. Let $e_j = (u, v)$, $e_{j+1} = (v, w)$ and let X_i be on the right of the edge e_j. There are four possibilities for the position of w with respect to (u, v): w is right of (u, v), $w \neq u$ is colinear with (u, v), w is left of (u, v), and $w = u$. For the last two cases define $e'_j = (v, v')$ to be the edge in $G_{-1} \setminus H$ that is colinear with e_j and set $e''_j = (v', v)$. See Figure 1.6.

(i). *w is right of* e_j. Then X_{i+1} is on the right of the edge e_{j+1}. Note that in this case, the walk has a repeated box since $X_i = X_{i+1}$.

(ii). *$w \neq u$ is collinear with* e_j. Then X_{i+1} is on the right of the edge e_{j+1}.

(iii). *w is left of* e_j. Then X_{i+1} is on the right of the edge e'_j and X_{i+2} is on the right of the edge e_{j+1}.

(iv). *$w = u$.* Then X_{i+1} is on the right of the edge e'_j and X_{i+2} is on the right of the edge e''_j and X_{i+3} is on the right of the edge e_{j+1}.

We say that the edge e_{j+1} *adds* to the H-walk $W_{\text{right}}(X_1, X_r)$ the box X_{i+1}, in cases (i) and (ii), boxes X_{i+1} and X_{i+2} in case (iii), and boxes X_{i+1}, X_{i+2} and X_{i+3} in case (iv). If e_s adds more than one box to $W_{\text{right}}(X_1, X_r)$ then we will adopt the convention that X_r is the last box added by the edge e_s and that all the boxes added by each edge of e_j, $j \in \{1, \ldots, s\}$ are on the right side of e_j. Note that this convention is necessary for the box X_{i+1} in Case (iv). We remark that the first edge e_1 can only add the single box X_1. The left H-walk $W_{\text{left}}(X_1, X_r)$ induced by $\overrightarrow{K}$ can be constructed analogously.

It is straightforward to extend FTW to the case where H is a cycle in G. We remark that in this case, no two directed edges of the H^*-trail may correspond to the same (undirected) edge of H and that, since cycles have no end-vertices, Case (iv) never occurs. In fact, we can orient any path (or cycle) H in G to obtain a directed trail $\overrightarrow{K}$ and then use FTW to obtain an H-walk from $\overrightarrow{K}$.

Let H be a path or cycle in G. Let $\mathcal{K}(H^*)$ be the set of all H^*-trails. We can view the FTW construction as a function Φ that assigns an H-walk to elements of $\mathcal{K}(H^*) \times \{\text{right, left}\}$. We will take a closer look at H^*-trails in the case where H is a Hamiltonian path or a Hamiltonian cycle of G.

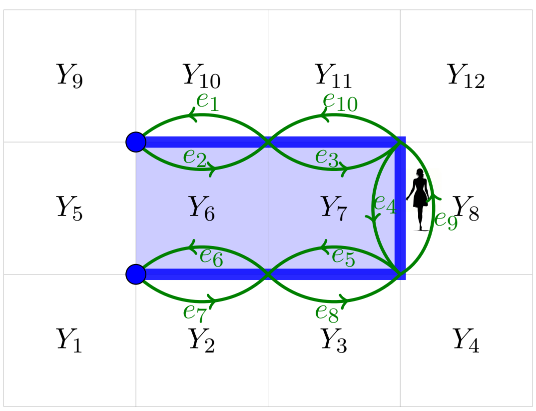

Fig. 1.7. G is on Y_6 and Y_7, shaded light blue. G_{-1} is the graph on $Y_1, \ldots, Y_{12}$. H in blue , H^* in green, $Y_1, \ldots, Y_{12}$ are boxes of G. $\overrightarrow{K}(e_{10})$ is flawed, $\overrightarrow{K}(e_1) = \overrightarrow{K}_H$ is not flawed. Let $\overrightarrow{K} = \overrightarrow{K}(e_3, e_5)$, $\overrightarrow{K}' = \overrightarrow{K}(e_8, e_{10})$. Then $\Phi\big(\overrightarrow{K}, \text{left}\big) = Y_{11}Y_{12}Y_8Y_4Y_3$ and $\Phi\big(\overrightarrow{K}', \text{right}\big) = Y_3Y_4Y_8Y_{12}Y_{11}$; silhouette in black following the wall along e_9.

Let $H = v_1, v_2, \ldots, v_r$ be a Hamiltonian path of G. Consider an Eulerian circuit (v_j, v_{j+1}), $(v_{j+1}, v_{j+2}), \ldots,$ (v_{r-1}, v_r), $(v_r, v_{r-1}), \ldots,$ (v_2, v_1), $(v_1, v_2), \ldots,$ (v_{j-1}, v_j) of H^*. Note that there are $2r - 2$ such distinct Eulerian circuits, one for each possible first edge. Observe that any subtrail of such a circuit is completely determined by its first and last edges. Therefore it will be fitting to use the notation $\overrightarrow{K}(e_s, e_t)$ to denote the unique subtrail starting at edge e_s and ending at edge e_t. We will use the notation $\overrightarrow{K}_{H^*}(e_j)$ to denote an Eulerian circuit of H^* starting at e_j and abbreviate to $\overrightarrow{K}_{H^*}$ whenever the first edge e_j is not relevant to our argument. Fix a side in $\{\text{right, left}\}$. We will use the notation W_{side,H^*} to denote $\Phi(\overrightarrow{K}_{H^*}, \text{side}) = W_{\text{side},H^*}$.

Consider an H-circuit $\overrightarrow{K}_{H^*}(e_{j+1}) = \overrightarrow{K}_{H^*}$, starting at e_{j+1} and ending at e_j. Fix a side in {right, left}. If e_j and e_{j+1} are as in cases (iii) and (iv) of the description of FTW, then $\Phi(\overrightarrow{K}_{H^*}, \text{side}) = W_{\text{side},H^*}$ might miss at least one box of G_{-1}. We will call such Eulerian circuits of H^* *flawed* and we avoid using them. We note that, for any path of length greater than two in a polyomino G with $G \notin \mathcal{P}_1$, it is possible to choose a starting edge for $\overrightarrow{K}_{H^*}$ so that $\overrightarrow{K}_{H^*}$ is not flawed. From here on, all Eulerian circuits of H^* we consider will be assumed to be not flawed. See Figure 1.7.

Observation 1.2.4. Let G be a polyomino and let H be a path or cycle in G.

(a) Let the H^*-trail $\overrightarrow{K}'$ be a subtrail of the H^*-trail $\overrightarrow{K}$ and fix a side in {right, left}. Then $\Phi(\overrightarrow{K}', \text{side})$ is an H-subwalk of $\Phi(\overrightarrow{K}, \text{side})$.

(b) Assume that H is a Hamiltonian path of G and fix a side in {right, left}. Then every box of G_{-1} is added to W_{side,H^*} by an edge of $\overrightarrow{K}_{H^*}$.

(c) Let $\overrightarrow{K} = (v_s, v_{s+1}), \ldots, (v_{t-1}, v_t)$ and $\overrightarrow{K}' = (v_t, v_{t-1}), \ldots, (v_{s+1}, v_s)$ be H trails. Then $\text{Boxes}(\Phi(\overrightarrow{K}, \text{right})) = \text{Boxes}(\Phi(\overrightarrow{K}', \text{left}))$ and $\text{Boxes}(\Phi(\overrightarrow{K}, \text{left})) = \text{Boxes}(\Phi(\overrightarrow{K}', \text{right}))$. See Figure 1.8.

(d) We may orient a subpath $P = v_s, v_{s+1}, \ldots, v_t$ of H to obtain a directed H^*-trail $\overrightarrow{K}_P = (v_s, v_{s+1}), \ldots, (v_{t-1}, v_t)$ of H^*. Similarly, an H^*-trail $\overrightarrow{K} = (v_s, v_{s+1}), \ldots, (v_{t-1}, v_t)$ not containing an end-vertex of H determines an oriented subpath $P = v_s, v_{s+1}, \ldots, v_t$ of H.

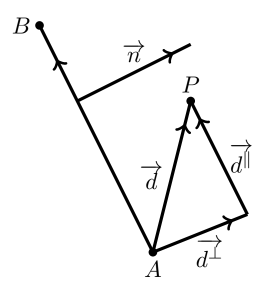

Fig. 1.8. P on the right of $\overrightarrow{AB}$.

Note 1.2.5. Let $A = (x_1, y_1)$, $B = (x_2, y_2)$ and $P = (x, y)$ be points in the plane that are not collinear. We define $(x_2 - x_1, y_2 - y_1)$ to be the direction of the vector $\overrightarrow{AB}$. Then the direction of the normal $\overrightarrow{n}$ to $\overrightarrow{AB}$, obtained by rotating $\overrightarrow{AB}$ by $-\frac{\pi}{2}$, is $(y_2 - y_1, x_1 - x_2)$. We want to know whether the point P is on the side of $\overrightarrow{AB}$ toward which $\overrightarrow{n}$ is pointing. Let $\overrightarrow{d} = \overrightarrow{AP}$. Let $\overrightarrow{d^{\perp}}$ be the component of $\overrightarrow{d}$ that is perpendicular to $\overrightarrow{AB}$ and let $\overrightarrow{d^{\parallel}}$ be the component of $\overrightarrow{d}$ that is parallel to $\overrightarrow{AB}$. See Figure 1.8. Note that:

$$\begin{aligned}
\overrightarrow{d} \cdot \overrightarrow{n} &= (\overrightarrow{d^{\parallel}} + \overrightarrow{d^{\perp}}) \cdot \overrightarrow{n} = \overrightarrow{d^{\perp}} \cdot \overrightarrow{n} \\
&= (x - x_1, y - y_1) \cdot (y_2 - y_1, x_1 - x_2) \\
&= (x - x_1)(y_2 - y_1) + (y - y_1)(x_1 - x_2).
\end{aligned}$$

We say that P is on the *right* of $\overrightarrow{AB}$ if $\overrightarrow{d} \cdot \overrightarrow{n} > 0$ and we say that P is on the *left* $\overrightarrow{AB}$ if $\overrightarrow{d} \cdot \overrightarrow{n} < 0$.

Let $e = (u, v)$ be an edge of a lattice animal G, where $u = v(k_1, l_1)$, $v = v(k_2, l_2)$. Let X be a box of the square lattice that is incident on (u, v). We say that X is on the *right* of the edge (u, v) if there is a vertex $w = v(k, l)$ in $V(X) \setminus V(e)$ such that $(k-k_1, l_2-l_1) \cdot (l-l_1, k_1-k_2) = 1$ and we say that X is on the *left* of the edge (u, v) if there is a vertex $w = v(k, l)$ in $V(X) \setminus V(e)$ such that $(k - k_1, l_2 - l_1) \cdot (l - l_1, k_1 - k_2) = -1$.

1.3 The structure of H-components

In this section we prove several structural properties of H-components. Key results include: H-components in simply connected polyominoes have unique necks (Corollary 1.3.2), and each H-component has a unique main trail - a subtrail of H^* that determines the component's structure (Proposition 1.3.9, Corollary 1.3.11). When H is a cycle, Corollary 1.3.15 proves properties of H-components that lie in the exterior of H (which we call cookies, and are central to the reconfiguration of Hamiltonian cycles in Chapter 2).

Lemma 1.3.1. Let G be a polyomino, let H be a Hamiltonian path of G, and let J be an H-component of G. Then J has a neck. Furthermore, if the boundary of G has q boundary components then J has at most q necks.

Proof. If J only has one box, we're done, so assume that J has more than one box. Let Z be a box of J. Without loss of generality, we may assume that Z is on the left of $e_z \in \overrightarrow{K}_H$. If no edge is incident on Z, we choose one of the four neighbours beside it. We claim that there exists a subtrail $\overrightarrow{K}(e_z, e_{t+1})$ of $\overrightarrow{K}_H$ such that $\Phi\big(\overrightarrow{K}(e_z, e_t), \text{left}\big)$ is contained in J but $\Phi\big(\overrightarrow{K}(e_z, e_{t+1}), \text{left}\big)$ is not. For a contradiction, assume that for every subtrail of $\overrightarrow{K}_{H^*}$ starting at e_z, $\Phi\big(\overrightarrow{K}(e_z, e_j), \text{left}\big)$ is contained in J, where $j \in \{z+1, \ldots, z\}$.

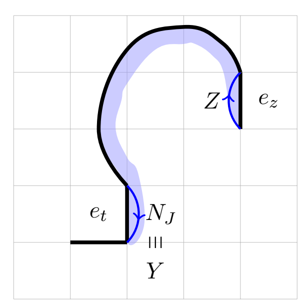

Fig. 1.9. $\Phi\big(\overrightarrow{K}(e_z, e_t), \text{left}\big)$ shaded in blue.

But then $\Phi\big(\overrightarrow{K}(e_z), \text{left}\big) = \Phi\big(\overrightarrow{K}_{H^*}, \text{left}\big)$ is contained in J, contradicting that $\Phi\big(\overrightarrow{K}_{H^*}, \text{left}\big)$ contains the boxes of $G_{-1} \setminus G$. It follows that e_{t+1} adds the first box Y of $\Phi\big(\overrightarrow{K}(e_z, e_{t+1}), \text{left}\big)$ that is not contained in J. Note that, by definition of H-component, Y must belong to $G_{-1} \setminus G$. (Since Y is H-adjacent to the box X preceding it, but Y does not belong to J, it must be the case that Y is not in G). See Figure 1.9. Let X be the box of J preceding Y in

$\Phi\big(\overrightarrow{K}(e_z, e_{t+1}), \text{left}\big)$. We have that X and Y are H-adjacent and share a boundary edge e_J of G that is not in H. By definition of neck of an H-component, $X = N_J$.

Now we show that J has at most q necks. For a contradiction, assume that J has $q' > q$ necks. By the pigeonhole principle there is a boundary component that has at least two necks $N_{J,1}$ and $N_{J,2}$ of J incident on it. Let $N'_{J,1}$ and $N'_{J,2}$ be the boxes in BBoxes(G) incident on $N_{J,1}$ and $N_{J,2}$, respectively. Note that Outer(G) is H-path-connected and so is every hole of G. Then there is an H-cycle $N'_{J,1}, \ldots, N'_{J,2}$ and $N_{J,2}, \ldots, N_{J,1}, N'_{J,1}$, containing the box $N_{J,1}$ of G, which contradicts Corollary 1.2.3. □

Corollary 1.3.2. Let G be a simply connected polyomino, let H be a Hamiltonian path of G and let J be an H-component of G. Then J has a unique neck. □

Definitions. Let G be a polyomino and let B_0 be the cycle bounding its outer face. Recall that the enclosure of G is the set of all boxes contained in the region of the plane bounded by B_0.

Let Q be a cycle of vertices in G consisting of a path $P(v_1, v_s) \neq \{v_1, v_s\}$ and the edge $\{v_1, v_s\}$. Q bounds a region U of Encl(G). Let N_U be the box incident on $\{v_1, v_s\}$ that is contained in U. We call N_U *the neck* of U and we call $\{v_1, v_s\}$ the *neck edge* of U.

Let $H = v_1, \ldots, v_s, \ldots, v_t, \ldots, v_r$ be any path of G such that $\{v_s, v_t\} \in E(\text{Encl}(G)) \setminus E(H)$. Let $P_1 = P(v_1, v_s)$, $P_2 = P(v_s, v_t)$ and $P_3 = P(v_t, v_r)$ be a partitioning of H into subpaths. We denote the H^*-trails $(v_s, v_{s-1}), \ldots, (v_{s-1}, v_s)$ and $(v_t, v_{t+1}), \ldots, (v_{t+1}, v_t)$ by $\overrightarrow{K}_{P_1}$ and $\overrightarrow{K}_{P_3}$, respectively.

Lemma 1.3.3. Let G be a polyomino. Let $H = v_1, \ldots, v_s, \ldots, v_t, \ldots, v_r$ be a Hamiltonian path of G [3] such that $\{v_s, v_t\} \in E(\text{Encl}(G)) \setminus E(H)$. Let $P_1 = P(v_1, v_s)$, $P_2 = P(v_s, v_t)$ and $P_3 = P(v_t, v_r)$ be a partitioning of H into subpaths. Let U be the region of the plane bounded by the polygon Q consisting of the subpath P_2 of H and the edge $\{v_s, v_t\}$. Let $\overrightarrow{K}_Q$ be the directed circuit obtained from orienting Q. Then:

(a) $U \subseteq \text{Encl}(G)$

(b) Boxes($\Phi(\overrightarrow{K}_Q, \text{right})$) $\subset U$ and Boxes($\Phi(\overrightarrow{K}_Q, \text{left})$) $\subset G_{-1} \setminus U$ iff $\Phi((v_t, v_s), \text{right})$ is a box of U.

(c) For $i \in \{1, 3\}$, P_i is contained in U or P_i is contained in $(G_{-1} \setminus U)$.

(d) For a side in $\{\text{right}, \text{left}\}$ and $i \in \{1, 3\}$, if $P_i \subset U$ then $\Phi(\overrightarrow{K}_{P_i}, \text{side}) \subset U$ and, if $P_i \subset G_{-1} \setminus U$ then $\Phi(\overrightarrow{K}_{P_i}, \text{side}) \subset G_{-1} \setminus U$.

[3] This result holds more generally when H is not Hamiltonian, but we state it this way as the general case is not needed for our purposes.

Proof. Part (a) follows from the fact that the boundary of U is contained in Encl(G), and so we must have $U \subseteq \text{Encl}(G)$ as well. Part (b) follows by definition of FTW and induction on the edges of Q. Part (d) follows by definition of $\overrightarrow{K}_{P_i}$ and FTW and induction on the edges of $\overrightarrow{K}_{P_i}$.

For part (c), we assume, for definiteness, that $v_1 \in U$. Since H is a path, $V(P_1) \cap V(Q) = v_s$. Thus $(P_1 \setminus v_s) \subset U$. The proofs for the other inclusions are similar. □

The next three Lemmas will only be used by Lemma 4.12 in Chapter 4.

Lemma 1.3.4. If a graph has a closed walk with a non-repeated edge, then the graph has a cycle[30]. □

The following lemma is a proof of the fact that around each hole of a polyomino, there is a cycle of boxes.

Lemma 1.3.5. Let G be a polyomino and let $\mathcal{O}$ be a hole of G with boundary $B(\mathcal{O})$. Let $\overrightarrow{K}$ be the circuit obtained by orienting the edges of $B(\mathcal{O})$ (Recall Lemma 1.1.7 (f)). Then it is possible to choose a starting edge e_1 for the $B(\mathcal{O})^*$-circuit $\overrightarrow{K}$ and a side in $\{\text{right}, \text{left}\}$ so that the $B(\mathcal{O})$-walk $\Phi(\overrightarrow{K}, \text{side})$ is contained in $G_{-1} \setminus \mathcal{O}$ and it contains a cycle of boxes, or the $B(\mathcal{O})$-walk $\Phi(\overrightarrow{K}, \text{side}), \Phi(e_1, \text{side})$ is contained in $G_{-1} \setminus \mathcal{O}$ and it contains a cycle of boxes.

Proof. By Lemma 1.3.3 (b), for any edge e of $\overrightarrow{K}$, there is a side in $\{\text{right}, \text{left}\}$ such that $\Phi(e, \text{side}) \subset G_{-1} \setminus \mathcal{O}$ implies $\Phi(\overrightarrow{K}, \text{side}) \subset G_{-1} \setminus \mathcal{O}$. For definiteness, let $\Phi(\overrightarrow{K}, \text{left}) \subset G_{-1} \setminus \mathcal{O}$. Let $\overrightarrow{K} = e_1, \ldots, e_t$ be such that e_1, e_t are as e_{j+1}, e_j, respectively, in Case (i) or Case (ii) of the definition of FTW. Let $W = W_{\text{left}}(X_1, X_s) = \Phi(\overrightarrow{K}, \text{left})$. If e_1, e_t are as e_{j+1}, e_j, respectively, in Case (i) of the definition of FTW then $X_1 = X_s$ (Figure 1.10 (a)); and if e_1, e_t are as e_{j+1}, e_j, respectively, in Case (ii) of the definition of FTW then X_1 is adjacent to X_s (Figure 1.10 (b)). If the latter, we may add X_1 to W after X_s to obtain a closed H-walk. From here on we will assume that e_1, e_t are as e_{j+1} and e_j, respectively, in Case (i) of FTW, so that W is the closed walk starting and ending at X_1. The case where e_1, e_t are as e_{j+1} and e_j, respectively, in Case (ii) of FTW is very similar, so we omit it. It remains to check that W contains a cycle.

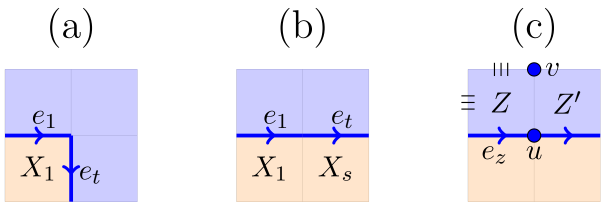

Fig. 1.10. $\overrightarrow{K}$ in blue, boxes of $G_{-1} \setminus \mathcal{O}$ shaded blue, boxes of $\mathcal{O}$ shaded orange.

Mod out W by consecutively repeated boxes to obtain the H-walk W'. That is, if the string of boxes XY^kZ appears in W, substitute it by the string XYZ. Note that W' is an H-walk that retains the order in which boxes of G_{-1} appear in W. Now we can view W' as an alternating sequence of boxes and gluing edges starting and ending at the box X_1. If we can find a non-repeated gluing edge of W', then Lemma 1.3.4 implies that W contains a cycle of boxes of G_{-1}.

Let e_z be a horizontal northernmost edge in $\overrightarrow{K}$. Let $Z = \Phi(e_z, \text{left})$. Note that $Z \in W$, $Z' = Z + (1, 0)$ is an H-neighbour of Z, and that no other edge of $\overrightarrow{K}$ is incident on Z. We claim that the vertical gluing edge $f = \{u, v\}$ between Z and Z', where $u \in \overrightarrow{K}$ and $v \notin \overrightarrow{K}$, is a non-repeated gluing edge of W'. See Figure 1.10 (c). For a contradiction, assume that there is a second occurrence of f in W'. This means that the sequence of boxes Z, Z' or the sequence of boxes $Z'Z$ occurs again in W'. By definition of FTW, and the fact that the edges of $\overrightarrow{K}$ are unique, such a sequence must be added to W by edges of $\overrightarrow{K}$ incident on v, of which there are none.

Lemma 1.3.6. Let H be a Hamiltonian path of a polyomino G. Assume that J is a non-self-adjacent H-subtree of an H-component of G, and that J has no polyking junctions. Then J is a simply connected polyomino.

Proof. Since J is an H-subtree, J is box-path-connected, so it is a polyomino. For a contradiction, assume that J is not simply connected. Then J has a hole $\mathcal{O}$ with boundary $B(\mathcal{O})$. Orient the edges of $B(\mathcal{O})$ to obtain a directed circuit $\overrightarrow{K}$. By Lemma 1.3.5, we may choose a side and a starting edge for $\overrightarrow{K}$ such that $\Phi(\overrightarrow{K}, \text{side}) \subseteq J_{-1} \setminus \mathcal{O}$ and such that $\Phi(\overrightarrow{K}, \text{side})$ contains a cycle C of boxes in $J_{-1} \setminus \mathcal{O}$. For definiteness, we assume that $\Phi(\overrightarrow{K}, \text{left}) \subseteq J_{-1} \setminus \mathcal{O}$.

Suppose that $\Phi(\overrightarrow{K}, \text{left})$ is contained in $J \setminus \mathcal{O}$. Then, either at least one of the gluing edges of C belongs to H, or no gluing edge of C is in H. The former contradicts the assumption that J is non-self-adjacent. The latter contradicts Proposition 1.2.1. It remains to show that $\text{Boxes}(\Phi(\overrightarrow{K}, \text{left})) \subseteq J$.

We use induction on the edges of $\overrightarrow{K}$ to show that $\text{Boxes}(\Phi(\overrightarrow{K}, \text{left})) \subseteq J$. Let Z and Z' be the boxes in G_{-1} incident on an edge $e \in \overrightarrow{K}$. By Lemma 1.1.7 (b), one of Z and Z' belongs to $\mathcal{O}$ and the other belongs to J. Let this be observation $(*)$. Note that by $(*)$ and the assumption that $\Phi(\overrightarrow{K}, \text{left}) \subseteq J_{-1} \setminus \mathcal{O}$, the base case holds: $\Phi(e_1, \text{left})$ is a box of J.

For the inductive case, assume that for $i \in \{2, \ldots, i_0\}$, $\text{Boxes}(\Phi(\overrightarrow{K}(e_1, e_{i_0}), \text{left})) \subset J$. For definiteness, assume that $e_{i_0} = (a, a+1; b) = (v(a, b), v(a+1; b))$ and let $X = R(a, b)$. Then $X \in J$ and $X + (0, -1) \in \mathcal{O}$. We want to check that $\text{Boxes}(\Phi(\overrightarrow{K}(e_1, e_{i_0+1}), \text{left})) \subset J$.

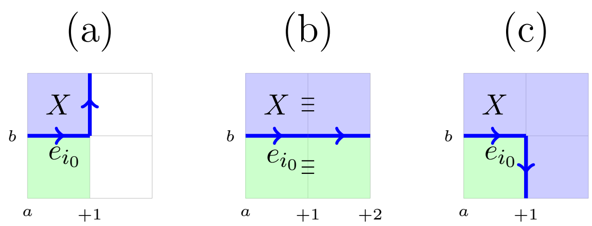

Fig. 1.11. (a) $e_{i_0+1} = e(a+1; b, b+1)$.
(b) $e_{i_0+1} = e(a+1, a+2; b)$.
(c) $e_{i_0+1} = e(a+1; b-1, b)$.

There are three possibilities: $e_{i_0+1} = e(a+1; b, b+1)$, $e_{i_0+1} = e(a+1, a+2; b)$, and $e_{i_0+1} = e(a+1; b-1, b)$. See Figure 1.11. If $e_{i_0+1} = e(a+1; b, b+1)$, then $\text{Boxes}(\Phi(\overrightarrow{K}(e_1, e_{i_0+1}), \text{left}))= \text{Boxes}(\Phi(\overrightarrow{K}(e_1, e_{i_0}), \text{left}))$ and we are done by the inductive hypothesis. If $e_{i_0+1} = e(a+1, a+2; b)$, then Observation 1.1.9 implies that $e(a+1; b, b+1) \notin B(J)$ and $e(a+1; b-1, b) \notin B(J)$. Then observation $(*)$ implies that one of $X + (1, 0)$ and $X + (1, -1)$ belongs to J, and the other belongs to $\mathcal{O}$. The assumption that J has no polyking junction implies that we must have $X + (1, 0) \in J$ and $X + (1, -1) \in \mathcal{O}$. If $e_{i_0+1} = e(a+1; b-1, b)$, then by $(*)$, $Z + (1, -1)$ belongs to J. Our assumption that J has no polyking junctions implies that $Z + (1, 0)$ must belong to J as well. Thus, $(\Phi(\overrightarrow{K}, \text{left}))$ is contained in J. □

Proposition 1.3.7. Let G, H, P_i for $i \in \{1, 2, 3\}$, $\{v_s, v_t\}$, Q and U be as in Lemma 1.3.3. In addition, assume that H is a Hamiltonian path of G, that $\{v_s, v_t\}$ is in $B(G)$, that J is an H-component of G that has $\{v_s, v_t\}$ as one of its neck-edges, and that the neck N_U of U coincides with a neck N_J of J[4]. Then $J \subseteq U$. Moreover, if G is a simply connected polyomino, then $J = U$.

Proof. For a contradiction, assume that there is a box Z in $J \setminus U$. Then there is an H-path $P(X_1, X_r)$ in J, where $X_1 = N_J = N_U$ and $X_r = Z$. Let Y be the other box incident on $\{v_s, v_t\}$ and let $c_1, c_2, \ldots, c_r = c_1$ be the centers of the boxes of $P(X_1, X_r)$. Then for each $j \in \{1, 2, \ldots, r-1\}$, $[c_j, c_{j+1}]$ intersects the gluing edge of X_j and X_{j+1} and $[c_j, c_{j+1}]$ intersects no other edge of G. By JCT, since $c_1 \in U$ and $c_r \notin U$, $P(c_1, c_r)$ intersects Q at some edge e. Since $Y \notin G$, $e \neq \{v_s, v_t\}$. Then e must be some other edge of Q. But all other edges of Q belong to H, contradicting that e is a gluing edge.

[4] This can be made more general by only assuming that $U \cap J \neq \emptyset$, but it makes the proof longer and we don't need it.

Assume that G is a simply connected polyomino. We will show that $U \subseteq J$. Since G is a simply connected polyomino, we have that N_U is the unique neck of J. For a contradiction, assume that there is a box Z in $U \setminus J$. For definiteness, by Observation 1.2.4 (b), we may assume that Z is on the left of an edge e_z of $\overrightarrow{K}_{H^*}$. (If no edge of H is incident on Z, as in Case (iii) of FTW, we replace Z with the second box added by e_z).

Let $\overrightarrow{K}(e_z, e_{y-1})$ be a subtrail of $\overrightarrow{K}_{H^*}$ such that $\Phi(\overrightarrow{K}(e_z, e_{y-1}), \text{left})$ is contained in $U \setminus J$ but $\Phi(\overrightarrow{K}(e_z, e_y), \text{left})$ is not. Let Y be the first box not in $U \setminus J$ added by e_y and let Y' be the box preceding Y in $\Phi(\overrightarrow{K}(e_z, e_y), \text{left})$. See Figure 1.12. This means that $Y' \in U \setminus J$. Let f be the edge that Y and Y' share. Since $J \subset U$, either $Y \in J$ or $Y \in G_{-1} \setminus U$.

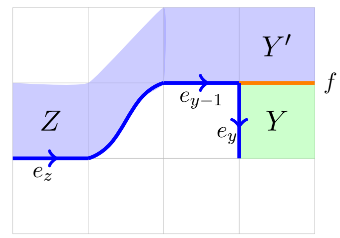

Fig. 1.12. $\Phi\big(\overrightarrow{K}(e_z, e_t), \text{left}\big)$ shaded in blue; f in orange.

Suppose that $Y \in J$. Since G is simply connected, by Corollary 1.3.2, J has a unique neck. Since Y' is not in J, f must be the neck-edge of J. By assumption, $Y = N_J = N_U$. But then, by definition of the neck of U, $Y' \notin U$, contradicting our assumption that $Y' \in U \setminus J$.

Suppose that $Y \in G_{-1} \setminus U$. Since $Y' \in U$, f must be a boundary edge of U, so $f \in Q$. Since $\Phi(\overrightarrow{K}(e_z, e_y), \text{left})$ is an H-walk, f is a gluing edge of $\Phi(\overrightarrow{K}(e_z, e_y), \text{left})$ and so $f \notin H$. It follows that $f = \{v_s, v_t\}$. Since $Y' \in U$ and $Y \in G_{-1} \setminus U$, we have $Y' = N_U = N_J \in J$. But this contradicts our assumption that $Y' \in U \setminus J$. □

Corollary 1.3.8. Let G be a simply connected polyomino and let H, P_i for $i \in \{1, 2, 3\}$, $\{v_s, v_t\}$, Q, U and J be as in Proposition 1.3.7, and let $\overrightarrow{K}_{P_i}$ be as in Lemma 1.3.3. Then:

(a) Exactly one of the following is true:

(I) $\text{Boxes}(\Phi(P_2, \text{right})) \subseteq J$ and $\text{Boxes}(\Phi(P_2, \text{left})) \subseteq G_{-1} \setminus J$, and

(II) $\text{Boxes}(\Phi(P_2, \text{left})) \subseteq J$ and $\text{Boxes}(\Phi(P_2, \text{right})) \subseteq G_{-1} \setminus J$.

(b) For a side in $\{\text{right}, \text{left}\}$ and $i \in \{1, 3\}$, $\text{Boxes}(\Phi(\overrightarrow{K}_{P_i}, \text{side})) \subseteq J$ or $\text{Boxes}(\Phi(\overrightarrow{K}_{P_i}, \text{side})) \subseteq G_{-1} \setminus J$. Furthermore, $\text{Boxes}(\Phi(\overrightarrow{K}_{P_i}, \text{side})) \subseteq J$ if and only if P_i is contained in J and $\text{Boxes}(\Phi(\overrightarrow{K}_{P_i}, \text{side})) \subseteq G_{-1} \setminus J$ if and only if P_i is contained in $G_{-1} \setminus J$. □

Proposition 1.3.9. Let G be a simply connected polyomino and let H, P_i for $i \in \{1, 2, 3\}$, $\{v_s, v_t\}$, Q, and J be as in Corollary 1.3.8. Then there are exactly two maximal and unique elements $(\overrightarrow{K}_J, \text{right})$ and $(\overrightarrow{K}'_J, \text{left})$ of $\mathcal{K} \times \{\text{right}, \text{left}\}$ such that $\text{Boxes}(\Phi(\overrightarrow{K}_J, \text{right})) = J$ and $\text{Boxes}(\Phi(\overrightarrow{K}'_J, \text{left})) = J$.

Proof. We may assume that J has more than one box. Let $\overrightarrow{K}_{P_1}$ and $\overrightarrow{K}_{P_3}$ be as in Lemma 1.3.3. By Observation 1.2.4 (d), P_2 determines the trail $\overrightarrow{K}_{P_2} = (v_s, v_{s+1}), \ldots, (v_{t-1}, v_t)$. Let $\overrightarrow{K}'_{P_2}$ be the trail $= (v_t, v_{t-1}), \ldots, (v_{s+1}, v_s)$. Note that $\overrightarrow{K}_{H^*} = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}, \overrightarrow{K}_{P_3}, \overrightarrow{K}'_{P_2}$ is an Eulerian circuit of H^*. We will construct $\overrightarrow{K}_J$ and $\overrightarrow{K}'_J$ explicitly and then check maximality and uniqueness.

Note that J may contain both end-vertices of H, exactly one end-vertex or neither end-vertex, so there are three cases to consider.

CASE 1. J contains no end-vertices. By Corollary 1.3.8 (a) $\mathrm{Boxes}(\Phi(P_2, \mathrm{right})) \subseteq J$ or $\mathrm{Boxes}(\Phi(P_2, \mathrm{left})) \subseteq J$. For definiteness, assume that $\mathrm{Boxes}(\Phi(P_2, \mathrm{right})) \subseteq J$. Then $\mathrm{Boxes}(\Phi(P_2, \mathrm{left})) \subseteq G_{-1} \setminus J$.

First we will show that $\overrightarrow{K}_J = \overrightarrow{K}_{P_2}$ satisfies $\mathrm{Boxes}(\Phi(\overrightarrow{K}_J, \mathrm{right})) = J$. We already have that $\mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}, \mathrm{right})) \subseteq J$. We need to check that $J \subset \mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}, \mathrm{right}))$. Let Z be a box of J. Note that at least one vertex of Z, say v_z, is neither an end-vertex of H nor incident on the neck of J. Either at least one of the two edges of Z incident on v_z is in H, or neither is.

CASE 1.1: At least one of the two edges of Z incident on v_z, say e_z, is in H. If we can show that $e_z \in \overrightarrow{K}_{P_2}$ then, by Observation 1.2.4 (a), we have that $Z = \Phi(e_z, \mathrm{right})$ is contained in $\Phi(\overrightarrow{K}_{P_2}, \mathrm{right})$ and we're done. Now, Corollary 1.3.8 (b) implies that $e_z \notin \overrightarrow{K}_{P_1}$ and $e_z \notin \overrightarrow{K}_{P_3}$. BWOC assume that $e_z \in \overrightarrow{K}'_{P_2}$. Then, by Observation 1.2.4 (a), $Z = \Phi(e_z, \mathrm{right})$ is contained in $\mathrm{Boxes}(\Phi(\overrightarrow{K}'_{P_2}, \mathrm{right}))$. By Observation 1.2.4 (c), $\mathrm{Boxes}(\Phi(\overrightarrow{K}'_{P_2}, \mathrm{right}))= \mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}, \mathrm{left}))$ and $\mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}, \mathrm{left}))$ is contained in $G_{-1} \setminus J$. But this contradicts that $Z \in J$. Thus we must have that $e_z \in \overrightarrow{K}_{P_2}$. End of Case 1.1.

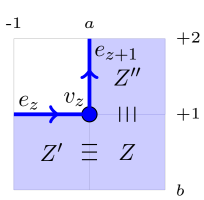

Fig. 1.13. Case 1.2.

CASE 1.2: Neither of the two edges of Z incident on v_z is in H. For definiteness, let $Z = R(a, b)$ and $v_z = v(a, b+1)$. Let $Z' = Z + (-1, 0)$ and $Z'' = Z + (0, 1)$. Then we have that $e(a-1, a; b+1) \in H$, $e(a; b+1, b+2) \in H$, $Z' \in J$ and $Z'' \in J$. Let $(v(a-1, b+1), v(a, b+1)) = e_z$ and $(v(a, b+1), v(a, b+2)) = e_{z+1}$. By Case 1.1, $e_z \in \overrightarrow{K}_{P_2}$ and $e_{z+1} \in \overrightarrow{K}_{P_2}$. Then $Z \in \mathrm{Boxes}(\Phi(\overrightarrow{K}(e_z, e_{z+1}), \mathrm{right})) \subset \mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}), \mathrm{right})$. End of Case 1.2. See Figure 1.13.

Cases 1.1 and 1.2 showed that $\mathrm{Boxes}(\Phi(\overrightarrow{K}_{P_2}, \mathrm{right})) = J$ and that the edges of $\overrightarrow{K}_J$ cannot belong to $\overrightarrow{K}_{P_1}$, $\overrightarrow{K}'_{P_2}$ or $\overrightarrow{K}_{P_3}$. It follows that $\overrightarrow{K}_J$ cannot be extended and thus is maximal.

To see that (K_J, right) is unique, assume toward a contradiction, that there exists an element $(\overrightarrow{K}^{\dagger}_J, \mathrm{right})$ in $\mathcal{K} \times \{\mathrm{right}\}$, distinct from $\overrightarrow{K}_J$, such that $\mathrm{Boxes}(\Phi(\overrightarrow{K}^{\dagger}_J, \mathrm{right})) = J$.

Then there is an edge $e \in \overrightarrow{K}^{\dagger}_J \setminus \overrightarrow{K}_J$. Since $\overrightarrow{K}_J = \overrightarrow{K}_{P_2}$, $e \in \overrightarrow{K}_{P_1}$ or $e \in \overrightarrow{K}_{P_3}$ or $e \in \overrightarrow{K}'_{P_2}$. By Corollary 1.3.8(b), $e \notin \overrightarrow{K}_{P_1}$ and $e \notin \overrightarrow{K}_{P_3}$. By Cases 1.1 and 1.2, $e \notin \overrightarrow{K}'_{P_2}$. But then e must belong to $\overrightarrow{K}_{P_2} = \overrightarrow{K}_J$, contradicting that $e \in \overrightarrow{K}^{\dagger}_J \setminus \overrightarrow{K}_J$. Thus, $(\overrightarrow{K}_J, \text{right})$ is unique.

By Observation 1.2.4 (c) we can see that $\text{Boxes}(\Phi(\overrightarrow{K}'_{P_2}, \text{left})) = J$. The proof that $\overrightarrow{K}'_{P_2} = \overrightarrow{K}'_J$ is maximal and unique is the same as the proof for $\overrightarrow{K}_J$, so we omit it. End of Case 1.

CASE 2. J contains exactly one end-vertex. For definiteness, assume that J contains v_1. By Corollary 1.3.8 (a) $\text{Boxes}(\Phi(P_2, \text{right})) \subseteq J$ or $\text{Boxes}(\Phi(P_2, \text{left})) \subseteq J$. For definiteness, assume that $\text{Boxes}(\Phi(P_2, \text{right})) \subseteq J$. We will show that $\overrightarrow{K}_J = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}$. By Corollary 1.3.8 (b), $\text{Boxes}(\Phi(\overrightarrow{K}_{P_1}, \text{right})) \subseteq J$. Then we have that $\text{Boxes}(\Phi(\overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}, \text{right})) \subseteq J$. We need to check that $J \subseteq \text{Boxes}(\Phi(\overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}, \text{right}))$.

Let Z be a box of J and let v_z be as in Case 1. By Case 1, it is sufficient to check the case where at least one of the two edges of Z incident on v_z, say e_z, is in H. By Corollary 1.3.8(b), $e_z \notin \overrightarrow{K}_{P_3}$ and by Case 1, $e_z \notin \overrightarrow{K}'_{P_2}$. Then $e_z \in \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}$.

Let $\overrightarrow{K}'_J = \overrightarrow{K}'_{P_2}, \overrightarrow{K}_{P_1}$. By Observation 1.2.4 (c) we can see that $\text{Boxes}(\Phi(\overrightarrow{K}'_J, \text{left})) = J$ as well. Proofs for maximality and uniqueness are similar to those in Case 1, so we omit them. End of Case 2.

CASE 3. J contains both end-vertices. As in previous cases, we may assume, for definiteness, that $\text{Boxes}(\Phi(P_2, \text{right})) \subseteq J$. Let $\overrightarrow{K}_J = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}, \overrightarrow{K}_{P_3}$ and $\overrightarrow{K}'_J = \overrightarrow{K}_{P_3}, \overrightarrow{K}'_{P_2}, \overrightarrow{K}_{P_1}$. Using the same arguments as above we can see that the elements of $\mathcal{K} \times \{\text{right}, \text{left}\}$ that satisfy the conclusion are $(\overrightarrow{K}_J, \text{right})$ and $(\overrightarrow{K}'_J, \text{left})$. End of Case 3. □

Corollary 1.3.10. Let G be a simply connected polyomino, let H be a Hamiltonian path of G, and let J be an H-component of G. If some edge $e_z \in \overrightarrow{K}_{H^*}$ adds a box Z in J to $W_{H^*,\text{right}}$, then $e_z \in \overrightarrow{K}_J$; and if $e_z \in \overrightarrow{K}_{H^*}$ adds a box Z in J to $W_{H^*,\text{left}}$, then $e_z \in \overrightarrow{K}'_J$. □

Definitions. Let G be a simply connected polyomino, and let J be an H-component of G. We call the elements $(\overrightarrow{K}_J, \text{right})$ and $(\overrightarrow{K}'_J, \text{left})$ of $\mathcal{K} \times \{\text{right}, \text{left}\}$ *the right main trail of J* and *the left main trail of J*, respectively. We call $\Phi((\overrightarrow{K}_J, \text{right}))$ and $\Phi((\overrightarrow{K}'_J, \text{left}))$ and *the right main walk of J* and the *left main walk of J*, respectively.

Corollary 1.3.11 Let G be a simply connected polyomino and let H, P_i for $i \in \{1, 2, 3\}$, $\{v_s, v_t\}$, Q, and J be as in Corollary 1.3.8. For $i \in \{1, 2, 3\}$, let $\overrightarrow{K}_{P_i}$ and $\overrightarrow{K}'_{P_i}$ be as in Proposition 1.3.9. Then:

(a) If J contains no end-vertices, then $\overrightarrow{K}_J = \overrightarrow{K}_{P_2}$.

(b) If J contains exactly one end-vertex, then $\overrightarrow{K}_J = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}$, or $\overrightarrow{K}_J = \overrightarrow{K}_{P_2}, \overrightarrow{K}_{P_3}$.

(c) If J contains both end-vertices, then $\overrightarrow{K}_J = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2}, \overrightarrow{K}_{P_3}$.

(d) J is a simply connected polyomino with boundary $B(J) = P_2, \{v_s, v_t\}$. $\square$

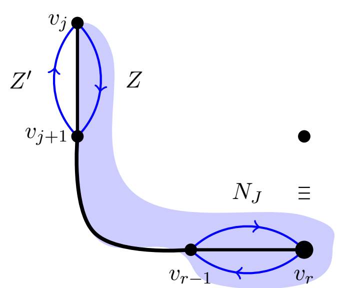

Fig. 1.14. $\Phi\big(\overrightarrow{K}((v_j, v_{j+1}), (v_r, v_{r-1}), \text{left}\big)$, shaded in blue.

Lemma 1.3.12. Let H be a Hamiltonian path of a simply connected polyomino G and let J be an H-component of G. Then J is self-adjacent if and only if an edge of a main trail of J has an end-vertex of H incident on it in $G \setminus B(G)$

Proof. Note that if the condition is satisfied, then the four boxes incident on the end-vertex belong to J and are H-path connected, and so J is self-adjacent. It remains to prove the converse.

Suppose that J is a self-adjacent H-component of G. We will show that some edge of the main left trail $\overrightarrow{K}'_J$ of J has an end-vertex of H incident on it in $G \setminus B(G)$. Since J is self-adjacent, there are H-path-connected boxes Z and Z' in J sharing an edge $e = \{v_j, v_{j+1}\}$ of H. Without loss of generality, we may assume that Z is on the left of (v_j, v_{j+1}). Then (v_j, v_{j+1}) adds Z to $\Phi((\overrightarrow{K}'_J, \text{left}))$. See Figure 1.14. By Corollary 1.3.10, $(v_j, v_{j+1}) \in \overrightarrow{K}'_J$. Similarly, $(v_{j+1}, v_j) \in \overrightarrow{K}'_J$. The maximality of $\overrightarrow{K}'_J$ implies that exactly one of $\overrightarrow{K}\big((v_j, v_{j+1}), (v_{j+1}, v_j)\big)$ and $\overrightarrow{K}\big((v_{j+1}, v_j), (v_j, v_{j+1})\big)$ is a subtrail of $\overrightarrow{K}'_J$. For definiteness, assume that $\overrightarrow{K}\big((v_j, v_{j+1}), (v_{j+1}, v_j)\big)$ is a subtrail of $\overrightarrow{K}'_J$. Note that the end-vertex v_r of H is incident on the edge (v_{r-1}, v_r) of $\overrightarrow{K}\big((v_j, v_{j+1}), (v_{j+1}, v_j)\big)$ and so v_r is incident on the edge (v_{r-1}, v_r) of $\overrightarrow{K}'_J$.

To see that $v_r \in G \setminus B(G)$, BWOC assume that $v_r \in B(G)$. By definition of FTW, (see Figure. 1.6 (iv)) all four boxes incident on v_r must belong to J. But then at least one box incident on v_r belongs to $G_{-1} \setminus G$, and so $(G_{-1} \setminus G) \cap J \neq \emptyset$, contradicting the definition of J. $\square$

Hamiltonian e-cycles, cycles, and cookies. Let G be a polyomino, let H be a Hamiltonian cycle of G and let e be an edge of H in the boundary of G. We call the path $H \setminus e$ a Hamiltonian *e-cycle* of G with H-components $J_0, \ldots, J_s$, where J_0 is the H-component that contains both end-vertices of H. By Lemma 1.3.12, J_0 and all other H-components of G are non-self-adjacent. We define $\text{int}(H)$ to be the H-component J_0 and $\text{ext}H$ to be the complement of J_0 in G. We will call the H-components $J_1, \ldots, J_s$ *cookies*[5]. If a cookie J

[5] This definition is similar to the one given in [17].

consists of exactly one box, we call that J a *small* cookie. Otherwise we call J a *large* cookie See Figure 1.15.

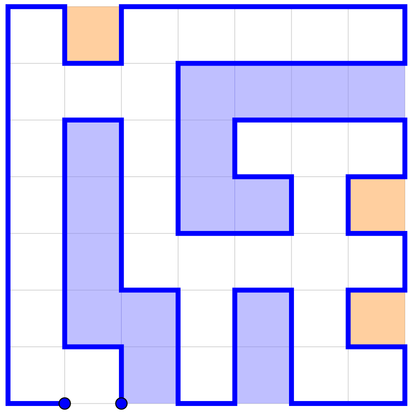

Fig. 1.15. An 8×8 grid graph with a Hamiltonian e-cycle in blue. Large cookies shaded blue. Small cookies shaded orange. $J_0 = \text{int}(H)$ in white.

Let H be a Hamiltonian path of G. A box of G with vertices a, b, c, d is *switchable* in H if it has exactly two edges in H and those edges are parallel to each other.

Lemma 1.3.14. Let G be a simply connected polyomino, let H be an e-cycle of G and let $J_0, \ldots, J_s$ be the H-components of G. Then every edge of H is incident on a box of J_0 and a box of $G_{-1} \setminus J_0$.

Proof. Let $H = v_1, \ldots, v_r$. We adopt here the partitioning of the Hamiltonian path into three subpaths and all relevant notation from Corollary 1.3.8. Note that $\{v_s, v_t\}$ in Corollary 1.3.8 corresponds to $\{v_1, v_r\}$ here. Then $P_1 = v_1$, $P_3 = v_r$ and $P_2 = H$. For definiteness, by Corollary 1.3.8 (a), we may assume that Boxes($\Phi(H, \text{right})$) $\subseteq J_0$ and Boxes($\Phi(P_2, \text{left})$) $\subseteq G_{-1} \setminus J_0$. Then every edge of H is incident on a box of J_0 and a box of $G_{-1} \setminus J_0$. $\square$

Corollary 1.3.15. Let G, H and $J_0, \ldots, J_s$ be as in Lemma 1.3.14. Then:

(a) Boxes of distinct cookies cannot be adjacent to one another.

(b) A large cookie has exactly one box incident on a boundary edge of G, namely its neck. Furthermore, if G is an $m \times n$ grid graph, then the neck of each large cookie is switchable.

Proof of (a). BWOC assume there are boxes $X \in J_i$, $Y \in J_j$ with $1 \leq i, j \leq s$, $i \neq j$ such that X and Y are adjacent sharing an edge e of H. Then, by Lemma 1.3.14, at least one of X and Y must belong to J_0, contradicting that $i, j \geq 1$.

Proof of (b). First we show that a large cookie has exactly one box incident on a boundary edge of G, namely its neck. Let J_i be a cookie. By Corollary 1.3.2, J_i has a neck N_{J_i} and

N_{J_i} is incident on $B(G)$. Suppose that there is another box X of J_i that is incident on $B(G)$. Note that $X \in G_{-1} \setminus J_0$. Let e be the boundary edge of X and let Y be the box in $G_{-1} \setminus G$ that is incident on e. Then either $e \in H$ or $e \notin H$. Note that $e \notin H$ contradicts Corollary 1.3.2 so we only need to check the case where $e \in H$. Suppose that $e \in H$. $X \in G_{-1} \setminus J_0$ and Lemma 1.3.14 imply that Y is in $J_0 \subset G$. But then $Y \in G$ and $Y \in G_{-1} \setminus G$ is a contradiction.

Now we show that in an $m \times n$ grid graph, the neck of each large cookie is a switchable box. Let $X = R(k,l)$ be the neck of a cookie J_i. Let $v(k,l) = a$, $v(k+1,l) = b$, $v(k+1,l+1) = c$ and $v(k,l+1) = d$. For definiteness, assume that $\{a,b\}$ is the neck edge of J_i. Since $i \neq 0$, neither a nor b is an end-vertex of H. This, together with the assumption that G is an $m \times n$ grid graph, implies that $\{a,d\} \in H$ and $\{b,c\} \in H$. Since J_i is a large cookie, $\{d,c\} \notin H$. Thus X is switchable. $\square$

Remark. All the definitions made for the case where H is a Hamiltonian e-cycle of a simply connected polyomino G, as well as Lemma 1.3.14 and Corollary 1.3.15, translate immediately to the case where H is a Hamiltonian cycle of G. Let H' be a Hamiltonian e-cycle of G. We may just add the edge e incident on the end-vertices of the H' to obtain the cycle H. J_0 no longer has a neck, and all other H-components and the properties we have observed remain unchanged.

1.4 Moves

Let G be a polyomino and let H be a Hamiltonian path of G. We will define here three types of moves we can apply to H in order to obtain a new Hamiltonian cycle H': switches, double-switches and backbites. In Chapter 2 we describe an algorithm that can reconfigure any two Hamiltonian cycles H and H' of an $m \times n$ grid graph into one another by using only double-switch moves. In Chapter 5 we describe an algorithm that can reconfigure any two Hamiltonian paths H and H' of an $m \times n$ grid graph into one another by using switch, double-switch and backbite moves.

Switch and double-switch moves. Let H be a Hamiltonian path of a polyomino G. Let $abcd$ be a switchable box with edges ab and cd in H. We define a *switch move* on the box $abcd$ in H as follows: remove edges ab and cd and add edges bc and ad. Let $X \in G$ be a switchable box in H. We denote a switch move by $\mathrm{Sw}(X)$.

A *double-switch move* is a pair of switch operations where we first switch X and then find a switchable Y and switch it. We denote a double-switch move by $X \mapsto Y$. See Figure 1.14.

If after a move, be it switch or double switch, we get a new Hamiltonian path, then we call the move a *valid move*. From here on we will often say "move" to mean "valid move". If X and Y share an edge and $X \mapsto Y$ is a valid move, we call $X \mapsto Y$ a *flip* move.

It will be useful to consider H as a directed path, so we assign an orientation $v_1, \dots, v_r$ to H. Let X be a switchable box in H with edges $e_1 = (v_s, v_{s+1})$ and $e_2 = (v_t, v_{t+1})$. We say e_1 and e_2 are *parallel* if v_s is adjacent to v_t in G, and *anti-parallel* if v_s is adjacent to v_{t+1}. Similarly, we call the box X a *parallel (anti-parallel)* box if its edges are parallel (anti-parallel).

Let $P(v_x, v_y)$ be a subpath of H and assume that μ is a valid move that gives a Hamiltonian path $H' = v'_1, \dots, v'_r$, where $v'_1 = v_1$ and $v'_r = v_r$. If there is a subpath $P(v'_x, v'_y)$ of H' such that $P(v'_x, v'_y) = P(v_x, v_y)$, then we say that μ *fixes* $P(v_x, v_y)$; and if there is a subpath $P(v'_x, v'_y)$ of H' such that $P(v'_x, v'_y) = P(v_y, v_x)$, then we say that μ *reverses* $P(v_x, v_y)$.

Let $H = v_1, \dots, v_r$. We will show below that if we switch a switchable box of H with anti-parallel edges, we get a path H_p and a cycle H_c. We define a (H_p, H_c)*-port* to be a switchable box in $H_p \cup H_c$ that has one edge in H_p and the other in H_c. See Figure 1.16

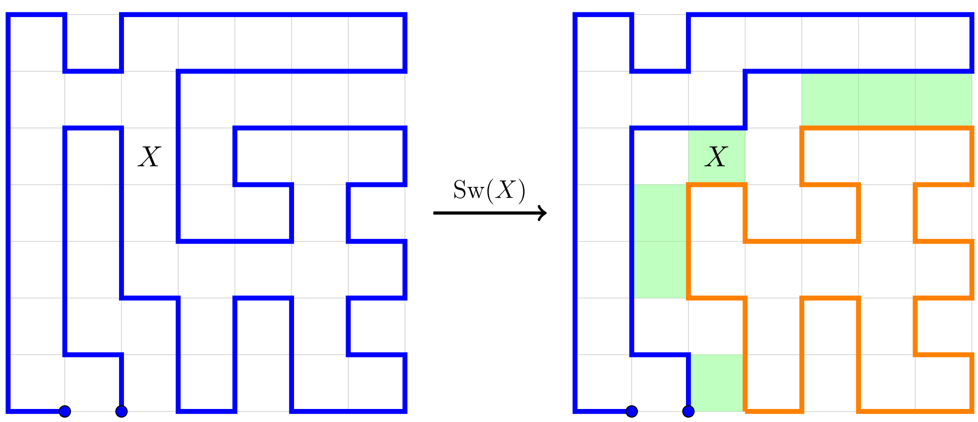

Fig. 1.16. Left. An 8×8 grid graph with a Hamiltonian path in blue. Right. The cycle H_c in orange and the path H_p in blue, after we apply the move Sw(X) on the switchable box X. (H_p, H_c)-ports shaded green.

Lemma 1.4.1. Let G be a polyomino. Let $H = v_1, \dots, v_r$ be a Hamiltonian path of G and let $s + 1 < t$. Let X be a switchable box of H with edges $e_1 = (v_s, v_{s+1})$ and $e_2 = (v_t, v_{t+1})$. Let $P_1 = P(v_1, v_s)$, $P_2 = P(v_{s+1}, v_t)$ and $P_3 = P(v_{t+1}, v_r)$.

(i) If e_1 and e_2 are parallel, Sw(X) is a valid move that fixes P_1 and P_3 and reverses P_2.

(ii) If e_1 and e_2 are anti-parallel Sw(X) splits H into a cycle H_c and a path H_p.

(iii) Suppose e_1 and e_2 are anti-parallel and assume we apply Sw(X). If Y is an

(H_p, H_c)-port then $X \mapsto Y$ is a valid double-switch move.

Proof. Let X be a switchable box of H. Without loss of generality we may assume that $s + 1 < t$. If we remove e_1 and e_2, H splits into three disjoint sub-paths: $P_1 = P(v_1, v_s)$, $P_2 = P(v_{s+1}, v_t)$ and $P_3 = P(v_{t+1}, v_r)$.

Proof of (i). Suppose that e_1 and e_2 are parallel. Then v_s is adjacent to v_t, and v_{s+1} is adjacent to v_{t+1}. Now, adding $e'_1 = (v_s, v_t)$ and $e'_2 = (v_{s+1}, v_{t+1})$ gives a new Hamiltonian path H' on $v_1, \ldots, v_s, v_t, \ldots, v_{s+1}, v_{t+1}, \ldots, v_r$. See Figure 1.17. End of proof for (i).

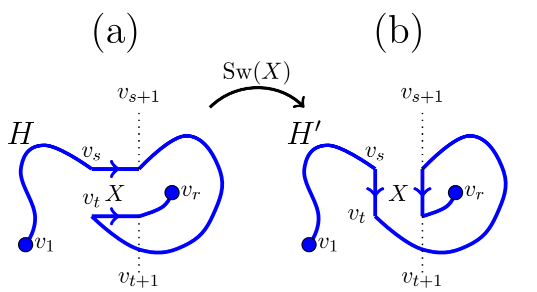

Fig. 1.17. (a) A directed Hamiltonian path H, X is switchable, e_s and e_t are parallel. (b) A directed Hamiltonian path H', obtained from H after $Sw(X)$.

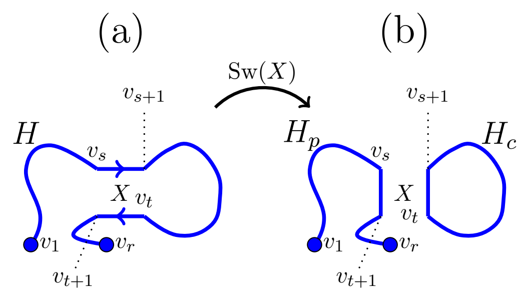

Fig. 1.18. (a). A directed Hamiltonian path H, X is switchable, e_1 and e_2 are anti-parallel. (b). A path H_p and a cycle H_c, obtained from H after $Sw(X)$.

Proof of (ii). Suppose that e_1 and e_2 are anti-parallel. Then we have that v_s is adjacent to v_{t+1} and v_{s+1} is adjacent to v_t. Now adding $e'_1 = (v_s, v_{t+1})$ gives path H_p on vertices $v_1, \ldots, v_s, v_{t+1}, \ldots, v_r$ that joins P_1 with P_3; and adding $e'_2 = (v_{s+1}, v_t)$ joins the endpoints of P_2 to give a cycle H_c on $v_{s+1}, v_{s+2}, \ldots, v_t, v_{s+1}$. See Figure 1.18. End of proof for (ii).

The proof of (iii) is essentially the same as the proof for (ii). □

Corollary 1.4.2. Let $H = v_1, \ldots, v_r$ be a Hamiltonian path of a polyomino G. Let X be an anti-parallel switchable box with edges $e_1(X)$ and $e_2(X)$, and let Y be a parallel switchable box that is edge-disjoint from X, on edges (v_s, v_{s+1}) and (v_t, v_{t+1}). Let $P_1 = P(v_1, v_s)$, $P_2 = P(v_{s+1}, v_t)$ and $P_3 = P(v_{t+1}, v_r)$, and assume that $s + 1 < t$. Then:

(i) If $e_1(X) \in P_i$ and $e_2(X) \in P_j$ where $(i, j) \in \{(1, 1), (2, 2), (3, 3), (1, 3), (3, 1)\}$, then after Sw($Y$), X remains anti-parallel.

(ii) If $e_1(X) \in P_i$ and $e_2(X) \in P_j$ where $(i, j) \in \{(1, 2), (2, 1), (2, 3), (3, 2)\}$, then after Sw($Y$), X becomes parallel. □

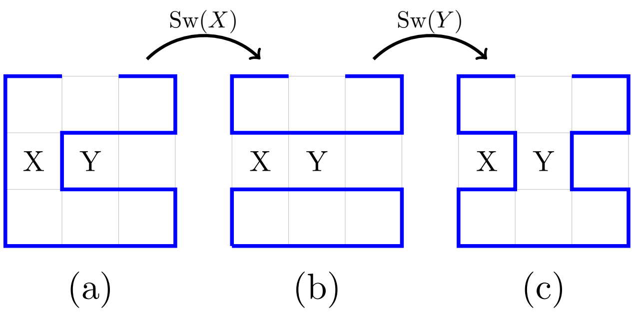

Fig. 1.19. (a) A Hamiltonian e-cycle H on a 4×4 grid; X is switchable. (b) H_{cycle} and H_{path} after switching X. Note that Y is an $(H_{\text{cycle}}, H_{\text{path}})$-port. (c) A Hamiltonian e-cycle H' after switching Y.

Corollary 1.4.3. Let H be a Hamiltonian cycle or e-cycle of an $m \times n$ grid graph G, and let X be a switchable box with edges e_1 and e_2. By Lemma 1.5.1 in the next section, all switchable boxes in a Hamiltonian cycle or e-cycle are anti-parallel. Thus, e_1 and e_2 are anti-parallel.

(i) Sw(X) splits H into the cycles H_1 and H_2.

(ii) Assume we apply Sw(X). If Y is an (H_1, H_2)-port then $X \mapsto Y$ is a valid double-switch move. See Figure 1.19. □

Backbite moves. A *backbite move* consists of adding an edge e of $G \setminus H$ incident on an endpoint of H and simultaneously removing an edge e' of H adjacent to e such that the resulting graph H' is a Hamiltonian path. We will use the notation $e' \mapsto e$ to indicate that we are removing the edge $e' \in H$ and adding the edge $e \in G \setminus H$.

Here is a more detailed description of a backbite move. Let $H = v_1, v_2, \ldots, v_r$ be a Hamiltonian path of G. First we choose a vertex $v_s \in G$ adjacent to v_1 such that $\{v_1, v_s\} \in G \setminus H$. We will add to H the edge $\{v_1, v_s\}$ and remove from H the edge (v_{s-1}, v_s). This is a valid backbite move that gives a Hamiltonian path H' distinct from H. See Figure 1.20.

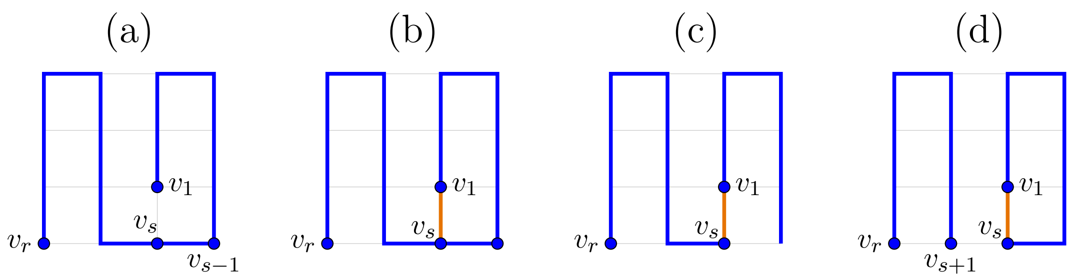

Fig. 1.20. A Hamiltonian path H on 4×4 grid. (b) H with $\{v_1, v_s\}$ added. (c) H' after the valid backbite move $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$ is applied in H. (d) The result of the invalid backbite move $(v_s, v_{s+1}) \mapsto \{v_1, v_s\}$.

Lemma 1.4.4. Let $H = v_1, v_2, \ldots, v_r$ be a Hamiltonian path of a graph G and assume that $\{v_1, v_s\} \in G \setminus H$. Then $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$ is a valid backbite move and $(v_s, v_{s+1}) \mapsto \{v_1, v_s\}$ is not. Furthermore, $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$ reverses $P(v_1, v_{s-1})$ and fixes $P(v_s, v_r)$, and the end-vertex in the resulting Hamiltonian path is located on the vertex of G where v_{s-1} was located in H.

Proof. Suppose we add the edge $\{v_1, v_s\}$ to H. Note that if we remove the edge $\{v_s, v_{s+1}\}$ we get a cycle H_c consisting of $P(v_1, v_s)$ and $\{v_1, v_s\}$ and a path $H_p = P(v_{s+1}, v_r)$ on the vertices $v_{s+1}, v_{s+2}, \ldots, v_r$. Thus $(v_s, v_{s+1}) \mapsto \{v_1, v_s\}$ is not a valid backbite move. If we remove the edge (v_{s-1}, v_s) we get a new Hamiltonian path $H' = v_{s-1}, v_{s-2}, \ldots, v_1, v_s, v_{s+1}, \ldots, v_r$, where we see that $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$ fixes $P(v_s, v_r)$ and reverses $P(v_1, v_{s-1})$. See Figure 1.20. □

Corollary 1.4.5. Let $H = v_1, \ldots, v_r$ be a Hamiltonian path of a polyomino G. Let X be an anti-parallel switchable box with edges $e_1(X)$ and $e_2(X)$, such that neither $\{v_1, v_s\}$ nor $\{v_{s-1}, v_s\}$ is an edge of X. Let $P_1 = P(v_1, v_{s-1})$ and $P_2 = P(v_s, v_r)$. Then:

(i) If $e_1(X) \in P_i$ and $e_2(X) \in P_j$ where $(i, j) \in \{(1, 1), (2, 2)\}$, then after $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$, X remains anti-parallel.

(ii) If $e_1(X) \in P_i$ and $e_2(X) \in P_j$ where $(i, j) \in \{(1, 2), (2, 1)\}$, then after $(v_s, v_{s-1}) \mapsto \{v_1, v_s\}$, X becomes parallel. □

Let $e \mapsto f$ be a backbite move. Note that if the end-vertices of H are not adjacent then $e \mapsto f$ is fully characterized by the edge f that we choose to add to H; and if the end-vertices of H are adjacent then $e \mapsto f$ is fully characterized by the edge f that we choose to add to H and the end-vertex to which the edge e is added. Whenever the edge that a backbite move removes is irrelevant to the argument, we will use the notation $bb_{v_1}(north)$ to indicate that the backbite move we're applying is the one that adds to H the edge of G on $\{v_1, v_{\text{north}}\}$, where v_{north} is the vertex of G north of v_1. The notation for backbite moves in other directions is defined analogously.

Definitions. We define the *inverse* of a move μ to be the move μ^{-1}, such that after applying μ and μ^{-1}, we get back the original Hamiltonian path. We call a move $X \mapsto X$ a *trivial* move. We define a *cascade* to be a sequence of moves $\mu_1, \ldots, \mu_r$ such that for $0 \leq j \leq r-1$:

1) μ_1 is valid,
2) if $\mu_1, \ldots, \mu_j$ have been applied then μ_{j+1} is valid, and
3) the sequence does not create any new cookies.

Let H be a Hamiltonian e-cycle of an $m \times n$ grid graph G and let J be a cookie of H with neck N_J. Consider a cascade $\mu_1, \ldots, \mu_r$ where μ_r is the nontrivial move $Z \mapsto N_J$. We say that the cascade $\mu_1, \ldots, \mu_r$ *collects* the cookie J. Note that switches, double-switches and backbites are all invertible moves. For non-adjacent boxes X and Y, the moves $X \mapsto Y$ and

$Y \mapsto X$ yield the same result. When X and Y are adjacent with X switchable and Y a leaf (i.e. $X \mapsto Y$ is a flip move), X must be switched first before Y becomes switchable, so the order matters. See Figure I.3.

1.5 Supporting lemmas

The Lemmas in this section are used in Chapters 4 and 5.

Lemma 1.5.1. Let $H = v_1, \ldots, v_r$ be a Hamiltonian cycle or e-cycle of a polyomino G. Then all switchable boxes of G are anti-parallel.

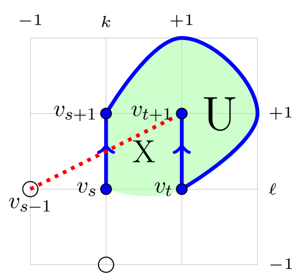

Fig. 1.21. Case 1: $X \in U$.

Proof. Assume that H is a Hamiltonian cycle or e-cycle of G. For a contradiction, assume that there is a parallel box X with edges (v_s, v_{s+1}) and (v_t, v_{t+1}), and without loss of generality, that $s < t+1$. For definiteness, let $X = R(k, l)$, $v_s = v(k, l)$, and $v_{s+1} = v(k, l+1)$. Then $v_t = v(k+1, l)$ and $v_{t+1} = v(k+1, l+1)$. Let Q be the cycle $P(v_s, v_t), \{v_t, v_s\}$ and let U be the region of G bounded by Q. Either $X \in U$ or $X \notin U$.

CASE 1: $X \in U$. Then, by Lemma 1.3.3 (b) $v_{t+1} \in U \setminus Q$. By Lemma 1.3.3 (c), $P(v_{t+1}, v_r)$ is contained in $U \setminus Q$ as well, so $v_r \in U \setminus Q$. Now, $v_{s-1} = v(k-1, l)$, or $v_{s-1} = v(k, l-1)$, or $v_s = v_1$. Note that if $v_s = v_1$, then all neighbours of v_s must belong to Q. But then none of them can be v_r which contradicts that H is a Hamiltonian cycle or e-cycle.

Then, either $v_{s-1} = v(k-1, l)$, or $v_{s-1} = v(k, l-1)$. See Figure 1.21. Without loss of generality, assume that $v_{s-1} = v(k-1, l)$. Since the segment $[v(k-1, l), v_{t+1}]$ (in red in Figure 1.21) satisfies the hypotheses of Corollary 1.1.5, $v(k-1, l)$ must belong to $G \setminus (U \cup Q)$. By JCT, $P(v_1, v(k-1, l))$ is contained in $G \setminus (U \cup Q)$ as well, so $v_1 \in G \setminus (U \cup Q)$. But then JCT implies that the path $v_{t+1}, \ldots, v_r, v_1$ intersects Q, contradicting that H is a Hamiltonian cycle or e-cycle. End of Case 1.

CASE 2: $X \notin U$. This is similar to Case 1, so we omit the proof. □

Corollary 1.5.2. Let G be an $m \times n$ grid graph and let H be a Hamiltonian path of G. If both end-vertices of H are in B_0, then every switchable box is anti-parallel. □

Proof. Let u and v be the end-vertices of H. Then u and v are on the same side of B_0, or on adjacent sides, or on opposite sides, so there are three cases to check.

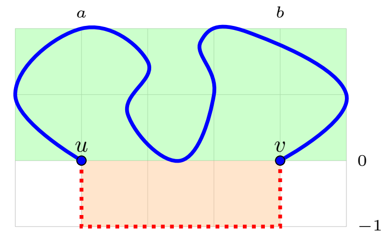

Fig. 1.22. Case 1: G shaded green, $G' \setminus G$ shaded orange, H in blue, $H' \setminus H$ in red.

CASE 1. u and v are on the same side of B_0. Without loss of generality assume that $u = v(a, 0)$ and $v = v(b, 0)$ are on the southern side of B_0. Let G' be the simply connected polyomino obtained by adding to G the boxes $R(a, -1), R(a, -2), \ldots, R(b-1, -1)$, and let H' be the Hamiltonian cycle of G' obtained by adding to H the edges $e(a; 0, -1), e(a, a+1; -1), e(a+1, a+2; -1), \ldots, e(b-1, b; -1), e(b; -1, 0)$. See Figure 1.22. By Lemma 1.5.1, G' has no boxes that are parallel in H'. It follows that G has no boxes that are parallel in H. End of Case 1.

The same argument works for the other two cases, so we omit the proof. $\square$

Corollary 1.5.3. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path of G, and let J be an H-component of G with switchable neck N_J such that no end-vertices of H are incident on $J \setminus B(J)$. Then the box N_J is anti-parallel.

Proof. By Corollary 1.3.11, J is a simply connected polyomino, and $\overrightarrow{K}_J$ is a Hamiltonian e-cycle of J. Then, by Lemma 1.5.1, N_J is anti-parallel. $\square$

Lemma 1.5.4. Let $H = v_1, \ldots, v_r$ be a path in a polyomino G such that the box X with edges (v_s, v_{s+1}) and (v_t, v_{t+1}) in H is anti-parallel. Let $s > 1$, $t < r$, and $s + 1 < t$. Let Q be the cycle $P(v_s, v_{t+1}), \{v_{t+1}, v_s\}$ and let U be the region of G bounded by Q. Then, either both end-vertices of H are incident on $U \setminus Q$, or neither is.

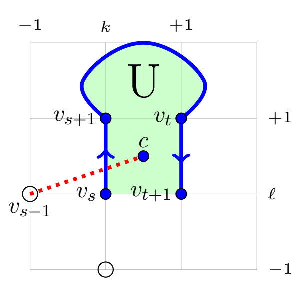

Fig. 1.23. Case 1: $X \in U$.

Proof. For definiteness, let $X = R(k, l)$, $v_s = v(k, l)$, and $v_{s+1} = v(k, l+1)$. Then $v_t = v(k+1, l+1)$ and $v_{t+1} = v(k+1, l)$. Let $c = (k + \frac{1}{2}, l + \frac{1}{2})$ be the center of the box X. Either $X \in U$ or $X \notin U$.

CASE 1: $X \in U$. Note that $v_{s-1} = v(k-1, l)$ or $v_{s-1} = v(k, l-1)$. In either case, the segment $[v_{s-1}, c]$ satisfies the premises of Corollary 1.1.5, and so v_{s-1} must belong to $G \setminus (U \cup Q)$. By JCT, v_1 must be in $G \setminus (U \cup Q)$. The same argument can be used to show that v_{t+2}, and consequently, v_r, must also belong to $G \setminus (U \cup Q)$. End of Case 1. See Figure 1.23.

CASE 2: $X \notin U$. A similar argument to the one in Case 1 can be used to show that in this case, both end-vertices must be incident on $U \setminus Q$. End of Case 2. □

Corollary 1.5.5. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path of G, and let J be an H-component of G with neck N_J. If N_J is anti-parallel, then either both end-vertices of H are incident on $J \setminus B(J)$, or neither is. □

1.6 Summary

This chapter proved structural properties of rectangular grid graphs, their Hamiltonian paths and cycles, and their H-components that are used throughout the dissertation. Key definitions include the Follow-The-Wall construction, H-components, cookies, and necks. Key results include Corollary 1.2.2, Corollary 1.3.11, and the move characterization results in Section 1.4. See Flowchart 1 on the next page for an illustration of how Chapter 1 results are used by Chapters 3–5.

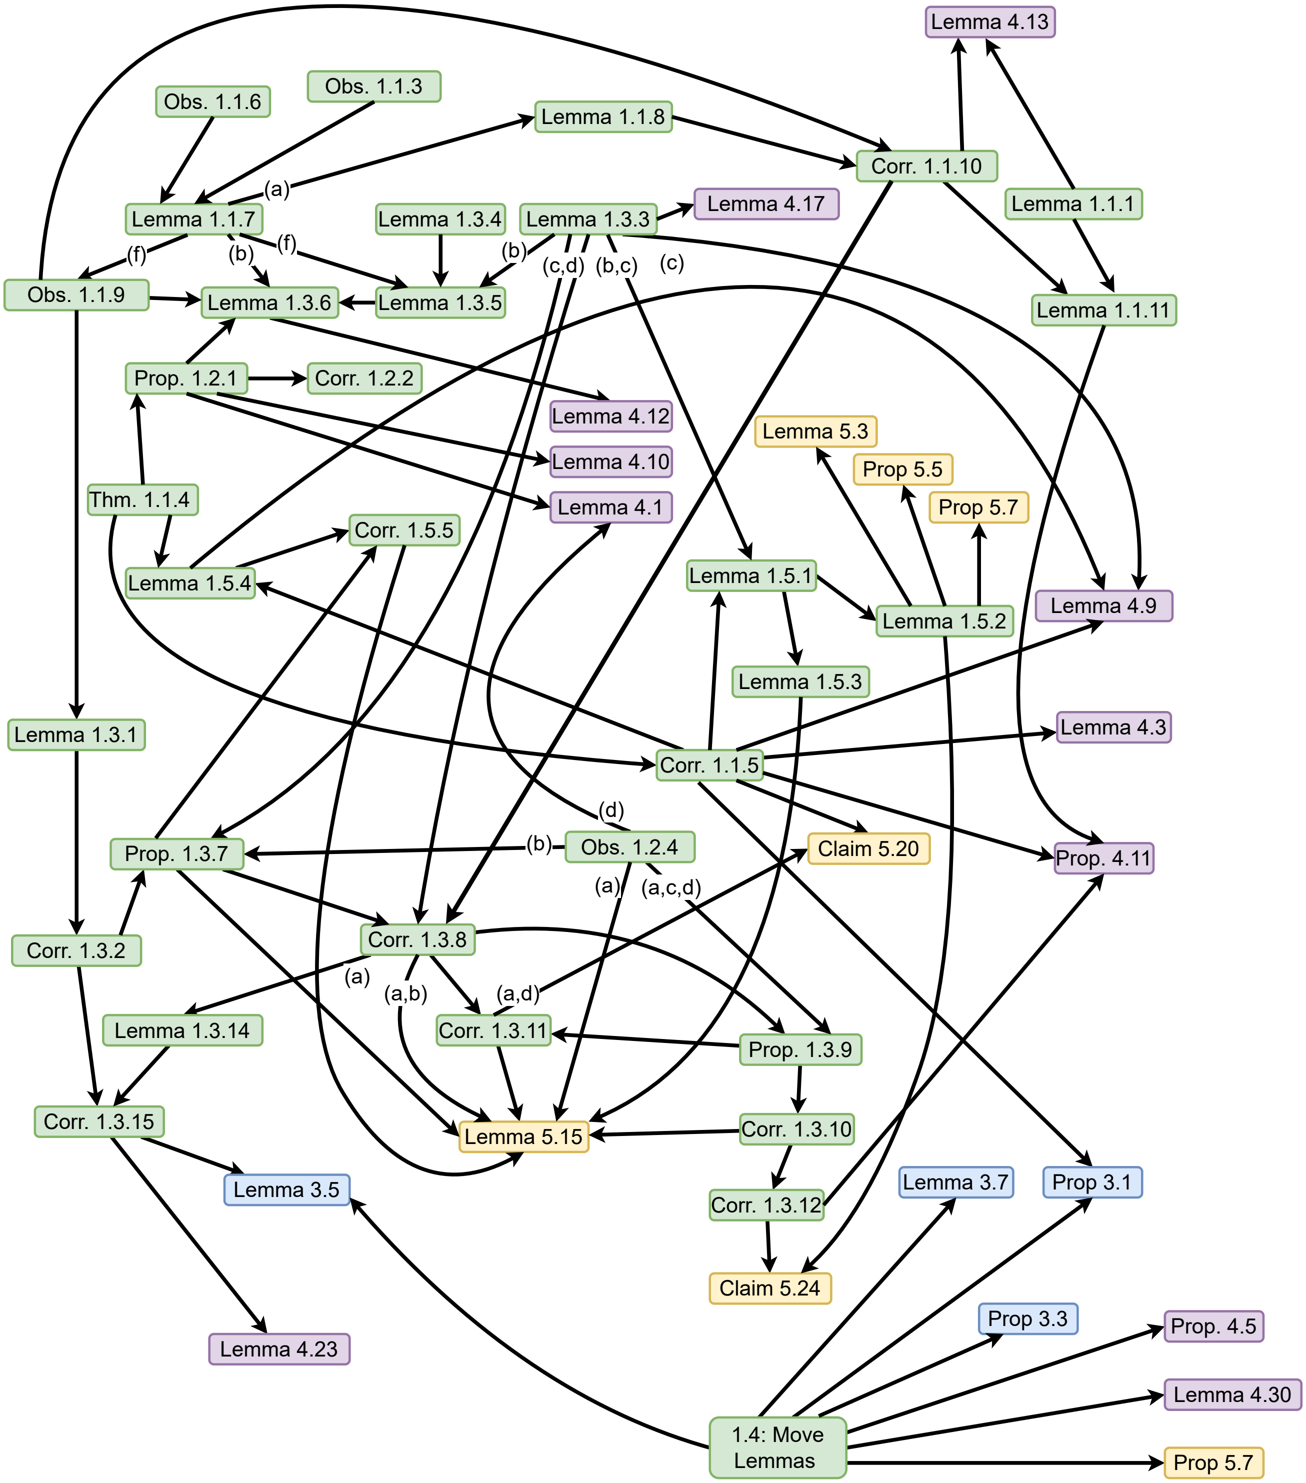

Flowchart 1. Dependencies within Chapter 1 and connections to later chapters. Arrows indicate prerequisite relationships between results.

2 Reconfiguration algorithm for Hamiltonian cycles

In this chapter we give the algorithms required for the reconfiguration of Hamiltonian cycles on an $m \times n$ grid graph G. The $3 \times n$ and $4 \times n$ cases were done by Nishat in [18], so from here on, we will assume that $m, n \geq 5$. Furthermore we assume that m and n are not both odd, since in that case G does not have a Hamiltonian cycle. See Note 2.3 at the end of the chapter.

We start with the statement of the main theorem and a few definitions, and follow by an overview of the general reconfiguration strategy.

Definitions. Let G be an $m \times n$ grid graph. Denote by G_s the induced subgraph of G on all the vertices with distance s or greater from the boundary of G. Denote by R_s the rectangular induced subgraph on vertices of G with distance s from the boundary. Then R_s is the boundary of G_s and the edges of R_0 are the boundary edges of G. We define (R_i, R_{i+1}) to be the set of all the boxes of G adjacent to both R_i and R_{i+1}. See Figure 2.1.

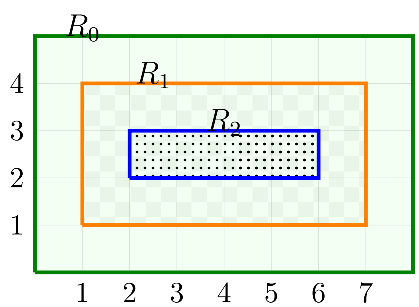

Fig. 2.1. R_0 in green, $G = G_0$ is a 9×6 grid graph on all light green boxes; R_1 in orange, G_1 is the rectangle $R(1, 7; 1, 4)$ on all checkered boxes; R_2 in blue, G_2 is the rectangle $R(2, 6; 2, 3)$ on all dotted boxes. (R_1, R_2) consists of all checkered but not dotted boxes.

Theorem 2.1 below is the main result of this dissertation. Its proof will require Chapters 3 and 4.

Theorem 2.1. Let G be an $m \times n$ grid graph with $n \geq m$. Let H and K be two Hamiltonian cycles or Hamiltonian e-cycles of G with the same edge e. Then there is a sequence of at most n^2m valid double-switch moves that reconfigures H into K.

2.1 Canonical forms

Overview of reconfiguration strategy. Let G be an $m \times n$ grid graph. We denote the set of all Hamiltonian cycles and e-cycles on an G by $\mathcal{H}(m, n)$. We shall define a subset

$\mathcal{H}_{\text{can}}(m,n)$ of Hamiltonian cycles of G, which we call canonical forms. Then we check that reconfiguration between canonical forms is easy. The problem then reduces to reconfiguring an arbitrary Hamiltonian cycle into a canonical form. Roughly, canonical forms have the following structure: the subpath $H_i = G_i \cap H$ of H on each of the nested rectangles $G_0 = G, G_1, G_2, \dots$ down to the rectangle before the central rectangle is a Hamiltonian e-cycle of that rectangle with exactly one large cookie and no small cookies. See Figure 2.2.

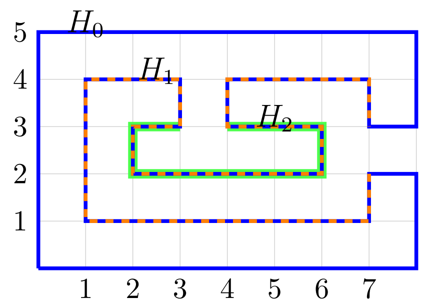

Fig. 2.2. $H = H_0$ in blue is a Hamiltonian e-cycle of a 9×6 grid graph G. The subpath H_1 of H in G_1 is in orange; it is a Hamiltonian e-cycle of G_1. The subpath H_2 of H in G_2 is highlighted in green; it is a Hamiltonian e-cycle of G_2. H_0 and H_1 have exactly one cookie.

Description of canonical forms. A Hamiltonian cycle H of an $m \times n$ grid graph G belongs to $\mathcal{H}_{\text{can}}(m,n)$ if and only if H can be constructed by the "Canonical Form Builder" algorithm described below.

Let $t = \left\lfloor \frac{\min(m,n)-4}{2} \right\rfloor$. Let $k_1 = |m-n|+2$ and $k_2 = |m-n|+3$. If $\min(m,n)$ is even, let D be the Hamiltonian cycle of the $2 \times k_1$ grid graph G_{t+1}. If $\min(m,n)$ is odd, let D be any Hamiltonian cycle of the $3 \times k_2$ grid graph G_{t+1}. Let $U = D \cup \bigcup_{i=0}^{t} R_i$. See Figure 2.3.

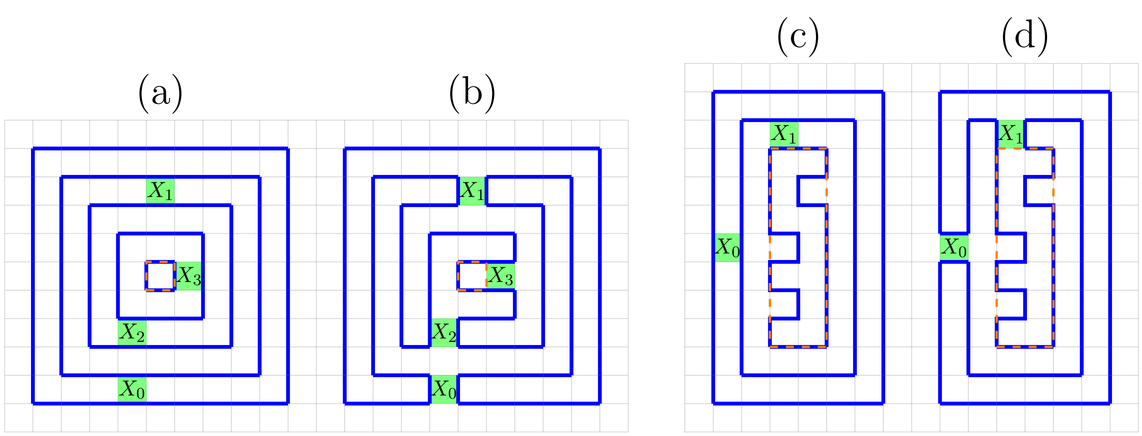

Fig. 2.3. (a) U_1 with $\min(m,n)$ even. D_1 in dashed orange. (b) A canonical form from U_1. (c) U_2 with $\min(m,n)$ odd. (d) A canonical form from U_2. D_2 in dashed orange.

Now we can state the Canonical Form Builder algorithm (CFB) that takes as inputs m and n and outputs an element of $\mathcal{H}_{\text{can}}(m,n)$.

Step 1. Set $i = 0$. Switch one of the $2(m-3)+2(n-3)$ switchable boxes of (R_0, R_1) of the graph U. This switch removes some edge, say e_1, from $E(R_1)$. If $t = 0$, stop. If $t > 0$, go to Step 2.

Step 2. Increase i by 1. Switch one of the switchable boxes of $(R_i, R_{i+1}) \setminus e_i$.

Step 3. If $i \leq t$, go to Step 2. If $i = t+1$, then stop.

We have arrived at a canonical form H. Record the switched boxes $X_0, \ldots, X_t$ in a list *List(H)*. So List(H) $= (X_0, \ldots, X_t)$ consists of the faces of G that were chosen to be switched to make U into a canonical form, listed in order.

We observe that the CFB algorithm above works just as well for e-cycles if we remove e from U.

Reconfiguration between canonical forms. Let $H, K \in \mathcal{H}_{\text{can}}(m,n)$. Let List(H) $= (X_0, \ldots, X_t)$ and List(K) $= (Y_0, \ldots, Y_t)$ be the switched boxes of H and K respectively. We will reconfigure H into K, so the algorithm will run on H.

Let $D(H) = H \cap G_{t+1}$ and note that $D(H)$ is an e-cycle of G_{t+1}. Using the result of Nishat in [18], $D(H)$ can be reconfigured into $D(K)$ by a sequence of valid moves.

Step 1. Set $i = 0$. If $t = 0$, go to Step 3. If $t > 0$, go to Step 2.

Step 2. If Y_i is switchable after switching X_i, switch both X_i and Y_i.

If Y_i is not switchable after switching X_i, switch X_{i+1} and any other switchable box in (R_{i+1}, R_{i+2}), say X'_{i+1}, such that $X_{i+1} \mapsto X'_{i+1}$ is valid. We remark that the only case where Y_i would not be switchable after switching X_i occurs when Y_i is adjacent to X_{i+1}. Note that there are many possible choices for X'_{i+1}. Now Y_i is switchable. Switch both X_i and Y_i. Update List(H) by setting the $(i+1)^{st}$ slot to X'_{i+1}. Increase i by 1.

Step 3. If $i < t$, go to Step 2.

If $i = t$ and $\min(m,n)$ is even, switch X_i and Y_i, and then stop.

If $i = t$ and $\min(m,n)$ is odd, go to Step 3.1.

Step 3.1. Switch X_t and any one of the four switchable boxes, say X'_t, located on the short sides of D. Run NRI's algorithm to reconfigure $D(H)$ into $D(K)$. Switch X'_t and Y_t. Stop.

2.2 The RtCF algorithm

Reconfiguration of a cycle into a canonical form (RtCF). Let $H = H_0$ be a Hamiltonian cycle of an $m \times n$ grid graph $G = G_0$. The RtCF algorithm takes as input a Hamiltonian e-cycle and outputs a canonical e-cycle. It works iteratively to reduce the number of cookies over each of the subgraphs $G_0, G_1, ...G_t$. The existence of required moves is guaranteed by the MLC and 1LC algorithms, stated in Proposition 2.2 below, and proved in Chapter 3. Once it reduces the number of cookies of $G = G_0$ to one, it moves on to G_1. Now, consider the resulting Hamiltonian cycle H_0'. The subpath $H_1 = H_0' \cap G_1$ is a Hamiltonian e-cycle of G_1, reducing the problem to a smaller grid. This process continues until reaching the central rectangle, at which point H has been reconfigured into canonical form. See Figure 2.4 on page 42 for an illustration of a full execution of the RtCF algorithm.

Proposition 2.2. Let $H \in \mathcal{H}(m, n)$.

(a) If H has more than one large cookie, then there is a cascade that reduces the number of large cookies of H by one. This is the ManyLargeCookies (MLC) algorithm.

(b) If H has exactly one large cookie and at least one small cookie, then there is a cascade that reduces the number of small cookies of H by one and such that it does not increase the number of large cookies. This is the OneLargeCookie (1LC) algorithm.

Now we can describe the RtCF algorithm. Without loss of generality, assume H is a Hamiltonian e-cycle of G. Suppose $H = H_0$ has $c_{0;\text{large}}$ large cookies and $c_{0;\text{small}}$ small cookies. We run MLC $c_{0;\text{large}} - 1$ times and then run 1LC $c_{0;\text{small}}$ times to reconfigure H_0 into H_0', where H_0' has exactly one (necessarily large) cookie C_0. We define $H_1 = (G_1 \cap H_0')$ and observe that H_1 is a Hamiltonian e-cycle of G_1. This is the first iteration of (RtCF). Now we describe the j^{th} iteration. We run MLC $c_{j;\text{large}} - 1$ times and then run 1LC $c_{j;\text{small}}$ times to reconfigure H_j into H_j', where H_j' has exactly one (necessarily large) cookie C_j. The RtCF algorithm stops when $j = t$, having completed $t + 1$ iterations. We give a summary of the algorithm below.

Step 1. Set $j = 0$. Run MLC $c_{0;\text{large}} - 1$ times and then 1LC $c_{0;\text{small}}$ times on H_0 to reconfigure H_0 into H_0'.

Step 2. Increase j by 1. Set $H_j = G_j \cap H_{j-1}'$ and note that H_j is a Hamiltonian e-cycle in G_j. Run MLC $c_{j;\text{large}} - 1$ times and then 1LC $c_{j;\text{small}}$ times on H_j to reconfigure H_j into H_j'.

Step 3. If $j < t$, go to Step 2. If $j = t$, stop.

Proof of the RtCF algorithm. Let N_j be the neck of the only cookie C_j of H'_j in G_j. Define $e_1(N_j) = N_j \cap R_j$, $e_2(N_j) = N_j \cap R_{j+1}$ and $\{e_3(N_j), e_4(N_j)\} = N_j \cap H'_j$. We observe that when the RtCF algorithm stops, we have reconfigured H into

$$H_c = D(H) \cup \bigcup_{j=0}^{t} \big(R_j \cup e_3(N_j) \cup e_4(N_j)\big) \setminus \big(e_1(N_j) \cup e_2(N_j)\big).$$

Now we can see that H_c is an element of $\mathcal{H}_{\text{can}}(m,n)$ by setting $X_j = N_j$ for $j = 0, 1, \ldots, t$ and running CFB.

Bound for Theorem 2.1. Recall that $n \geq m$. Note that it takes at most $2m$ [6] moves to reconfigure canonical forms into one another. Now we count the moves required for RtCF to terminate. Observe that for each $j \in \{0, \ldots, t-1\}$, H_j has at most $2\big(\frac{n-2j}{2} + \frac{m-2j}{2}\big) = n+m-4j$ cookies. This is the number of iterations of MLC or 1LC required for each j. It will follow from the proofs in Sections 3 and 4 that each application of MLC or 1LC in H_j requires at most $\frac{1}{2}n + m - 3j + 3$ moves. So, RtCF requires at most:

$$\begin{aligned}
S &= \sum_{j=1}^{\lfloor \frac{1}{2}(m-2) \rfloor} \Big(n + m - 4j\Big)\Big(\frac{n}{2} + m - 3j + 2\Big) \\
&= \sum_{j=1}^{\lfloor \frac{1}{2}(m-2) \rfloor} 12j^2 + (-5n - 7m - 8)j + \frac{n^2}{2} + \frac{3nm}{2} + 2n + m^2 + 2m \\
&\leq \frac{m^3}{2} - \frac{3m^2}{2} + m + (-5n - 7m - 8)\left(\frac{m^2}{8} - \frac{m}{4}\right) + \frac{\left(\frac{n^2}{2} + \frac{3nm}{2} + 2n + m^2 + 2m\right)(m-2)}{2} \\
&= \frac{n^2 m}{4} + \frac{nm^2}{8} + \frac{3nm}{4} + \frac{m^3}{8} - \frac{3m^2}{4} - \frac{n^2}{2} + m - 2n \\
&= \frac{n^2 m}{2} + \frac{nm}{4}\big(3 - \frac{3m}{n} - \frac{2n}{m}\big) + m - 2n.
\end{aligned}$$

Let $x = \frac{m}{n}$. Then $\frac{3m}{n} + \frac{2n}{m} = 3x + \frac{2}{x}$. Using calculus, we find that it attains a minimum of $2\sqrt{6}$ at $x = \frac{\sqrt{6}}{3}$. Then $\big(3 - \frac{3m}{n} - \frac{2n}{m}\big)$ can be at most $3 - 2\sqrt{6} \leq -1$. It follows that RtCF requires at most $\frac{n^2 m}{2} - \frac{nm}{4} + m - 2n$ to terminate. For a complete reconfiguration we need to run RtCF once for each e-cycle and reconfigure the resulting canonical forms. So, we need at most $2\big(\frac{n^2 m}{2} - \frac{nm}{4} + m - 2n\big) + m = n^2 m - \frac{nm}{2} - 4n + 3m < n^2 m$ moves for a complete reconfiguration. We remark that this is a worst case scenario and conjecture that the typical

[6] If $\min(m,n)$ is odd, add another $(n-m)/2$ moves to reconfigure D.

number of moves required is of the order of n^2.

Note 2.3. An $m \times n$ rectangular grid graph does not admit a Hamiltonian cycle when both m and n are odd. To see this, observe that the grid graph is bipartite: color a vertex (i, j) red if $i + j$ is even and blue if $i + j$ is odd. Any cycle in a bipartite graph must alternate between the two color classes, and therefore has even length. However, when both m and n are odd, the grid has mn vertices, which is an odd number. Thus, no Hamiltonian cycle can exist.

2.3 Summary

This chapter defined canonical forms and the RtCF algorithm for transforming Hamiltonian cycles into canonical form. We observed that reconfiguration between canonical forms is straightforward, reducing the general reconfiguration problem to reaching a canonical form. The main result is Theorem 2.1. Proving that RtCF can always find its required moves is handled by the MLC and 1LC algorithms, in Chapter 3.

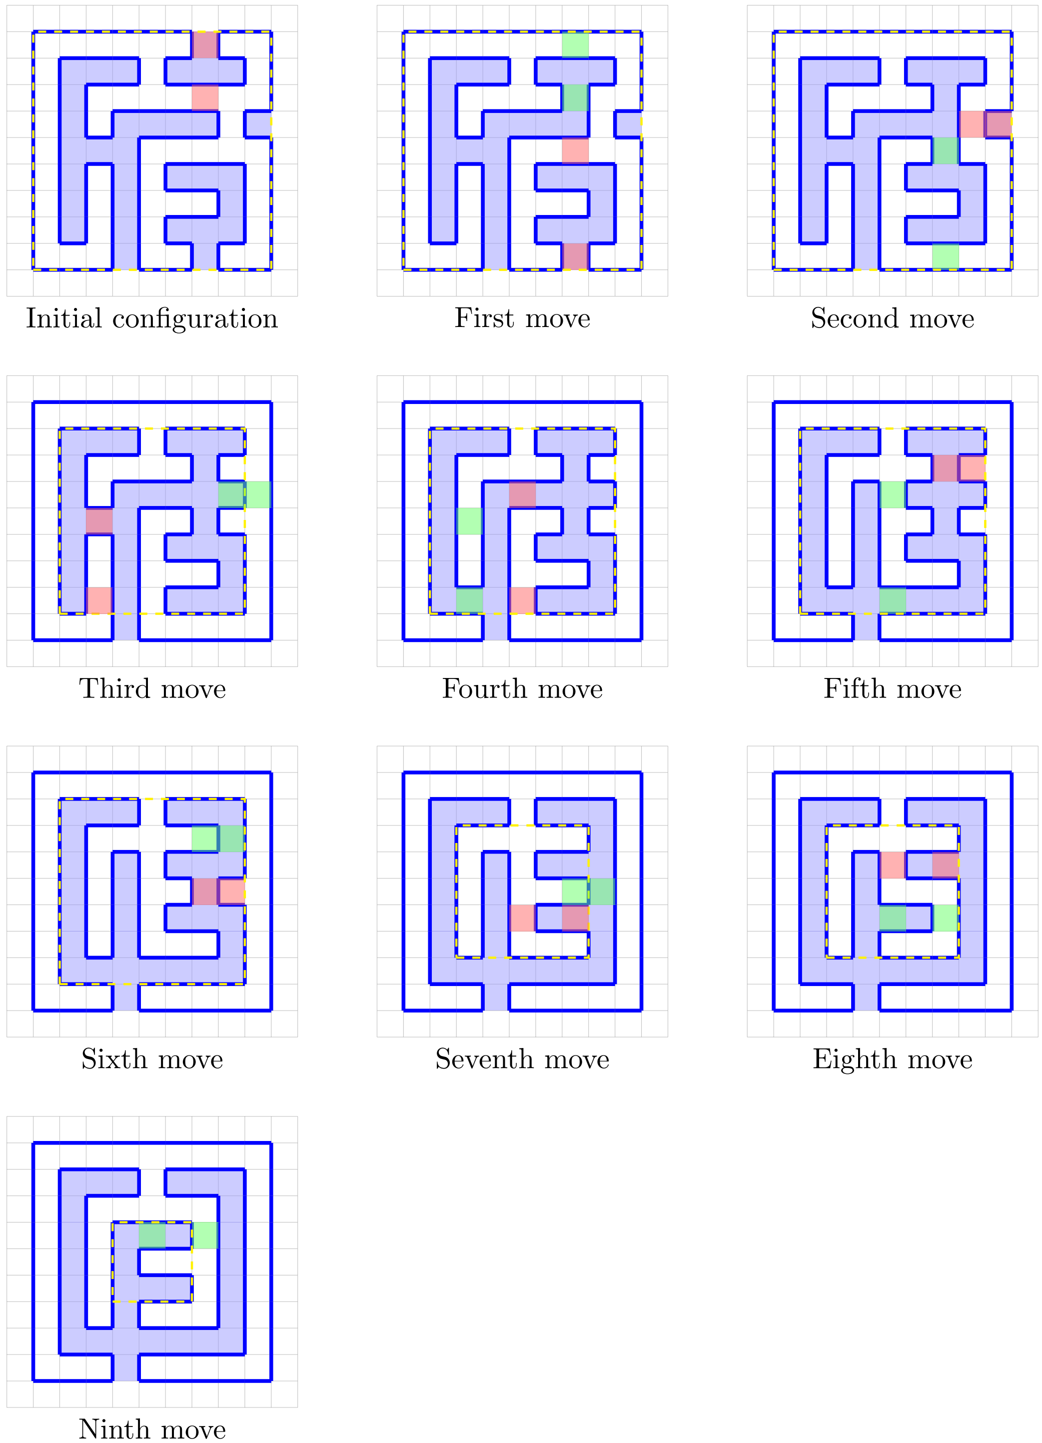

Figure 2.4. An illustration of RtCF on a 10×10 grid graph. The first three moves are in the $j = 0$ iteration of RtCF, the next three moves are in the $j = 1$ iteration, and the last three moves are in the $j = 2$ iteration. No moves are required for the $j = 3$ iteration, since H_3 has exactly one large cookie. Note that each move of the j^{th} iteration involves a box incident on the rectangle R_j, shown dashed in yellow.

3 Existence of the MLC and 1LC Algorithms

Recall that RtCF algorithm (Chapter 2) presupposes the existence of the moves required for its execution. The proofs of existence were deferred to the ManyLargeCookies (MLC) and OneLargeCookie (1LC) algorithms, which we prove in Sections 3.1 and 3.2, respectively. These algorithms ensure that whenever RtCF requires a particular move, either the move is immediately available, or else there exists a cascade after which the required move becomes available. Importantly, such cascades do not undo the progress already made: RtCF does not regress. The restrictions in the definition for cascades at the end of Section 1.4 were designed precisely for this reason.

Consider an iteration of RtCF on rectangle G_i with Hamiltonian e-cycle H_i. At this stage, H_i either has more than one large cookie, exactly one large cookie with at least one small cookie, or exactly one large cookie with no small cookies. In the last case, H_i is already in the desired form for this iteration, and RtCF proceeds to G_{i+1}. The MLC algorithm handles the first case by finding the required cascades to collect large cookies when multiple large cookies are present. The 1LC algorithm handles the second case by finding the required cascades to collect small cookies when exactly one large cookie remains.

Why do we need two separate algorithms for what appears to be the same task? This is because small cookies can be harder to collect than large ones. A second large cookie J always has a switchable neck N_J; to collect J we need only find another switchable box X such that $N_J \mapsto X$ is a valid move, or a cascade delivering such a switchable box. In Section 3.1, we show that it takes at most two moves to accomplish this (Proposition 3.8). Small cookies, by contrast, consist of a single non-switchable box. To collect a small cookie, either the box Y adjacent to it in (R_i, R_{i+1}) must be switchable, or we must find a cascade that makes Y switchable. The latter task can require much longer cascades, and it is more difficult to deal with. It requires Lemma 3.7, all of Section 3.2, most of Chapter 4, and several results from Chapter 1. Furthermore, the assumption that exactly one large cookie is present significantly shortens and simplifies the proofs of Proposition 3.10 and Lemmas 3.11–3.15 in Section 3.2, by precluding the possibility of several tedious cases.

Notation. In this chapter, we prove results for a Hamiltonian cycle H of an $m \times n$ grid graph G. These results apply verbatim to H_i on G_i at each iteration of RtCF, so we omit subscripts to simplify notation.

3.1 Existence of the MLC Algorithm

Let G be an $m \times n$ grid graph, let H be a Hamiltonian path or cycle of G, and let W be a switchable box in H. Let X and Y be the boxes adjacent to W that are not its H-neighbours, and assume that X and Y belong to the same H-component. By Corollary 1.2.2, the H-path $P(X,Y)$ is unique. We call $P(X,Y)$ *the looping H-path of W*. See Figure 3.1 for an illustration of the looping H-path of a switchable box W in a Hamiltonian cycle of a 6×6 grid graph.

We remark that all boxes Z for which $W \mapsto Z$ is a valid move (that is, in the terminology of Section 1.4, all (H_1, H_2)-ports that arise once W is switched) lie within the looping H-path $P(X,Y)$.

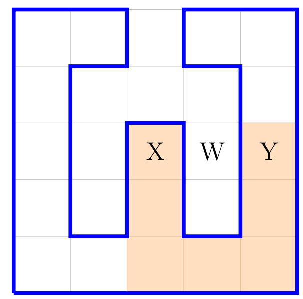

Fig. 3.1. The looping H-path of W shaded orange.

Outline of the MLC algorithm. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G with multiple large cookies. We first identify a large cookie J with switchable neck N_J. Consider what happens if N_J is switched: this would produce two cycles, H_1 and H_2. First we observe that there must be some edge $\{v_1, v_2\}$ in R_2 (recall the nested rectangles from Chapter 2) with $v_1 \in H_1$ and $v_2 \in H_2$ (Lemma 3.7). The proximity of $\{v_1, v_2\}$ to the boundary constrains the possible configurations of edges in its vicinity. We analyze those configurations (Lemma 3.5) and show that either an (H_1, H_2)-port already exists near $\{v_1, v_2\}$, or a single-move cascade on the original H yields a Hamiltonian cycle H' where such a port exists after switching N_J.

Proposition 3.1. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let $P(X,Y)$ be the looping H-path of a switchable box W. Let H' be the graph consisting of the cycles H_1 and H_2 obtained after switching W. Then a box Z of G belongs to the H'-cycle $P(X,Y), W, X$ if and only if Z is incident on a vertex of H_1 and on a vertex of H_2.

Proof. Orient H as a directed cycle $H = v_1, \ldots, v_r, v_1$. By Corollary 1.4.3, W is anti-parallel. Let the edges of W in H be $\{v_x, v_{x+1}\}$ and $\{v_{y-1}, v_y\}$. For definiteness, assume that X is adjacent to $\{v_x, v_{x+1}\}$, Y is adjacent to $\{v_{y-1}, v_y\}$ and that W is on the right of $\{v_x, v_{x+1}\}$. Then we have that $\Phi((v_x, v_{x+1}), \text{left}) = X$ and $\Phi((v_{y-1}, v_y), \text{left}) = Y$. Define $\overrightarrow{K}_1$ and $\overrightarrow{K}_2$ to be the subtrails $\overrightarrow{K}((v_x, v_{x+1}), (v_{y-1}, v_y))$ and $\overrightarrow{K}((v_{y-1}, v_y), (v_x, v_{x+1}))$ of $\overrightarrow{K}_H$, respectively. By Lemma 1.4.1 (ii), switching W splits H into two cycles H_1 and H_2, with $V(H_1) = V(\overrightarrow{K}_1) \setminus \{v_x, v_y\}$ and $V(H_2) = V(\overrightarrow{K}_2) \setminus \{v_{x+1}, v_{y-1}\}$. See Figure 3.2.

Proof of ($\Longrightarrow$). Since $P(X, Y)$ is unique, any H-walk of boxes between X and Y contains $P(X, Y)$. In particular, $\Phi(\overrightarrow{K}_1, \text{left})$ contains $P(X, Y)$ and $\Phi(\overrightarrow{K}_2, \text{left})$ contains $P(X, Y)$. Let Z be a box of $P(X, Y)$. Then Z is added to $\Phi(\overrightarrow{K}_1, \text{left})$ by an edge of $\overrightarrow{K}_1$ and Z is added to $\Phi(\overrightarrow{K}_2, \text{left})$ by an edge of $\overrightarrow{K}_2$. By definition of FTW, Z is incident on a vertex of $\overrightarrow{K}_1$ and a vertex of $\overrightarrow{K}_2$. Since for $i \in \{1, 2\}$, $V(\overrightarrow{K}_i) \supset V(H_i)$, we have that Z is incident on a vertex of H_1 and a vertex of H_2. Lastly, note that W is incident on $v_{x+1} \in H_1$ and $v_x \in H_2$. End of proof for ($\Longrightarrow$).

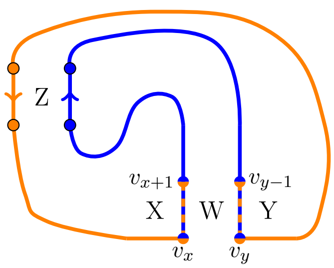

Fig. 3.2.

Proof of ($\Longleftarrow$). Suppose we switch W and obtain the graph H' consisting of the cycles H_1 and H_2. Observe that $P(X, Y), W, X$ is the only H'-cycle in G. We will say that a box Z of G satisfies $(*)$ if Z is incident on a vertex in H_1 and H_2. We will show that if a box Z of G satisfies $(*)$ then it must belong to an H'-cycle of boxes that satisfy $(*)$. Then, since there is only one H'-cycle in G, $Z = W$ or $Z \in P(X, Y)$.

Let Z be a box in G that satisfies $(*)$. For definiteness, assume that $Z = R(k, l)$. We will show that Z has exactly two neighbours in G that satisfy $(*)$ and that Z is H'-adjacent to those two neighbours. Since Z satisfies $(*)$, either Z has two vertices in H_1 and two vertices in H_2 or Z has one vertex in one of H_1 and H_2 and three vertices in the other.

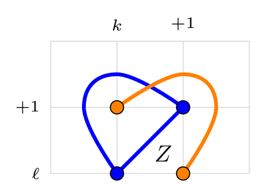

Fig. 3.3.

CASE 1: Z has two vertices in H_1 and two vertices in H_2. First we will check that the pair of vertices belonging to H_i, $i \in \{1, 2\}$ must be adjacent in Z. For contradiction, assume that $v(k, l)$ and $v(k+1, l+1)$ belong to H_1 and $v(k+1, l)$ and $v(k, l+1)$ belong to H_2. Let Q be the closed polygon consisting of the subpath $P(v(k, l), v(k+1, l+1))$ of H_1 and the segment $[v(k, l), v(k+1, l+1)]$. Then, by Corollary 1.1.5, $v(k+1, l)$ and $v(k, l+1)$ are on different sides of H_1. It follows that the subpath $P(v(k+1, l), v(k, l+1))$ intersects Q. Since $P(v(k+1, l), v(k, l+1))$ does not intersect Q at the segment $[v(k, l), v(k+1, l+1)]$, it must intersect Q at some vertex in $P(v(k, l), v(k+1, l+1))$. But this contradicts that H_1 and H_2 are disjoint. It follows that the pair of vertices belonging to H_i, $i \in \{1, 2\}$ must be adjacent in Z.

For definiteness, assume that $v(k, l)$ and $v(k+1, l)$ belong to H_1 and that $v(k+1, l+1)$ and $v(k, l+1)$ belong to H_2. Since H_1 and H_2 are disjoint $e(k; l, l+1) \notin H'$ and $e(k+1; l, l+1) \notin H'$ so $Z + (-1, 0)$ and $Z + (1, 0)$ are H'-adjacent to Z. Since $v(k, l) \in H_1 \cap V(Z + (-1, 0))$ and

$v(k, l+1) \in H_2 \cap V(Z + (-1, 0))$, $Z + (-1, 0)$ satisfies $(*)$. Similarly, $Z + (1, 0)$ satisfies $(*)$. It remains to check that $Z + (0, 1)$ and $Z + (0, -1)$ do not satisfy $(*)$.

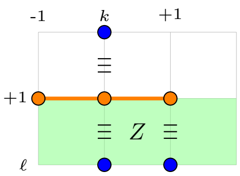

Fig. 3.4.

For contradiction, assume that one of $Z + (0, 1)$ and $Z + (0, -1)$ satisfies $(*)$. By symmetry we may assume WLOG that $Z + (0, 1)$ satisfies $(*)$. Then at least one of $v(k, l+2)$ and $v(k+1, l+2)$ belongs to H_1. By symmetry we may assume WLOG that $v(k, l+2) \in H_1$. It follows that $e(k; l+1, l+2) \notin H'$ and that $v(k-1, l+1) \in H_2$. Then we must have $e(k-1, k; l+1) \in H_2$ and $e(k, k+1; l+1) \in H_2$. But then by Corollary 1.1.5, one of $v(k, l)$ and $v(k, l+2)$ lies inside the region bounded by H_2 and the other lies outside it. It follows that the subpath $P(v(k, l), v(k, l+2))$ of H_1 intersects H_2, contradicting that H_1 and H_2 are disjoint. End of Case 1.

CASE 2: Z has one vertex in one of H_1 and H_2 and three vertices in the other. For definiteness, assume that $v(k, l)$, $v(k, l+1)$ and $v(k+1, l+1)$ belong to H_1 and that $v(k+1, l)$ belongs to H_2. Then $e(k, k+1; l) \notin H$ and $e(k+1; l, l+1) \notin H$, so $Z + (1, 0)$ and $Z + (0, -1)$ are H'-neighbours of Z. Since $v(k, l) \in H_1 \cap V(Z + (0, -1))$ and $v(k+1, l) \in H_2 \cap V(Z + (0, -1))$, $Z + (0, -1)$ satisfies $(*)$. Similarly, $Z + (1, 0)$ satisfies $(*)$. It remains to check that $Z + (0, 1)$ and $Z + (0, -1)$ do not satisfy $(*)$.

For a contradiction, assume that one of $Z + (-1, 0)$ and $Z + (0, 1)$ satisfies $(*)$. By symmetry we may assume WLOG that $Z + (0, 1)$ satisfies $(*)$. Then one of $v(k, l+2)$ and $v(k+1, l+2)$ belongs to H_2. Note that if $v(k+1, l+2) \in H_2$ we run into the same contradiction as in Case 1, so we only need to check the case where $v(k, l+2) \in H_2$. Now, either $e(k, k+1; l+1) \in H_1$, or $e(k, k+1; l+1) \notin H'$.

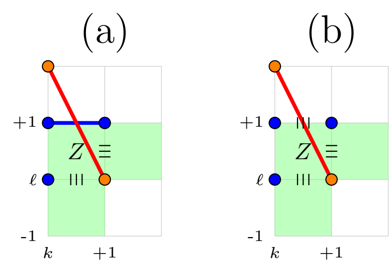

Fig. 3.5. (a) Case 2.1. (b) Case 2.2.

CASE 2.1: $e(k, k+1; l+1) \in H_1$. Then the segment $[v(k, l+2), v(k+1, l)]$ intersects H_1 at the point $(k\frac{1}{2}, l+1)$. by Corollary 1.1.5, the vertices $e(k, l+2)$ and $v(k+1, l)$ are on different sides of H_1, and we run into the same contradiction as in Case 1 again. See Figure 3.5 (a).

CASE 2.2: $e(k, k+1; l+1) \notin H_1$. Consider the polygon Q consisting of the segment $[v(k, l+2), v(k+1, l)]$ and the subpath $P(v(k, l+2), v(k+1, l))$ of H_2. By Corollary 1.1.5, the vertices $v(k, l+1)$ and $v(k+1, l+1)$ are on different sides of Q. By JCT the subpath $P(v(k, l+1), v(k+1, l+1))$ of H_1 intersects Q. Since $P(v(k, l+1), v(k+1, l+1))$ does not intersect Q at the segment $[v(k, l+2), v(k+1, l)]$, it must do so at some vertex of $P(v(k, l+2), v(k+1, l))$, contradicting that H_1 and H_2 are disjoint. See Figure 3.5 (b). End of Case 2.2. End of Case 2. □

Corollary 3.2 (i). Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let W be a switchable box in H. Let H' be the graph consisting of the cycles H_1 and H_2 obtained after switching W. Let a, b and c be colinear vertices such that b is adjacent to a and c. Then, for $i \in \{1, 2\}$, If a and c belong to H_i, so must b. See Figure 3.4. □

Corollary 3.2 (ii). Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let W be a switchable box in H. Let H' be the graph consisting of the cycles H_1 and H_2 obtained after switching W. Let Z be a box on vertices a, b, c, and d such that a and b belong to H_1, and c and d belong to H_2. Then a is adjacent to b, and c is adjacent to d. See Figure 3.3. □

Proposition 3.3. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G, let W be a switchable box in H and let $P(X, Y)$ be the looping H-path of W. If $P(X, Y)$ has a switchable box Z, then $Z \mapsto W$ is a valid move.

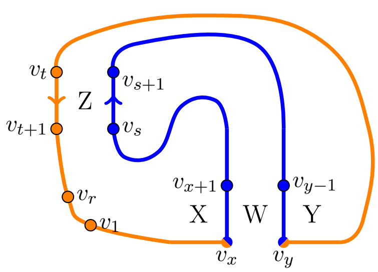

Fig. 3.6. P_1 in blue, P_2 in orange.

Proof. Let $H = v_1, v_2 \dots, v_r, v_1$. By Corollary 1.4.3, W is anti-parallel. Let W, $P(X, Y)$, $\{v_x, v_{x+1}\}$ and $\{v_{y-1}, v_y\}$ be as in Proposition 3.1. Let $P_1 = P(v_x, v_y)$ and let $P_2 = P(v_y, v_x)$. By Proposition 3.1, every box of $P(X, Y)$ is incident on a vertex of P_1 and a vertex of P_2.

Let Z be a switchable box of $P(X, Y)$. Let (v_s, v_{s+1}) and (v_t, v_{t+1}) be the edges of Z in H. For definiteness, assume $s + 1 < t$. Proposition 3.1 implies that exactly one of (v_s, v_{s+1}) and (v_t, v_{t+1}) is in P_1 and the other is in P_2. WLOG assume that (v_s, v_{s+1}) is in P_1 and that (v_t, v_{t+1}) is in P_2. Then we can partition the edges of H as follows: $P(v_1, v_s)$, (v_s, v_{s+1}), $P(v_{s+1}, v_t)$, (v_t, v_{t+1}), $P(v_{t+1}, v_r)$, $\{v_r, v_1\}$ where $1 < x < s < y < t < r$.

We check that $Z \mapsto W$ is a valid move. After removing the edges (v_s, v_{s+1}) and (v_t, v_{t+1}) we are left with two paths: $P(v_{t+1}, v_s)$ and $P(v_{s+1}, v_t)$. Note that adding the edge $\{v_s, v_{t+1}\}$ gives a cycle H_1 consisting of the path $P(v_s, v_{t+1})$ and the edge $\{v_s, v_{t+1}\}$, and adding the edge $\{v_{s+1}, v_t\}$ gives a cycle H_2 consisting of the path $P(v_{s+1}, v_t)$ and the edge $\{v_{s+1}, v_t\}$. Now $1 < x < s$ implies that $(v_x, v_{x+1}) \in H_1$ and $s < y < t$ implies that $(v_{y-1}, v_y) \in H_2$. It follows that W is now an (H_1, H_2)-port . By Lemma 1.4.1 (iii), $Z \mapsto W$ is a valid move. □

Remark. Proposition 3.3 is the primary tool for collecting cookies and characterizes when double-switch moves are valid. To collect a large cookie with switchable neck W, we search for a switchable box Z in the looping H-path of W. If such a Z exists, we can collect the cookie immediately. If not, this imposes strong structural constraints on the H-path. This dichotomy—either we find a switchable box and succeed, or we find none and exploit the resulting structure (and eventually find one anyway)—drives the case analyses of Lemma 3.5 below as well as much of Section 3.2 and Chapters 4 and 5. Beyond collecting large cookies, the proposition is essential for collecting small cookies (which requires finding a switchable box adjacent to the small cookie) and for characterizing valid double-switch moves more generally. In Chapter 4 we prove analogous results that are used in the Hamiltonian path reconfiguration algorithms of Chapter 5. Proposition 4.5 extends Proposition 3.3 and Proposition 4.4 extends Proposition 3.1.

Let H be a Hamiltonian cycle of an $m \times n$ grid graph G. Recall the definitions of the subgraphs G_i and R_i of G on page 37 in Chapter 2.

Observation 3.4. Let $X \mapsto Y$ be a valid move. If $X \in \text{ext}(H)$ and $Y \in \text{int}(H)$ then:

(i) $X \mapsto Y$ increases the total number of cookies if and only if $X \in G_1$ and $Y \in G_0 \setminus G_1$.

(ii) $X \mapsto Y$ increases the total number of large cookies, leaving the total number of cookies unchanged, if and only if $X \in G_1$, $Y \in G_1 \setminus G_2$ and Y is adjacent to a small cookie.

(iii) $X \mapsto Y$ decreases the total number of cookies if and only if $X \in G_0 \setminus G_1$, $Y \in G_1$.

Lemma 3.5. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G and let Z be a switchable box in $\text{ext}(H) \cap ((G_0 \setminus G_1) \cup G_3)$. Assume that switching Z splits H into the cycles H_1 and H_2 that are such that there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2. Then either $Z \mapsto Z'$ is a cascade, or there is a cascade $\mu, Z \mapsto Z'$ (of length two), with $Z \mapsto Z'$ nontrivial in either case.

Remark. For the MLC algorithm (Proposition 3.8) in this section, we only need the case where $Z \in \text{ext}(H) \cap (G_0 \setminus G_1)$, since cookie necks reside in $G_0 \setminus G_1$. However, the more general statement (including $Z \in G_3$) will be needed for the 1LC algorithm in Section 3.2, specifically in the proof of Lemma 3.12. We prove the general result here to avoid duplication.

In the figures illustrating the proof, vertices and edges of H_1 and H_2 are in blue and orange, respectively, and boxes of $\text{int}(H)$ are shaded in green.

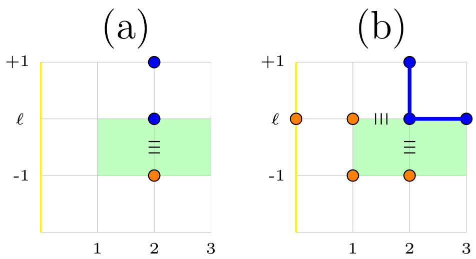

Fig. 3.7. (a). Initial configuration. (b). Case 1.

Proof. We will use the assumption that $Z \in (G_0 \setminus G_1) \cup G_3$ repeatedly and implicitly throughout the proof. Switch Z to obtain H' consisting of the disjoint cycles H_1 and H_2. Note that now, if a vertex belongs to H_i for $i \in \{1, 2\}$, both edges incident on it must also belong to H_i.

For definiteness, let $v_1 \in H_1 \cap R_2$ be the vertex $v(2, l)$ for some $l \in \{2, \ldots, n-3\}$ and let $v_2 = v(2, l-1) \in H_2 \cap R_2$. Then $e(2; l-1, l) \notin H$, and by Corollary 3.2 (i), $v(2, l+1) \in H_1$ as well. By Proposition 3.1, $R(1, l-1)$ and $R(2, l-1)$ belong to $\text{int}(H)$. Now, either $v(1, l) \in H_1$, or $v(1, l) \in H_2$. See Figure 3.7.

CASE 1: $v(1, l) \in H_2$. Then $e(1, 2; l) \notin H$. Corollary 3.2(i), $v(3, l) \in H_1$. It follows that $e(2; l, l+1) \in H'$ and $e(2, 3; l) \in H'$. Now, by Corollary 3.2 (ii), $v(1, l-1) \in H_2$ and by Corollary 3.2 (i), $v(0, l) \in H_2$. At this point we must either have, $e(1; l, l+1) \in H'$ or $e(1; l, l+1) \notin H'$. See Figure 3.7.

CASE 1.1: $e(1; l, l+1) \notin H'$. Then $e(1; l-1, l) \in H_2$ and $e(0, 1; l) \in H_2$. Since $Z \in \text{ext}(H)$, by Proposition 3.1, $R(1, l-1) \in \text{int}(H)$. Then $R(0, l-1)$ must be a small cookie of G, so we must have that $e(0, 1; l-1) \in H_2$ and that $e(1, 2; l-1) \notin H'$. It follows that $e(2, 3; l-1) \in H'$ and $e(3; l-1, l) \notin H'$. Now note that $R(2, l-1)$ is an (H_1, H_2)-port. Then, by Lemma 1.4.1 (iii) and Observation 3.4, $Z \mapsto R(2, l-1)$ is valid move that does not create new cookies. So, $Z \mapsto R(2, l-1)$ is the cascade we seek. End of Case 1.1.

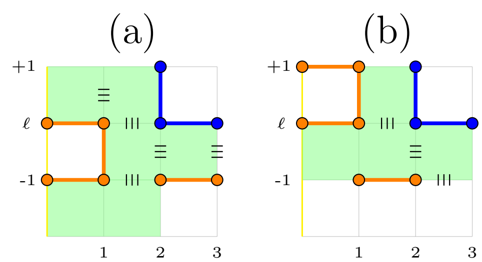

Fig. 3.8. (a) Case 1.1. (b) Case 1.2.

CASE 1.2: $e(1; l, l+1) \in H'$. Proposition 3.1, and the assumption that $Z \in \text{ext}(H)$ imply that $R(1, l) \in \text{int}(H)$. Then, Corollary 1.3.15 (b) implies that $R(0, l)$ is a small cookie of G. Note that if $e(2, 3; l-1) \in H'$, then we're back to Case 1.1, so we may assume that $e(2, 3; l-1) \notin H'$. It follows that $e(1, 2; l-1) \in H_2$ and that $R(0, l-1)$ is not a small cookie of H. Then, by Observation 3.4 and Proposition 3.3, $R(0, l) \mapsto R(1, l)$, $Z \mapsto R(1, l-1)$ is the cascade we seek. End of Case 1.2. End of Case 1.

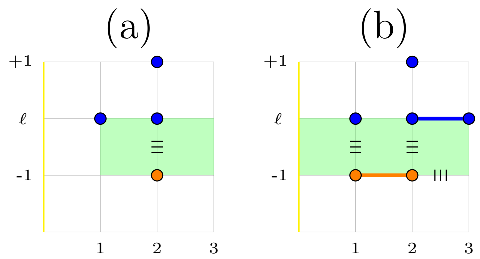

Fig. 3.9. (a) Case 2. (b) Case 2.1.

CASE 2: $v(1,l) \in H_1$. Then $e(2,3;l) \in H_1$ or $e(2,3;l) \notin H'$. See Figure 3.9.

CASE 2.1: $e(2,3;l) \in H_1$. Note that if $e(2,3;l-1) \in H_2$, then we're back to essentially the same scenario as Case 1.1, so we may assume that $e(2,3;l-1) \notin H_2$. Then $e(1,2;l-1) \in H_2$. It follows that $e(1;l-1,l) \notin H'$ and so $R(0,l-1)$ is not a small cookie of H. Now, either $e(1,2;l) \in H_1$ (Figure 3.10 (a)), or $e(1,2;l) \notin H'$ (Figure 3.10 (b)).

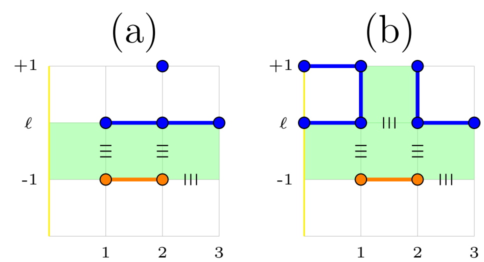

Fig. 3.10. (a) Case 2.1(a). (b) Case 2.1(b).

CASE 2.1 (a): $e(1,2;l) \in H_1$. By Observation 3.4 and Proposition 3.3, $Z \mapsto R(1,l-1)$ is the cascade we seek. End of Case 2.1(a).

CASE 2.1 (b): $e(1,2;l) \notin H'$. Then we have that $e(2;l,l+1) \in H_1$, that $e(1;l,l+1) \in H_1$ and that $R(1,l) \in \text{int}(H)$. It follows that $R(0,l) \in \text{ext}(H)$, so $R(0,l)$ must be a small cookie. Then, after $R(0,l) \mapsto R(1,l)$, we are back to Case 2.1 (a). End of Case 2.1(b). End of Case 2.1.

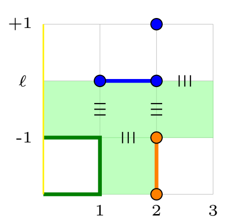

Fig. 3.11. Case 2.2.

CASE 2.2: $e(2,3;l) \notin H'$. Then we have that $e(1,2;l) \in H_1$. Note that if $e(1,2;l-1) \in H_2$, then we're back to Case 2.1, so we may assume that $e(1,2;l-1) \notin H'$. It follows that $e(2;l-2,l-1) \in H_2$. Note that if $e(1;l-1,l) \in H_1$, then $R(0,l-1) \in \text{ext}(H)$ and $R(0,l) \in \text{ext}(H)$, contradicting Corollary 1.3.15 (b), so we may assume that $e(1;l-1,l) \notin H'$. Then we must have that $e(1;l-2,l-1) \in H'$. This implies that $R(0,l-2)$ is a small cookie of H. Then after $R(0,l-2) \mapsto R(1,l-2)$, we're back to Case 2.1(a). End of Case 2.2. End of Case 2. □.

Observation 3.6. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G, with $m,n \geq 5$ and let J be a large cookie of G. Then $J \cap R_2 \neq \emptyset$.

Lemma 3.7. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G, with $m,n \geq 5$ and assume that G has at least two large cookies J_1 and J_2. Then switching the neck N_{J_1} of J_1 splits H into two cycles H_1 and H_2 such that there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with

v_1 adjacent to v_2.

Proof. Orient H. Let $\{v_x, v_y\}$ be the boundary edge of the neck N_{J_1} of J_1. Define $\overrightarrow{K}_1$ and $\overrightarrow{K}_2$ to be the subtrails $\overrightarrow{K}((v_x, v_{x+1}),\ (v_{y-1}, v_y))$ and $\overrightarrow{K}((v_y, v_{y+1}), (v_{x-1}, v_x))$ of $\overrightarrow{K}_H$, respectively. By Lemma 1.4.1 (ii), switching N_{J_1} gives two cycles H_1 and H_2, with $V(H_1) = V(\overrightarrow{K}_1 \setminus \{(v_x, v_y)\})$ and $V(H_2) = V(\overrightarrow{K}_2)$. By Corollary 1.3.11 (a) and (d), $V(J_1) = V(\overrightarrow{K}_1)$. By Observation 3.6, $V(J_1) \cap R_2 \neq \emptyset$. Then $V(H_1) = R_2 \neq \emptyset$. Since $V(J_1) = V(\overrightarrow{K}_1)$, we have that $V(J_2) \subseteq V(\overrightarrow{K}_2) = V(H_2)$. By Observation 3.6, $V(J_2) \cap R_2 \neq \emptyset$. It follows that $V(H_2) \cap R_2 \neq \emptyset$.

We have shown that $V(H_1) = R_2 \neq \emptyset$ and that $V(H_2) \cap R_2 \neq \emptyset$. It remains to check there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2. Let $v_1 \in H_1 \cap R_2$ and let $R_2 = w_1, \ldots, w_s$, with $v_1 = w_1$. Sweep R_2 starting at w_1. If there is $i \in \{1, \ldots, s-2\}$ such that $v_i \in H_1$ and $v_{i+1} \in H_2$, we are done. If there is no such i then $R_2 \cap H_2 = \emptyset$, contradicting that $R_2 \cup H_2 \neq \emptyset$. □

Proposition 3.8. (MLC Algorithm.) Let H be a Hamiltonian cycle of an $m \times n$ grid graph G. Assume that G has more than one large cookie. Then there is a cascade of length at most two that reduces the number of large cookies of G by one.

Proof. Let J be a large cookie of G with neck N_J. Switching N_J splits H into the cycles H_1 and H_2. By Lemma 3.7 there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2. By Lemma 3.5 there is a cascade of length at most two, whose last move is $N_J \mapsto N'_J$ with $N_J \neq N'_J$. By Observation 3.4, this cascade decreases the number of large cookies of G by one. □

Algorithm: ManyLargeCookies

Input: Hamiltonian cycle H on $m \times n$ rectangular grid graph.

Output: Hamiltonian cycle H' with all large cookies collected.

Assumptions and Notation: Without loss of generality, assume $m \leq n$. For a large cookie J with neck N_J, let (v_1, v_2) denote a pair of adjacent vertices belonging to cycles H_1 and H_2 (arising from switching N_J) that are incident on R_2 (guaranteed by Lemma 3.7). Let $\mathbf{B}$ denote the 4×4 rectangle centered at v_1 (or v_2).

Identify all large cookies `//` $O(m + n)$
Let numLargeCookie $\leftarrow$ number of large cookies;
while *numLargeCookie* ≥ 2 **:**
    Let J be any remaining large cookie;
    Let $N_J \leftarrow$ neck of J `//` $O(1)$
    Let $(v_1, v_2) \leftarrow$ vertex pair for J `// Lemma 3.7`
    Let $\mathbf{B} \leftarrow 4 \times 4$ box at v_1;
    `// Try direct move`
    for each *box* Z' *in region* $\mathbf{B}$ **:**
        if *move* $N_J \rightarrow Z'$ *is valid* **then**
            Execute($N_J \rightarrow Z'$);
            numLargeCookie $\leftarrow$ numLargeCookie -1;
            break;
        end
    end
    `// Direct move failed - find prep move`
    for each *box* Z'' *in region* $\mathbf{B}$ **:**
        for each *box* Z' *in region* $\mathbf{B}$ **:**
            if *both* $Z'' \rightarrow Z'$ *and* $N_J \rightarrow Z''$ *valid* **then**
                Execute($Z'' \rightarrow Z'$) `// Prep move`
                Execute($N_J \rightarrow Z''$) `// direct move`
                numLargeCookie $\leftarrow$ numLargeCookie -1;
                break;
            end
        end
    end
end
return H';

Time complexity: $O(m+n)$ total. The initial scan takes $O(m+n)$ time. Each cookie collection takes $O(1)$ time (search is localized to region **B**), and there are at most $O(m+n)$ cookies.

3.2 Existence of the 1LC Algorithm

In this section we prove the 1LC algorithm (Proposition 3.10). We need some definitions before we can state it.

Definitions. Let G be an $m \times n$ grid graph and H be a Hamiltonian cycle of G. We call a subpath of H on the edges $e(k;l,l+1), e(k,k+1;l+1)$ and $e(k+1;l,l+1)$ a *northern leaf.* We will often say that $R(k,l)$ is a northern leaf to mean that $e(k;l,l+1), e(k,k+1;l+1)$ and $e(k+1;l,l+1)$ belong to H. Southern, eastern and western leaves are defined analogously. We call the subgraph of H on the edges $e(k-1,k;l)$, $e(k;l,l+1)$, $e(k+1,k+2;l)$, and $e(k+1;l,l+1)$ a *northern A-type.* Suppose H has a northern A-type. We call the subgraph $A \cup e(k,k+1;l+1)$ of H a *northern A_0-type*, and we call the subgraph $A \cup e(k;l+1,l+2) \cup e(k+1;l+1,l+2)$ of H a *northern A_1-type.* We make analogous definitions for eastern, southern and western A-types. See Figure 3.12.

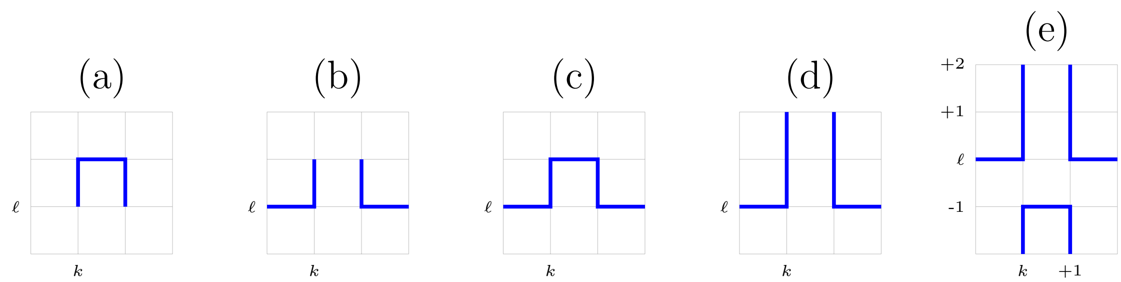

Fig. 3.12. (a) A northern leaf. (b) An A-type. (c) A_0-type. (d) A_1-type. (e) The edge $e(k,k+1;l)$ and the northern leaf $R(k,l-1)$ followed by a northern A_1-type.

Let $R(k,l-1)$ be a northern leaf. If H has a northern A-type on $e(k-1,k;l+1), e(k;l+1,l+2)$ and $e(k+1,k+2;l+1), e(k+1;l+1,l+2)$ then we say that A-type *follows* the northern leaf $R(k,l-1)$ *northward.* We call the boxes $R(k,l+1)$ and $R(k,l+2)$ the *middle-boxes* of the A_1-type. We call the box $R(k,l+1)$ the *switchable middle-box* of the A_1-type. Analogous definitions apply for other compass directions.

Let $e(k,k+1;l)$ be an edge in H, let A be a northern A_0-type in H on the edges $e(k-1,k;l+1)$, $e(k;l+1,l+2)$, $e(k,k+1,l+2)$, $e(k+1;l+1,l+2)$, and $e(k+1,k+2;l+1)$, and let $j \in \left\{1,2,\dots \left\lfloor \frac{n-1-l}{2} \right\rfloor\right\}$. We define a *northern j-stack of A_0s starting at A* to be a subgraph *stack*$(j;A_0)$ of H, where $\text{stack}(j;A_0) = \bigcup_{i=0}^{j-1}\Big(A_0 + (0,2i)\Big)$. If $j = \left\lfloor \frac{n-1-l}{2} \right\rfloor$, we call the j-stack a *full j-stack of A_0's.* We will often say that the j-stack of A_0's is *following* the edge $e(k, k+1; l)$ whenever this edge plays an important role. Eastern, southern, and western j-stacks are defined analogously. See Figure 3.14(a) on page 53.

We extend the notion of *collecting* to leaves. Let $R(k, l)$ be leaf of H. WLOG assume that $R(k, l)$ is northern. We say that the cascade $\mu_1, \ldots, \mu_r$ *collects* $R(k, l)$ if μ_r is the move $R(k, l+1) \mapsto R(k, l)$.

We denote the set of northern and southern small cookies by SmallCookies$\{N, S\}$ and the set of eastern and western small cookies by SmallCookies$\{E, W\}$. Let C be an easternmost or westernmost northern or southern small cookie in SmallCookies$\{N, S\}$. Then we say that C is an *outermost* small cookie in SmallCookies$\{N, S\}$. Outermost small cookies in SmallCookies$\{E, W\}$ are defined analogously.

Given a leaf L, we want to show that there is a cascade that collects L. For definiteness, assume that L is a northern leaf $R(k, l)$. Furthermore, assume that L is not incident on the northern side of the boundary, so $l+1 < n-1$. If $e(k, k+1; l+2) \in H$, then $L + (0, 1) \mapsto L$ is the cascade we seek, so we only need to consider the case where $e(k, k+1; l+1) \notin H$. Then we must have $e(k-1, k; l+2) \in H$, $e(k+1, k+2; l+2) \in H$, $e(k; l+2, l+3) \in H$ and $e(k+1; l+2, l+3) \in H$. Note that if $e(k-1, k; l+3) \in H$, then $L+(0, 2) \mapsto L+(-1, 2)$ followed by $L+(0, 1) \mapsto L$ is the cascade we seek, so we consider the case where $e(k-1, k; l+3) \notin H$ and, by symmetry, where $e(k+1, k+2; l+3) \notin H$. Now, we either have $e(k; l+3, l+4) \in H$ and $e(k+1; l+3, l+4) \in H$ or $e(k, k+1; l+3) \in H$. That is, L is followed northward by an A_0-type or by an A_1-type. See Figure 3.13. From this point onward, we will omit the compass direction when it does not introduce ambiguity. We coalesce this paragraph into the following lemma:

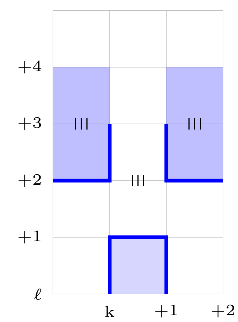

Fig. 3.13.

Lemma 3.9. Let $L = R(k, l)$ be a northern leaf with $l+1 < n-1$. If L is not followed by an A_0-type or by an A_1-type, then there is a cascade of length at most two that collects L.

Proposition 3.10. (1LC Algorithm). Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G. If H has exactly one large cookie and at least one small cookie, then there is a cascade of length at most $\frac{1}{2}\max(m, n) + \min(m, n) + 2$ moves that reduces the number of small cookies of H by one and such that it does not increase the number of large cookies.

The proof of Proposition 3.10 requires the following two lemmas.

Lemma 3.11. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let $C \in \text{SmallCookies}\{N, S\}$ be an easternmost small cookie. Assume that G has only one large cookie. Then there cannot be a full j-stack of A_0s starting at the A_0-type that contains C.

Lemma 3.12. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let $C \in \text{SmallCookies}\{N, S\}$ be an easternmost small cookie. Assume that there is a j-stack of A_0 starting at the A_0-type containing C. Let L be the leaf in the top (j^{th}) A_0 of the stack. Assume that L is followed by an A_1-type and that G has only one large cookie. Then there is a cascade of at most $\min(m, n) + 3$ moves that collects L.

Proof of Proposition 3.10. Since there is at least one small cookie, one of SmallCookiesN, S or SmallCookiesE, W must be nonempty. Without loss of generality we may assume that SmallCookies$\{N, S\}$ is nonempty. Let $C \in \text{SmallCookies}\{N, S\}$ be an easternmost small cookie. For definiteness, assume that C is a small northern cookie on the southern boundary. Let $Q(j)$ be the statement "There is a j-stack of A_0's starting at the A_0-type containing C". Note that C is contained in an A_0-type, so $Q(1)$ is true. By Lemma 3.11, there is a $j_0 \in \left\{2, 3, \ldots \left\lfloor \frac{n}{2} \right\rfloor \right\}$ such that for each $j \in \{2.., j_0 - 1\}$, $Q(j)$ is true for each $j < j_0$ but $Q(j_0)$ is not true.

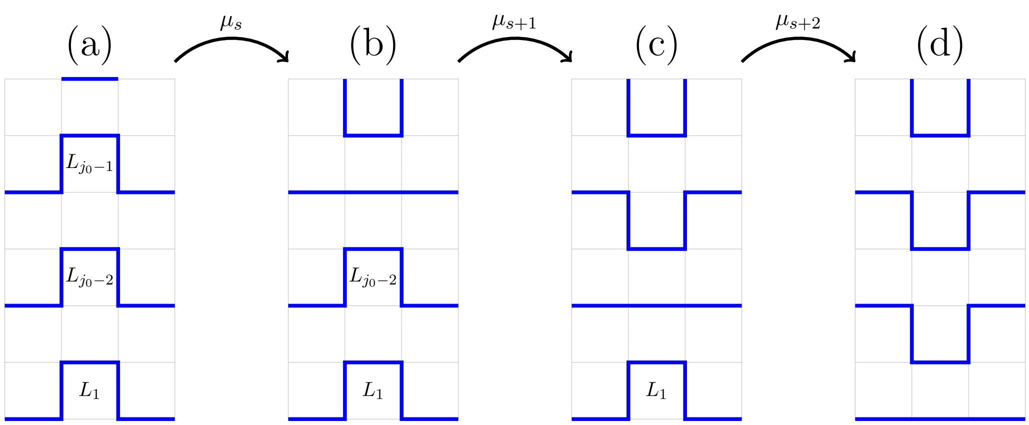

Fig. 3.14. (a) A 3-stack of A_0s starting at L_1. (b) The resulting graph after the application of μ_s. (c) The resulting graph after the application of μ_{s+1}. (d) The resulting graph after the application of μ_{s+2}.

For $j \in \{1.., j_0 - 1\}$ let L_j be the northern leaf of the j^{th} A_0-type in the stack. Note that $L_1 = C$. Lemma 3.9 implies that L_{j_0-1} is either followed by an A_1-type or there is a cascade that collects L_{j_0-1}. If L_{j_0-1} is followed by an A_1-type, then by Lemma 3.12, we can find a cascade that collects it, so we only need to check the case in which there is a cascade $\mu_1, \ldots, \mu_s$ that collects L_{j_0-1}. Note that μ_s must be the move $L_{j_0-1} + (1, 0) \mapsto L_{j_0-1}$. Then

$\mu_1, \ldots, \mu_s, L_{j_0-2}+(0,1) \mapsto L_{j_0-2}, \ldots, L_1+(0,1) \mapsto L_1$ is a cascade that collects C. Note that $j \leq \frac{n}{2}$, and that by Lemma 3.12, there are at most $\min(m,n)+3$ moves required to collect L. After that, we need at most another $j-1$ flips to collect C, so C can be collected after at most, $\frac{1}{2}\max(m,n)+\min(m,n)+2$ moves to collect C. See Figure 3.14 for an illustration with $j_0 - 1 = 3$. □

We prove Lemma 3.11 next. The proof of Lemma 3.12 takes up the remainder of the section.

Proof of Lemma 3.11. Assume for a contradiction that C is in a full stack of A_0's starting at the A_0 that contains C. For definiteness, assume that $C = R(k,0)$ is a small northern cookie on the southern boundary. First we check that $m-1 > k+2$. If $m-1 = k+2$, then we must have $e(k+2;0,1) \in H$ and $e(k+2;1,2) \in H$. But then H misses $v(k+2,3)$ (in green in Figure 3.15 (a)). There fore, we must have that $m-1 > k+2$.

The number j of A_0's in the full stack is even or an odd so there are two cases to check. Note that for each odd $i \in \{1,2,\ldots,j\}$, the i^{th} A_0 belongs to $\text{ext}(H)$ and for each even $i \in \{1,2,\ldots,j\}$, the i^{th} A_0 belongs to $\text{int}(H)$.

CASE 1: j is even. Note that the top leaf of the stack is in $\text{int}(H)$. Now, $n-1$ is either even or odd.

CASE 1.1: $n-1$ is even. We have that $R(k,n-3) \in \text{int}(H)$. But then we must have $R(k,n-2) \in \text{ext}(H)$ and $R(k+1,n-2) \in \text{ext}(H)$, contradicting Lemma 1.14. End of Case 1.1. See Figure 3.15 (a).

CASE 1.2: $n-1$ is odd. Then we must have that $R(k+1,n-2)$ is a small southern cookie. But this contradicts our assumption that C is the easternmost small cookie in SmallCookies$\{N,S\}$. End of Case 1.2. End of Case 1. See Figure 3.15 (b).

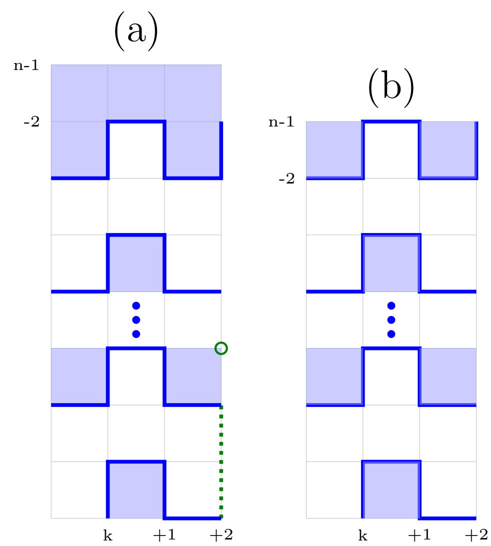

Fig 3.15. (a) Case 1.1. (b) Case 1.2.

CASE 2: j is odd. Note that the top leaf of the stack is in $\text{ext}(H)$. Again, $n-1$ is either even or odd.

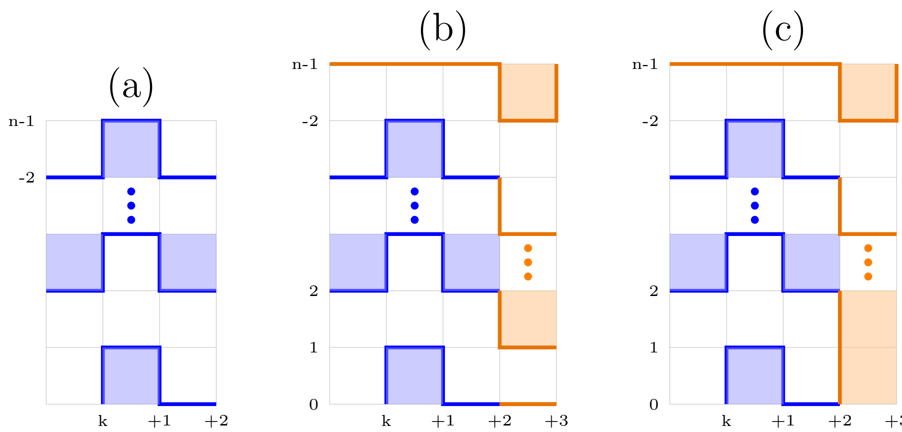

Fig 3.16. (a) Case 2.1. (b) Case 2.2(a). (c) Case 2.2(b).

CASE 2.1: $n-1$ is odd. We have that $R(k, n-2) \in \text{ext}(H)$. But then, the fact that $e(k, k+1; n-1) \in H$ implies that $R(k, n-2)$ is not a cookie neck, contradicting Lemma 1.14. End of Case 2.1. See Figure 3.16 (a).

CASE 2.2: $n-1$ is even. We have that $e(k+1, k+2; 0) \in H$. Then, either $e(k+2, k+3; 0) \in H$, or $e(k+2; 0, 1) \in H$.

CASE 2.2(a): $e(k+2, k+3; 0) \in H$. Then we must have $e(k+2; 1, 2) \in H$. Note that for $i \in \{1, 3, \ldots, n-2\}$, $e(k+2; i, i+1) \in H$ implies $e(k+2; i+2, i+3) \in H$. Then, for $i \in \{1, 3, \ldots, n-2\}$, we have that $e(k+2; i, i+1) \in H$. Note that we must also have $e(k+2, k+3; n-1) \in H$. Then $R(k+2, n-2)$ must be a southern small cookie, contradicting the eastmost assumption. See Figure 3.16 (b) End of Case 2.2(a).

CASE 2.2(b): $e(k+2; 0, 1) \in H$. Note that if $e(k+2, k+3; 1) \in H$, then $R(k+2, 0)$ must be a small cookie, contradicting the eastmost assumption. Then $e(k+2, k+3; 1) \notin H$. But then we have $e(k+2; 1, 2) \in H$, and we are back to Case 2.2(a). See Figure 3.16 (c). End of Case 2.2(b). End of Case 2.2. End of Case 2. □

It remains to prove Lemma 3.12. Its proof will require Lemmas 3.13-3.16.

Lemma 3.13. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian cycle of G. Let C be a small cookie of G. Assume that G has only one large cookie, and that there is a j-stack of A_0 starting at the A_0-type containing C. Let L be the leaf in the top (j^{th}) A_0 of the stack, and assume that L is followed by an A_1-type. Let X and Y be the boxes adjacent to the switchable middle-box of the A_1-type that are not its H-neighbours. If $P(X, Y)$ has no switchable boxes, then either:

(i) there is a cascade of length at most $\min(m, n)$, which avoids the stack of A_0's, and after which $P(X, Y)$ gains a switchable box, or

(ii) there is a cascade of length at most $\min(m, n) + 1$, that collects L and avoids the stack of A_0's.

We postpone the proof of Lemma 3.13 to Chapter 4. For now, we assume Lemma 3.13 and use it to prove Lemma 3.12. Recall the setup: we have a j-stack of A_0-types followed by an A_1-type with middle-box X'. A consequence of the assumption that there is only one large cookie is that X' must determine an H-path. That is, if X and Y are the boxes adjacent to X' that are not its H-neighbors, then X and Y lie in the same H-component, so the H-path $P(X, Y)$ is well-defined. Lemma 3.13 guarantees that either $P(X, Y)$ already has a switchable box or there exists a cascade that produces one.

Lemma 3.14 shows that X' must lie in G_2. Lemma 3.15 handles the case $X' \in G_2 \backslash G_3$. For the remaining case, $X' \in G_3$, we use the existence of a switchable box in $P(X, Y)$ (guaranteed by Lemma 3.13) to prove Lemma 3.16. This lemma plays a role analogous to Lemma 3.7 in Section 3.1: it shows that switching X' yields two cycles with adjacent vertices on R_2. We then can apply the part of Lemma 3.5 covering the case $Z \in G_3$ to find either a switchable box or a cascade of length at most two that delivers the required move.

Lemma 3.14. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let $C \in \text{SmallCookies}\{N, S\}$ be an easternmost small cookie. Assume that G has only one large cookie, and that there is a j-stack of A_0 starting at the A_0-type containing C. Let L be the leaf contained in the top (j^{th}) A_0 of the stack. Assume that L is followed by an A_1-type with looping H-path $P(X, Y)$. Let X' be the box of G that shares edges with X and Y. Then $X' \in G_2$.

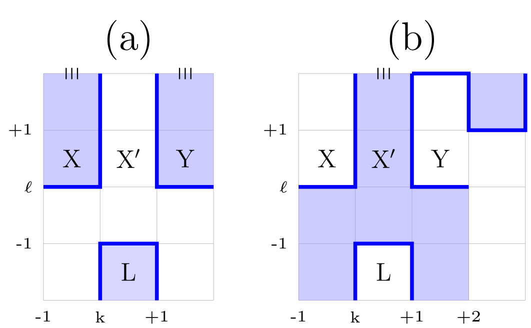

Fig 3.17. (a) Case 1. (b) Case 2.

Proof. For definiteness, assume that L is the northern leaf $R(k, l-2)$, and that $X = R(k-1, l)$. Then $X' = R(k, l)$ and $Y = R(k+1, l)$. Note that $l - 2 \geq 0$ and $l + 2 \leq n - 1$. Either $P(X, Y)$ is contained in $\text{ext}(H)$, or $P(X, Y)$ is contained in $\text{int}(H)$, so there are two cases to check.

CASE 1: $P(X, Y) \subset \text{ext}(H)$. By Lemma 1.14, we must have that $m - 1 > k + 2$ and $k - 1 > 0$. To see that $n - 1 > l + 2$, assume for a contradiction that $n - 1 = l + 2$. By Lemma 1.14, $X + (1, 0)$ and $Y + (1, 0)$ are cookie necks. But this contradicts the assumption that there is only one large cookie in G. See Figure 3.17 (a). End of Case 1.

CASE 2: $P(X, Y) \subset \text{int}(H)$. By Lemma 1.14, we must have that $m-1 > k+2$ and $k-1 > 0$. To see that $n-1 > l+2$, assume for a contradiction that $n-1 = l+2$. Lemma 1.14 implies that $X' + (0, 1)$ is the neck of the large cookie of G. But now $X' + (2, 1)$ must be a small cookie of G, contradicting the easternmost assumption. See Figure 3.17 (b).

Lemma 3.15. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let $C \in \text{SmallCookies}\{N, S\}$ be an easternmost small cookie. Assume that G has only one large cookie, and that there is a j-stack of A_0 starting at the A_0-type containing C. Let L be the leaf in the top (j^{th}) A_0 of the stack. Assume that $L \in \text{int}(H)$ and that L is followed by an A_1-type with looping H-path $P(X, Y)$. Let X' be the box of G that shares edges with X and Y. If X' is not in G_3 then, either $X' \mapsto W$ is a cascade, or there is a cascade $\mu, X' \mapsto W$, of length two, with $X' \mapsto W$ nontrivial in either case.

Proof. Suppose that X' is not in G_3. By Lemma 3.14, $X' \in G_2 \setminus G_3$. For definiteness, assume that L is a northern leaf, and let $X' = R(k, l)$. The assumption that $L \in \text{int}(H)$ implies that $l - 2 > 0$. See Figure 3.18.

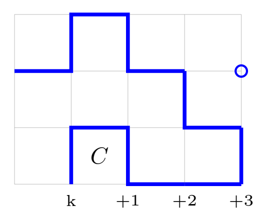

Fig 3.18.
$m - 1 = k + 3$.

Now we check that $m-1 > k+3$. By Lemma 3.14, $m-1 > k+2$. For a contradiction, assume that $m - 1 = k + 3$. Note that we must have $e(k+1, k+2; 0) \in H$, $e(k+2, k+3; 0) \in H$, and $e(k+3; 0, 1) \in H$. This implies that we must have $e(k + 2; 1, 2) \in H$ and $e(k + 2, k + 3; 1) \in H$. But now H misses $v(k+3; 2)$. It follows that we must have $m-1 > k+3$. See Figure 3.18. By symmetry, $0 < k - 2$. It follows that $l + 3 = n - 1$.

The same argument used in Case 2.2 of Lemma 3.11 (see Figure 3.16 (b) and (c)) shows that we have $e(k+2; l+1, l+2) \in H$ and $e(k+2, k+3; l+1) \in H$. Now either $e(k, k+1; l+2) \in H$ or $e(k, k+1; l+2) \notin H$. See Figure 3.19.

CASE 1: $e(k, k + 1; l + 2) \in H$. By Lemma 1.16, $X' \mapsto X' + (1, 1)$ is a valid move, and by Observation 3.4, $X' \mapsto X' + (1, 1)$ creates no new cookies. End of Case 1.

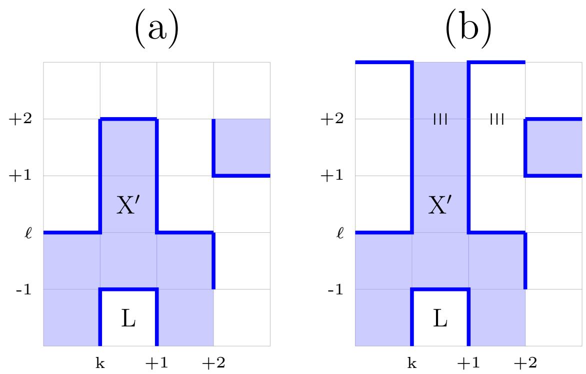

Fig 3.19. (a) Case 1. (b) Case 2.

CASE 2: $e(k, k+1; l+2) \notin H$. Lemma 1.14 implies that $e(k+1, k+2; l+2)$ cannot be in H either. It follows that $X' + (0, 2)$ must be the neck of the large cookie. The assumption that there is only one large cookie implies that $e(k+2; l+2, l+3) \notin H$. Then we must have $e(k+2, k+3; l+2) \in H$. Then, by Lemma 1.16, $X' + (1, 1) \mapsto X' + (2, 1)$, $X' \mapsto X' + (1, 0)$ is the cascade we seek. □

Lemma 3.16. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G. Let $X' \in G_3 \cap \text{ext}(H)$ be a switchable box, and let $P(X, Y)$ be the looping H-path of X'. Assume that $P(X, Y)$ has a switchable box in $G_0 \setminus G_2$. Then switching X' splits H into two cycles H_1 and H_2 such that there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2.

Proof. Let Z be a switchable box of $P(X, Y)$ in $G_0 \setminus G_2$. Orient H. Let (v_x, v_{x+1}) and (v_{y-1}, v_y) be the edges of X' in H. Define $\overrightarrow{K}_1$ and $\overrightarrow{K}_2$ to be the subtrails $\overrightarrow{K}((v_x, v_{x+1}), (v_{y-1}, v_y))$ and $\overrightarrow{K}((v_y, v_{y+1}), (v_{x-1}, v_x))$ of $\overrightarrow{K}_H$, respectively. By Lemma 1.16 (i), switching X' gives two cycles H_1 and H_2, with $V(H_1) = V(\overrightarrow{K}_1 \setminus \{v_x, v_y\})$ and $V(H_2) = V(\overrightarrow{K}_2)$. By Proposition 3.1, Z has a vertex in H_1 and another in H_2, and the same holds for X'. Since $X' \in G_3$ and $Z \in G_0 \setminus G_2$, by JCT, $H_1 \cap R_2 \neq \emptyset$. Similarly, $H_2 \cap R_2 \neq \emptyset$. Now the argument in the last paragraph of Lemma 3.7 shows that there must be $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2. □

Proof of Lemma 3.12. For definiteness, assume that L is the northern leaf $R(k, l)$. Let $P(X, Y)$ be the looping H-path following L, with $X = R(k-1, l+2)$ and $Y = R(k+1, l+2)$. By Lemma 3.13, either there is a cascade that collects L, or a cascade after which $P(X, Y)$ gains a switchable box, with each cascade having length at most $\min(m, n) + 1$, and both avoiding the j-stack of A_0's starting at C. If the former, we are done, so may assume that $P(X, Y)$ has a switchable box Z. Let J be the large cookie of G and let N_J be the neck of J. Note that N_J cannot be a box of $P(X, Y)$. If it were, since the neck has exactly one H-neighbour in G, $N_J = X$ or $N_J = Y$, but neither box is switchable . Now, $P(X, Y)$ is either contained in $\text{ext}(H)$ or $\text{int}(H)$.

CASE 1: $P(X, Y) \subseteq \text{ext}(H)$. Then $X' \subset \text{int}(H)$. By Lemma 3.14, $X' \in G_2$. By Proposition

3.3, $Z \mapsto X'$ is a valid move. By Observation 3.4, $Z \mapsto X'$ does not create additional cookies. Then $Z \mapsto X'$, $L + (0,1) \mapsto L$ is a cascade that collects L. End of Case 1.

CASE 2: $P(X,Y) \subseteq \text{int}(H)$. Then $X' \subset \text{ext}(H)$. If $Z \subset G_2$, by Proposition 3.3, $Z \mapsto X'$ is a valid move, and by Observation 3.4, $Z \mapsto X'$ does not create additional cookies. Then $Z \mapsto X'$, $L \mapsto L + (0,1)$ is a cascade that collects L.

Suppose then that $Z \subset G_0 \setminus G_2$. By Lemma 3.15, we only need to check the case where $X' \in G_3$. Note that switching X' splits H into two cycles H_1 and H_2. By Lemma 3.16 there is $v_1 \in H_1 \cap R_2$ and $v_2 \in H_2 \cap R_2$ with v_1 adjacent to v_2. By Lemma 3.5 there is a cascade $\mu, X' \mapsto W$, or a cascade $X' \mapsto W$ with $X' \neq W$. Note that here, X' plays the role that Z played in Lemma 3.5. Then $\mu, X' \mapsto W, L + (0,1) \mapsto L$ or $X' \mapsto W, L + (0,1) \mapsto L$ is a cascade that collects L. End of Case 2.

We have just shown that if $P(X,Y)$ has a switchable box, then the cascade required to collect L has length at most three. By Lemma 3.13, the cascade after which $P(X,Y)$ gains a switchable box has length at most $\min(m,n)$. Thus, at most $\min(m,n) + 3$ moves are required to collect L. □

3.3 Summary

In this chapter we proved the MLC and 1LC algorithms. The proof of the MLC algorithm is fully contained here, while the proof of the 1LC algorithm depends on Lemma 3.13, whose proof is given in Chapter 4.

Proposition 3.3 characterizes when double-switch moves are valid and serves as the primary tool for both algorithms.

The MLC algorithm handles the case where H has multiple large cookies. To collect a large cookie J with switchable neck N_J, we look for a switchable box Z in the looping H-path of N_J. Proposition 3.8 shows that either such a Z already exists, or a there is single preparatory move that produces one.

The 1LC algorithm handles the case where H has exactly one large cookie and at least one small cookie. It collects outermost small cookies. Suppose that C is an outermost small cookie. Either C can be collected immediately by a single move, or C is followed by a j-stack of A_0-types and an A_1-type with switchable middle-box X'. If the latter, let $P(X,Y)$ be the H-path determined by X'. If $P(X,Y)$ contains a switchable box Z, then C can be collected by either switching X' directly (if $Z \in G_2$) or by using Lemma 3.5 to find a cascade of length at most two that enables switching X' (if $Z \in G_0 \setminus G_2$). In both cases, a cascade of flips then collects C. The existence of such a switchable box Z is guaranteed by Lemma 3.13, whose proof takes up much of Chapter 4.

4 Fat Paths

In this chapter we prove Lemma 3.13 and some related results that will be needed for the reconfiguration of Hamiltonian paths in Chapter 5. Recall the setup of Lemma 3.13: we have a Hamiltonian cycle H in an $m \times n$ grid graph G with a j-stack of A_0-types followed by an A_1-type with middle-box W. We must show that if the H-path $P(X, Y)$ determined by W contains no switchable boxes, then there exists a cascade that produces one. The absence of switchable boxes in $P(X, Y)$ constrains the possible edge configurations of H in the vicinity of $P(X, Y)$, making it highly structured. We will be able to use this structure to find the cascade we seek.

We begin with the necessary definitions. Among these, *turn* and *looping fat path* are used most extensively. We then sketch the proof of Lemma 3.13 and give an outline of the chapter.

Definitions. Let G be a simply connected polyomino and let H be a Hamiltonian path or cycle of G. Let $e(k, k+1; l) \in H$. If H has a northern A-type on $e(k-1, k; l+1), e(k; l+1, l+2)$ and $e(k+1, k+2; l+1), e(k+1; l+1, l+2)$ then we say that A-type *follows* the edge $e(k, k+1; l)$ *northward*. Similarly, if H has a southern A-type on $e(k-1, k; l-1), e(k; l-1, l-2)$ and $e(k+1, k+2; l-1), e(k+1; l-1, l-2)$ then we say that A-type *follows* the edge $e(k, k+1; l+1)$ *southward*. We make analogous definitions for the other directions. See Figure 4.1.

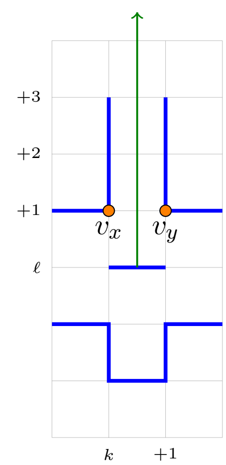

Fig. 4.1. Edge followed by a northern A_1-type and a southern A_0-type.

Let $e(k, k+1; l) \in H$ be followed by a northern A_1-type and let $PBR_{\text{North}}(e(k, k+1; l))$ denote the northward perpendicular bisector ray $PBR_{\text{North}}(e(k, k+1; l))$ of the edge $e(k, k+1; l)$ originating from the point $v(k + \frac{1}{2}, l)$ (green in Figure 4.1).

We call the vertices $v(k, l+1)$ and $v(k+1, l+1)$ *the corner vertices of* A_1. More specifically, we call the vertex $v(k, l+1)$ of A_1, located on the left of $PBR_{\text{North}}(e(k, k+1; l))$ the *left corner of* A_1 and denote it by $v_{\text{left}}(A_1)$. Similarly, we call the vertex of $v(k+1, l+1)$ of A_1 located on the right of $PBR_{\text{North}}(e(k, k+1; l))$ the *right corner of* A_1, and denote it by $v_{\text{right}}(A_1)$ (orange in Figure 4.1).

We call the edges $e(k; l+1, l+2)$ and $e(k; l+2, l+3)$ the *left colinear edges of the* A_1*-type.* Similarly, the edges $e(k+1; l+1, l+2)$ and $e(k+1; l+2, l+3)$ are called the *right colinear* *edges of the A_1-type*. We refer to these in Proposition 4.14. We make analogous definitions for the other directions. See Figure 4.1. Note that if the A_1-type is southern, then the left corner of the A_1-type is on the right of the image in our page; and similarly if the A_1-type is western (eastern), then the left corner of the A_1-type is south (north) of the perpendicular bisector ray.

Suppose that the edge $e(k,k+1,l)$ is followed by an A_1-type southward. Define an *A_1-partitioning of H* to be the partition of H into the sub-paths $P_0 = P(v_x, v_y)$, $P_1 = P(v_1, v_x)$ and $P_2 = P(v_y, v_r)$, where $v_x = v_{\text{left}}(A_1)$ and $v_y = v_{\text{right}}(A_1)$ or $v_x = v_{\text{right}}(A_1)$ and $v_y = v_{\text{left}}(A_1)$.

Let $W = R(k, l-2)$, $X = R(k+1, l-2)$, and $Y = R(k-1, l-2)$, and assume that the looping H-path $P(X,Y)$ of W is contained in an H-component of G. Then we say that $P(X,Y)$ is a *southern looping H-path* and that $P(X,Y)$ *follows* $e(k,k+1,l)$ southward. Similarly, if $R(k,l)$ is a southern leaf, we say that $P(X,Y)$ *follows* $R(k,l)$ southward. We make analogous definitions for western, northern and eastern looping H-paths. From here on, assume that every looping H-path follows an edge of H and that its end-boxes are incident on an A_1-type, as described above.

Let $J = \{X_1, X_2, \ldots, X_r\}$ be a collection of boxes of a polyomino G and let H be a Hamiltonian path or cycle of G. We will use the notation $G\langle J\rangle$ to denote the subgraph of G with vertex set $V(G\langle J\rangle) = V(J)$ and edge set $E(G\langle J\rangle) = E(J) \cap E(H)$. The boxes of $G\langle J\rangle$ are the boxes of J. We call $G\langle J\rangle$ the *subgraph of G induced by J*.

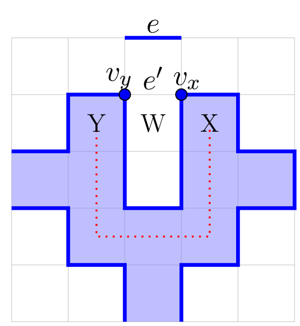

Fig. 4.2. A southern looping fat path $G\langle N[P(X,Y]\rangle$. $N[P(X,Y]$ shaded in light blue. $P(X,Y)$ traced in red.

Let e be an edge of H followed by an A_1-type with switchable middle-box W. Let X and Y be the boxes adjacent to W that are not its H-neighbours, and assume that X and Y belong to the same H-component. We will call the set of boxes consisting of the boxes of the H-path $P(X,Y)$ together with all the H-neighbours of $P(X,Y)$ *the H-neighbourhood of $P(X,Y)$* and denote it by $N[P(X,Y)]$. Consider the subgraph $F = G\langle N[P(X,Y)]\rangle$ of G induced by $N[P(X,Y)]$. We call F a *looping fat path* if:

1. no end-vertex of H is incident on $P(X,Y)$,
2. W is anti-parallel, and
3. $P(X,Y)$ has no switchable boxes.

If furthermore, F is such that:

4. There is no switch or flip or transpose (recall definitions from Introduction) move after

which W is switched, or $P(X,Y)$ gains a switchable box, or $W \mapsto X$ or $W \mapsto Y$ is valid, then we say that F is a *sturdy looping fat path.* We call these conditions the *fat path conditions* (FPC), and say FPC-j when referring to condition j, for $j \in \{1, \ldots, 4\}$. The role of FPC-1 in handling end-vertices is addressed in Sections 4.3 (Remark after Observation 4.20) and later sections. Roughly, our goal will be to find a cascade after which the edge $e' = \{v_x, v_y\}$ (Figure 4.2.) is in the resulting Hamiltonian path. If any of the conditions FPC-j fail, we can find such a cascade immediately. In that sense, a sturdy looping fat path represents an unyielding configuration – the kind of subgraph where there are no obvious good moves, and we have to look further into structure to obtain the required cascade.

Recall the definition of an outermost small cookie from the beginning of Section 3.2. Suppose G has only one large cookie, and that there is a j-stack of A_0's starting at the A_0-type containing an outermost southern small cookie C. Let L be the leaf in the top (j^{th}) A_0 of the stack, and assume that L is followed by an A_1-type with looping H-path $P(X,Y)$. Let $F = G\langle N[P(X,Y)]\rangle$, and assume that F is a fat path. We say that F is a *southern* looping fat path *anchored* at the outermost southern small cookie C. Analogous definitions apply for northern, eastern, and western looping fat paths.

From here on, whenever H is a cycle and we are discussing a looping fat path F, we will assume that G has exactly one large cookie and that F is anchored at an outermost small cookie.

Assume that H is a Hamiltonian path of G, and that $F = G\langle N[P(X,Y)]\rangle$ is a looping fat path such that $P(X,Y)$ follows an edge of H southward. Then we say that F is a *southern* looping fat path. Analogous definitions apply for northern, eastern, and western looping fat paths.

Let H be a Hamiltonian path of an $m \times n$ grid graph G. In Chapter 5 we will define an appropriate canonical form and show that H can be reconfigured into such a form. If H has adjacent end-vertices, the reconfiguration process shortens and simplifies. The more involved cases, requiring some of the tools we build in this chapter, arise when H does not have adjacent end-vertices. Therefore, from here on, we will often assume that the Hamiltonian path we are considering does not have adjacent end-vertices. We call this assumption the non-adjacency assumption (NAA).

We define below a subgraph of G consisting of the union of translations of two adjacent and perpendicular edges of G. Let $r \in \mathbb{N}$ and let *the stairs from (k,l) to (k+r,l-r) east* be denoted by $S_{\rightarrow}(k,l;k+r,l-r)$ and be defined as:

$$S_{\rightarrow}(k,l;k+r,l-r) = \bigcup_{j=0}^{r-1} \Big(e(k,k+1;l)+(j,-j)\Big) \cup \Big(e(k+1;l-1,l)+(j,-j)\Big).$$

We define $d(S) = r$ to be the *length* of $S_{\rightarrow}(k,l;k+r,l-r)$. We say that $S_{\rightarrow}(k,l;k+r,l-r)$ starts at $v(k,l)$ and ends at $v(k+r,l-r)$. The subscripted arrow indicates the direction from $v(k,l)$ of the first edge of the subgraph. By choosing an "up", "down", "left" or "right" arrow for direction and a sign for the third and fourth arguments of $S_{\square}(k,l;k\pm r,l\pm r)$ we may describe any of the eight possible steps subgraphs starting at the vertex $v(k,l)$. See Figure 4.3 (a).

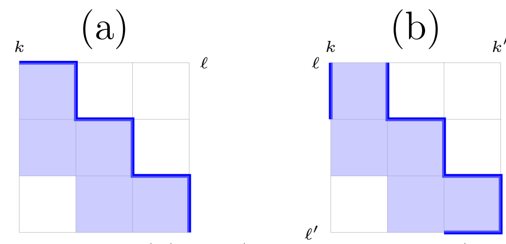

Fig. 4.3. (a) $S_{\rightarrow}(k,l;\ k+3,l-3)$. (b) A half-open northeastern turn.

Let H be a Hamiltonian path or cycle of a polyomino G. Let T be the subgraph of H on the edges $S_{\downarrow}(k+1,l;k',l'+1)$, $e(k;l-1,l)$ and $e(k'-1,k';l')$, where $k' = k+d(T)$, $l' = l-d(T)$, where $d(T) = d(S)+1$ is *the length of* T and $d(T) \geq 2$. We call T a *northeastern turn*. If both $e(k,k+1;l)$ and $e(k';l',l'+1)$ belong to $G \setminus H$, call T an *open northeastern turn*. If exactly one of $e(k,k+1;l)$ and $e(k';l',l'+1)$ is in H, then T is a *half-open northeastern turn*. See Figure 4.3 (b). If both $e(k,k+1;l) \in H$ and $e(k';l',l'+1) \in H$, then T is a *closed northeastern turn*. For any northeastern turn T, we say that $R(k,l-1)$ is the *northern leaf* of T and $R(k'-1,l')$ is the *eastern leaf* of T. If $e(k,k+1;l) \notin H$ we call $R(k,l-1)$ an *open northern leaf of T* and if $e(k,k+1;l) \in H$ we call $R(k,l-1)$ a *closed northern leaf of T*. We note that the two leaves of a turn will determine its "leaf prefix": If a turn has a northern leaf and an eastern leaf then the turn is a northeastern turn.

We will say that a looping fat path F *has a turn* (open, half-open or closed) to mean that there exists some turn T of H such that $E(F) \supset E(T)$.

Sketch of proof of Lemma 3.13. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G. Assume that $P(X,Y)$ is a looping H-path with no switchable boxes, following a leaf L. It follows that $P(X,Y)$ is contained in a looping fat path F. In Section 4.2, we show that every looping fat path has a turn. In Sections 4.3-4.6 we show that given a turn, we can find a cascade we call a *weakening* (precise definition in Section 4.3) that collects one of its leaves. It then follows that after this cascade, either $P(X,Y)$ gains a switchable box, or we can extend the cascade by a single move to collect L.

The rest of the chapter is organized as follows. Section 4.1 proves structural properties of fat paths, which the later sections build on. Section 4.2 shows that every fat path contains a turn (Proposition 4.18). Section 4.3 defines weakenings and proves some basic properties. Sections 4.4-4.6 prove that turns have weakenings: Section 4.4 handles the case where H is a Hamiltonian cycle (Proposition 4.24), while Sections 4.5 and 4.6 handle the case where H is a Hamiltonian path (Propositions 4.27 and 4.28). The proof of Lemma 3.13 appears at the end of Section 4.6. Before we proceed to Section 4.1 we prove results 4.1-4.5, which generalize Propositions 3.1 and 3.3 to Hamiltonian paths in polyominoes.

Lemma 4.1. Let G be a polyomino and let $H = v_1, \dots, v_r$ be a Hamiltonian path of G. Let $P(X,Y)$ be the looping H-path of a switchable box W. Assume that $P(X,Y)$ is contained in an H-component of G and that the edges of W are anti-parallel. Then the A_1-partitioning of H is such that every box of $P(X,Y)$ is incident on a vertex of P_0.

Proof. For definiteness, assume that $P(X,Y)$ is a southern looping H path following the edge $e(k,k+1,l)$ southward, that $v_x = v_{\text{left}}(A_1)$, and $v_y = v_{\text{right}}(A_1)$. Then $v_x = v(k+1,l-1)$ and $v_y = v(k,l-1)$. Now, $v_{x+1} = v(k+1,l-2)$ or $v_{x+1} = v(k+2,l-1)$.

Assume that $v_{x+1} = v(k+1,l-2)$. Then $v_{y-1} = v(k,l-2)$. Then $\Phi((v_x,v_{x+1}),\text{left}) = X$ and $\Phi((v_{y-1},v_y),\text{left}) = Y$. By Observation 1.2.4 (d), we can identify $\overrightarrow{K}((v_x,v_{x+1}),(v_{y-1},v_y))$ with P_0. Since the end-boxes X and Y of $P(X,Y)$ are contained in G, it follows from Proposition 1.2.1 that $P(X,Y)$ is not contained in any cycle of boxes in G_{-1}. Thus, $P(X,Y)$ is unique. It follows that any H-walk of boxes between X and Y contains $P(X,Y)$. In particular, $\Phi(P_0,\text{left})$ contains $P(X,Y)$. Then, by definition of FTW, any box of $P(X,Y)$ is incident on a vertex of P_0.

The case where $v_{x+1} = v(k+2,l-1)$ uses a similar argument, so we omit the proof. □

Corollary 4.2. Let G, H, $P(X,Y)$, P_0, P_1, and P_2 be as in Lemma 4.1. Then:

(a) For every box Z of $P(X,Y)$ there is an edge $e_z \in P_0$, such that e_z adds Z to $\Phi(P_0,\text{left})$.

(b) The middle-boxes of the A_1-type following $e(k,k+1,l)$ are not in $P(X,Y)$.

Proof. Part (a) follows from the definition of FTW. It remains to prove part (b). Either $v_{x+1} = v(k+1,l-2)$ or $v_{x+1} = v(k+2,l-1)$. Assume that $v_{x+1} = v(k+1,l-2)$. For a contradiction, assume that $X + (-1,0)$ is in $P(X,Y)$. Then $X + (-1,0)$ must be added to $\Phi(P_0,\text{left})$ by (v_{x+1},v_x) or by (v_y,v_{y-1}), but neither edge belongs to P_0. Similarly if $X + (-1,-1)$ were in $P(X,Y)$, it would have to be added to $\Phi(P_0,\text{left})$ by (v_{x+2},v_{x+1}) or by (v_{y-1},v_{y-2}), or, if $\{v_{y-2},v_{x+2}\} \in H$, by (v_{y-2},v_{x+2}) but none of those edge belong to P_0.

The case where $v_{x+1} = v(k+2, l-1)$ uses a similar argument, so we omit the proof. $\square$.

Lemma 4.3. Let G be a polyomino and let $H = v_1, \ldots, v_r$ be a Hamiltonian path of G. Let $P(X,Y)$ be a looping H-path. Let P_0, P_1 and P_2 be the partitioning of H given in Lemma 4.1. Then, no edge of P_0 has both of its incident boxes belonging to $P(X,Y)$.

Proof. For contradiction, assume that there is an edge e of P_0 with both its incident boxes Z and Z' belonging to $P(X,Y)$. Let Q be the cycle consisting of the edges of P_0 and the edge $\{v_x, v_y\}$ of $G \setminus H$. Consider the H-subpath $P(Z_1, Z_s)$ of $P(X,Y)$, where $Z_1 = Z$ and $Z_s = Z'$. Let c_j be the center of the box Z_j, for each $j \in \{1, \ldots, s\}$. Note that each segment $[c_j, c_{j+1}]$ intersects one gluing edge of $P(Z_1, Z_s)$ and no other edge or vertex of G, and that $[c_1, c_2], \ldots, [c_{s-1}, c_s]$ is a path $P(c_1, c_s)$. By Corollary 1.1.5, c_1 and c_s are on different sides of Q. By JCT, $P(c_1, c_s)$ intersects Q. Since the edges of $P(c_1, c_s)$ only intersects gluing edges of $P(Z_1, Z_s)$, some edge $\{c_i, c_{i+1}\}$ of $P(c_1, c_s)$, must intersect Q at $\{v_x, v_y\}$. Now, Z_i and Z_{i+1} belong to $P(X,Y)$, and, since they are incident on $\{v_x, v_y\}$, at least one of them must be a middle-box of F, which contradicts Corollary 4.2 (b). $\square$

Proposition 4.4. Let G be a polyomino and let $H = v_1, \ldots, v_r$ be a Hamiltonian path of G. Let $P(X,Y)$ be a looping H-path in G and assume that the switchable box between X and Y is anti-parallel. Then the partitioning of H into the subpaths $P_0 = P(v_x, v_y)$, $P_1 = P(v_1, v_x)$ and $P_2 = P(v_y, v_r)$, where v_x and v_y are the corners of the A_1-type, is such that every box of $P(X,Y)$ is incident on a vertex of P_0 and a vertex of P_1 or P_2.

Proof. Let v_x be incident on X and v_y be incident on Y. By Lemma 4.1, every box of $P(X,Y)$ has a vertex incident on P_0. To finish the proof, we proceed by induction. Note that, since $v_x \in V(P_0) \cap V(P_1)$ and $v_y \in V(P_0) \cap V(P_2)$, X and Y satisfy the conclusion. Let $P(X,Y)$ be the path $X = X_1, \ldots, X_t = Y$. Assume that the box X_j is incident on a vertex in P_0 and a vertex of P_1 or P_2. Let $\{v_p, v_q\}$ be the gluing edge between X_j and X_{j+1}. By Lemma 4.1, at least one vertex of X_{j+1} belongs to $V(P_0)$. For a contradiction, assume that no other vertex of X_{j+1} belongs to $V(P_0)$.

Without loss of generality we may assume that $p < q$. Note that if $x = p$ or $y = q$ then X_{j+1} satisfies the conclusion of the proposition, so we may assume that $x < p$ and $q < y$. Then we have that $1 < x < p < q < y < r$. Let U be the region of G bounded by the polygon Q consisting of the subpath $P(v_p, v_q)$ of P_0 and the edge $\{v_p, v_q\}$ of $G \setminus H$. Note that X_j and X_{j+1} are on distinct sides of Q. Assume for now that this implies that X and Y are on distinct sides of Q as well. Then v_x and v_y must be on distinct sides of Q (v_x, v_y cannot be in Q since x is too small, and y is too large. Therefore, if X and Y are on distinct sides of Q, so must v_x and v_y). But this contradicts the fact that v_x and v_y are adjacent. It remains to show that X and Y are on distinct sides of Q. Either $X_{j+1} \in U$ or $X_{j+1} \notin U$.

CASE 1: $X_{j+1} \in U$. We check that $Y \in U$. For a contradiction, assume $Y \notin U$. Consider the H-path $P(X_{j+1}, Y)$. For each $i \in \{j+1, j+2, \ldots, t\}$, let c_i be the center of the boxes X_i. Observe that for each $i \in \{j+1, j+2, \ldots, t\}$, the segment $[c_i, c_{i+1}]$ intersects the gluing edge of X_i and X_{i+1} and $[c_i, c_{i+1}]$ intersects no other edge of G. By JCT, any path between a point in U and a point in $G \setminus U$ intersects Q. Since $X_{j+1} \in U$ and $Y \notin U$, we have that $c_{j+1} \in U$ and $c_t \notin U$. In particular, the path $P(c_{j+1}, c_t)$ intersects Q. By the ordering of $P(X, Y)$, $X_j \notin P(X_{j+1}, Y)$. Then the segment $[c_j, c_{j+1}]$ is not in $P(c_{j+1}, c_t)$, so $P(c_{j+1}, c_t)$ does not intersect Q at the edge $\{v_p, v_q\}$. It follows that $P(c_{j+1}, c_t)$ intersects Q at some other edge e of $P(v_p, v_q)$. But then e is not a gluing edge of $P(X_{j+1}, Y)$, contradicting our observation above. So we must have that $Y \in U$.

A similar argument can be used to show that X must belong to $G \setminus U$. End of Case 1.

CASE 2: $X_{j+1} \notin U$. A similar argument can be used to show that X and Y are on distinct sides of Q. End of Case 2. $\square$.

Proposition 4.5. Let H be a Hamiltonian path of a polyomino G with end-vertices u and v, let W be a switchable box in H, and let $P(X, Y)$ be the looping H-path of W, with $v_x = v_{\text{left}}(A_1)$ and $v_y = v_{\text{right}}(A_1)$. Assume that $P(X, Y)$ is contained in an H-component of G, and that W is anti-parallel. If $P(X, Y)$ has a switchable box Z then either $\text{Sw}(Z)$, $\text{Sw}(W)$ is a cascade, or $Z \mapsto W$ is a valid move.

Proof. Without loss of generality, let $x < y$. By Lemma 4.1, we may partition H into the subpaths $P_0 = P(v_x, v_y)$, $P_1 = P(v_1, v_x)$ and $P_2 = P(v_y, v_r)$ such that every box of $P(X, Y)$ is incident on a vertex of P_0 and a vertex of P_1 or P_2. Then the edges of W in H are either (v_x, v_{x+1}) and (v_{y-1}, v_y), or (v_{x-1}, v_x) and (v_y, v_{y+1}).

CASE 1: The edges of W in H are (v_x, v_{x+1}) and (v_{y-1}, v_y). Suppose that $P(X, Y)$ has a switchable box Z. By Corollary 4.2 (b), $Z \neq W$. Since Z is switchable, Z is not any of the boxes incident on v_x or v_y. Let (v_s, v_{s+1}) and (v_t, v_{t+1}) be the edges of Z in H. Then exactly one of (v_s, v_{s+1}) and (v_t, v_{t+1}) is in P_0 and the other is in P_1 or P_2. WLOG assume that (v_s, v_{s+1}) is in P_0 and that (v_t, v_{t+1}) is in P_2. Then we can partition the edges of H as follows: $P(v_1, v_s)$, (v_s, v_{s+1}), $P(v_{s+1}, v_t)$, (v_t, v_{t+1}), $P(v_{t+1}, v_r)$, where $1 < x < s < y < t < r$.

Now, (v_s, v_{s+1}) and (v_t, v_{t+1}) are either parallel or anti-parallel.

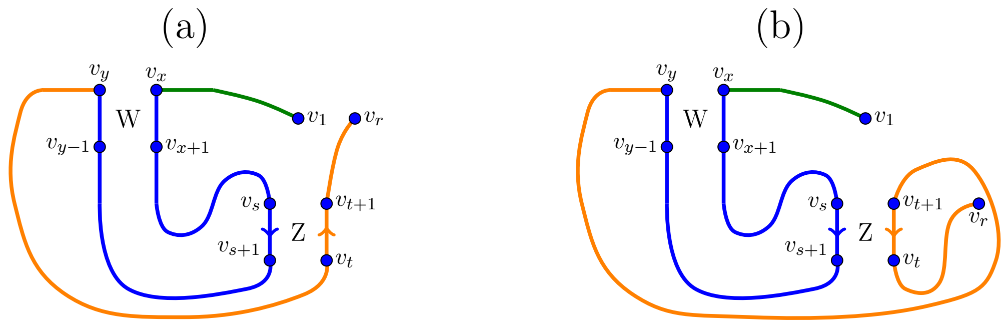

Fig. 4.4. (a) Case 1.1. Edges of Z are anti-parallel. P_0 in blue, P_1 in green, P_2 in orange. (b) Case 1.2. Edges of Z are parallel; $\overrightarrow{K}_0$ in blue, P_1 in green, P_2 in orange.

CASE 1.1: Z is anti-parallel. We show that $Z \mapsto W$ is a valid move. After removing the edges (v_s, v_{s+1}) and (v_t, v_{t+1}) we are left with three paths: $P(v_1, v_s), P(v_{s+1}, v_t)$ and $P(v_{t+1}, v_r)$. Note that adding the edge $\{v_{s+1}, v_t\}$ gives a cycle H_c consisting of the path $P(v_{s+1}, v_t)$ and the edge $\{v_{s+1}, v_t\}$, and adding the edge $\{v_s, v_{t+1}\}$ gives a path $H_p = P(v_1, v_s), \{v_s, v_{t+1}\}, P(v_{t+1}, v_r)$. Now $1 < x < s$ implies that $e_x \in H_p$ and $s < y < t$ implies that $e_y \in H_c$. Thus we have that W is now an (H_c, H_p)-port . By Lemma 1.4.1 (iii), $Z \mapsto W$ is a valid move. See Figure 4.4 (a). End of Case 1.1.

CASE 1.2: Z is parallel. We show that Sw(Z), Sw(W) is a valid sequence of moves. Since Z is parallel, by Lemma 1.4.1(i), Sw(Z) is a valid move. After the removal of the edges (v_s, v_{s+1}) and (v_t, v_{t+1}) we are left with three paths: $P(v_1, v_s), P(v_{s+1}, v_t)$ and $P(v_{t+1}, v_r)$. Adding $\{v_s, v_t\}$, $\{v_{s+1}, v_{t+1}\}$ gives a new Hamiltonian path $H' = v_1, \ldots, v_x, \ldots, v_s, v_t, \ldots, v_y, \ldots, v_{s+1}, v_{t+1}, \ldots, v_r$. Note that now the edges (v_x, v_{x+1}) and (v_y, v_{y+1}) of H' are parallel so Sw(W) is a valid move. See Figure 4.4 (b). End of Case 1.2. End of Case 1.

CASE 2: The edges of W in H are (v_{x-1}, v_x) and (v_y, v_{y+1}). This follows by an argument similar to that of Case 1, so we omit the proof. End of Case 2. □.

4.1 Properties of looping fat paths

Section 4.1 derives and formalizes the structural properties of looping fat paths. There are three important results here: Lemmas 4.6 and 4.13, and Proposition 4.14. Lemma 4.6 gives some preliminary properties. Lemma 4.13 shows that the boundary $B(F)$ of a looping fat path F is a Hamiltonian cycle of F. It builds on the facts that looping fat paths are non-self-adjacent (Proposition 4.7) and have no polyking junctions (Proposition 4.11). Proposition 4.14 shows that fat paths have no consecutive colinear edges other than those belonging to the A_1-type configuration. These three results are used in Section 4.2, where we scan looping fat paths for turns.

Lemma 4.6. Let $F = G\langle N[P(X,Y)]\rangle$ be a looping fat path. Suppose that $W = R(k,l)$ in P has an H-neighbour Z southward in $N[P] \setminus P$. Then:

(a) Z has exactly one H-neighbour in P and W has no other H-neighbour in $N[P] \setminus P$.

(b) If W is not an end-box of P, then the H-neighbours of W in P are $W + (-1,0)$ and $W + (1,0)$. Furthermore, $S_{\rightarrow}(k-1,l;k,l-1) \in H$, $S_{\uparrow}(k+1,l-1;k+2,l) \in H$, and $e(k,k+1;l+1) \in H$.

(c) If W is an end-box of P, then $e(k,k+1;l+1) \in H$ and exactly one of $e(k;l,l+1)$ and $e(k+1;l,l+1)$ belong to H.

(d) Z is a leaf or Z is a switchable box in H.

Analogous statements apply when Z is west, north or east of W.

Proof of (a). Note that if Z has more than one H-neighbour in P then we can make an H-cycle, which contradicts Proposition 1.2.1. To see that W has no other H-neighbour in $N[P] \setminus P$, assume, BWOC, that W has at least two H-neighbours in $N[P] \setminus P$.

If W is an end-box of P, then, by definition of A_1, W has at most two H-neighbours, and at least one of them must belong to P, contradicting our assumption that W has at least two H-neighbours in $N[P] \setminus P$.

If W is not an end-box, then W must have four H neighbours: two in $N[P] \setminus P$ and two in P. By definition of a looping fat path, at least one neighbour of W, say W', is not an end-box. But then W' must be switchable, which conflicts with the definition of a looping fat path. End of proof for (a).

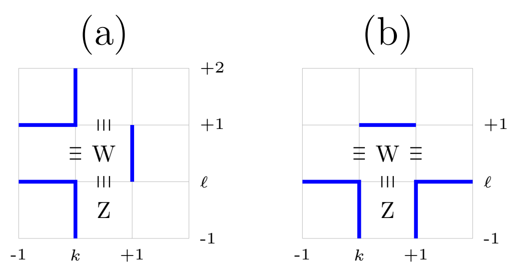

Fig. 4.5. (a) $W{+}(1,0)$ is not an H-neighbour of W. (b) $W{+}(1,0)$ is an H-neighbour of W.

Proof of (b). First, we show that the H-neighbours of W are $W + (1,0)$ and $W + (-1,0)$. BWOC, assume that $W + (1,0)$ is not an H-neighbour of W (See fig). Then the H-neighbours of W in P must be $W + (-1,0)$ and $W + (0,1)$. It follows that $S_{\rightarrow}(k-1,l;k,l-1) \in H$ and $S_{\rightarrow}(k-1,l+1;k,l+2) \in H$. Note that, by definition of A_1 and looping fat paths,

$W+(-1,0)$ is not an end-box of P. But then $W+(-1,0)$ is a switchable box of P, which conflicts with the definition of a looping fat path. Therefore the H-neighbours of W in P are $W+(-1,0)$ and $W+(1,0)$. It follows that $S_{\to}(k-1,l;k,l-1) \in H$ and $S_{\uparrow}(k+1,l-1;k+2,l) \in H$, and by part (a), $(k,k+1;l+1) \in H$. See Figure 4.5. End of proof for (b).

Proof of (c). By part (a) and the assumption that W is an end-box of P, W has exactly one H-neighbour in P and no other H-neighbours in $N[P] \setminus P$. It follows that W has exactly two edges in H and two edges not in H. BWOC, assume that the other edge of W not in H is $e(k,k+1;l+1)$. But then $e(k;l-1,l) \in H$ and $e(k+1;l-1,l) \in H$, and W is switchable, which conflicts with the definition of a looping fat path. It follows that $e(k,k+1;l+1) \in H$. See Figure 4.6.

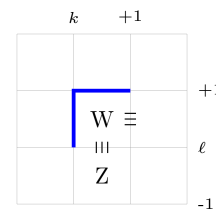

Fig. 4.6.

That exactly one of $e(k;l,l+1)$ and $e(k+1;l,l+1)$ belong to H, follows from the fact that W has exactly one H-neighbour in P and exactly one H-neighbour, Z, in $N[P] \setminus P$. End of proof for (c).

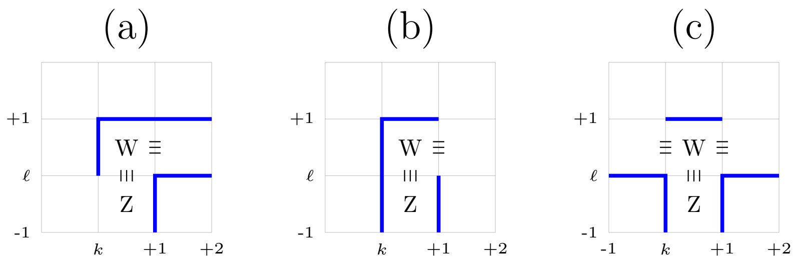

Fig. 4.7. (a) Case 1: F is eastern. (b) Case 1: F is southern. (c) Case 2.

Proof of (d). W is either an end-box of P or it is not.

CASE 1: W is an end-box of P. By part (c), we may assume WLOG that $e(k;l,l+1) \in H$ and $e(k+1;l,l+1) \notin H$. Then F is eastern or southern. Suppose that F is eastern. Then $e(k+1,k+2;l+1) \in H$. It follows that $S_{\uparrow}(k+1,l-1;k+2,l) \in H$. But then $W+(1,0) \in P$ is switchable, which conflicts with the definition of a looping fat path. So F must be southern. Then $e(k;l-1,l) \in H$. It follows that $S_{\uparrow}(k+1,l-1;k+2,l) \in H$. Now, either $e(k,k+1;l-1) \in H$ or $e(k,k+1;l-1) \notin H$. Either way, (d) is satisfied. See Figure 4.7 (a) and (b).

CASE 2: W is not an end-box of P. By part (b), the H-neighbours of W in P are $W+(-1,0)$ and $W+(1,0)$ and we have that $S_{\to}(k-1,l;k,l-1) \in H$, $S_{\uparrow}(k+1,l-1;k+2,l) \in H$. Then, either $e(k,k+1;l-1) \in H$ or $e(k,k+1;l-1) \notin H$. Either way, (d) is satisfied. See Figure 4.7 (c).□

Let $F = G\langle N[P(X,Y)]\rangle$ be a looping fat path following a leaf L. We will often write "the middle-boxes of F" to refer to the middle-boxes of the A_1-type following L.

Proposition 4.7. Let G be a polyomino and let H be a Hamiltonian path of G. Then every sturdy looping fat path of G is non-self-adjacent.

We will need the following lemmas and observations before we prove Proposition 4.7.

Observation 4.8. Let $F = G\langle N[P(X,Y)]\rangle$ be a southern looping fat path. Let P_0, P_1 and P_2 be the partitioning of H given in Lemma 4.1. Then (v_x, v_{x+1}) and (v_{y-1}, v_y) are either anti-parallel and incident on the box $X + (-1, 0)$ (as in Figure 4.8 (a)) or collinear with the edge $\{v_x, v_y\}$ of $G \setminus H$ (as in Figure 4.8 (b)). Analogous observations apply to eastern, northern and western looping fat paths. □

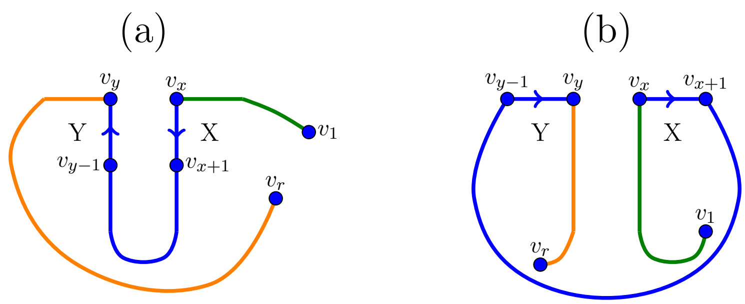

Fig. 4.8. P_0 in blue, P_1 in green, P_2 in orange.

Lemma 4.9. Let G, H, $P(X,Y)$, P_0, P_1 and P_2 be as in Lemma 4.1. Let Q be the cycle $P_0, \{v_x, v_y\}$ in G. Let the vertices v and v' belong to $V((P_1 \setminus \{v_x\}) \cup (P_2 \setminus \{v_y\}))$. Then the segment $[v, v']$ does not intersect Q exactly once at point p of an edge $\{u, u'\}$ of Q with $p \neq u$ and $p \neq u'$.

Proof. Let U be the region bounded by Q. Let the vertices v and v' belong to $V((P_1 \setminus \{v_x\}) \cup (P_2 \setminus \{v_y\}))$. Then $v \in V(U)$ or $v \in V(G) \setminus V(U)$. For definiteness, assume that $v \in V(U)$. For a contradiction, assume that the segment $[v, v']$ that intersects Q exactly once at a point p of an edge $\{u, u'\}$ of Q with $p \neq u$ and $p \neq u'$. Now, Corollary 1.1.5 implies that $v' \in V(G) \setminus V(U)$. Lemma 1.3.3 (c) implies that $v_1 \in V(U)$. By Lemma 1.5.4 in Chapter 1, $v_r \in V(U)$ as well. By Lemma 1.3.3 (c) again, $v' \in V(U)$ contradicting that $v' \in V(G) \setminus V(U)$ □.

Lemma 4.10. Let G be a polyomino and let $H = v_1, \ldots, v_r$ be a Hamiltonian path of G. Let $F = G\langle N[P(X,Y)]\rangle$ be a sturdy looping fat path. Let P_0, P_1 and P_2 be the partitioning of H given in Lemma 4.1. Then, no edge of P_1 or P_2 has both of its incident boxes belonging to F.

Remark. The non-self-adjacency of sturdy looping fat paths is a natural consequence of their local structure but the proof requires lengthy and detailed case analysis. In the proofs of Lemma 4.10 and Proposition 4.7 we will use repeatedly and implicitly Proposition 4.5 (a), Lemma 4.6 (a), and the facts that a box of P that is not an end-box has exactly two H-neighbours in P while an end-box of P has exactly one H-neighbour in P.
Lastly, the following type of scenario occurs several times. Let F be a looping fat path. Suppose $Z = R(a,b)$ is an end-box of P and $Z + (1,0)$ is its H-neighbour in $N[P] \setminus P$. Assume that $e(a;b,b+1) \in H$, $e(a,a+1;b) \notin H$, $e(a,a+1;b+1) \in H$ $e(a+1,a+2;b+1) \in H$. See Figure 4.9 (a). Then the F can be eastern or southern. But if F is southern then $e(a;b-1,b) \in H$ and $S_{\uparrow}(a+1,b-1;a+2,b) \in H$, and then $Z+(0,-1)$ is in P and it is switchable. Hence, F must be eastern. See Figure 4.9 (b). The cases where the H-neighbour of Z is north, west, or south of Z are analogous. We denote the argument in this paragraph $(**)$ for reference.

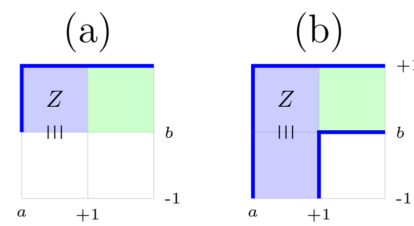

Fig. 4.9.

Proof of Lemma 4.10. Let F be a sturdy looping fat path. Assume that some edge $e_j = (v_j, v_{j+1})$ of P_1 or P_2 has both its incident boxes in F. Without loss of generality assume that e_j is an edge of P_2. We will use induction to show that now we can find a box of P that has an end-vertex of H incident on it, which contradicts the assumption that F is a sturdy looping fat path.

Let e_j be an edge of P_2. Then e_{j+1} is one edge closer to v_r then e_j. We denote this remark by $(*)$ for reference.

Let $Q(j)$ be the statement "Both boxes incident on e_j belong to F". Assume that $Q(j)$ implies $Q(j+1)$. BWOC assume $Q(j_0)$ for some edge e_{j_0} of P_2. By induction, we have $Q(r-1)$. For definiteness let $W_1 = R(a,b)$ and let $W_2 = R(a,b-1)$ be the boxes of F incident on e_{r-1}, and let $v_r = v(a+1,b)$. If any of the boxes incident to v_1 belong to P, then this conflicts with the definition of F. Then we must have that W_1 and W_2 belong to $N[P] \setminus P$. By Lemma 4.6 (a), $W_1 + (1,0)$ and $W_2 + (1,0)$ do not belong to F. Let W_1'

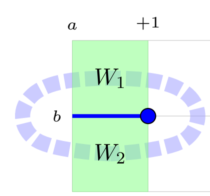

Fig. 4.10.

and W_2' be the H-neighbours of W_1 and W_2 in P, respectively. Then there is an H-cycle $W_2, W_2', \dots, W_1', W_1, W_1 + (1,0), W_2 + (1,0), W_2$, contradicting Proposition 1.2.1. See Figure 4.10.

It remains to show that $Q(j)$ implies $Q(j+1)$. For definiteness, let $e_j = (v_j, v_{j+1}) = (v(a,b), v(a+1,b))$. We remark that assuming $e_{(j)}$ to be horizontal and fixing its orientation comes at the expense of allowing F to have any direction; however, this trade-off is reasonable as it simplifies and shortens the proof. Assume $Q(j)$ is true, and let Z_1 and Z_2 be the boxes of F incident on e_j. There are three possibilities: both Z_1 and Z_2 belong to P, or both Z_1 and Z_2 belong to $N[P] \setminus P$, or one of Z_1 and Z_2 belongs to P and the other belongs to $N[P] \setminus P$. WLOG let Z_1 be on the left of e_j and let Z_2 be on the right of e_j.

CASE 1: Both Z_1 and Z_2 belong to P. There are three possibilities: $e_{j+1} = e(a+1; b, b+1)$, $e_{j+1} = e(a+1, a+2; b)$, or $e_{j+1} = e(a+1; b-1, b)$. By symmetry of the first and third, we only need to check the first two.

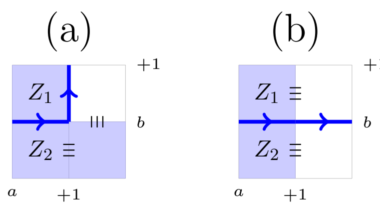

Fig. 4.11. (a) Case 1.1. (b) Case 1.2.

CASE 1.1: $e_{j+1} = e(a+1; b, b+1)$. Then $e(a+1; b-1, b) \notin H$ and $e(a+1, a+2; b) \notin H$. Lemma 4.6 (d) implies that $Z_2 + (1,0) \in P$. Then $Z_1 + (1,0) \in P$ or $Z_1 + (1,0) \in N[P] \setminus P$. Either way, $Q(j+1)$ holds. End of Case 1.1.

CASE 1.2: $e_{j+1} = e(a+1, a+2; b)$. Then $e(a+1; b-1, b) \notin H$ and $e(a+1, b; b+1) \notin H$. Then $Z_1 + (1,0)$ and $Z_2 + (1,0)$ either belong to P or to $N[P] \setminus P$. Either way $Q(j+1)$ holds. End of Case 1.2. End of Case 1.

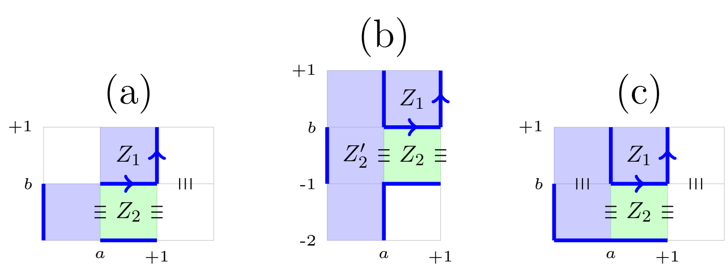

Fig. 4.12. (a) Case 2.1. (b) Case 2.1(a). (c) Case 2.1(b_1).

CASE 2: One of Z_1 and Z_2 belongs to P and the other belongs to $N[P] \setminus P$. For definiteness, let $Z_1 \in P$ and $Z_2 \in N[P] \setminus P$. Then $e_{j+1} = e(a+1; b, b+1)$, $e_{j+1} = e(a+1, a+2; b)$ or $e_{j+1} = e(a+1, b-1; b)$. Note that the case where $e_{j+1} = e(a+1, b-1; b)$ is the same as Case 1.1, so we only need to check the cases where $e_{j+1} = e(a+1; b, b+1)$ and where $e_{j+1} = e(a+1, a+2; b)$.

CASE 2.1: $e_{j+1} = e(a+1; b, b+1)$. Then $e(a+1; b-1, b) \notin H$. By Lemma 4.6 (d), $e(a, a+1; b-1) \in H$. By Lemma 4.6 (a), $Z_2 + (1,0)$ either belongs to P or it is not a box of F. If $Z_2 + (1,0) \in P$, then $Z_2 + (1,1) \in P$ or $Z_2 + (1,1) \in N[P] \setminus P$. Either way, $Q(j+1)$ holds. We check that $Z_2 + (1,0) \notin P$ is impossible.

Toward a contradiction, assume that $Z_2+(1,0) \notin P$. Then we have that $Z_2+(-1,0) \in P$, and Lemma 4.6 (b) and (c) imply that $e(a-1; b-1, b) \in H$. Now, $Z_2+(-1,0)$ is an end-box of P or it is not. See Figure 4.12 (a).

CASE 2.1(a): $Z_2 + (-1,0)$ is not an end-box of P. By Lemma 4.6 (b), $Z_2 + (-1,-1) \in P$, $Z_2+(-1,1) \in P$, $S_{\downarrow}(a, b+1; a+1, b) \in H$, $S_{\uparrow}(a, b-2; a+1, b-1) \in H$ and $e(a-1; b-1, b) \in H$. Then, after $Z_2 \mapsto Z_1$, $Z_2 + (-1,0) \in P$ is switchable, contradicting FPC-4. See Figure 4.12 (b). End of Case 2.1(a).

CASE 2.1(b): $Z_2+(-1,0)$ is an end-box of P. By Lemma 4.6(c) exactly one of $e(a-1, a; b-1)$ and $e(a-1, a; b)$ belong to H.

CASE 2.1(b_1): $e(a-1, a; b-1) \in H$ and $e(a-1, a; b) \notin H$. Then we have that $e(a; b, b+1) \in H$. But this implies that Z_1 is an end-box of P as well, which is impossible. See Figure 4.12 (c). End of Case 2.1(b_1).

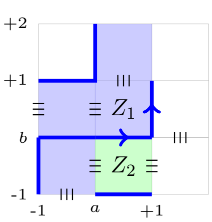

Fig. 4.13. Case 2.1(b_2).

CASE 2.1(b_2): $e(a-1, a; b) \in H$ and $e(a-1, a; b-1) \notin H$. Then $e(a; b, b+1) \notin H$ and $e(a-1; b, b+1) \notin H$. Note that if $e(a, a+1; b+1) \in H$, then Z_1 is the other end-box of P, which does not agree with the assumption that $Z_2 + (-1,0)$ is an end-box of P. So we only need to check the case where $e(a, a+1; b+1) \notin H$. In this case, we must have $S_{\rightarrow}(a-1, b+1; a, b+2) \in H$. Now, the fact that Z_1 is not an end-box of P implies that both H-neighbours $Z_1 + (-1,0)$ and $Z_1 + (0,1)$ of Z_1 are boxes of P. But then $Z_1 + (-1,0)$ is switchable. See Figure 4.13. End of Case 2.1(b_2). End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $e_{j+1} = e(a+1, a+2; b)$. Then $e(a+1; b-1, b) \notin H$, $e(a+1; b, b+1) \notin H$ and $Z_1 + (1,0)$ must be a box of F. Now, $Z_2 + (1,0)$ either belongs to P or it is not a box of F. If $Z_2 + (1,0) \in P$, then $Q(j+1)$ holds. We will show that $Z_2 + (1,0) \notin P$ is impossible.

Toward a contradiction, assume that $Z_2 + (1,0)$ is not a box of F. Then we have that $Z_2 + (-1,0) \in P$, and, by Lemma 4.6 (d), $e(a, a+1; b-1) \in H$. Lemma 4.6 (b) and (c) imply that $e(a-1; b-1, b) \in H$. See Figure 4.14. Now, either $Z_2 + (-1,0)$ is an end-box of P or it is not.

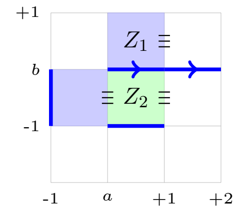

Fig. 4.14. Case 2.2.

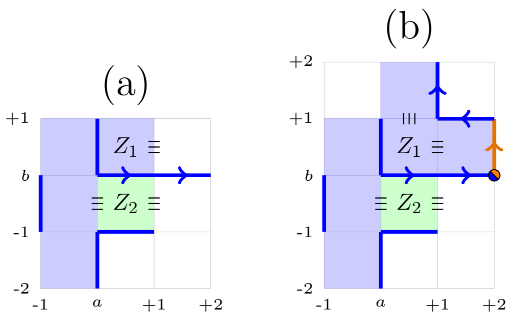

Fig. 4.15. (a) Case 2.2(a). (b) Case 2.2(a_1), e_{j+2} in orange.

CASE 2.2(a): $Z_2 + (-1,0)$ is not an end-box of P. By Lemma 4.6 (b) we have that $Z_2 + (-1,-1) \in P$, $Z_2 + (-1,1) \in P$, $e(a; b-2, b-1) \in H$ and $e(a; b, b+1) \in H$. See Figure 4.15 (a). Now, either Z_1 is an end-box of P or it is not.

CASE 2.2(a_1): Z_1 is not an end-box of P. Then $e(a, a+1; b+1) \notin H$. It follows that $S_{\downarrow}(a+1, b+2; a+2, b+1) \in H$. Note that if $e(a+2; b, b+1) \notin H$, then $Z_1 + (1,0)$ is switchable, so we only need to check the case where $e(a+2; b, b+1) \in H$. This implies that $Z_1 + (1,0)$ is an end-box of P. It follows that F is eastern and that $v(a+2, b)$ is a corner of A_1. But now, the A_1-partitioning of H implies that $e(a+2; b, b+1) = e_{j+2} \in P_0$, contradicting $(*)$ in the beginning of the lemma. See Figure 4.15 (b). End of Case 2.2(a_1).

CASE 2.2(a_2): Z_1 is an end-box of P. Then $e(a; b+1, b+2) \in H$ or $e(a; b+1, b+2) \notin H$. If $e(a; b+1, b+2) \in H$, then we must have $S_{\downarrow}(a+1, b+2; a+2, b+1) \in H$. But this implies that at least one of $Z_1 + (0,1)$ and $Z_1 + (1,0)$ belongs to P and is switchable. So we only need to check the case where $e(a; b+1, b+2) \notin H$. In this case, we observe that F must be western. This implies that one of $v(a,b)$ and $v(a, b+1)$ is a corner of A_1. If the former, then $e(a; b, b+1) \in P_0$, which contradicts Lemma 4.3. See Figure 4.16 (a). If the latter then the other end-box of P is $Z_1 + (0,2)$ and F follows the edge $e(a-1; b+1, b+2)$. This means that $e(a, a+1; b+1) \in H$, and that $e(a+1, a+2; b+1) \in H$. But then $Z_1 + (1,0) \in H$ is switchable. See Figure 4.16 (b). End of Case 2.2(a_2).

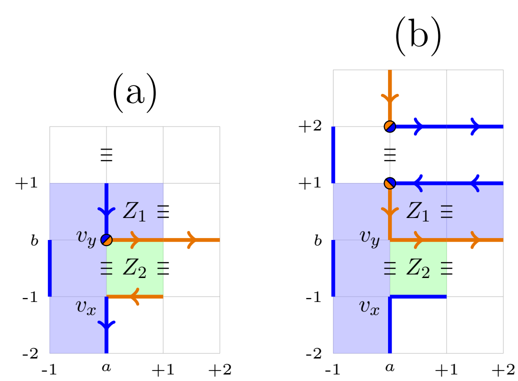

Fig. 4.16. (a) Case 2.2(a_2). $v(a,b)$ is a corner of A_1. (b) Case 2.2(a_2). $v(a, b+1)$ is a corner of A_1.

CASE 2.2(b): $Z_2 + (-1, 0)$ is an end-box of P. By Lemma 4.6 (c), exactly one of $e(a - 1, a; b - 1)$ and $e(a - 1, a; b)$ belongs to H.

CASE 2.2(b_1): $e(a-1, a; b-1) \in H$ and $e(a-1, a; b) \notin H$. Then $e(a; b, b+1) \in H$. Since Z_1 cannot be the other end-box of P, its H-neighbours in P must be $Z_1 + (0, 1)$ and $Z_1 + (1, 0)$, so $e(a, a+1; b+1) \notin H$. Then we must have that $e(a+1, a+2; b+1) \in H$. Since $Z_1 + (1, 0)$ cannot be the other end-box of P either, it must be switchable, contradicting FPC-3. See Figure 4.17 (a). End of Case 2.2(b_1).

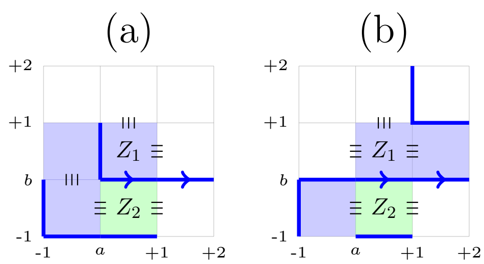

Fig. 4.17. (a) Case 2.2(b_1). (b) Case 2.2(b_2).

CASE 2.2(b_2): $e(a - 1, a; b - 1) \notin H$ and $e(a - 1, a; b) \in H$. Then $e(a; b, b + 1) \notin H$. Note that Z_1 and $Z_1 + (1, 0)$ are not end-boxes. If $e(a, a + 1; b + 1) \in H$, then Z_1 is switchable, so we only need to check the case where $e(a, a+1; b+1) \notin H$. Then we have that $S_{\downarrow}(a+1, b+2; a+2, b+1) \in H$. Lemma 4.6(b) implies $Z_1 + (1, 0) \in P$. Now, either $e(a+2; b, b+1) \in H$ or $e(a+2; b, b+1) \notin H$. The former implies that $Z_1 + (1, 0)$ is an end-box of P, which is impossible; and if the latter, then $Z_1 + (1, 0)$ is switchable, contradicting FPC-3. See Figure 4.17 (b). End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2.

CASE 3: both Z_1 and Z_2 belong to $N[P] \setminus P$. There are three possibilities: $e_{j+1} = e(a + 1; b, b + 1)$, $e_{j+1} = e(a + 1, a + 2; b)$ or $e_{j+1} = e(a + 1; b - 1, b)$. By symmetry of the first and third, we only need to check the first two.

CASE 3.1: $e_{(j+1)} = e(a + 1; b, b + 1)$. Then $e(a + 1; b - 1, b) \notin H$. By Lemma 4.6 (d), $e(a, a + 1; b - 1) \in H$. Now the H-neighbour of Z_2 in P is $Z_2 + (-1, 0)$ or $Z_2 + (1, 0)$. If the latter, then we're done by Case 1.1. We will show the former is impossible.

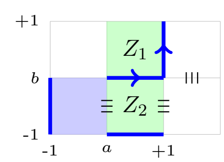

Fig. 4.18.

Assume that the H-neighbour of Z_2 in P is $Z_2 + (-1, 0)$. It follows that $e(a; b - 1, b) \notin H$ and, by Lemma 4.6 (b) and (c), that $e(a - 1; b - 1, b) \in H$. Now, $Z_2 + (-1, 0)$ is an end-box of P or it is not. Note that if $Z_2 + (-1, 0)$ is not an end-box of P, then, as in Case 2.1(a), after $Z_2 \mapsto Z_1$, $Z_2 + (-1, 0) \in P$ is switchable, contradicting FPC-4. Therefore, we only need to check the case where $Z_2 + (-1, 0)$ is an end-box of P. See Figure 4.18. By Lemma 4.6(c), exactly one of $e(a - 1, a; b - 1)$ and $e(a - 1, a; b)$ belongs to H.

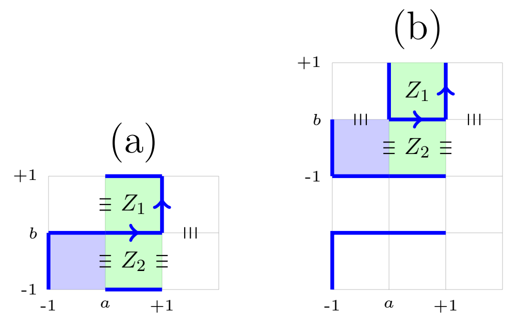

Fig. 4.19. (a) Case 3.1(a). (b) Case 3.1(b).

CASE 3.1(a): $e(a-1,a;b) \in H$ and $e(a-1,a;b-1) \notin H$. Then $e(a;b,b+1) \notin H$. By Lemma 4.6 (d), it follows that $e(a,a+1;b+1) \in H$. Then we have that $Z_1+(1,0) \in P$. Note that Lemma 4.6(b) implies that $Z_1+(1,0)$ is an end-box of P, but this disagrees with the assumption that $Z_2+(-1,0)$ is the other end-box of P. See Figure 4.19 (a). End of Case 3.1(a).

CASE 3.1(b): $e(a-1,a;b-1) \in H$ and $e(a-1,a;b) \notin H$. Then $e(a;b,b+1) \in H$. Note that F can be northern or eastern. If F is northern, then $e(a-1;b,b+1) \in H$, and the box $Z_2+(-1,1) \in P$ is switchable; and if F is eastern, then the switchable middle-box of F is $W = Z_2+(-1,-1)$. But then, after $Z_2 \mapsto Z_1$, $W \mapsto Z_2+(-1,0)$ is valid, contradicting FPC-4. See Figure 4.19 (b). End of Case 3.1(b). End of Case 3.1

CASE 3.2: $e_{(j+1)} = e(a+1,a+2;b)$. If both $Z_1+(1,0)$ and $Z_2+(1,0)$ belong to F we're done, so we assume that at least one of $Z_1+(1,0)$ and $Z_2+(1,0)$ is not a box of F. By symmetry, we may assume WLOG that $Z_2+(1,0)$ is not a box of F. By Lemma 4.6(d) we have that $e(a,a+1;b-1) \in H$, $Z_2+(-1,0) \in P$ and $e(a;b-1,b) \notin H$. By Lemma 4.6 (b) and (c), $e(a-1;b-1,b) \in H$. See Figure 4.20. Now, $Z_2+(-1,0)$ is an end-box of P or it is not.

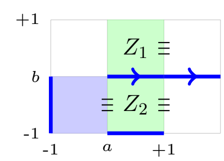

Fig. 4.20. Case 3.2.

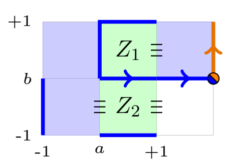

Fig. 4.21. Case 3.2(a).

CASE 3.2(a): $Z_2+(-1,0)$ is not an end-box of P. Then, by Lemma 4.6 (b), $Z_1+(-1,0) \in P$ and $e(a;b,b+1) \in H$. By Lemma 4.6 (d) we have that $e(a,a+1;b+1) \in H$. It follows that the H-neighbour of Z_1 in P is $Z_1+(1,0)$. Lemma 4.6 (b) and (c) imply that $e(a+2;b,b+1) \in H$. Then it must be the case that $Z_1+(1,0)$ is an end-box of P. Observe that this implies that $v(a+2,b)$ is a corner vertex of A_1. Now, the A_1-partitioning of H implies that $e(a+2;b,b+1) \in P_0$, while our assumption that $e(a,a+1;b) \in E(P_1) \cup E(P_2)$ implies that $e(a+2;b,b+1) = e_{(j+2)}$. But then $e_{(j+2)} \in P_0$, contradicting $(*)$. See Figure 4.21. End of Case 3.2(a).

CASE 3.2(b): $Z_2+(-1,0)$ is an end-box of P. By Lemma 4.6(c), exactly one of $e(a-1,a;b-1)$ and $e(a-1,a;b)$ belongs to H.

CASE 3.2(b_1): $e(a-1,a;b-1) \in H$ and $e(a-1,a;b) \notin H$. It follows that $Z_1 + (-1,0) \in P$ and $e(a;b,b+1) \in H$. As in Case 3.2(a), $Z_1 + (1,0)$ must be an end-box of P. But this conflicts with our assumption that $Z_2 + (-1,0)$ is an end-box of P. End of Case 3.2(b_1)

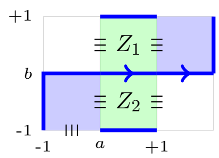

Fig. 4.22. Case 3.2(b_2).

CASE 3.2(b_2): $e(a-1,a;b) \in H$ and $e(a-1,a;b-1) \notin H$. Then $e(a;b,b+1) \notin H$ and, by Lemma 4.6 (d), $e(a,a+1;b+1) \in H$. Lemma 4.6 (b) and (c) imply that $e(a+2;b,b+1) \in H$. Then it must be the case that $Z_1 + (1,0)$ is an end-box of P, but again, this conflicts with the assumption that $Z_2 + (-1,0)$ is an end-box of P. See Figure 4.22. End of Case 3.2(b_2). End of Case 3.2(b). End of Case 3.2. End of Case 3. □

Proof of Proposition 4.7. Let $F = G\langle N[P(X,Y)]\rangle$ be a sturdy looping fat path of G. Let P_0, P_1 and P_2 be the partitioning of H given in Lemma 4.1. For a contradiction, assume that there is an edge e of H such that both its incident boxes Z_1 and Z_2 belong to F. By Lemma 4.10, we may assume that e is an edge of P_0. By Lemma 4.3, Z_1 and Z_2 cannot both be boxes of $P(X,Y) = P$. Then, either both Z_1 and Z_2 are boxes of $N[P] \setminus P$, or exactly one of Z_1 and Z_2 is a box of P, and the other is a box of $N[P] \setminus P$. For definiteness, assume that e is the horizontal edge $e(a,a+1;b)$, that Z_1 is north of $e(a,a+1;b)$ and that Z_2 is south of $e(a,a+1;b)$,

CASE 1: Z_1 and Z_2 belong to $N[P] \setminus P$. Let Z_1' and Z_2' be the H-neighbours of Z_1 and Z_2 in P, respectively. Note that it is not possible for both Z_1' and Z_2' to be end-boxes of P. Then, either exactly one of Z_1 and Z_2 is adjacent to an end-box of P or neither is.

CASE 1.1: Neither Z_1 nor Z_2 is adjacent to an end-box of P. There are three possibilities: $Z_2' = Z_2 + (0,-1)$; $Z_2' = Z_2 + (-1,0)$; and $Z_2' = Z_2 + (1,0)$. Since the second and third are symmetric, we only need to check the first two.

CASE 1.1(a): $Z_2' = Z_2 + (0,-1)$. By Lemma 4.6 (b), $Z_2' + (-1,-1) \in P$, $Z_2' + (1,-1) \in P$, $S_{\rightarrow}(a-1,b-1;a,b) \in H$, $S_{\downarrow}(a+1,b;a+2,b-1) \in H$ and $e(a,a+1;b-2) \in H$. It follows that $e(a;b,b+1) \notin H$ and $e(a+1;b,b+1) \notin H$. By Lemma 4.6(d), $e(a,a+1;b+1) \in H$. Then, after $Z_1 \mapsto Z_2$, the box $Z_2 + (0,-1)$ is in P and is switchable, contradicting FPC-4. See Figure 4.23 (a). End of Case 1.1(a).

CASE 1.1(b): $Z_2' = Z_2 + (-1, 0)$. By Lemma 4.6 (b), $Z_2' + (0, -1) \in P$, $Z_2' + (0, 1) \in P$, $S_{\downarrow}(a, b + 1; a + 1, b) \in H$, $S_{\uparrow}(a, b - 2; a + 1, b - 1) \in H$ and $e(a - 1; b - 1, b) \in H$. Now $Z_1' = Z_1 + (0, 1)$ or $Z_1' = Z_1 + (1, 0)$. See Figure 4.23 (b).

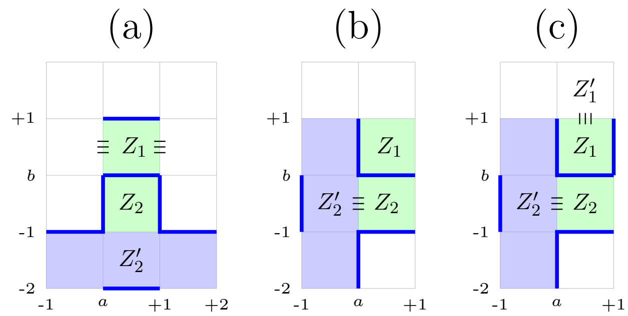

Fig. 4.23. (a) Case 1.1(a). (b) Case 1.1(b). (c) Case 1.1(b_1).

CASE 1.1(b_1): $Z_1' = Z_1 + (0, 1)$. By Lemma 4.6 (d), $e(a + 1; b, b + 1) \in H$. Then, after $Z_2 \mapsto Z_1$, $Z_2 + (-1, 0) \in P$ is switchable, contradicting FPC-4. See Figure 4.23 (c). End of Case 1.1(b_1).

CASE 1.1(b_2): $Z_1' = Z_1 + (1, 0)$. By Lemma 4.6 (b), $Z_1' + (0, -1) \in P$, $Z_1' + (0, 1) \in P$, $S_{\rightarrow}(a, b + 1; a + 1, b + 2) \in H$, $e(a + 1; b - 1, b) \in H$ and $e(a + 2; b, b + 1) \in H$. Note if $Z_1 + (2, 0) \in P$ then, $e(a + 2; b, b + 1) \in E(P_1) \cup E(P_2)$ contradicts Lemma 4.10; and $e(a+2; b, b+1) \in P_0$ contradicts Lemma 4.3. Therefore we may assume that $Z_1 + (2, 0) \notin P$. If $e(a + 2; b - 1, b) \in H$ then $Z_2 + (1, 0) \in P$ is switchable, so we may also assume that $e(a+2; b-1, b) \notin H$. It follows that $S_{\uparrow}(a+2, b-2; a+3, b-1) \in H$ and $e(a+2, a+3; b) \in H$. By symmetry, we may assume that $e(a - 1; b, b + 1) \notin H$, $e(a - 2, a - 1; b) \in H$ and $S_{\rightarrow}(a - 2, b + 1; a - 1, b + 2) \in H$.

Note that, by definition of A_1, none of the vertices of Z_1 and Z_2 can be corners of A_1. If one of $v(a + 2, b)$ and $v(a + 2, b + 1)$ is a corner of A_1, then F has to be eastern. But this implies that $Z_1 + (2, 0) \in P$, contradicting our assumption in the previous paragraph. So we may assume that $v(a+2, b)$ and $v(a+2, b+1)$ are not corners of A_1. By symmetry $v(a - 1, b - 1)$ and $v(a - 1, b)$ are not corners of A_1 either. By Proposition 4.4, $v(a+2, b)$ and $v(a+2, b+1)$ must belong to $(V(P_1) \cup V(P_2)) \setminus \{v_x, v_y\}$. By symmetry $v(a-1, b-1)$ and $v(a-1, b)$ must belong to $(V(P_1) \cup V(P_2)) \setminus \{v_x, v_y\}$ as well. But now, the existence of the segment $[v(a - 1, b - 1), v(a + 2, b + 1)]$ contradicts Lemma 4.9. See Figure 4.24. End of Case 1.1(b_2). End of Case 1.1(b). End of Case 1.1.

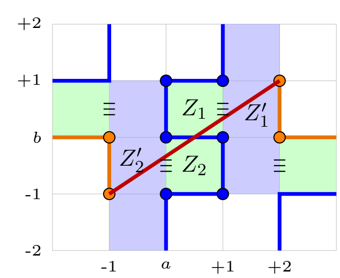

Fig. 4.24. Case 1.1(b_2).

CASE 1.2: Exactly one of Z_1 and Z_2 is adjacent to an end-box of P. WLOG we may assume that Z_1 is adjacent to the end-box Z_1'. Then $Z_1' = Z_1 + (0, 1)$, $Z_1' = Z_1 + (-1, 0)$ or $Z_1' = Z_1 + (1, 0)$. Since the second and third are symmetric, we only need to check the first two.

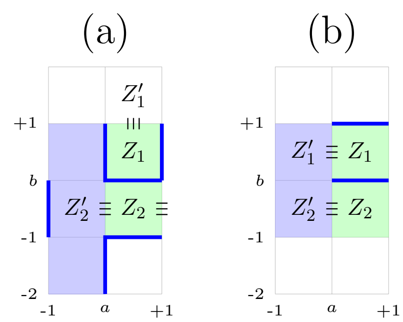

Fig. 4.25. (a) Case 1.2(a). (b) Case 1.2(b_1).

CASE 1.2(a): $Z_1' = Z_1 + (0,1)$. By Lemma 4.6 (d), $e(a;b,b+1) \in H$ and $e(a+1;b,b+1) \in H$. Then $e(a;b-1,b) \notin H$ and $e(a+1;b-1,b) \notin H$. By Lemma 4.6 (d), $e(a,a+1;b-1) \in H$. Now, $Z_2' = Z_2 + (-1,0)$ or $Z_2' = Z_2 + (1,0)$. By symmetry, we may assume WLOG that $Z_2' = Z_2 + (-1,0)$. But now we're back to Case 1.1(b_1). See Figure 4.25 (a). End of Case 1.2(a).

CASE 1.2(b): $Z_1' = Z_1 + (-1,0)$. Then $e(a;b,b+1) \notin H$, and by Lemma 4.6 (d), $e(a,a+1;b+1) \in H$. Now there are three possibilities for the H-neighbour Z_2' of Z_2: $Z_2' = Z_2 + (-1,0)$, $Z_2' = Z_2 + (0,-1)$, or $Z_2' = Z_2 + (1,0)$.

CASE 1.2(b_1): $Z_2' = Z_2 + (-1,0)$. Lemma 4.6 (b) implies that $e(a;b,b+1) \in H$. But this contradicts the assumption of Case 1.2(b). See Figure 4.25 (b). End of Case 1.2(b_1).

CASE 1.2(b_2): $Z_2' = Z_2 + (0,-1)$. We note that this is the same as Case 1.1(a). See Figure 4.23 (a). End of Case 1.2 (b_2).

CASE 1.2(b_3): $Z_2' = Z_2 + (1,0)$. By Lemma 4.6 (b), $Z_2' + (0,1) \in P$, $Z_2' + (0,-1) \in P$, $e(a+1;b,b+1) \in H$, $S_{\rightarrow}(a,b-1;a+1,b-2) \in H$ and $e(a+2;b-1,b) \in H$. By Lemma 4.6(c) we have that $e(a-1;b,b+1) \in H$ and that exactly one of $e(a-1,a;b)$ and $e(a-1,a;b+1)$ belongs to H. See Figure 4.26.

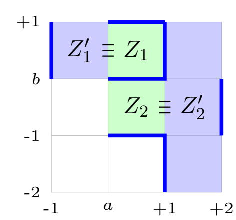

Fig. 4.26. Case 1.2(b_3).

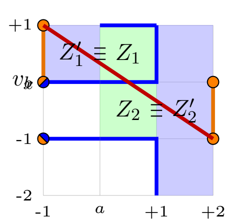

Fig. 4.27. Case 1.2(b_3).(i).

CASE 1.2(b_3).(i): $e(a-1,a;b) \in H$. By ((∗∗)), F is eastern. Note that the corner vertices of A_1 must be $v(a-1,b)$ and $v(a-1,b-1)$. It follows that $V(Z_2')$ contains no corners of A_1 and that $v(a-1,b+2) \in (V(P_1) \cup V(P_2))$. Our assumption that $e(a,a+1;b) \in P_0$ and Observation 4.8, now imply that $v(a+1,b-1)$ belongs to $V(P_0) \setminus (V(P_1) \cup V(P_2))$. By Proposition 4.4, we have that $v(a+2;b-1)$ and $v(a+2;b)$ belong to $(V(P_1) \cup V(P_2)) \setminus V(P_0)$. But then, the existence of the segment $[v(a-1,b+1), v(a+2,b-1)]$ contradicts Lemma 4.9. See Figure 4.27. End of Case 1.2(b_3).(i).

CASE 1.2$(b_3).(ii)$*:* $e(a-1,a;b+1) \in H$. Then we have that $e(a;b-1,b) \in H$, and, by ((∗∗)), that F is eastern. Using the same arguments as in Case 1.2$(b_3).(i)$, we have that $e(a;b-1,b) \in (P_0)$ and that $v(a-1,b)$ and $v(a+2,b-1)$ belong to $(V(P_1) \cup V(P_2)) \setminus V(P_0)$. But then the existence of the segment $[v(a-1,b), v(a+2,b-1)]$ contradicts Lemma 4.9. See Figure 4.28. End of Case 1.2$(b_3).(ii)$. End of Case 1.2(b_3). End of Case 1.2(b). End of Case 1.2. End of Case 1.

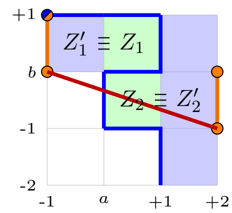

Fig. 4.28. Case 1.2$(b_3).(ii)$.

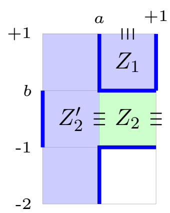

Fig. 4.29. Case 2.1(a).

CASE 2: Exactly one of Z_1 *and* Z_2 *is a box of* P*, and the other is a box of* $N[P] \setminus P$. For definiteness, assume that $Z_1 \in P$ and $Z_2 \in N[P] \setminus P$. Either Z_1 is an end-box of P or it is not.

CASE 2.1: Z_1 *is an end-box of* P. Let Z_2' be the H-neighbour of Z_2 in P. Then $Z_2' = Z_2 + (0,-1)$, $Z_2' = Z_2 + (-1,0)$, or $Z_2' = Z_2 + (0,1)$. Since the second and third are symmetric, we only need to check the first two.

CASE 2.1(a): $Z_2' = Z_2 + (0,-1)$. The assumption of Case 2.1 implies that Z_2' is not an end-box of P. By Lemma 4.6(b), $Z_2' + (0,1) \in P$ and $e(a;b,b+1) \in H$. But now, $e(a;b,b+1) \in P_0$, contradicts Lemma 4.3, and $e(a;b,b+1) \in E(P_1) \cup E(P_2)$ contradicts Lemma 4.10. See Figure 4.29. End of Case 2.1(a).

CASE 2.1(b): $Z_2' = Z_2 + (0,-1)$. This is the same as Case 1.1(a). End of Case 2.1(b). End of Case 2.1.

CASE 2.2: Z_1 *is not an end-box of* P. Let Z_2' be the H-neighbour of Z_2 in P. Then $Z_2' = Z_2 + (0,-1)$ or $Z_2' = Z_2 + (-1,0)$ or $Z_2' = Z_2 + (0,1)$. Since the second and third are symmetric, we only need to check the first two.

CASE 2.2(a): $Z_2' = Z_2 + (0,-1)$. By Lemma 4.6 (b) and (c) we have that $e(a,a+1;b-2) \in H$. By Lemma 4.6(d), $e(a;b-1,b) \in H$ and $e(a+1;b-1,b) \in H$. Then $e(a;b,b+1) \notin H$ and $e(a+1;b,b+1) \notin H$. If $e(a,a+1;b+1) \in H$, then Z_1 is a switchable box of P so we may assume that $e(a,a+1;b+1) \notin H$. Then $S_{\rightarrow}(a-1,b+1;a,b+2) \in H$ and $S_{\downarrow}(a+1,b+2;a+2,b+1) \in H$. Note that, by definition of A_1, $v(a+1,b)$ is not a corner of A_1. It follows that

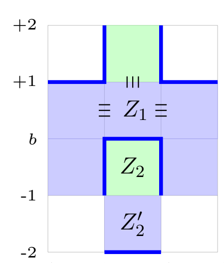

Fig. 4.30. Case 2.2(a).

$e(a+1; b-1, b) \in E(P_0)$. Now, if $Z_2 + (1,0) \in N[P] \setminus P$, then we're back to Case 1, so we may assume that $Z_2 + (1,0) \in P$.

In the same way we find that $v(a,b)$ is not a corner of A_1, $e(a; b-1, b) \in E(P_0)$ and $Z_2 + (-1,0) \in P$. See Figure 4.30. Now, either Z_2' is an end-box of P or it is not.

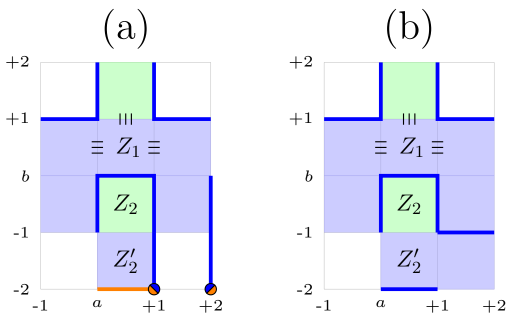

Fig. 4.31. (a) Case 2.2(a_1). (b) Case 2.2(a_2).

CASE 2.2(a_1): Z_2' is an end-box of P. By Lemma 4.6(c), exactly one of $e(a; b-2, b-1)$ and $e(a+1; b-2, b-1)$ belong to H. By symmetry, we may assume WLOG that $e(a+1; b-2, b-1) \in H$. By ((**)), we have that F is northern with $v(a+1, b-2)$ and $v(a+2, b-2)$ being the corners of A_1. But then $Z_2 + (1,0)$ is a switchable box of P. See Figure 4.31 (a). End of Case 2.2(a_1).

CASE 2.2(a_2): Z_2' is not an end-box of P. By Lemma 4.6(b), $Z_2' + (1,0) \in P$ and $e(a+1, a+2; b-1) \in H$. By definition of A_1, $v(a+1, b-1)$, cannot be a corner of A_1. This implies that $e(a+1, a+2; b-1) \in E(P_0)$. But this contradicts Lemma 4.3. See Figure 4.31 (b). End of Case 2.2(a_2). End of Case 2.2(a).

CASE 2.2(b): $Z_2' = Z_2 + (0,-1)$. By Lemma 4.6(d), $e(a, a+1; b-1) \in H$. By Lemma 4.6 (b) and (c), $e(a-1; b-1, b) \in H$. Now, either Z_2' is an end-box of P or it is not.

CASE 2.2(b_1): Z_2' is not an end-box of P. By Lemma 4.6(b), $Z_2' + (0,1) \in P$ and $e(a; b, b+1) \in H$. The assumption of Case 2.2 and the definition of A_1 imply that $v(a,b)$ is not a corner of A_1. But then $e(a; b, b+1) \in E(P_0)$, which contradicts Lemma 4.3. See Figure 4.32 (a). End of Case 2.2(b_1).

CASE 2.2(b_2): Z_2' is end-box of P. By Lemma 4.6(c), exactly one of $e(a-1, a; b-1)$ and $e(a-1, a; b)$ belongs to H. Note that if the former, then we arrive at the same contradiction as we did in Case 2.2(b_1), so we may assume the latter. Now, by ((**)), F is eastern. An argument similar to that in Case 2.2 (a_2) in Lemma 4.10 shows that the corners of A_1 must be $v(a-1, b)$ and $v(a-1, b+1)$.

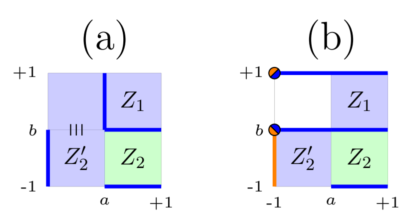

Fig. 4.32. (a) Case 2.2(b_1). (b) Case 2.2(b_2).

But then Z_1 is a middle-box of A_1 in P, contradicting Corollary 4.2 (b). See Figure 4.32 (b). End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2. $\square$

Proposition 4.11. Let H be a Hamiltonian path of a polyomino G and let F be a looping fat path in G. Then F has no polyking junctions whenever:

(i) G is a simply connected polyomino and both end-vertices of H are on $B(G)$, or

(ii) F is sturdy.

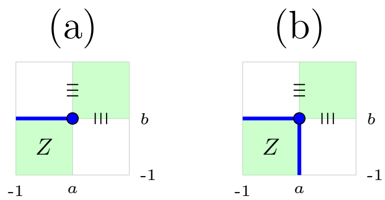

Fig. 4.33. (a) $v(a,b)$ is an end-vertex of H. (b) $S_{\rightarrow}(a-1,b;a,b-1) \in H$.

Proof. The proofs of (i) and (ii) overlap a lot so we combine them. We will prove the contrapositive. Assume that F has a polyking junction. We will show that this assumption leads to a contradiction under (i) and do the same for (ii). We remark that we won't have to use (i) and (ii) until Case 2.

Let $v(a,b)$ be a polyking junction of F with $Z = R(a-1,b-1) \in F$ and $Z+(1,1) \in F$. Either $v(a,b)$ is an end-vertex of H or it is not.

Suppose that $v(a,b)$ is an end-vertex of H. By symmetry, we may assume that $e(a-1,a;b) \in H$. Then $e(a;b,b+1) \notin H$ and $e(a,a+1;b) \notin H$. But now, $Z+(1,1) \in P$ conflicts with FPC-1 and $Z+(1,1) \in N[P] \setminus P$ contradicts Lemma 4.6 (d). See Figure 4.33 (a). It remains to check the case where $v(a,b)$ is not an end-vertex of H.

By symmetry, there are three possibilities to consider: $S_{\rightarrow}(a-1,b;a,b-1) \in H$, $S_{\rightarrow}(a-1,b;a,b+1) \in H$, and the case where $e(a-1,a;b)$ and $e(a,a+1;b)$ belong to H. We note that if $S_{\rightarrow}(a-1,b;a,b-1) \in H$ then $e(a;b,b+1) \notin H$ and $e(a,1+;b) \notin H$. See Figure 4.33 (b). Now, $Z+(1,1) \in N[P] \setminus P$ contradicts Lemma 4.6(d); and if $Z+(1,1) \in P$, then, by Lemma 4.6(a), at least one of $Z+(1,0)$ and $Z+(0,1)$ belongs to P, contradicting that $v(a,b)$ is a polyking junction. It remains to check the other two possibilities.

CASE 1: $S_{\rightarrow}(a-1,b;a,b+1) \in H$. It follows that $e(a;b-1,b) \notin H$ and $e(a,a+1;b) \notin H$. Since $Z+(1,0)$ is not in F, and it is H-adjacent to Z, Z must be in $N[P] \setminus P$. Similarly, $Z+(1,1) \in N[P] \setminus P$. Now Lemma 4.6 (d) implies that $e(a-1,a;b-1) \in H$ and $e(a+1;b,b+1) \in H$. Then we must have $Z+(-1,0) \in P$ and $Z+(1,2) \in P$. But then there is an H-cycle $Z+(-1,0),\ldots,Z+(1,2),Z+(1,1),Z+(1,0),Z,Z+(-1,0)$ which contradicts Proposition 1.2.1. See Figure 4.34. End of Case 1.

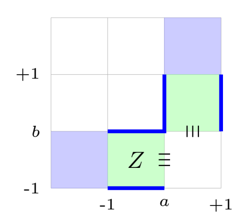

Fig. 4.34. Case 1.

CASE 2: $e(a-1,a;b)$ *and* $e(a,a+1;b)$ *belong to* H. Then $e(a;b-1,b) \notin H$ and $e(a;b,b+1) \notin H$. As in Case 1, we must have that $Z \in N[P] \setminus P$ and $Z+(1,1) \in N[P] \setminus P$. By Lemma 4.6 (d), $e(a-1,a;b-1) \in H$ and $e(a,a+1;b+1) \in H$. Then the H-neighbour of Z in P must be $Z+(-1,0)$ and the H-neighbour of $Z+(1,1)$ in P must be $Z+(2,1)$.

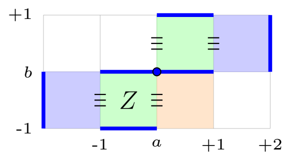

Fig. 4.35.

Assume (i) holds. By Lemma 1.1.11, since G is a simply connected polyomino, G has no polyking junctions. Then at least one of $Z+(1,0)$ and $Z+(0,1)$ must belong to G. By symmetry we may assume without loss of generality that $Z+(1,0)$ belongs to G. Let J_F be the H-component of G that contains F. Then both $Z+(1,1)$ and $Z+(1,0)$ must belong to J_F. Then J_F is self-adjacent. By Lemma 1.3.12, some edge of a main trail of J_F has an end-vertex of H incident on it in $G \setminus B(G)$. But this contradicts the assumption that both end-vertices of H are on $B(G)$. See Figure 4.35. End of proof for (i).

Orient H. Let $v(a+1,b+1) = v_s$ and $v(a+1,b) = v_t$. Without loss of generality assume that $s < t$. Either $v(a,b) = v_{t+1}$ or $v(a,b) = v_{t-1}$.

CASE 2.1: $v(a,b) = v_{t+1}$. Let Q be the cycle in G consisting of the subpath $P(v_s, v_t)$ of H and the edge $\{v_s, v_t\} \in G \setminus H$, and let U be the region bounded by Q. Now, either $Z+(2,1) \in U$, or $Z+(2,1) \notin U$.

CASE 2.1(a): $Z+(2,1) \in U$. Consider the H-subpath $P(Z_1, Z_r)$ of $P(X,Y)$, where $Z_1 = Z+(2,1)$ and $Z_r = Z+(-1,0)$. For $j \in \{1,\dots,r\}$ let c_j be the center of the box Z_j. Note that the edges (v_t, v_{t+1}) and (v_{t+1}, v_{t+2}) of H do not belong to Q. Then the segment $[c_1, c_r]$ (red in Figure 4.36) intersects Q exactly once at the edge $\{v_s, v_t\}$. By Corollary 1.1.5, c_1 and c_r are on different sides of Q. This implies that Z_r is in $G \setminus U$. Consider the polygonal path $P(c_1, c_r)$ in the plane that has as edges the segments $[c_1, c_2], \dots, [c_{r-1}, c_r]$. Note that each edge $[c_j, c_{j+1}]$ of $P(c_1, c_r)$ bisects the gluing edge between Z_j and Z_{j+1}, and intersects no other edge of G. Since c_1 and c_r are on different sides of Q, by JCT, $P(c_1, c_r)$ intersects Q. Then $P(c_1, c_r)$ either intersects Q at an edge of $P(v_s, v_t)$ or at $\{v_s, v_t\}$. The former is not possible, because gluing edges of $P(Z_1, Z_r)$ are not in H. The latter implies that $Z+(1,1) \in P$, contradicting the assumption that $Z+(1,1) \in N[P] \setminus P$. End of Case 2.1(a).

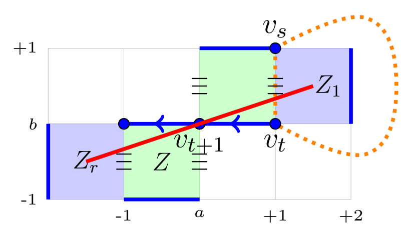

Fig. 4.36. Case 2.1(a). Q in orange; the segment $[c_1, c_r]$ in red.

CASE 2.1(b): $Z+(2,1) \notin U$. The proof uses the same argument as above, so we omit it. End of Case 2.1(b). End of Case 2.1.

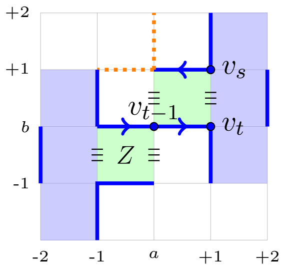

Fig. 4.37. Case 2.2(a).

CASE 2.2: $v(a,b) = v_{t-1}$. Assume that (ii) holds. Note that at least one of $Z+(-1,0)$ and $Z+(2,1)$ is not an end-box of P. By symmetry we may assume WLOG that $Z+(-1,0)$ is not an end-box of P. By Lemma 4.6 (b) and (c), $Z+(-1,1) \in P$, $Z+(-1,-1) \in P$, $e(a-2;b-1,b) \in H$, and $e(a+2;b,b+1) \in H$. Now, either $Z+(2,1)$ is an end-box of $P(X,Y)$, or it is not.

CASE 2.2(a): $Z+(2,1)$ is not an end-box of $P(X,Y)$. Then $Z+(2,2) \in P$ and $Z+(2,0) \in P$. If $Z+(1,1)$ is parallel, then, after $\text{Sw}(Z+(1,1))$, $Z+(2,1) \in P$ is switchable, violating FPC-4, so we may assume that $Z+(1,1)$ is anti-parallel. Note that this implies that $v(a,b+1)$ is not an end-vertex of H. Then exactly one of $e(a-1,a;b+1)$ and $e(a;b+1,b+2)$ belongs to H. If the former, then $Z+(1,1) \mapsto Z$ is a valid double-switch move after which $Z+(-1,0) \in P$ is switchable; and if the latter, then $Z+(1,1) \mapsto Z+(1,2)$, is a valid flip move after which $Z+(2,1) \in P$ is switchable. See Figure 4.37. End of Case 2.2(a).

CASE 2.2(b): $Z+(2,1)$ is an end-box of $P(X,Y)$. By Lemma 4.6(c), exactly one of $e(a+1,a+2;b)$ and $e(a+1,a+2;b+1)$ belongs to H.

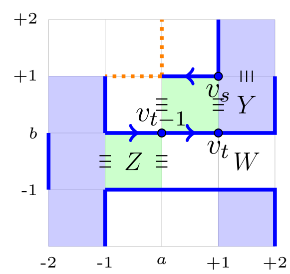

Fig. 4.38. Case 2.2(b_1).

CASE 2.2(b_1): $e(a+1,a+2;b) \in H$ and $e(a+1,a+2;b+1) \notin H$. Then F can be northern or western. By (∗∗) we may assume that F is western. Then the switchable middle-box of F is $W = Z+(2,0)$. Let $Z+(2,1)$ be the end-box Y of $P(X,Y)$. Note that $e(a+1,a+2;b+1) \notin H$, so $e(a+1;b+1,b+2) \in H$. If $Z+(1,1)$ is parallel then after $\text{Sw}(Z+(1,1))$, $W \mapsto Y$ is a valid flip move, violating FPC-4. So we may assume that $Z+(1,1)$ is anti-parallel. See Figure 4.38. As before, this implies that $v(a,b+1)$ is not an end-vertex of H. Then exactly one of $e(a-1,a;b+1)$ and $e(a;b+1,b+2)$ belongs to H. If the former, then $Z+(1,1) \mapsto Z$ is a valid move after which $Z+(-1,0) \in P$ is switchable; and if the latter, then after $Z+(1,1) \mapsto Z+(1,2)$, $W \mapsto Y$ is a valid flip move, violating FPC-4. End of Case 2.2(b_1).

CASE 2.2(b_2): $e(a+1,a+2;b) \notin H$ and $e(a+1,a+2;b+1) \in H$. Then F can be northern or western. By (∗∗) we may assume that F is western. Then the switchable middle-box of F is $W = Z+(2,2)$. Let $Z+(2,1)$ be the end-box X of $P(X,Y)$. If $Z+(1,1)$ is parallel then after $\text{Sw}(Z+(1,1))$, $W \mapsto X$ is a valid flip move, violating FPC-4.

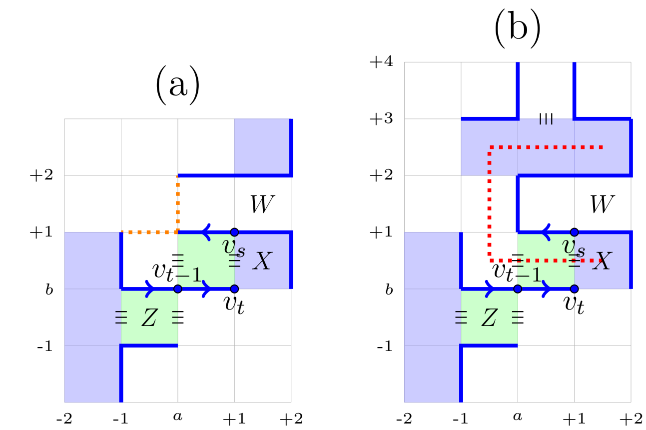

Fig. 4.39. Case 2.2(b_2). (a) $e(a{-}1, a; b{+}1) \in H$. (b) $e(a; b{+}1, b{+}2) \in H$.

So we may assume that $Z + (1, 1)$ is anti-parallel. Once again, this implies that $v(a, b+1)$ is not an end-vertex of H. Then exactly one of $e(a-1, a; b+1)$ and $e(a; b+1, b+2)$ belongs to H.

If the former, then $Z + (1, 1) \mapsto Z$ is a valid double-switch move after which $Z + (-1, 0) \in P$ is switchable (Figure 4.39 (a)), so we may assume the latter (Figure 4.39 (b)). Now, if $e(a, a+1; b+3) \in H$, then $W \mapsto W + (-1, 1)$ is a valid transpose move, violating FPC-4 so we may assume that $e(a, a+1; b+3) \notin H$. Then we must have that $S_{\downarrow}(a+1, b+4; a+2, b+3) \in H$ and that $W + (-1, 1) \in P$. Then $S_{\rightarrow}(a-1, b+3; a, b+4)$ is in H. The H-neighbour of $W + (-1, 1)$ in P cannot be $W + (-1, 2)$, since then $W + (-1, 2)$ would have to be switchable, so the H-neighbour of $W + (-1, 1)$ in P must be $W + (-2, 1)$. Observe that this implies that $W + (-2, 0)$ belongs to G, since all four vertices of $W + (-2, 0)$ are in G. Similarly, $Z + (0, 1)$ must belong to G as well. But then P must be the H-path $X, X + (-1, 0), X + (-2, 0), X + (-2, 1), X + (-2, 2), X + (-1, 2), X + (0, 2)$, contradicting that $Z + (-1, 0) \in P$. See Figure 4.39(b). End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2. □

Definitions. Let G be a polyomino, let H be a Hamiltonian path or cycle of G and let J be an H-subtree of an H-component of G. We define a *shadow edge* of J to be a boundary edge of J that is not in H. We denote the set of all shadow edges of J by $E_{\text{sh}}(J)$.

Suppose that J is a looping fat path $G\langle N[P(X, Y)]\rangle$ that is also non-self adjacent, and has no polyking junctions (but is not necessarily sturdy). Then we say that J is a *standard looping fat path.* We note that shadow edges of J cannot be incident on boxes of P.

Corollary 4.12. Let G be a polyomino, let H be a Hamiltonian path or cycle of G and let F be a looping fat path of G.

(a) If F is a sturdy looping fat path, then F is a standard looping fat path.

(b) If F is a standard looping fat path, then F is a simply connected polyomino.

Proof of (a). Let F be a sturdy looping fat path. By Propositions 4.7 and 4.11, F is non-self-adjacent and has no polyking junctions, so it is a standard looping fat path.

Proof of (b). Let F be a standard looping fat path. By definition, F is an H-subtree in an H-component of G, that is non-self adjacent and has no polyking junctions. By Lemma 1.3.6, it follows that F is a simply connected polyomino.

Lemma 4.13. Let G be a polyomino, let H be a Hamiltonian path of G and let $F = G\langle N[P(X,Y)]\rangle$ be a standard looping fat path in G. Then $E_{\text{sh}}(F) \cup (E(F) \cap E(H))$ is a Hamiltonian cycle of F and the boundary of F.

Proof. Let $E' = E_{\text{sh}}(F) \cup (E(F) \cap E(H))$. It follows from Lemma 4.6 (b), (c) and (d) that each vertex of F must be incident on some edge of $E(F) \cap E(H)$. So, it is enough to show that E' is a cycle.

By definition, F is not self-adjacent and has no polyking junctions. Corollary 4.12 (b), F is a simply connected polyomino. By Corollary 1.1.10, $B(F) = B_0(F)$. By Lemma 1.1.1, $B(F)$ is a cycle. Then the conclusion will follow if we can show that $B(F) = E'$.

Suppose that $e \in B(F)$ and assume that e is incident on boxes $Z \in \text{Boxes}(F)$ and $Z' \in \text{Boxes}(G_{-1}) \setminus \text{Boxes}(F)$. Then $e \in H$ or $e \notin H$. If $e \in H$, since e is incident on the box Z of F, we have that $e \in E(F) \cap E(H) \subset E'$; and if $e \notin H$, by definition, $e \in E_{\text{sh}}(F) \subset E'$.

Suppose that $e \in E'$, and assume that e is incident on boxes Z and Z' of G_{-1}. For definiteness, assume that $Z \in \text{Boxes}(F)$. If $e \notin H$, then e is a shadow edge of F, so by definition, $e \in B(F)$. If $e \in H$, then, by Proposition 4.7, F is non-self-adjacent, so we must have $Z' \in \text{Boxes}(G_{-1}) \setminus \text{Boxes}(F)$. Thus $e \in B(F)$. ☐

Proposition 4.14. Let G be a polyomino, let H be a Hamiltonian path of G and let $F = G\langle N[P(X,Y)]\rangle$ be a standard looping fat path in G following an edge e_F. Then $B(F)$ does not have consecutive colinear edges other than the left and right collinear edges in the A_1-type of F following e_F.

Proof. We need to check that $B(F)$ does not have consecutive colinear edges in the case where one of those colinear edges is one of the left or right colinear edges of the A_1-type of F and in the case where neither of those colinear edges is one of the left or right colinear edges of the A_1-type of F. We divide the proof into Lemmas 4.15 and 4.16.

Lemma 4.15. The shadow of F does not have a pair of consecutive colinear edges in the case where one edge of the pair is one of the left or right colinear edges of the A_1-type that follows e_F.

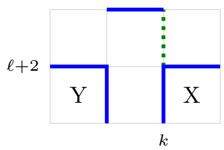

Fig. 4.40. The pair of consecutive colinear edges is $e(k; l+1, l+2)$ and $e(k; l+2, l+3)$.

Proof. For definiteness, assume that F is a standard southern looping fat path following the edge $e(k-1, k; l+3) = e_F$ southward, and let $X = R(k, l+1)$. BWOC assume that $B(F)$ does have consecutive colinear edges and one of those edges is one of the right or left colinear edges of the A_1-type that follows e_F. For definiteness, assume that one of those consecutive colinear edges of $B(F)$ is one of the right colinear edges of the A_1-type that follows e_F. Then either $e(k; l+1, l+2)$, $e(k; l+2, l+3)$ is a pair of consecutive colinear edges or $e(k; l-1, l)$, $e(k; l, l+1)$ is a pair of consecutive colinear edges. If the former, we note that $e(k; l+2, l+3) \in B(F)$ contradicts Lemma 4.13, so we only need to check the latter. See Figure 4.40.

Suppose then that $e(k; l-1, l)$, $e(k; l, l+1)$ is a pair of consecutive colinear edges of $B(F)$. Note that X and $X+(0,-1)$ belong to F. If the edge $e(k; l-1, l)$ is in $B(F)$ then it is either a shadow edge or it belongs to H. We will show that both cases lead to contradictions.

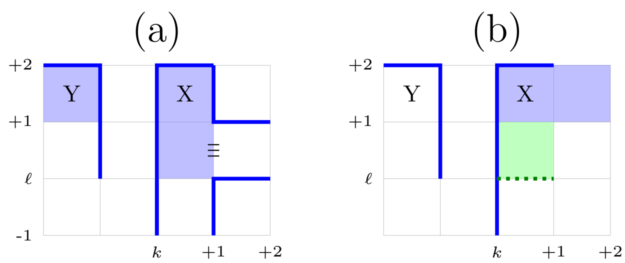

Fig. 4.41.(a) Case 1.1. (b) Case 1.2.

CASE 1: $e(k; l-1, l) \in H$. Then exactly one of $X+(0,-1)$ and $X+(1,0)$ must belong to P.

CASE 1.1: $X+(0,-1) \in P$. Note that if $e(k+1; l, l+1) \in H$, then $X+(0,-1)$ is switchable, so we only need to check the case where $e(k+1; l, l+1) \notin H$. Then $S_{\uparrow}(k+1, l-1; k+2, l) \in H$ and $S_{\downarrow}(k+1, l+2; k+2, l+1) \in H$. Now exactly one of $X+(0,-2)$ and $X+(1,-1)$ belong to P. But, since neither can be the other end-box of P, either one would have to be a switchable box. See Figure 4.41 (a). End of Case 1.1

CASE 1.2: $X+(1,0) \in P$. Then $X+(0,-1) \in N[P] \setminus P$ and, by Lemma 4.6(a), $X+(0,-2) \notin F$. Then $e(k, k+1; l)$ is a shadow edge of F. But then $\deg_{B(F)}(v(k,l)) = 3$, contradicting Lemma 4.13. See Figure 4.41 (b). End of Case 1.2.

CASE 2: $e(k; l-1, l) \in E_{sh}(F)$. Then exactly one of $X-(1,2)$ and $X-(0,2)$ belongs to $N[P] \setminus P$.

If $X-(1,2)\in N[P]\setminus P$ then Lemma 4.13 implies that $e(k-1,k;l)\notin B(F)$, but this is a contradiction to Lemma 4.6(d); and if $X-(0,2)\in N[P]\setminus P$, by Lemma 4.6(d), $e(k,k+1;l)\in H$. But then, since $X-(0,1)$ must belong to F, this contradicts the assumption that F is non-self-adjacent. See Figure 4.42. End of Case 2. □

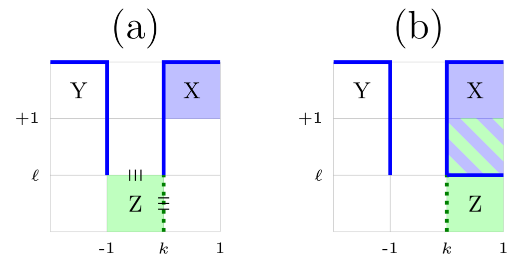

Fig. 4.42. Case 2. (a) $X-(1,2)\in N[P]\setminus P$. (b) $X-(0,2)\in N[P]\setminus P$.

Lemma 4.16. The shadow of F does not have a pair of consecutive colinear edges in the case where neither edge of the pair is one of the left or right colinear edges of the A_1-type that follows e_F.

Proof. For a contradiction, assume that there is a pair of consecutive collinear edges in $B(F)$ where neither edge of the pair is a left or right colinear edge of the A_1-type that follows e_F. For definiteness, we may assume that these edges are the horizontal edges $e(a,a+1;b)$ and $e(a+1,a+2;b)$, and that $Z=R(a,b)$ belongs to F. By Lemma 4.13, $e(a+1;b,b+1)\notin B(F)$, otherwise $\deg_{B(F)}(v(a+1,b))=3$. Now, either $e(a,a+1;b)\in E_{\text{sh}}(F)$, or $e(a,a+1;b)\in H$.

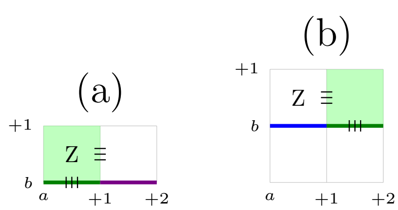

Fig. 4.43. (a) Case 1. (b) Case 2 with $e(a+1,a+2;b)\in E_{\text{sh}}(F)$.

CASE 1: $e(a,a+1;b)$ belongs to $E_{sh}(F)$. Since Z is incident on a shadow edge, Z belongs to $N[P]\setminus P$, and $Z+(0,-1)\notin F$. Since $e(a+1;b,b+1)\notin B(F)$, $e(a+1;b,b+1)$ is not in H either. But then Z can be neither switchable nor a leaf, which contradicts Lemma 4.6 (d). See Figure 4.43 (a). End of Case 1.

CASE 2. $e(a,a+1;b)$ belongs to H. Since F is non-self-adjacent, $Z+(0,-1)\notin F$. Since F is non-self-adjacent, exactly one of $Z+(1,0)$ and $Z+(1,-1)$ belongs to G. Since F has no polyking junctions, it must be the case that $Z+(1,0)\in F$ and $Z+(1,-1)\notin F$. Note that if $e(a+1,a+2;b)\in E_{\text{sh}}(F)$, then $Z+(1,0)$ must belong to $N[P]\setminus P$. But then, as in Case 1, $Z+(1,0)$ can be neither a leaf nor a switchable box, contradicting Lemma 4.6(d). See Figure 4.43 (b). Therefore we may assume that $e(a+1,a+2;b)$ also belongs to H.

Now, if Z and $Z+(1,0)$ both belong to $N[P]\setminus P$, then there is an H-cycle on $Z,Z'\dots,Z'',Z+(1,0)$, where Z' and Z'' are the H-neighbours in P of Z and $Z+(1,0)$, respectively, which contradicts Proposition 1.2.1. It follows, either Z and $Z+(1,0)$ both belong to P, or exactly one of them belongs to P and the other belongs to $N[P]\setminus P$.

CASE 2.1. Z and $Z + (1,0)$ both belong to P. We have that at least one of $e(a, a+1; b+1)$ and $e(a+1, a+2; b+1)$ belongs to H. For definiteness, assume that $e(a, a+1; b+1) \in H$. Now, either $e(a; b, b+1) \notin H$ or $e(a; b, b+1) \in H$. If $e(a; b, b+1) \notin H$ then Z is switchable (Figure 4.44 (a)), so we only need to check the case where $e(a; b, b+1) \in H$. In that case (Figure 4.44 (b)), Z is an end-box of P, and F must be eastern. Note that if $e(a+1, a+2; b+1) \in H$, then $Z + (1,0) \in P$ is switchable, so we may assume that $e(a+1, a+2; b+1) \notin H$. But then, $e(a, a+1; b)$ and $e(a+1, a+2; b)$ must be the right colinear edges of the A_1-type that follows e_F, which contradicts the initial assumption of the Lemma. End of Case 2.1

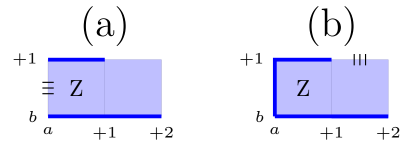

Fig. 4.44. Case 2.1. (a) $e(a; b, b+1) \notin H$. (b) $e(a; b, b+1) \in H$

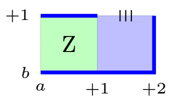

Fig. 4.45. Case 2.2.

CASE 2.2. Exactly one of Z and $Z + (1,0)$ belongs to $N[P] \setminus P$. For definiteness, assume that $Z \in N[P] \setminus P$. By Lemma 4.6 (d), $e(a, a+1; b+1) \in H$. Lemma 4.6 (b) implies that $Z + (1,0)$ is an end-box of P. By Lemma 4.6 (c), $e(a+2; b, b+1) \in H$, and $e(a+1, a+2; b+1) \notin H$. Then F must be western. But then $e(a, a+1; b)$ and $e(a+1, a+2; b)$ must be the left colinear edges of the A_1-type that follows e_F, contradicting the initial assumption of the lemma. See Figure 4.45. $\square$

This completes the proof of Proposition 4.14. An immediate consequence of it is that the A_1-type following e_F is the only A_1-type in F, so we can refer to it as *the* A_1-type of F.

4.2 Turns

In this section we show that every standard looping fat path F must have a turn . We do this by showing that the boundary $B(F)$ of a looping fat path F must have a (necessarily closed) turn and note that this would immediately imply that F must have a turn. We will often use the definition of the shadow of standard southern looping fat path, Proposition 4.14, and the fact that the A_1-type of F is unique in F, and write (DsFP) whenever we appeal to them.

Lemma 4.17. Let H be a Hamiltonian path or cycle of a polyomino G and let $B(F)$ be the boundary of a standard looping fat path F of G. Then $B(F)$ has at least one turn T_1 such that both leaves of T_1 belong to F.

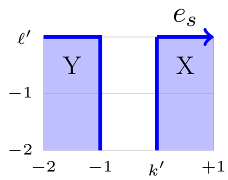

Fig. 4.46. Boxes(F) shaded blue.

Proof. For definiteness, assume that $F = G\langle N[P(X,Y)]\rangle$ is a standard southern looping fat path following $e(k'-1, l'; l'+1)$ southward, with $X = R(k', l'-1)$ and $Y = R(k'-2, l'-1)$. By Lemma 4.13, $B(F)$ is a cycle. Orient $B(F)$ into a directed trail $\overrightarrow{K}$ so that the first edge of $\overrightarrow{K}$ is $(v(k', l'), v(k'+1, l'))$. With this orientation we can give a direction - N, S, E or W - to edges of $\overrightarrow{K}$, defined as the position of the head of an edge relative to its tail. Our choice of direction for the first edge and Lemma 1.3.3 (b) imply that Boxes($\Phi(\overrightarrow{K}, \text{right})$) $\subset$ Boxes(F) so the boxes of F are on the right side of the oriented edges of K. We call this fact (RSK) for reference. See Figure 4.46. We sweep the edges of K in the direction of the orientation starting at $v(k', l')$. We observe that we must encounter at least one west edge e_W, since Y is west of X. Let $e_W = e_0 = e(k-1, k; l)$ be the first west edge encountered, let e_1 be the edge preceding e_0 in the sweep, let e_j be the edge preceding the edge e_{j-1} in the sweep and let $(v(k', l'), v(k'+1, l')) = e_s$.

In this proof we will use the fact that e_W is the first west edge encountered (1stW) often.

By (DsFP), e_0 was immediately preceded by a south edge or a north edge.

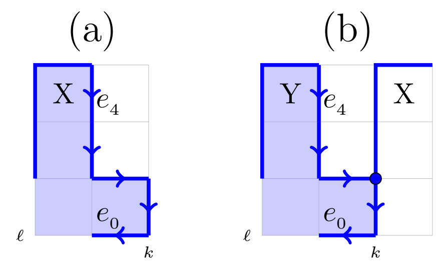

Fig. 4.47. Case 1.1.
(a) $\Phi(e_4, \text{right}) = X$
(b). $\Phi(e_4, \text{right}) = Y$.

CASE 1: e_1 is southern. We shall find a northeastern turn. By (DsFP) and (1stW), the preceding edge e_2 must be eastern; By (DsFP) and (1stW), e_3 has to be southern. By (1stW) e_4 cannot be western. Then e_4 is southern or e_4 is eastern.

CASE 1.1: e_4 is southern. Then, by (DsFP) and (RSK), we have that $\Phi(e_4, \text{right}) = X$ or $\Phi(e_4, \text{right}) = Y$. But the former contradicts Proposition 4.14, and the latter implies that $\deg_{B(F)}(v(k, l+1) = 3$, contradicting Lemma 4.13. See Figure 4.47. Thus, e_4 must be eastern. End of Case 1.1.

CASE 1.2: e_4 is eastern. By (DsFP), e_5 is not eastern. Then e_5 is northern or e_5 is southern.

CASE 1.2(a): e_5 is northern. Then there is a northeastern turn T_1 on the edges $e_0, \ldots, e_5$ with both leaves contained in F. See Figure 4.48 (a). End of Case 1.2(a).

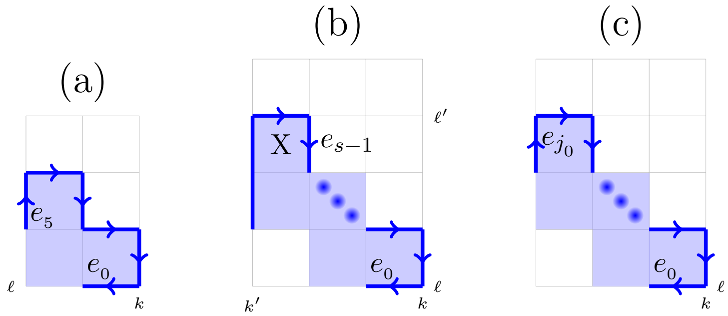

Fig. 4.48. (a) Case 1.2(a). (b) Case 1.2(b_1). (c) Case 1.2(b_2).

CASE 1.2 (b): e_5 is southern. Let $Q(j)$ be the statement: "e_j is southern and e_{j+1} is eastern". Now, either $Q(j)$ is true for each $j \in \{1, 3, \ldots, s-1\}$ (Case 1.2(b_1)), or there is some $j_0 \in \{5, 9, \ldots, s-1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true (Case 1.2(b_2)).

CASE 1.2(b_1). Then we have a northeastern turn T_1 on the edges $e_0, \ldots, e_s, e(k'; l'-1, l')$ with both leaves contained in F. Figure 4.48 (b). End of Case 1.2(b_1).

CASE 1.2(b_2). By (DsFP), e_{j_0} is not eastern. If e_{j_0} is southern, then we run into the same contradiction as in Case 1.1; and if e_{j_0} is northern then we have a northeastern turn T_1 on the edges $e_0, \ldots, e_{j_0}$ with both leaves contained in F. See Figure 4.48 (c). End of Case 1.2(b_2). End of Case 1.2(b). End of Case 1.2.

CASE 2: e_1 is northern. We shall find a southeastern turn with both leaves contained in F. By (1stW) and (DsFP), e_2 is eastern. By (DsFP), e_3 is northern. Note that (RSK) and Lemma 4.6 (d) imply that $\Phi(e_1, \text{right}) \in P$ and $e_1 \in H$. By Lemma 4.13, $e(k, k+1; l-1) \notin H$ and $e(k, k+1; l) \notin H$. If $e(k+2; l-1, l) \in H$, then $\Phi(e_1, \text{right})$ is switchable and in P, contradicting FPC-3, so we may assume that $e(k+2; l-1, l) \notin H$. Then we must have that $S_{\downarrow}(k+1, l+1; k+2, l) \in H$, $S_{\uparrow}(k+1, l-2; k+2, l-1) \in H$ and that $e(k+2; l-1, l) \in B(F)$.

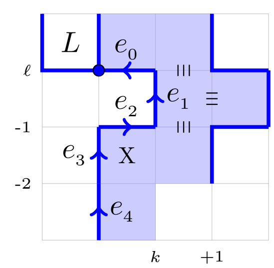

Fig. 4.49. Case 2.

By (1stW), e_4 is not western. If e_4 is northern then (DsFP) and (RSK) imply that $\Phi(e_3, \text{right}) = X$. But then $L = R(k-1, l)$ and then $\deg_H(v(k-1, l)) = 3$, contradicting that H is Hamiltonian. Then e_4 must be eastern. By (DsFP), e_5 is not eastern. Then e_5 is southern or northern. See Figure 4.49.

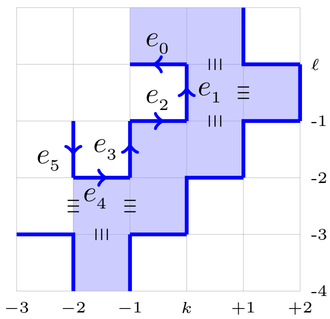

Fig. 4.50. Case 2.1.

CASE 2.1: e_5 is southern. By (RSK) and Lemma 4.6 (d), we have that $\Phi(e_4, \text{right}) \in P$, and that $e_4 \in H$. By Lemma 4.13, $e(k-2; l-3, l-2) \notin H$ and $e(k-1; l-3, l-2) \notin H$. If $e(k-2, k-1; l-3) \in H$, then $\Phi(e_4, \text{right})$ is switchable and in P, contradicting FPC-3, so we may assume that $e(k-2, k-1; l-3) \notin H$. It follows that $S_{\rightarrow}(k-3, l-3; k-2, l-4) \in H$ and that $S_{\uparrow}(k-1, l-4; k+2, l-1) \in H$. Then there is a southeastern turn on $e(k-2; l-4, l-3)$, $S_{\uparrow}(k-1, l-4; k+2, l-1)$, $e(k+1, k+2; l)$ with both leaves in F. See Figure 4.50 (b). End of Case 2.1

CASE 2.2: e_5 is northern. Let $Q(j)$ be the statement: "e_j is northern and e_{j+1} is eastern". Now, either $Q(j)$ is true for each $j \in \{1, 3, \ldots, s-1\}$ (Case 2.2(a)), or there is some $j_0 \in \{5, 7, \ldots, s-1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true (Case 2.2(b)).

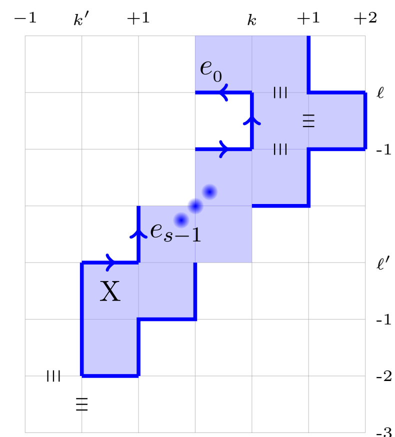

Fig. 4.51. Case 2.2(a).

CASE 2.2(a). Then (DsFP) and (RSK) imply that $\Phi(e_s, \text{right}) = X$. This means that $e(k'; l'-2, l'-1) \in H$, $e(k'; l'-1, l') \in H$, and that $\Phi(e_{s-1}, \text{right}) \in F$. By Proposition 4.14, $e(k'; l'-3, l'-2) \notin H$. Note that if $e(k'-1, k'; l'-2) \in H$, then $P(X, Y)$ is the H-path $X, X+(0,-1), X+(0,-2), X+(-1,-2), X+(-2,-2), X+(-2,-1), Y$. This contradicts our finding that $\Phi(e_{s-1}, \text{right}) \in F$. Then it must be the case that $e(k'-1, k'; l'-2) \notin H$. Then we must have $e(k', k'+1; l'-2) \in H$. It follows that $S_{\uparrow}(k'+1, l'-2; k+2, l-1) \in H$. See Figure 4.51. Then there is a southeastern turn T_1 on $e(k'; l'-2, l'-1)$, $S_{\rightarrow}(k'+1, l'-2; k+2, l-1)$, $e(k+1, k+2; l)$ with both leaves contained in F (see Figure 4.72). End of Case 2.2(a).

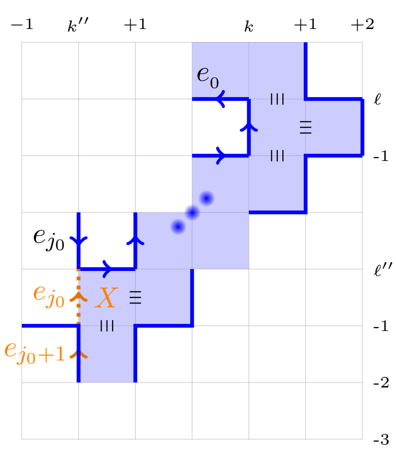

Fig. 4.52. Case 2.2(b).

CASE 2.2(b). By (DsFP), e_{j_0} is not eastern. Suppose that e_{j_0} is northern (in orange in Figure 4.25). The assumption that $Q(j_0)$ is false implies that e_{j_0+1} is not eastern and (1stW) implies that e_{j_0+1} is not western. It must be the case that e_{j_0+1} is also northern. By (DsFP) and (RSK), we have that $\Phi(e_{j_0}, \text{right}) = X$. But then $j_0 - 1 = s$, contradicting the assumption that $j_0 \in \{5, \ldots, s-1\}$ (in orange in Figure 4.52). Thus e_{j_0}

cannot be northern. It follows that e_{j_0} is southern. Let $e_{j_0} = e(k''; l'', l''+1)$.

Using the same arguments as in Case 2.1, we find that $e(k''; l''-1, l'') \notin H$, $e(k''+1; l''-1, l'') \notin H$, $e(k'', k''+1; l''-1) \notin H$, and that $\Phi(j_0 - 1, \text{right}) \in P$. It follows that $S_{\rightarrow}(k''-1, l''-1; k'', l''-2) \in H$ and that $S_{\uparrow}(k''+1, l''-2; k+2, l-1) \in H$. Then there is a southeastern turn T_1 on $e(k''; l''-2, l''-1)$, $S_{\uparrow}(k''+1, l''-2; k+2, l-1)$, $e(k+1, k+2; l)$ with both leaves contained in F (in blue in Figure 4.52). End of Case 2.2(b). End of Case 2.2. End of Case 2. □

Recall once more the setup from the sketch of the proof of Lemma 3.13 at the start of this chapter: there is an H-path $P(X, Y)$ following a leaf L contained in a looping fat path F. We then promised to show that F must have a turn T, and that we can find a cascade (called a *weakening*) that collects one of its leaves. As we shall see in Sections 4.3-4.6 we cannot control which leaf of the turn is collected. In some cases this leads to unwanted behavior. Recall that one of the steps toward our goal is to make $P(X, Y)$ gain a switchable box. In particular, if one of the turn's leaves happens to be an end-box of $P(X, Y)$, then the weakening does not give a useful outcome: instead of allowing us to extend the weakening and collect L (the original goal), it can cement L into a stair subgraph. To avoid such situations, we restrict attention to well-behaved turns, called *admissible turns*, defined below.

Definition. Let G be an $m \times n$ grid graph, let H be a Hamiltonian cycle of G, and let F be a standard looping fat path in G. We say that a turn T of $B(F)$ is *admissible* if:

(i) no leaf of T is an end-box of F, and

(ii) both leaves of T belong to F.

Lemma 4.18. Let H be a Hamiltonian path or cycle of a polyomino G and let $B(F)$ be the boundary of a standard looping fat path of G. Then $B(F)$ has an admissible turn.

Proof. Let F, X, Y, $\overrightarrow{K}$ and $e_W, e_1, \ldots e_s$ be as in Lemma 4.17, including the assumption that F is southern and (RSK).

CASE 1: e_1 is southern. By Case 1 in Lemma 4.17, there is a northeastern turn T_1. We continue sweeping $\overrightarrow{K}$, beginning from e_W, until we find the first northern edge e_N in the subtrail $\overrightarrow{K}(e_W, e_N)$ of $\overrightarrow{K}$, where $e_N = (\widehat{k}; \widehat{l}, \widehat{l}+1) = \widehat{e_0}$. We write (1stN) to refer to the fact that e_N is the first northern edge encountered after e_W, whenever we appeal to it. Let $\widehat{e_1}$ be the edge preceding $\widehat{e_0}$ in the sweep, let $\widehat{e_j}$ be the edge preceding the edge $\widehat{e_{j-1}}$ in the sweep and let $\widehat{e_{t+1}} = e_W$. Then $\widehat{e_1}$ is western or $\widehat{e_1}$ is eastern.

Before we consider each case, we will check that the subtrail $\overrightarrow{K}(\widehat{e_0}, \widehat{e_t})$ of $\overrightarrow{K}$ does not contain the right or left colinear edges of the A_1-type of F. To this end, we will translate H by $(-k', -l')$ to simplify calculations. (DsFP) and (1stW) imply that for every eastern edge in the subtrail $\overrightarrow{K}(e_s, e_1)$ of $\overrightarrow{K}$ there is at most one northern or southern edge. Denote by v_{end} the head of the edge e_W. The assumption that e_1 is southern and the fact that a shortest turn has length two imply that v_{end} is contained in the region $U_{1,\text{end}}$, determined by $x \geq 1$ and $|y+2| \leq x-1$ (Eq.1). Let $v_{\text{end}} = v(a,b)$. It follows that $\overrightarrow{K}(\widehat{e_t}, \widehat{e_0})$ is contained in the region U_2 bounded by $y \leq b+1$ and $|x-a| \leq b+1-y$. (Eq.2)

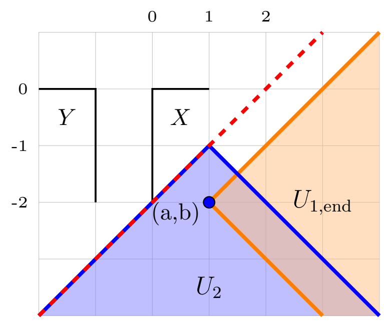

Fig. 4.53. Case 1. The line $y-x+2=0$ in red; $U_{1,\text{end}}$ shaded orange, U_2 shaded blue.

We will check that U_2 and the colinear edges of the A_1-type of F lie on two different sides of the line $y = x-2$. See Figure 4.53. By (Eq.1) we have that $b \leq a-3$ and by (Eq.2) we have that $y \leq x-a+b+1$. Let $(x,y) \in U_2$. Then $y-x+2 \leq -a+b+3 \leq 0$, so U_2 lies below the line $y-x+2$. Plugging in the values of the coordinates of the vertices of the A_1-type, we see that they lie above the line $y = x-2$. This shows that $\overrightarrow{K}(\widehat{e_t}, \widehat{e_0})$ does not contain colinear edges. We will write (NCE) whenever we appeal to this fact. Note that (NCE) implies (i).

CASE 1.1: $\widehat{e_1}$ is western. Note that $\widehat{e_1} \neq e_W$, otherwise we get a cycle on e_1, e_W, e_1, e_2. By (DsFP), $\widehat{e_2}$ is not western and by (1stN), $\widehat{e_2}$ is not northern, so $\widehat{e_2}$ must be southern. By (NCE) $\widehat{e_3}$ cannot be southern. Then $\widehat{e_3}$ must be western. If $\widehat{e_3} = e_W$, then we have a southeastern turn T_2 on $\widehat{e_0}, \ldots, \widehat{e_3}, e_1, e_2$, satisfying (i) and, by (RSK), (ii) (in orange in Figure 4.54), so we may assume that $\widehat{e_3} \neq e_W$. By (DsFP), $\widehat{e_4}$ is not western and by (1stN), $\widehat{e_4}$ is not northern. Then $\widehat{e_4}$ is southern.

Let $Q(j)$ be the statement: "$\widehat{e_j}$ is western and $\widehat{e_{j+1}}$ is southern". Now, either $Q(j)$ is true for each $j \in \{1, 3, \ldots, t+1\}$, or there is some $j_0 \in \{5, 7, \ldots, t+1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true. If the former, then we have a southeastern turn T_2 on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_t}, e_W, e_1, e_2$ satisfying (i) and (ii) (in blue in Figure 4.54), so assume the latter. By (NCE), $\widehat{e_{j_0}}$ is not southern. If $\widehat{e_{j_0}}$ is eastern then we have a southeastern turn on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{j_0}}$ satisfying (i) and (ii) (blue in Figure 4.75). Suppose then that $\widehat{e_{j_0}}$ is

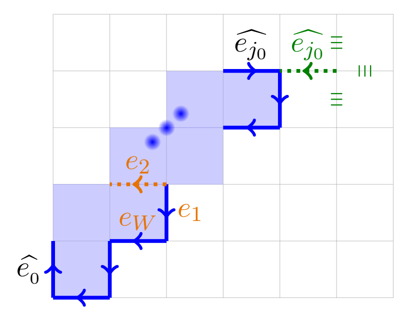

Fig. 4.54. Case 1.1.

western (in green in Figure 4.54). This is impossible: by (1stN). $\widehat{e_{j_0+1}}$ is not northern; since $Q(j_0)$ is false, $\widehat{e_{j_0+1}}$ is not southern; and by (DsFP), $\widehat{e_{j_0+1}}$ is not western. End of Case 1.1

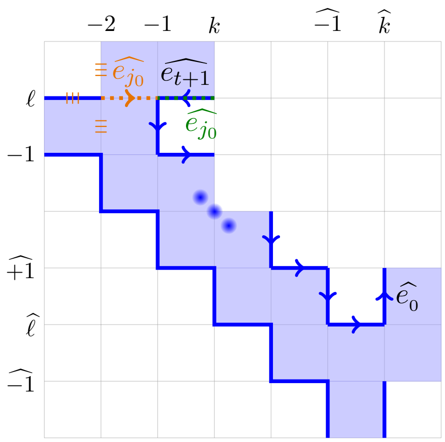

Fig. 4.55. Case 1.2(a) and (b).

CASE 1.2: $\widehat{e_1}$ is eastern. By (DsFP), $\widehat{e_2}$ is not eastern and by (1stN), $\widehat{e_2}$ is not northern, so $\widehat{e_2}$ must be southern. By (NCE) $\widehat{e_3}$ is not southern and by (DsFP), $\widehat{e_3}$ is not western. Then $\widehat{e_3}$ must be eastern. (DsFP) and (1stN) imply that $\widehat{e_4}$ must be southern.

Let $Q(j)$ be the statement: "$\widehat{e_j}$ is eastern and $\widehat{e_{j+1}}$ is southern". Now, either $Q(j)$ is true for each $j \in \{1, 3, \ldots, t-1\}$, or there is some $j_0 \in \{5, 7, \ldots, t-1\}$ such that $Q(j)$ is true for each odd $j < j_0$, but $Q(j_0)$ is not true.

CASE 1.2(a): $Q(j)$ is true for each $j \in \{1, 3, \ldots, t-1\}$. Then we have a southwestern turn on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{t-1}}, \widehat{e_t}, e_W$. Recall that $e_W = e(k-1, k; l)$. Observe that $R(k-2; l-1) \in P$, so $e(k-2; l-1, l) \notin H$.
Similarly, $R(\widehat{k}-1, \widehat{l}-1) \in P$ and $e(\widehat{k}-1, \widehat{k}; \widehat{l}-1) \notin H$. It follows that $R(k-3, l-1) \in F$, $R(\widehat{k}-1, \widehat{l}) \in F$ and that there is a southeastern
turn T_2 on $e(k-3, k-2; l)$, $S_{\rightarrow}(k-3, l-1; \widehat{k}-1, \widehat{l}-2)$, $e(\widehat{k}; \widehat{l}-2; \widehat{l}-1)$ satisfying (i) and (ii) (in blue in Figure 4.55). End of Case 1.2(a).

CASE 1.2(b): There is some $j_0 \in \{5, 7, \ldots, t-1\}$ such that $Q(j)$ is true for each odd $j < j_0$, but $Q(j_0)$ is not true. If $\widehat{e_{j_0}}$ is western then we have a southeastern turn on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{j_0}}$. As in Case 1.2 (a), this turn satisfies (i) and (ii) (in blue in Figure 4.55, with e_{j_0} dotted green).

By (DsFP), $\widehat{e_{j_0}}$ is not southern. Suppose then $\widehat{e_{j_0}}$ is eastern. This is impossible: by (1stN), $\widehat{e_{j_0+1}}$ is not northern; since $Q(j_0)$ is false, $\widehat{e_{j_0+1}}$ is not southern; and by (DsFP), $\widehat{e_{j_0+1}}$ is not eastern (dotted orange in Figure 4.55). End of Case 1.2(b). End of Case 1.2.

CASE 2: e_1 is northern. By Case 2 in Lemma 4.17, $\overrightarrow{K}$ has a southeastern turn T_1. We continue sweeping K, beginning from e_W, until we find the first southern edge e_S in the subtrail $\overrightarrow{K}(e_W, e_S)$ of $\overrightarrow{K}$, where $e_S = (\widehat{k}; \widehat{l}, \widehat{l}+1) = \widehat{e_0}$. We write (1stS) to refer to the fact that e_S is the first southern edge encountered after e_W, whenever we appeal to it. Let $\widehat{e_1}$ be the edge preceding $\widehat{e_0}$ in the sweep, let $\widehat{e_j}$ be the edge preceding the edge $\widehat{e_{j-1}}$ in the sweep and let $\widehat{e_{t+1}} = e_W$. Then $\widehat{e_1}$ is western or $\widehat{e_1}$ is eastern.

CASE 2.1: $\widehat{e_1}$ is western. Note that the assumption that e_1 is northern implies that $\widehat{e_1} \neq e_W$, otherwise there is a cycle $e_2, e_1, e_W, \widehat{e_0}$. By (DsFP), $\widehat{e_2}$ is not western and by 1stS, $\widehat{e_2}$ is not southern, so $\widehat{e_2}$ must be northern. By (DsFP) $\widehat{e_3}$ is not northern or eastern. Then $\widehat{e_3}$ must be western. By 1stS $\widehat{e_4}$ is not southern, and by (DsFP), $\widehat{e_4}$ is not western. Then $\widehat{e_4}$ must be northern. Let $Q(j)$ be the statement: "$\widehat{e_j}$ is western and $\widehat{e_{j+1}}$ is northern. Then either $Q(j)$ is true for each $j \in \{1, 3, \ldots, t-1\}$ or there is some $j_0 \in \{5, 7, \ldots, t-1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true.

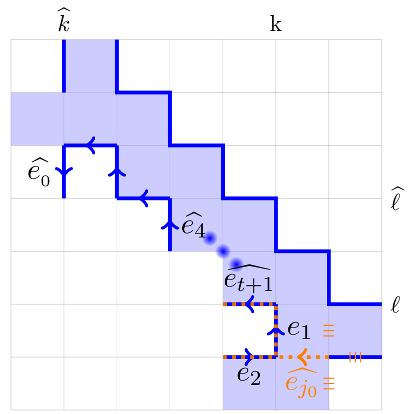

Fig. 4.56. Case 2.1(a) and (b).

CASE 2.1(a): $Q(j)$ is true for each $j \in \{1, 3, \ldots, t-1\}$. Then there is a northeastern turn on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{t-1}}, \widehat{e_t}, e_W, e_1, e_2$. As in Case 1.2 (a), we have that $R(\widehat{k}, \widehat{l}+1) \in P$, $R(k, l-1) \in P$, $e(\widehat{k}, \widehat{k}+1; \widehat{l}+2) \notin H$ and $e(k+1; l-1, l) \notin H$. Then there is a northeastern turn T_2 on $e(\widehat{k}; \widehat{l}+2, \widehat{l}+3)$, $S_{\downarrow}(\widehat{k}+1, \widehat{l}+3; k+2, l)$, $e(k+1, k+2; l-1)$ satisfying (i) and (ii) (in blue in Figure 4.56). End of Case 2.1(a).

CASE 2.1(b): There is some $j_0 \in \{5, 7, \ldots, t-1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true. If $\widehat{e_{j_0}}$ is eastern then we have a northeastern turn on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{j_0}}$. Then, as in Case 2.1(a), there is a northeastern turn T_2 satisfying (i) and (ii).

By (DsFP), $\widehat{e_{j_0}}$ is not southern. Suppose then that $\widehat{e_{j_0}}$ is western. This is impossible: by 1stS. $\widehat{e_{j_0+1}}$ is not southern; since $Q(j_0)$ is false, $\widehat{e_{j_0+1}}$ is not northern; and by (DsFP), $\widehat{e_{j_0+1}}$ is not western (in orange in Figure 4.56). End of Case 2.1. End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $\widehat{e_1}$ is eastern. By 1stS and (DsFP), $\widehat{e_2}$ is northern. By (DsFP), $\widehat{e_3}$ is not western. An argument analogous to (NCE-1) in Case 1 can be used to show that T_2 and the A_1-type lie on two different sides of the line $y = 2 - x$. In this case, we have that the region $U_{1,\text{end}}$ containing v_{end} is determined by $x \geq 1$ and $|y-2| \leq x-1$, and the region U_2 containing T_2, as defined in the next paragraph, is determined by $y \geq b-1$ and $|x-a| \leq y-b+1$. We will refer to this argument as (NCE-2). Note that by (NCE-2), $\widehat{e_3}$ is not northern. Then $\widehat{e_3}$ must be eastern. By (DsFP) and (1stS) $\widehat{e_4}$ is not southern or eastern. Then $\widehat{e_4}$ must be northern.

Let $Q(j)$ be the statement: "$\widehat{e_j}$ is eastern and $\widehat{e_{j+1}}$ is northern". Now, either $Q(j)$ is true for each $j \in \{1, 3, \ldots, t-1\}$, or there is some $j_0 \in \{5, 7, \ldots, t-1\}$ such that $Q(j)$ for each odd $j < j_0$, but $Q(j_0)$ is not true. If the former then we have a northwestern turn T_2 on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{t-1}}, \widehat{e_t}, e_W$ satisfying (i) and (ii), (in blue in Figure 4.57 with $\widehat{e_{t-1}}, \widehat{e_t}, e_W$ dotted orange) so assume the latter. If $\widehat{e_{j_0}}$ is western then again we have a northwestern turn T_2 on $\widehat{e_0}, \widehat{e_1}, \ldots, \widehat{e_{j_0}}$ satisfying (i) and (ii) (in blue in Figure 4.57). By (NCE-2), $\widehat{e_{j_0}}$ is not southern. Then, suppose that $\widehat{e_{j_0}}$ is western (green in Figure 4.57). This is impossible: by 1stS. $\widehat{e_{j_0+1}}$ is not southern; since $Q(j_0)$ is false, $\widehat{e_{j_0+1}}$ is not northern; and by (DsFP), $\widehat{e_{j_0+1}}$ is not eastern. End of Case 2.2. □

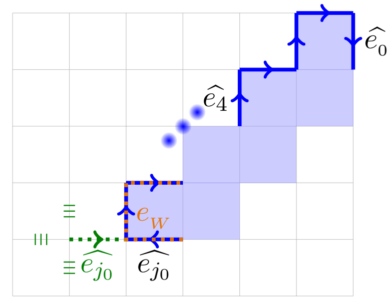

Fig. 4.57. Case 2.2.

Corollary 4.19. Let G be a polyomino, let H be a Hamiltonian cycle or path of G, let $F = G\langle N[P(X,Y)]\rangle$ be a standard looping fat path in G, let T be an admissible turn of F, let L be a leaf of T, and let L' be the H-neighbour of L in F. Then:

(a) $d(T) \geq 3$, and

(b) $L \in N[P] \setminus P$ and $L' \in P$.

Proof of (a). We prove the contrapositive. For definiteness, assume that T is northeastern with northern leaf $L_N = R(k, l-1)$. Suppose that $d(T) < 3$. Then $d(T) = 2$ and the eastern leaf of T must be $L_E = R(k+1, l-2)$. Note that $L_N + (0, -1) \in F$, otherwise $L_E, \ldots, L_N, L_N + (0, -1)$ is an H-cycle of boxes in G, contradicting Proposition 1.2.1.

Now, by Proposition 4.14, $e(k; l-2, l-1)$ and $e(k, k+1; l-2)$ cannot both belong to H. Then, either exactly one of $e(k; l-2, l-1)$ and $e(k, k+1; l-2)$ belongs to H, or neither does.

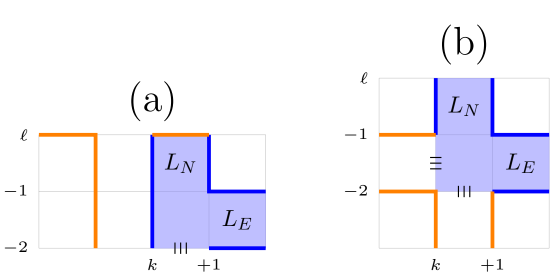

Fig. 4.58. (a) Case 1. (b) Case 2.

CASE 1: exactly one of $e(k; l-2, l-1)$ *and* $e(k, k+1; l-2)$ *belongs to* H. By symmetry, we may assume WLOG that $e(k; l-2, l-1) \in H$ and $e(k, k+1; l-2) \notin H$. Note that the assumption that $e(k, k+1; l-2) \notin H$ implies that F is northern. It follows that L_N is an end-box of $P(X,Y)$. Therefore T is not admissible. See Figure 4.58 (a). End of Case 1.

CASE 2: neither $e(k; l-2, l-1)$ *nor* $e(k, k+1; l-2)$ *belongs to* H. Then $e(k-1, k; l-1)$,

$e(k+1;l-3,l-2)$ and $S_{\rightarrow}(k-1,l-2;k,l-3)$ belong to H. By Lemma 4.6(d), $L_N+(0,-1)$ must belong to $P(X,Y)$. It follows that at least one of L_N, $L_N+(0,-2)$, L_E and $L_E+(-2,0)$ belongs to $P(X,Y)$ and is switchable, contradicting the assumption that F is a looping fat path. See Figure 4.58 (b). End of Case 2. End of proof for (a).

Proof of (b). Let T be an admissible turn. For definiteness, assume that T is northeastern with northern leaf $L_N = R(a,b)$. By Corollary 4.19 (a), $d(T) \geq 3$. Then we have that $e(a,a+1;b-1) \in H$, and that $e(a+1;b-1,b) \notin H$. By (RSK), $L_N+(0,-1)$ and $L_N+(1,-1)$ belong to F. This means that $L_N+(0,-1) = L'_N$ is the H-neighbour of L_N in F. Now, either $e(a,a+1;b+1) \in H$ or $e(a,a+1;b+1) \notin H$. If $e(a,a+1;b+1) \in H$ (Figure 4.59 (a)), then, since L_N is not an end-box of P, $L_N \in N[P] \setminus P$. Then, by Lemma 4.6 (a), $L'_N \in P$. And if $e(a,a+1;b+1) \notin H$ (Figure 4.59 (b)), then L_N is switchable, so $L_N \in N[P] \setminus P$. Since $L'_N \in F$, by Lemma 4.6 (a), $L'_N \in P$. Either way we have that $L'_N \in P$ and and $L_N \in N[P] \setminus P$. End of proof for (b). □

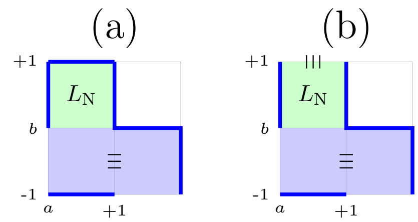

Fig. 4.59. (a) $e(a,a+1;b+1) \in H$. (b) $e(a,a+1;b+1) \notin H$.

In this section we have shown that every standard looping fat path contains at least one admissible turn.

4.3 Turn weakenings.

Definitions. Let H be a Hamiltonian path or cycle of an $m \times n$ grid graph G. Let T be a turn of H on $\{e(k;l-1,l), S_{\downarrow}(k+1,l;k',l'+1), e(k'-1,k';l')\}$ and let L_N and L_E be the northern and eastern leaves of T, respectively. Define the *sector of* T to be the induced subgraph of G bounded by $e(k,k+1;l)$, $S_{\downarrow}(k+1,l;k',l')$, $e(k'-1,k';l')$, and the segments $[(k',l'),(m-1,l')]$, $[(m-1,l'),(m-1,n-1)]$, $[(m-1,n-1),(k,n-1)]$, $[(k,n-1),(k,l)]$, and denote it by $\text{Sector}(T)$. See Figure 4.60. Analogous definitions apply to sectors of southeastern, southwestern and northwestern turns. Define the *end-vertex weakening terminal of* T to be the set of vertices $V(S_{\rightarrow}(k+1,l-1;k'-1,l'+1)) \cup \{v(k,l-1), v(k'-1,l')\}$ and denote it by $\text{ew-Terminal}(T)$. We will call the edges $e(k,k+1;l-1)$ and $e(k'-1;l',l'+1)$ of $G \setminus H$ the *northern* and *eastern terminal edge of* T, respectively.

Define a *weakening* of T to be a cascade $\mu_1, \ldots, \mu_s$ such that μ_s is the first move after which one of the following holds:

1. One of the terminal edges of T is in the resulting Hamiltonian path of G, in which case we say that T has an *edge* weakening, or

2. An end-vertex of the resulting Hamiltonian path of G is incident on ew-Terminal(T), in which case we say that T has an *end-vertex* weakening.

We call a weakening of T consisting of three or less moves a *short weakening of* T. We remark that if H is a Hamiltonian path of G, then cookies are not defined and we need not be concerned with preserving their count after applying $\mu_1, \ldots, \mu_s$. We call the subgraph $S_{\downarrow}(k+1, l; k', l'+1)$ *the stairs-part of* T and denote it by stairs(T). We say that T has a *lengthening* T' if T' is a turn of H such that:

a) $d(T') \geq d(T)$ and

b) stairs(T') $\supseteq$ stairs(T) + (1, 1).

Define a set of lengthenings $\mathcal{T}(T_0)$ as follows:

1. The turn $T_0 \in \mathcal{T}(T_0)$.
2. The turn $T_j \in \mathcal{T}(T_0)$ if and only if T_j is a lengthening of the turn T_{j-1}.

Note that by definition, $\mathcal{T}(T_0)$ is maximal.

Define the *flank of* T to be the set of vertices $V(\text{stairs}(T)+(1,1)) \cup V(L_N+(0,1)) \cup V(L_E+(1,0))$. We call $V(L_N+(0,1)) \cup V(L_E+(1,0))$ the *leaf flank of* T and denote it ℓ-Flank(T), and we call $V(\text{stairs}(T)+(1,1))$ the *stairs flank of* T and denote it s-Flank(T). See Figure 4.60.

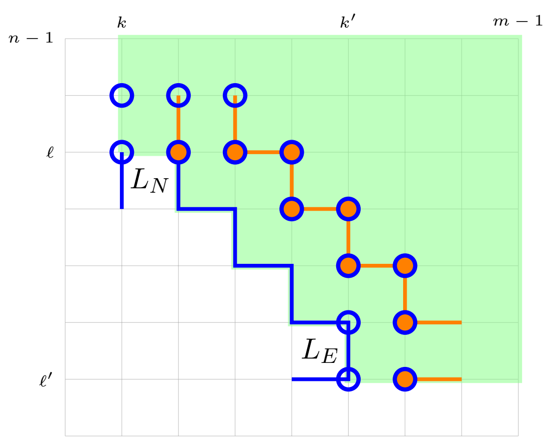

Fig. 4.60. A half-open turn T in blue, its lengthening T' in orange, Flank(T) marked with blue circles, ew-Terminal(T') marked with circles filled in orange, Sector(T) shaded green. s-Flank(T)=$V(S_{\downarrow}(k+2, l+1; k'+1, l'+2))$.

Observation 4.20. Let $F = G\langle N[P(X,Y)]\rangle$ be a standard looping fat path, let T be an admissible turn of F, and let T' be a lengthening of T. Then:

(a) Every vertex of ew-Terminal(T) is incident on a box of $P(X,Y)$.

(b) ew-Terminal(T') $\subset$ Flank(T).

(c) Flank(T) $\subset$ Sector(T).

(d) $\mathcal{T}(T) \subset$ Sector(T).

Remark. Recall from the definition of looping fat paths at the beginning of this chapter that our goal, roughly, is to find a cascade after which the edge e' is in the resulting Hamiltonian path (see Figure 4.3). What happens if there is an end-vertex weakening? By Observation 4.20(a), every vertex of ew-Terminal(T) is incident on a box of $P(X, Y)$. So, after applying the end-vertex weakening, FPC-1 fails in the resulting Hamiltonian path. We show in Chapter 5 (Claims 5.18-5.21) that when FPC-1 fails, we can find a cascade after which e' is in the resulting Hamiltonian path. This explains why attaining an end-vertex weakening is sufficient for our purposes. This type of weakening is an effective way to deal with the Section 4.6 case, that considers Hamiltonian paths with Sector(T) containing at most one end-vertex u with $u \notin R_0$.

Lemma 4.21. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian cycle of G. Let F be a looping fat path of G, anchored at some outermost small cookie C. Then F has an admissible turn T such that Sector(T) and the j-stack of A_0's following C are disjoint.

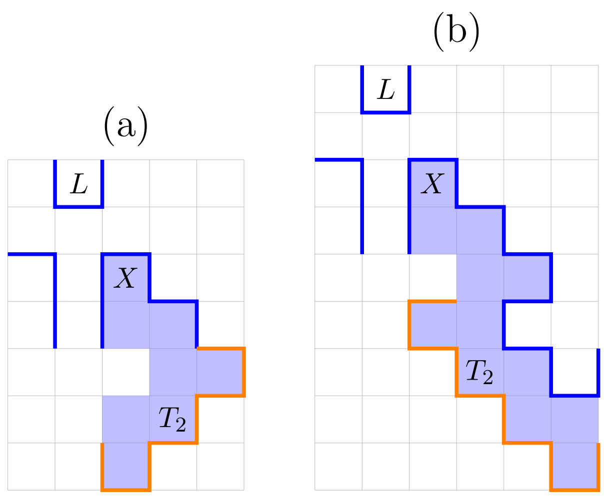

Fig. 4.61. An illustration of Case 1 with T_2: (a) southeastern; (b) southwestern.

Proof. For definiteness, assume that C is a small southern cookie followed by a j-stack of A_0's, which is then followed by the southern looping fat path $F = G\langle N[P(X, Y)]\rangle$, and let $X = R(k', l' - 1)$. Let $\overrightarrow{K}$ and $e_W, e_1, \ldots e_s$ be as in Lemma 4.17. By Lemma 4.17, F has a turn T_1. Either T_1 is a northeastern turn with X as its northern leaf, or it is not.

CASE 1: T_1 is a northeastern turn with X as its northern leaf. Then e_1 is southern. By the proof of Lemma 4.18, T_2 is either southeastern or southwestern. In either case, we note that Sector(T_2) is south of the stack of A_0's. By Lemma 4.18, T_2 is admissible. See Figure 4.61. End of Case 1.

CASE 2: T_1 is not a northeastern turn with X as its northern leaf. By the proof of Lemma 4.17, T_1 is a northeastern turn or T_1 is a southeastern turn.

CASE 2.1: T_1 is a northeastern turn. It follows from the proof of Lemma 4.17 that Sector(T_1) is east of the stack of A_0's, that both leaves of T_1 are in F, and that neither leaf of T_1 is an end-box of $P(X, Y)$. End of Case 2.1. See Figure 4.62 (a).

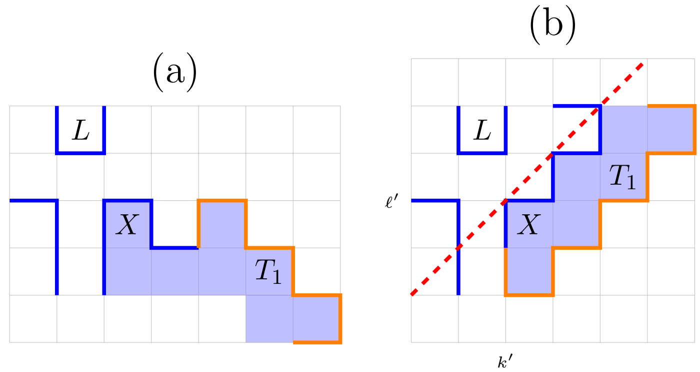

Fig. 4.62. (a) An illustration of Case 2.1. (b) An illustration of Case 2.2; $y = x - k' + l'$ in red.

CASE 2.2: T_1 is a southeastern turn. It follows from the proof of Lemma 4.17 that Sector(T_1) is southeast of the stack of A_0's. More precisely, we can check that Sector(T_1) is below the line $y = x+(l'-k')$, and the j-stack of A_0's is above the line $y = x + (l' - k')$. By the proof of Lemma 4.17, we have that both leaves of T_1 are in F, and that no leaf of T_1 is an end-box of $P(X, Y)$. See Figure 4.62 (b). End of Case 2.2. End of Case 2. □

Corollary 4.22. Let H be a Hamiltonian path of an $m \times n$ grid graph G, and let $e_F = e(k, k+1; l)$. Assume that e_F is followed by a southern j-stack of A_0's, which is then followed by a standard southern looping fat path $F = G\langle N[P(X, Y)]\rangle$. Then F has an admissible turn T such that Sector(T) avoids the j-stack of A_0's as well as all boxes incident on e_F. Furthermore, T may be chosen to lie below the line $y = x - k + l - 2j - 2$, or below the line $y = -x + k + l - 2j - 1$.

Proof. One such turn T lying below the line $y = x-k+l-2j-2$ can be found by inspecting the turns identified in Lemma 4.21. By the same argument and symmetry, we may find another turn that lies below the line $y = -x + k + l - 2j - 1$. □

In the next three sections we show that every looping fat path contains an admissible turn T that has a weakening. We will need three separate versions of this fact addressing three distinct cases:

1. H is a Hamiltonian cycle (Section 4.4);
2. H is a Hamiltonian path and Sector(T) contains at most one end-vertex u with $u \in R_0$ (Section 4.5); and
3. H is a Hamiltonian path and Sector(T) contains at most one end-vertex u with $u \notin R_0$ (Section 4.6).

The first will be needed for the proof of Lemma 3.13 while thee second and third are used in Chapter 5.

4.4 Weakenings when H is a Hamiltonian cycle of G

Lemma 4.23. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G, and let T be a turn in H with $d(T) \geq 3$. Then:

I. T has a short weakening or T has a lengthening.

II. If T' is a lengthening of T and T' has a weakening of length at most s, then T has a weakening of length at most $s+1$, with $s+1 \leq \min(m,n)$.

We prove Lemma 4.23 after we use it to prove Proposition 4.24.

Proposition 4.24. Let H be a Hamiltonian cycle of an $m \times n$ grid graph G, and let T be a turn in H with $d(T) \geq 3$. Then T has a weakening of length at most $\min(m,n)$.

Proof. Let $T = T_0$ be a turn of H with $d(T) \geq 3$. If T_0 has a short weakening, then we're done, so we assume T_0 has no short weakening. By I in Lemma 4.23, T_0 has a lengthening T_1. So, $T_1 \in \mathcal{T}$, where $\mathcal{T} = \mathcal{T}(T_0)$. Since $m, n < \infty$, we have that $|\mathcal{T}| < \infty$. Let $\mathcal{T} = \{T_0, T_1, \ldots, T_j\}$. Then T_j has no lengthening; thus, by I of Lemma 4.23, it must have a short weakening. Then, by induction and II on Lemma 4.23, T_0 has a weakening. The bound follows immediately. □

Proof of Lemma 4.23. We first remark that none of the moves we use throughout this proof fit the description of the moves in Observation 3.4 (i) and (ii) in Section 3. We will use this fact repeatedly and implicitly.

Let H be a Hamiltonian cycle of G and let T be a turn of H with $d(T) \geq 3$. For definiteness, assume that T is northeastern and that T is on $\{e(k; l-1, l), S_{\downarrow}(k+1, l; k', l'+1), e(k'-1, k'; l')\}$. Let L_{N} be the northern leaf of T and let L_{E} be the eastern leaf of T. Since $d(T) \geq 3$, $m-1 \geq k+3$ and $0 \leq l-3$. L_{N} can be open or closed, so there are two cases to check.

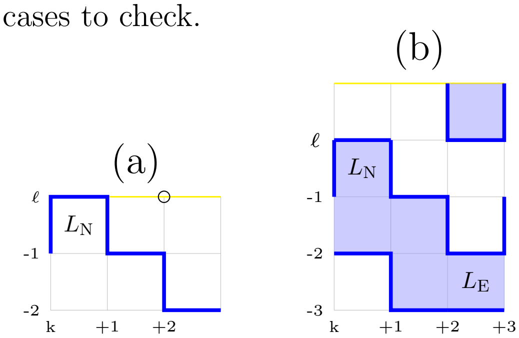

Fig. 4.63. (a) Case 1. (b) Case 1.1.

CASE 1: L_N is closed. Proof of I. First we note $n-1 \neq l$, otherwise H misses $v(k+2, l)$. Then we must have $S_{\downarrow}(k+2, l+1; k+3, l) \in H$. See Figure 4.63 (a). Now, $n-1 = l+1$, $n-1 = l+2$, or $n-1 \geq l+3$.

CASE 1.1: n-1=l+1. By Corollary 1.3.15(b), $L_N + (0,1) \in \mathrm{int}(H)$. This implies that $L_N + (2,1)$ is a small cookie of H, so $e(k+3; l, l+1) \in H$. Then $e(k+3; l-2, l-1) \in H$. It follows that $L_N + (2,-2) = L_E$. But then $L_E \mapsto L_E + (0,1)$ is a short weakening of T. See

Figure 4.63 (b). End of Case 1.1.

CASE 1.2: n-1= l+2. Either $e(k, k+1; l+1) \in H$ or $e(k, k+1; l+1) \notin H$.

CASE 1.2(a): $e(k, k+1; l+1) \in H$. Either $L_N + (0,2) \in \text{int}(H)$ or $L_N + (0,2) \in \text{ext}(H)$. If $L_N + (0,2) \in \text{int}(H)$, then $L_N + (0,1) \mapsto L_N$ is a short weakening of T. See Figure 4.64 (a). Suppose then that $L_N + (0,2) \in \text{ext}(H)$. This implies that $L_N + (0,2)$ is a small cookie. Then $L_N + (0,1) \mapsto L_N + (0,2)$, $L_N \mapsto L_N + (1,1)$ is a short weakening. See Figure 4.64 (b). End of Case 1.2(a).

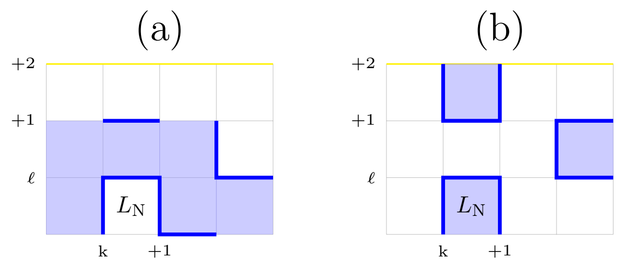

Fig. 4.64. Case 1.2(a). (a) $L_N + (0,2) \in \text{int}(H)$. (b) $L_N + (0,2) \in \text{ext}(H)$.

CASE 1.2(b): $e(k, k+1; l+1) \notin H$. Then $S_{\rightarrow}(k-1, l+1; k, l+2) \in H$ and $S_{\downarrow}(k+1, l+2; k+2, l+1) \in H$. Note that if $e(k, k+1; l+2) \in H$, then H misses $v(k+2, l+2)$ (Figure 4.65 (a)), so we may assume that $e(k, k+1; l+2) \notin H$ (Figure 4.65 (b)). It follows that $e(k+1, k+2; l+2) \in H$. Then $L_N + (0,2) \mapsto L_N + (1,2)$[1], $L_N + (0,1) \mapsto L_N$ is a short weakening. End of Case 1.2(b). End of Case 1.2.

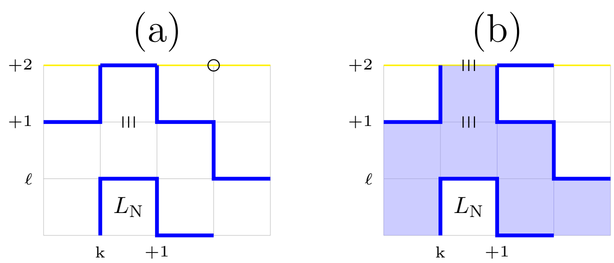

Fig. 4.65. Case 1.2(b). (a) $e(k, k+1; l+2) \in H$. (b) $e(k, k+1; l+2) \notin H$.

CASE 1.3: $n - 1 \geq l + 3$. By Case 1.2, we may assume that $e(k, k+1; l+1) \notin H$, $S_{\rightarrow}(k-1, l+1; k, l+2) \in H$ and that $S_{\downarrow}(k+1, l+2; k+2, l+1) \in H$. Now L_E is either open or closed.

CASE 1.3(a): L_E is closed. By previous cases and symmetry we may assume that $m - 1 \geq k' + 3$. Using symmetry once more, we may assume that $e(k'+1; l', l'+1) \notin H$. Then the turn $\hat{T}$ on $\{e(k; l+1, l+2), S_{\downarrow}(k+1, l+2; k'+2, l'+1), e(k'+1, k'+2; l')\}$ is in H and it is a lengthening of T. See Figure 4.66. End of proof of I for Case 1.3(a).

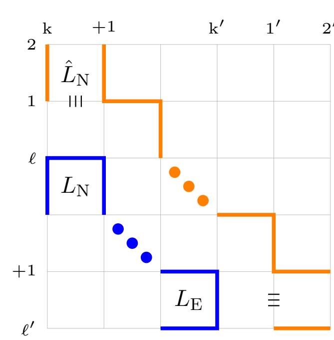

Fig. 4.66. Case 1.3 (a).

Proof of II for Case 1.3 (a). WLOG assume that the last

[1] This move will be referenced in Chapter 5, in the proof of Proposition 5.4.

move μ_s of a weakening $\mu_1, \ldots, \mu_s$ of $\hat{T}$ is $Z \mapsto \hat{L}_{\mathrm{N}}$, where $\hat{L}_{\mathrm{N}}$ is the northern leaf of $\hat{T}$. Then $\mu_1, \ldots, \mu_s, L_N + (0,1) \mapsto L_N$, is a weakening of T.
It remains to check that $s+1 \leq \min(m,n)$. Since the j^{th} lengthening T_j in $\mathcal{T}(T)$ is j units north and east of T, and $d(T) \geq 3$, there can be at most $\min(m,n) - 3$ such lengthenings. Since a turn with no lengthening has a short weakening, $s+1 \leq 3 + \min(m,n) - 3 = \min(m,n)$. End of proof of II for Case 1.3(a). End of Case 1.3(a).

CASE 1.3(b): L_E is open. Either $m-1 < k'+2$ or $m-1 \geq k'+2$. It will follow from Case 2 that if a turn has an open leaf adjacent to the boundary or at distance one away from the boundary, then we can find a weakening outright. Therefore, we may assume that $m-1 \geq k'+2$.

If $e(k'; l'+1, l'+2) \in H$, then there is a weakening $L_E \mapsto L_E + (0,1)$, so we may assume that $e(k'; l'+1, l'+2) \notin H$. Then the turn $\hat{T}$ on $\{e(k; l+1, l+2), S_{\downarrow}(k+1, l+2; k'+1, l'+2), e(k', k'+1; l'+1)\}$ is in H and it is a lengthening of T. See Figure 4.67. End of proof of I for Case 1.3.

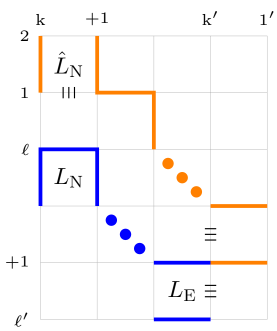

Fig. 4.67. Case 1.3 (b).

Proof of II for Case 1.3(b). Let $\hat{L_N}$ and $\hat{L_E}$ be the northern and eastern leaves of $\hat{T}$ respectively, and let $\mu_1, \ldots, \mu_s$ be a weakening of $\hat{T}$. If μ_s is the move $X \mapsto \hat{L}_N$, then, as in Case 1.3(a), $\mu_1, \ldots, \mu_s, L_N + (0,1) \mapsto L_N$, is a weakening of T. Suppose then that μ_s is the move $Z' \mapsto \hat{L}_E$. Then $\mu_1, \ldots, \mu_s, L_E \mapsto L_E + (0,1)$, is a weakening of T. The argument that $s+1 \leq \min(m,n)$ is the same as the one in Case 1.3(a), so we omit it. End of proof of II for Case 1.3(b) End of Case 1.3(b). End of Case 1.3. End of Case 1.

CASE 2: L_N is open. Proof of I. If $n-1 = l$ then we must have $e(k+1, k+2; l) \in H$. Then $L_{\mathrm{N}} \mapsto L_{\mathrm{N}} + (1,0)^1$ is a weakening. See Figure 4.68. Therefore, we may assume that $n-1 > l$.

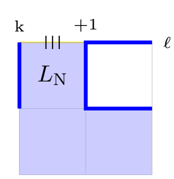

Fig. 4.68. Case 2, $n-1 = l$.

CASE 2.1: $n-1 = l+1$. Either $e(k+1, k+2; l) \in H$ or $e(k+1, k+2; l) \notin H$.

CASE 2.1(a): $e(k+1, k+2; l) \in H$. Then $e(k, k+1; l+1) \in H$ and $e(k+1, k+2; l+1) \in H$. Then by Corollary 1.3.15 (b), $L_N + (0,1) \in \mathrm{int}H$, and so $L_N \in \mathrm{int}H$ and $L_N + (1,0) \in \mathrm{ext}H$. Either $k = 0$, or $k > 0$.

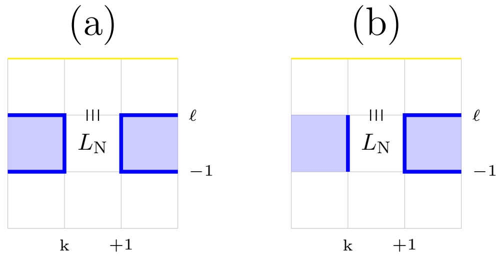

Fig. 4.69. Case 2.1(a_1). (a) $L_N + (-1,0)$ is a small cookie. (b) $L_N + (-1,0)$ is not a small cookie.

CASE 2.1(a_1): $k > 0$. Then $L_N+(-1,0) \in \text{ext}H$, and $L_N + (-1,0)$ is either a small cookie of H or it is not. If the former, then $L_N \mapsto L_N + (-1,0)$ is a short weakening; and if the latter then $L_N \mapsto L_N + (1,0)$ is a short weakening. See Figure 4.69. End of Case 2.1(a_1).

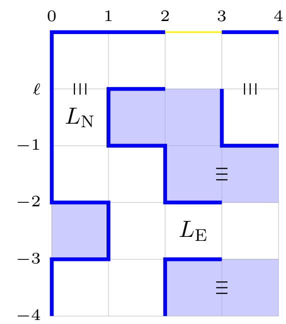

Fig. 4.70. Case 2.1(a_2).

CASE 2.1(a_2): $k = 0$. Then $e(0;l,l+1) \in H$, $e(0;l-2,l-1) \in H$, and $S_{\rightarrow}(0,l-2;1,l-3) \in H$. This implies that $0 \leq l-4$ and that $S_{\leftarrow}(1,l-3;0,l-4) \in H$. Then we must have $S_{\uparrow}(2,l-4;3,l-3) \in H$ as well, that $L_E = L_N+(2,-2)$ and that $L_E + (0,-1) \in \text{ext}H$. See Figure 4.70. Note that if $L_E + (0,-1)$ is a small cookie, then $L_E \mapsto L_E + (0,-1)$ is a short weakening, so we may assume that $L_E + (0,-1)$ is not a small cookie. Note that this implies that $0 \leq l-5$. If $e(3;l-4,l-3) \in H$, then again $L_E \mapsto L_E + (0,-1)$ is a short weakening. Similarly, if $e(3;l-2,l-1) \in H$, then $L_E \mapsto L_E + (0,1)$ is a short weakening. Therefore we only need to check the case where $e(3;l-4,l-3) \notin H$ and $e(3;l-2,l-1) \notin H$. Then $L_E+(1,-1) \in \text{ext}H$, and by the assumption that H is Hamiltonian, $m-1 \geq 4$. Then we have $S_{\downarrow}(3,l;4,l-1) \in H$.

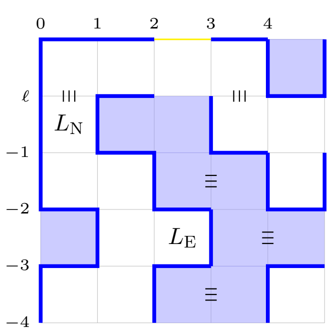

Fig. 4.71. Case 2.1(a_2).(i).

Note that either $e(2,3;l+1) \in H$ and $e(2,3;l) \in H$, or $e(2;l,l+1) \in H$ and $e(3;l,l+1) \in H$. Either way, we must have $e(3,4;l) \notin H$ and $e(3,4;l+1) \in H$. Now, either $e(3;l-3,l-2) \in H$ or $e(3;l-3,l-2) \notin H$.

CASE 2.1(a_2).(i): $e(3;l-3,l-2) \in H$. Then $L_E+(1,0) \in \text{ext}H$. By Corollary 1.3.15(b), this implies that $m-1 \geq 5$. If $e(4;l-3,l-2) \in H$, then $L_E + (1,0) \mapsto L_E$ is a short weakening, so we may assume that $e(4;l-3,l-2) \notin H$. Then $S_{\downarrow}(4,l-1;5,l-2) \in H$ and $S_{\uparrow}(4,l-4;5,l-3) \in H$. We must also have that $S_{\downarrow}(4,l+1;5,l) \in H$ and $e(5;l,l+1) \in H$. Then $e(5;l-2,l-1) \in H$ as well. See Figure 4.71. Then $L_E + (2,0) \mapsto L_E + (2,1)$, $L_E + (1,0) \mapsto L_E$ is a short weakening of T. End of Case 2.1(a_2).(i).

CASE 2.1(a_2).(ii): $e(3;l-3,l-2) \notin H$. Then $e(3,4;l-3) \in H$ and $e(3,4;l-2) \in H$. Now, either $e(2,3;l+1) \in H$ and $e(2,3;l) \in H$, or $e(2;l,l+1) \in H$ and $e(3;l,l+1) \in H$.

CASE 2.1(a_2).(ii)$_1$: $e(2,3;l+1) \in H$ and $e(2,3;l) \in H$. Then $e(4;l-2,l-1) \notin H$. Note that if $e(4;l-1,l) \in H$, then $L_E+(1,1) \mapsto L_E+(1,2)$, $L_E \mapsto L_E+(0,1)$ is a short weakening of T, so we may assume that $e(4;l-1,l) \notin H$. Then $e(4,5;l-1) \in H$, $S_\downarrow(4,l+1;5,l) \in H$, and $e(5;l,l+1) \in H$. Then $L_E+(2,2) \mapsto L_E+(2,3)$, $L_E+(1,1) \mapsto L_E+(1,2)$, $L_E \mapsto L_E+(0,1)$ is a short weakening of T. See Figure 4.72 (a). End of Case 2.1(a_2).(ii)$_1$.

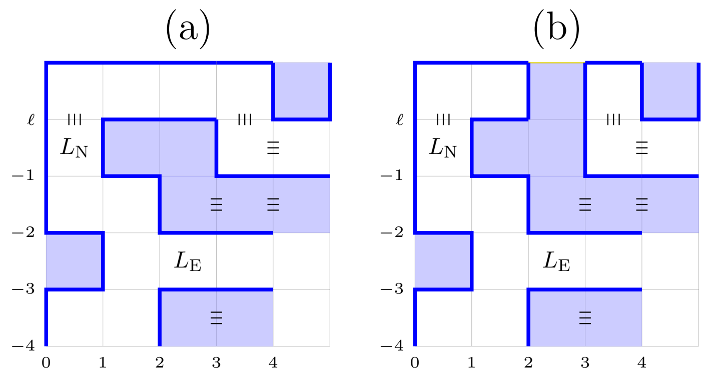

Fig. 4.72. (a) Case 2.1(a_2).(ii)$_1$. (b) Case 2.1(a_2).(ii)$_2$.

CASE 2.1(a_2).(ii)$_2$: $e(2;l,l+1) \in H$ and $e(3;l,l+1) \in H$. Now, if $e(4;l-2,l-1) \notin H$, then we can use the same argument and find the same cascades as in Case 2.1(a_2).(ii)$_1$ (Figure 4.72 (a)); and if $e(4;l-2,l-1) \in H$, then we must have that $e(5;l-2,l-1) \in H$ as well. Then $L_E+(2,1) \mapsto L_E+(1,1)$, $L_E \mapsto L_E+(0,1)$ is a short weakening of T. See Figure 4.72 (b). End of Case 2.1(a_2).(ii)$_2$. End of Case 2.1(a_2).(ii). End of Case 2.1(a_2). End of Case 2.1(a).

CASE 2.1(b): $e(k+1,k+2;l) \notin H$. Then we must have $e(k+1;l,l+1) \in H$ and $S_\downarrow(k+2,l+1;k+3,l) \in H$. Now, either $e(k+1,k+2;l+1) \notin H$ or $e(k+1,k+2;l+1) \in H$.

CASE 2.1(b_1): $e(k+1,k+2;l+1) \notin H$. Then we must have $e(k+2,k+3;l+1) \in H$ and that $L_N+(1,1) \in \text{ext}H$ is the neck of the large cookie. Then $k>0$, or $k=0$.

If $k>0$, then after $L_N+(1,1) \mapsto L_N+(2,1)$ [1], we are back to Case 2.1(a_1). And if $k=0$ then we are effectively in the same scenario as in Case 2.1(a_2), except that in this case, in the analogues of Cases 2.1(a_2).(i) and (ii), we get lengthenings instead of short weakenings. See Figure 4.73 (a). End of Case 2.1(b_1).

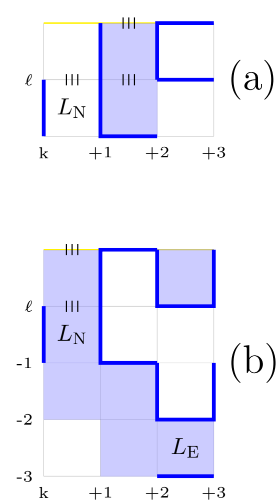

Fig. 4.73. (a) Case 2.1(b_1). (b) Case 2.1(b_2).

CASE 2.1(b_2): $e(k+1,k+2;l+1) \in H$. Then $e(k,k+1;l+1) \notin H$ and $S_\downarrow(k+2,l+1;k+3,l) \in H$. Then we must have that $e(k+3;l,l+1) \in H$ as well. This implies that $e(k+3;l-2,l-1) \in H$. It follows that $L_N+(2,-2) = L_E$, and that L_E is open. Then

[1] This move will be referenced in Chapter 5, in the proof of Proposition 5.4.

$L_E \mapsto L_E + (0, 1)$ is a short weakening. See Figure 4.73 (b). End of Case 2.1(b_2). End of Case 2.1(b). End of Case 2.1.

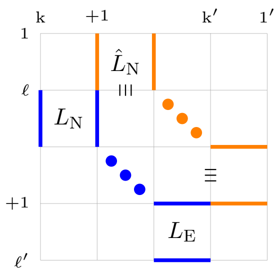

Fig. 4.74. Case 2.2.

CASE 2.2: $n-1 \geq l+2$. By previous cases we may assume that $m-1 \geq k'+2$, $e(k+1, k+2; l) \notin H$, $e(k+1; l, l+1) \in H$ and $S_{\downarrow}(k+2, l+1; k+3, l) \in H$. If L_E is closed, then we're done by Case 1, so we may assume that L_E is open. By Case 2.1, we may assume that $e(k'; l'+1, l'+2) \notin H$. Then the turn $\hat{T}$ on $\{e(k+1; l, l+1), S_{\downarrow}(k+2, l+1; k'+1, l'+2), e(k', k'+1; l'+1)\}$ is in H and it is a lengthening of T. See Figure 4.74. End of proof of I for Case 2.2.

The proof of II for Case 2.2 is the same as the proof of II for Case 1.3(a). End of proof for Case 2. □

Observation 4.25. All turn weakenings found in Lemma 4.23 are contained in Sector(T).

4.5 Weakenings when H is a Hamiltonian path of G with $u \in R_0$

Lemma 4.26. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path with end-vertices u and v, and let T be a turn in H with $d(T) \geq 3$. Assume that u and v are not adjacent, that $u \in R_0$, and that v is not in Sector(T). Then:

I. T has a short double-switch weakening[7] (SDSW) or T has a lengthening.

II. If T' is a lengthening of T and T' has a double-switch weakening of length at most s, then T has a double-switch weakening of length at most $s+1$, with $s+1 \leq \min(m, n)$.

Proposition 4.27. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path with end-vertices u and v, and let T be a turn in H with $d(T) \geq 3$. Assume that u and v are not adjacent, that $u \in R_0$, and that v is not in Sector(T). Then T has a double-switch weakening of length at most $\min(m, n)$ moves.

Proof. Let $T = T_0$ be a turn of H with $d(T) \geq 3$. Assume that $u \in R_0$, and that $v \notin$ Sector(T). If T_0 has a short double-switch weakening, we are done, so we may assume that T_0 does not have a short double-switch weakening. By I in Lemma 4.26, T_0 has a lengthening T_1. So, $T_1 \in \mathcal{T}$, where $\mathcal{T} = \mathcal{T}(T_0)$. Since $m, n < \infty$, we have that $|\mathcal{T}| < \infty$.

[7] That is, a weakening consisting entirely of double-switch moves.

Let $\mathcal{T} = \{T_0, T_1, \ldots, T_j\}$. Then T_j has no lengthening. So, by I in Lemma 4.26, T_j must have a short double-switch weakening. Then, by induction and II on Lemma 4.26, T_0 has a double-switch weakening. The bound follows immediately. □

Proof of Lemma 4.26. For definiteness, assume that T is the northeastern turn on $\{e(k; l-1, l), S_{\downarrow}(k+1, l; k', l'+1), e(k'-1, k', l')\}$. Let L_N and L_E be the northern and eastern leaves of T, respectively. We will consider three cases: L_N is closed and $n-1 < l+2$, L_N is open and $n-1 < l+2$, and $n-1 \geq l+2$.

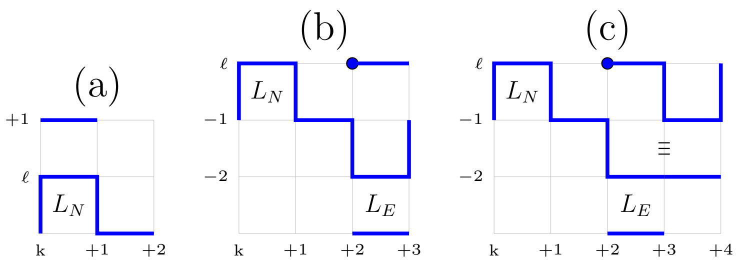

Fig. 4.75. (a) Case 1, $n-1 = l+1$. (b) $n-1 = l$, $e(k+2, k+3; l+1) \in H$. (c) $n-1 = l$. Case 1.1.

CASE 1. L_N is closed and $n-1 < l+2$. Suppose that $n-1 = l+1$. Observe that (NAA) implies that $e(k, k+1; l+1) \in H$. Then $L_N + (0,1) \mapsto L_N$ is a SDSW. See Figure 4.75 (a). It remains to check the case where $n-1 = l$.

Suppose that $n-1 = l$. Then $v(k+2, l) = u$, and $e(k+2, k+3; l) \in H$. Note that if $e(k+3; l-2, l-1) \in H$, then we must have that $L_E = L_N + (2, -2)$, and that L_E is open. Then $L_E \mapsto L_E + (0,1)$ is a SDSW. See Figure 4.75 (b). Therefore we may assume that $e(k+3; l-2, l-1) \notin H$. Since $v \notin \text{Sector}(T)$, this implies that $m-1 > k+3$. Then we have that $S_{\downarrow}(k+3, l; k+4, l-1) \in H$, and $e(k+4; l-1, l) \in H$. Now, $e(k+3, k+4; l-2) \in H$, or $e(k+3, k+4; l-2) \notin H$.

CASE 1.1: $e(k+3, k+4; l-2) \in H$. This also implies that $L_E = L_N + (2, -2)$, and that L_E is open. Then $L_E + (1,1) \mapsto L_E + (1,2)$, $L_E \mapsto L_E + (0,1)$ is a SDSW. See Figure 4.75 (c). End of Case 1.1.

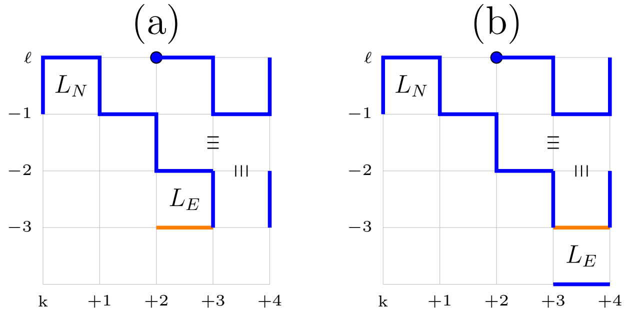

Fig. 4.76. $n-1 = l$. (a) Case 1.2, with $e(k+2, k+3; l-3) \in H$. (b) Case 1.2, with $e(k+2, k+3; l-3) \in H$.

CASE 1.2: $e(k{+}3, k{+}4; l{-}2) \notin H$. Then $e(k{+}3; l{-}3, l{-}2) \in H$ and $e(k{+}4; l{-}3, l{-}2) \in H$. Now, either $e(k+2, k+3; l-3) \in H$, or $e(k+3, k+4; l-3) \in H$. If the former, then $L_E = L_N + (2, -2)$, L_E is closed, and $L_E + (1, 0) \mapsto L_E$ is a SDSW. See Figure 4.76 (a). And if the latter, $L_E = L_N + (3, -3)$, L_E is open, and $L_E \mapsto L_E + (0, 1)$ is a SDSW. Figure 4.76 (b). End of Case 1.2. End of Case 1.

CASE 2. L_N is open and $n - 1 < l + 2$. Suppose that $n - 1 = l$. Either $e(k + 1, k + 2; l) \in H$, or $e(k + 1, k + 2; l) \notin H$.

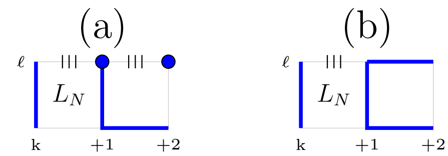

Fig. 4.77. (a) $e(k{+}1, k{+}2; l) {\in} H$. (b) $e(k{+}1, k{+}2; l) {\notin}) H$.

Note that the latter implies that the end-vertices of H are on $v(k + 1, l)$ and $v(k + 2, l)$, contradicting the assumption that $v \notin \text{Sector}(T)$. Then, it must be the case that $e(k + 1, k + 2; l) \in H$. Then $L_N \mapsto L_N + (1, 0)$ is a SDSW. See Figure 4.77. It remains to check the case where $n - 1 = l + 1$.

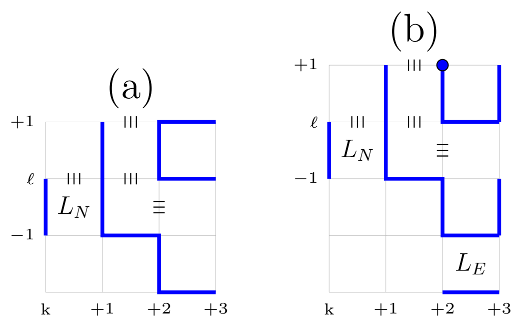

Fig. 4.78. $n{-}1{=}l{+}1$. (a) $e(k + 2, k + 3; l + 1) \in H$. (b) $u = v(k + 2, l + 1)$.

Suppose that $n - 1 = l + 1$. Note that if $e(k + 1, k + 2; l) \in H$ then $L_N \mapsto L_N + (1, 0)$ is a SDSW, so we may assume that $e(k{+}1, k{+}2; l) \notin H$. Then $e(k + 1; l, l + 1) \in H$. It follows that $S_{\downarrow}(k + 2, l + 1; k + 3, l) \in H$. There are three possibilities: $e(k + 2, k + 3; l + 1) \in H$ (Figure 4.78 (a)), $v(k + 2, l + 1) = u$ (Figure 4.78 (b)), and $e(k + 1, k + 2; l + 1) \in H$ (Figure 4.79 (a)). If the first, then $L_N + (1, 1) \mapsto L_N + (2, 1)$, $L_N \mapsto L_N + (1, 0)$ is a SDSW. If the second, then we must have $e(k + 3; l, l + 1) \in H$. It follows that $e(k + 3; l - 2, l - 1) \in H$ as well. It follows that $L_E = L_N + (2, -2)$, and that L_E is open. Then $L_E \mapsto L_E + (0, 1)$ is a SDSW. It remains to check the third.

Suppose that $e(k + 1, k + 2; l + 1) \in H$. As above, if $e(k + 3; l - 2, l - 1) \in H$, then $L_E = L_N + (2, -2)$, L_E is open, and $L_E \mapsto L_E + (0, 1)$ is a SDSW, so we may assume that $e(k + 3; l - 2, l - 1) \notin H$. It follows that $m - 1 \geq k + 4$, that $S_{\downarrow}(k + 3, l; k + 4; l - 1) \in H$, that $u = v(k + 3, l + 1)$, that $e(k + 3, k + 4; l + 1) \in H$. Now, $e(k + 4; l - 1, l) \in H$, or $e(k + 4; l - 1, l) \notin H$.

CASE 2.1. $e(k+4; l-1, l) \in H$. Now, either $e(k+3, k+4; l-2) \in H$ or $e(k+3, k+4; l-2) \notin H$.

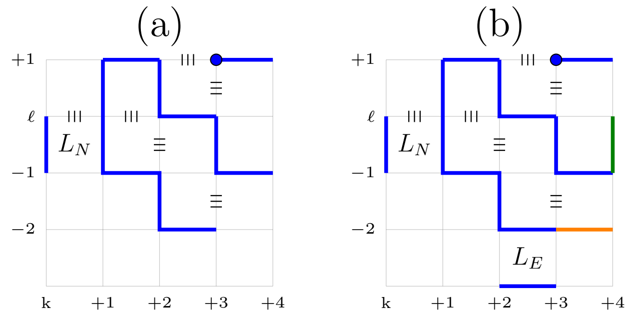

Fig. 4.79. (a) $n-1=l+1$, $e(k+1, k+2; l+1) \in H$. (b) Case 2.1(a).

CASE 2.1(a). $e(k+3, k+4; l-2) \in H$. Then $L_E = L_N + (2, -2)$, L_E is open, and $L_E + (1, 1) \mapsto L_E + (1, 2)$, $L_E \mapsto L_E + (0, 1)$ is a SDSW. See Figure 4.79 (b). End of Case 2.1(a).

CASE 2.1(b). $e(k+3, k+4; l-2) \notin H$. Then $e(k+3; l-3, l-2) \in H$, and $e(k+4; l-3, l-2) \in H$. Now either $e(k+2, k+3; l-3) \in H$, or $e(k+2, k+3; l-3) \notin H$.

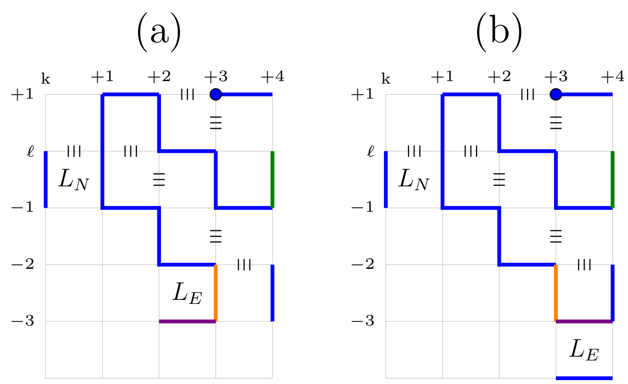

Fig. 4.80. (a) Case 2.1(b_1). (b) Case 2.1(b_2).

CASE 2.1(b_1). $e(k+2, k+3; l-3) \in H$. Then $L_E = L_N + (2, -2)$, L_E is closed, and $L_E + (1, 0) \mapsto L_E$ is a SDSW. See Figure 4.80 (a). End of Case 2.1(b_1).

CASE 2.1(b_2). $e(k+2, k+3; l-3) \notin H$. Then we must have $e(k+3, k+4; l-3) \in H$, that $L_E = L_N + (3, -3)$, that L_E is open, and that $L_E \mapsto L_E + (0, 1)$ is a SDSW. See Figure 4.80 (b). End of Case 2.1(b_2). End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $e(k+4; l-1, l) \notin H$. It follows that $m-1 \geq k+5$, that $S_{\downarrow}(k+4, l+1; k+5, l) \in H$, and that $e(k+5; l, l+1) \in H$. Note that if $e(k+4, k+5; l-1) \in H$, then, after $L_N + (4, 0) \mapsto L_N + (4, 1)$, we are back to Case 2.1. Note that the weakenings in Case 2.1 required at most two moves. So, including $L_N + (4, 0) \mapsto L_N + (4, 1)$, we would still need only three moves for a weakening. See Figure 4.81 (a).

Assume that $e(k+4, k+5; l-1) \notin H$. It follows that $e(k+4; l-2, l-1) \in H$, $e(k+5; l-2, l-1) \in H$, $e(k+3, k+4; l-2) \notin H$, and $e(k+3, l-3, l-2) \in H$. Now, either $e(k+4; l-3, l-2) \in H$ or $e(k+4; l-3, l-2) \notin H$.

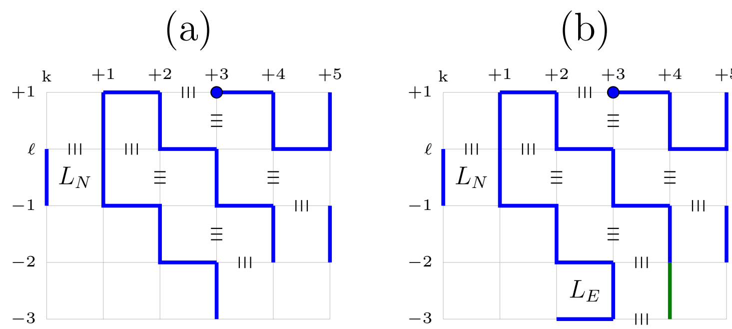

Fig. 4.81. (a) Case 2.2. (b) Case 2.2(a).

CASE 2.2(a): $e(k+4; l-3, l-2) \in H$. Then $e(k+3, k+4; l-3) \notin H$. It follows that $L_E = L_N + (2, -2)$, and that L_E is closed. Then $L_E + (1, 0) \mapsto L_E$ is a SDSW. See Figure 4.81 (b). End of Case 2.2(a)

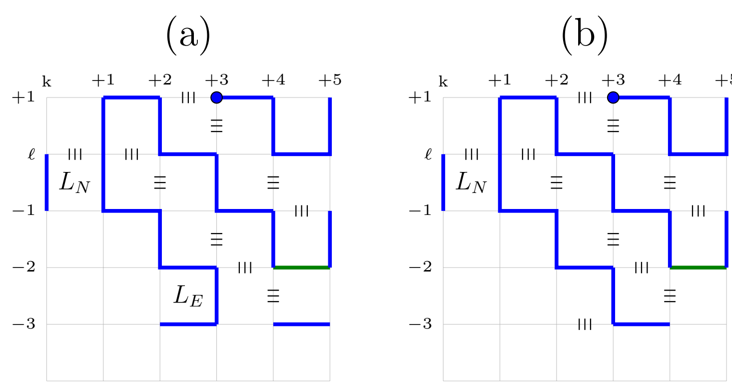

Fig. 4.82. Case 2.2(b) with: (a) $e(k+2, k+3; l-3) \in H$; and (b) $e(k+2, k+3; l-3) \notin H$.

CASE 2.2(b): $e(k+4; l-3, l-2) \notin H$. Then $e(k+4, k+5; l-2) \in H$. Note that if $e(k+2, k+3; l-3) \in H$, then $L_E = L_N + (2, -2)$, L_E is closed, and $e(k+4, k+5; l-3) \in H$. Then $L_E + (2, 0) \mapsto L_E + (2, 1)$, $L_E + (1, 0) \mapsto L_E$ is a SDSW. See Figure 4.82 (a). So, we may assume that $e(k+2, k+3; l-3) \notin H$ (Figure 4.82 (b)). Then $e(k+3, k+4; l-3) \in H$. Now, either $e(k+4, k+5; l-3) \in H$, or $e(k+4, k+5; l-3) \notin H$.

CASE 2.2(b_1): $e(k+4, k+5; l-3) \in H$. Then $L_E = L_N + (3, -3)$, and $L_E + (1, 1) \mapsto L_E + (1, 2)$, $L_E \mapsto L_E + (0, 1)$ is a SDSW. See Figure 4.83 (a). End of Case 2.2(b_1).

CASE 2.2(b_2): $e(k+4, k+5; l-3) \notin H$. It follows that $e(k+4; l-4, l-3) \in H$, and $e(k+5; l-4, l-3) \in H$. Now either $e(k+3, k+4; l-4) \in H$, or $e(k+4, k+5; l-4) \in H$. If the former, then $L_E = L_N + (3, -3)$, L_E is closed, and $L_E + (1, 0) \mapsto L_E$ is a SDSW. And if the latter, then $L_E = L_N + (4, -4)$, L_E is open, and $L_E \mapsto L_E + (0, 1)$ is a SDSW. See

Figure 4.83 (b) and (c). End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2.

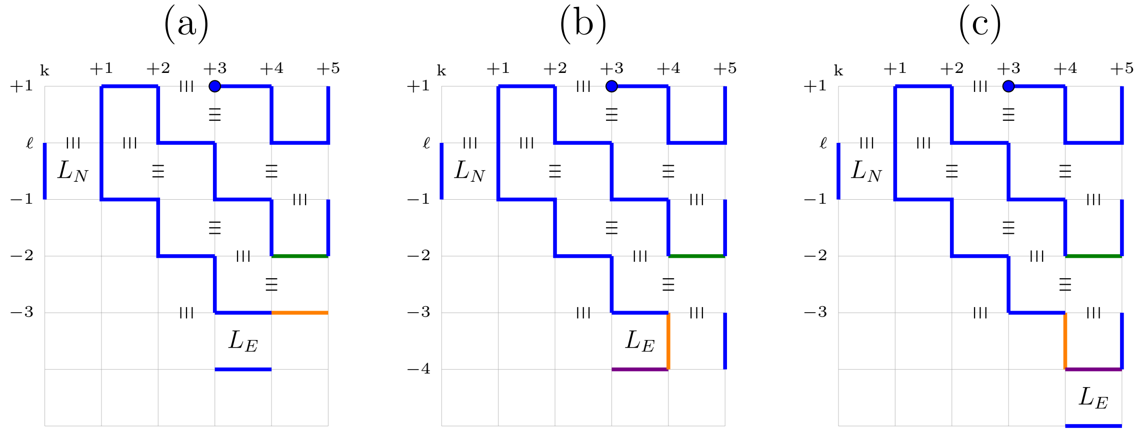

Fig. 4.83. (a) Case 2.2(b_1). (b) Case 2.2(b_2) with $e(k+3, k+4; l-4) \in H$. (c) Case 2.2(b_2) with $e(k+4, k+5; l-4) \in H$.

CASE 3: $n-1 \geq l+2$. This is very similar to Cases 1.3 and 2.2 in Lemma 4.23, so we omit it. End of Case 3. □

4.6 Weakenings when H is a Hamiltonian path of G with $u \notin R_0$

Turn weakenings are useful for the reconfiguration of Hamiltonian paths, which are discussed in Chapter 5. So far, we have found turn weakenings using only double-switch moves. When an end-vertex of H lies within the sector of a turn, we will find a weakening that also allows single-switches and backbites. While a double-switch weakening may still exist, it seems that it would require several additional pages of configuration checking, and we believe it would be no shorter than the one we find.

Proposition 4.28. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path with end-vertices u and v, and let T be a turn in H with $d(T) \geq 3$. Assume that u and v are not adjacent, that $u \notin R_0$, and that v is not in Sector(T). Then T has a weakening of length at most $\min(m, n)$ moves.

The proof of Proposition 4.28 requires the following lemma.

Lemma 4.29. Let G be an $m \times n$ grid graph, let H be a Hamiltonian path with end-vertices u and v, and let T be a turn in H with $d(T) \geq 3$. Assume that u and v are not adjacent, that $u \notin R_0$, and that v is not in Sector(T). Then:

I. T has a short weakening or T has a lengthening.

II. If T' is a lengthening of T and T' has a weakening of length at most s, then T has a weakening of length at most $s+1$, with $s+1 \leq \min(m,n)$.

Proof of Proposition 4.28. This is the same as the proof of Proposition 4.24, once Lemma 4.29 has been proved, so we omit it. □.

The proof of Lemma 4.29 requires Lemma 4.30 below. The definitions of edge weakening, end-vertex weakening, end-vertex weakening terminal, and Flank, introduced at the start of Section 4.3, will be used throughout the remainder of this section.

Lemma 4.30. Let H be a Hamiltonian path of G with non-adjacent end-vertices, and let T be a turn of H with $d(T) \geq 3$. Then:

(a) If Flank(T) contains an end-vertex of H, then T has a short weakening.

(b) If T' is a lengthening of T, and T' has an end-vertex weakening of length at most s, then T has a weakening of length at most $s+1$, with $s+1 \leq \min(m,n)$.

Proof of Lemma 4.29. For definiteness, let T be the northeastern turn of H on $\{e(k;l-1,l), S_{\downarrow}(k+1,l;k',l'+1), e(k'-1,k',l')\}$. Since $v \notin$ Sector(T), by Observation 4.20(c), $v \notin$ Flank(T). If $u \in$ Flank(T), by Lemma 4.30 (a), T has a short weakening, so we may assume that $u \notin$ Flank(T). We observe that the proof of Cases 1 and 2 in Lemma 4.26 show that if $n-1 < l+2$, or if $m-1 < k'+2$, then T has a weakening, so we may assume that $n-1 \geq l+2$ and that $m-1 \geq k'+2$. Let $L_N = R(k,l-1)$ and $L_E = R(k'-1,l')$. Then L_N is either open or closed.

CASE 1. L_N is closed. As in Lemma 4.23, if $e(k,k+1;l+1) \in H$, then T has a weakening, so we may assume that $e(k,k+1;l+1) \notin H$. Then $e(k;l+1,l+2) \in H$. Now, L_E is open, or L_E is closed.

CASE 1.1: L_E is open. If $e(k';l'+1,l'+2) \in H$, then $L_E \mapsto L_E + (0,1)$ is a short weakening, so we may assume that $e(k';l'+1,l'+2) \notin H$. Then $e(k',k'+1;l'+1) \in H$ and $S_{\downarrow}(k+1,l+2;k'+1,l'+2) \in H$. Now the turn T' on $e(k;l+1,l+2)$, $S_{\downarrow}(k+1,l+2;k'+1,l'+2)$ and $e(k',k'+1;l'+1)$ is a lengthening of T. See Figure 4.84. End of proof of I for Case 1.1.

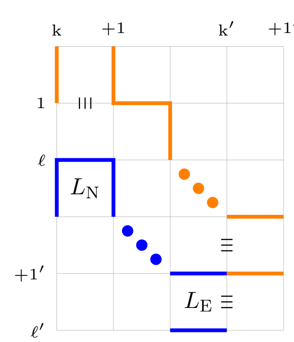

Fig. 4.84. Case 1.1.

Assume that T' has a weakening $\mu_1, \ldots, \mu_s$. Either it has an edge weakening or an end-vertex weakening. If $\mu_1, \ldots, \mu_s$ is an edge weakening, after applying it, either $e(k, k+1; l+1)$, or $e(k'; l'+1, l'+2)$ is in the resulting Hamiltonian path. If the former then $L_N + (0,1) \mapsto L_N$ is a weakening of T; and if the latter, then $L_E \mapsto L_E + (0,1)$ is a weakening of T. In either case, the argument that $s+1 \leq \min(m,n)$ is the same as the one used in Case 1.3(a) of Lemma 4.23, so we omit it.

If T' $\mu_1, \ldots, \mu_s$ is an end-vertex weakening, then, by Lemma 4.30 (b), T has a weakening of length $s+1 \leq \min(m,n)$. End of Case 1.1.

Case 1.2 and Case 2 are similar, so we omit the proofs. $\square$

Proof of Lemma 4.30(a). For definiteness, let T be the northeastern turn of H on $\{e(k; l-1, l), S_{\downarrow}(k+1, l; k', l'+1), e(k'-1, k', l')\}$ and assume that v is an end-vertex of H in Flank(T). Orient H so that $v = v_1$. If either leaf of T is parallel, then switching that leaf is an edge weakening, so we may assume that both leaves of T are anti-parallel. We will use Lemmas 1.4.1 and 1.4.4 repeatedly and implicitly. Now, $v \in s$-Flank(T) or $v \in \ell$-Flank(T), so there are two cases to check.

CASE 1: $v \in \ell$-Flank(T). For definiteness, assume that $v \in V(L_N + (0,1))$. Now, L_N is open or closed.

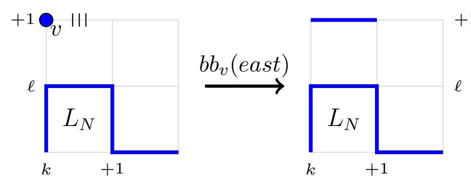

Fig. 4.85. Case 1.1 with $v = v(k, l+1)$.

CASE 1.1: L_N is closed. If $e(k, k+1; l+1) \in H$ then $L_N + (0,1) \mapsto L_N$ is a short edge weakening, so we may assume that $e(k, k+1; l+1) \notin H$. Then $v = v(k, l+1)$ or $v = v(k+1, l+1)$. If the former, then $bb_v(east)$, $L_N + (0,1) \mapsto L_N$ is a short edge weakening (see Figure 4.85); and if the latter, then $bb_v(west)$, $L_N + (0,1) \mapsto L_N$ is a short edge weakening. End of Case 1.1.

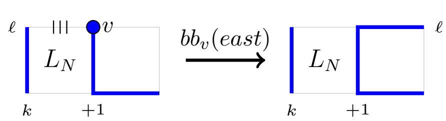

Fig. 4.86. Case 1.2(a).

CASE 1.2: L_N is open. Then $v = v(k+1, l)$, $v = v(k, l)$, $v = v(k+1, l+1)$, or $v = v(k, l+1)$.

CASE 1.2(a): $v = v(k+1, l)$. Then $bb_v(east)$, $L_N \mapsto L_N + (1,0)$ is a short edge weakening. See Figure 4.86.

End of Case 1.2+(a).

CASE 1.2 (b): $v = v(k,l)$. Then, after $bb_v(east)$, the end-vertex is on $v(k+1, l-1) \in$ ew-Terminal(T), so $bb_v(east)$ is a short end-vertex weakening. See Figure 4.87. End of Case 1.2(b).

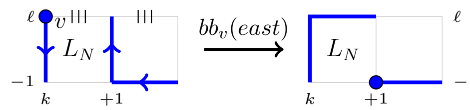

Fig. 4.87. Case 1.2(b).

CASE 1.2(c): $v = v(k+1, l+1)$. By (NAA), $v(k+2, l+1)$ and $v(k+1, l)$ are not end-vertices of H. If $e(k+1, k+2; l) \in H$, then $L_N \mapsto L_N + (1,0)$ is a short edge weakening, so we may assume that $e(k+1, k+2; l) \notin H$. Then $e(k+1; l, l+1) \in H$. Note that if $e(k+2; l-1, l) \in H$, then $L_N + (1,-1) = L_E$, which contradicts the assumption that $d(T) \geq 3$, so we must have that $e(k+2; l-1, l) \notin H$. It follows that $e(k+2; l, l+1) \in H$. Note that if $L_N + (1,1)$ is parallel, then $\text{Sw}(L_N + (1,1))$, $L_N \mapsto L_N + (1,0)$ is a short edge weakening, so we may assume that $L_N + (1,1)$ is anti-parallel.

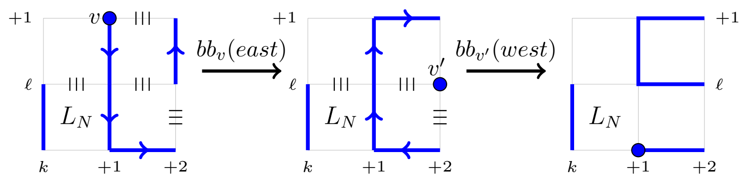

Fig. 4.88. Case 1.2(c).

Let H' be the Hamiltonian path obtained after $bb_v(east)$, and let v' be the end-vertex of H' to which v is relocated by the move. Note that $v' = v(k+2, l)$, and that $e(k+1; l-1, l)$ precedes $e(k+1; l, l+1)$ in H'. Now we can apply $bb_{v'}(west)$ to H', after which, v' is relocated to $v(k+1, l-1) \in$ ew-Terminal(T), so $bb_v(east)$ by $bb_{v'}(west)$ is a short end-vertex weakening. See Figure. 4.88. End of Case 1.2(c).

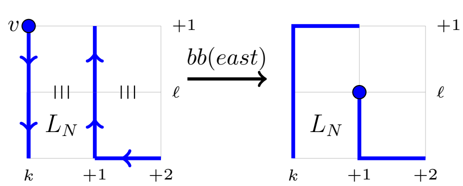

Fig. 4.89. Case 1.2(d_1).

CASE 1.2(d): $v = v(k, l+1)$. By case 1.2(a), we may assume that $v(k+1, l)$ is not an end-vertex of H, and that $e(k+1, k+2; l) \notin H$. Then $e(k+1; l, l+1) \in H$. Now, either $e(k; l, l+1) \in H$ or $e(k; l, l+1) \notin H$.

CASE 1.2(d_1): $e(k; l, l+1) \in H$. The assumption that L_N is anti-parallel implies that $bb_v(east)$ relocates v to $v(k+1, l)$, and we are back to Case 1.2(a). See Figure 4.89. End of Case 1.2(d_1).

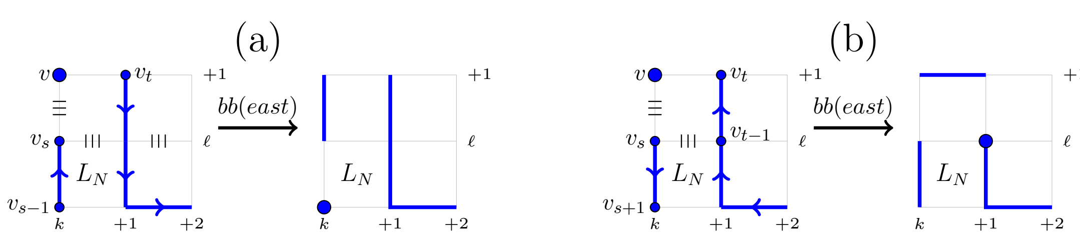

Fig. 4.90. Case 1.2 (d_2) with: (a) $v(k, l-1) = v_{s-1}$, and (b) $v(k, l-1) = v_{s+1}$.

CASE 1.2 (d_2): $e(k; l, l+1) \notin H$. By NAA, $v(k,l)$ and $v(k+1,l+1)$ are not end-vertices of H. Let $v_s = v(k,l)$ and $v_t = v(k+1,l+1)$. Now, $v(k,l-1) = v_{s-1}$ (Figure 4.90 (a)), or $v(k,l-1) = v_{s+1}$ (Figure 4.90 (b)). If the former, then after $bb_v(south)$, v is relocated to $v(k,l-1) \in$ ew-Terminal(T), so $bb_v(south)$ is a short end-vertex weakening. And if the latter, than the assumption that L_N is anti-parallel implies that $v_{t-1} = v(k+1,l)$. Then, after $bb_v(east)$, v is relocated to $v(k,l+1)$, and we are back to Case 1.2(a). End of Case 1.2(d_2). End of Case 1.2(d). End of Case 1.2. End of Case 1.

CASE 2: $v \in$ s-Flank(T). As we have seen before, if $e(k+1,k+2;l) \in H$, then $L_N \mapsto L_N + (1,0)$ is a short edge weakening, so we may assume that $e(k+1,k+2;l) \notin H$. Let:

$V_1 = \{v(k+2,l), v(k+2,l+1), v(k+3,l), v(k',l'+2), v(k'+1,l'+2), v(k',l'+3)\}$,

$V_2 = \{v(k+3,l-1), v(k+4,l-2), \ldots, v(k'-1,l'+3)\}$, and

$V_3 = \{v(k+4,l-1), v(k+5,l-2), \ldots, v(k'-1,l'+4)\}$,

Note that $\{V_1, V_2, V_3\}$ is a partition of s-Flank(T), that $V_3 = \emptyset$ if $d(T) = 4$, and that V_2 and V_3 are both empty if $d(T) = 3$. Then there are three cases to check: $v \in V_1$, $v \in V_2$ and $v \in V_3$.

CASE 2.1: $v \in V_1$. By symmetry we only need to check the cases where $v = v(k+2,l)$, $v = v(k+2,l+1)$ and $v = v(k+3,l)$.

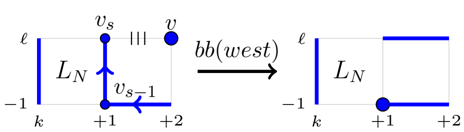

Fig. 4.91. Case 2.1(a_1).

CASE 2.1 (a): $v = v(k+2,l)$. If $e(k+2; l-1, l) \in H$ then $L_N + (1,-1) = L_E$, which contradicts the assumption that $d(T) \geq 3$, so we may assume that $e(k+2; l-1, l) \notin H$. Let $v(k+1,l) = v_s$. Then $v(k+1,l-1) = v_{s-1}$ or $v(k+1,l-1) = v_{s+1}$.

CASE 2.1 (a_1): $v(k+1,l-1) = v_{s-1}$. Then, after $bb_v(west)$, v is relocated to $v(k+1,l-1) \in$ ew-Terminal(T), so $bb_v(west)$ is a short end-vertex weakening. See Figure 4.91. End of Case 2.1 (a_1).

(a) (b)

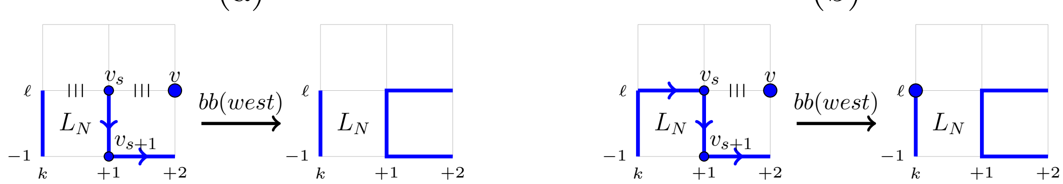

Fig. 4.92. Case 2.1 (a_2) with: (a) L_N open; (b) L_N closed.

CASE 2.1 (a_2): $v(k+1,l-1)=v_{s+1}$. If L_N is open (Figure 4.92 (a)), then $bb_v(west)$ is the move $e(k+1;l,l+1)\mapsto e(k+1,k+2;l)$. Then we can apply $L_N\mapsto L_N+(1,0)$, and see that $bb_v(west)$, $L_N\mapsto L_N+(1,0)$ is a short edge weakening. And if L_N is closed (Figure 4.92 (b)), then $bb_v(west)$ is the move $e(k,k+1;l)\mapsto e(k+1,k+2;l)$. Then we can apply $L_N\mapsto L_N+(1,0)$, and see that $bb_v(west)$, $L_N\mapsto L_N+(1,0)$ is a short weakening. Either way $bb_v(west)$, $L_N\mapsto L_N+(1,0)$ is a short edge weakening. End of Case 2.1 (a_2). End of Case 2.1(a).

CASE 2.1 (b): $v=v(k+2,l+1)$. By NAA $v(k+2,l)$ and $v(k+1,l+1)$ are not end-vertices of H. Since $d(T)\geq 3$, $e(k+2;l-1,l)\notin H$. Then we must have that $e(k+2;l,l+1)\in H$.

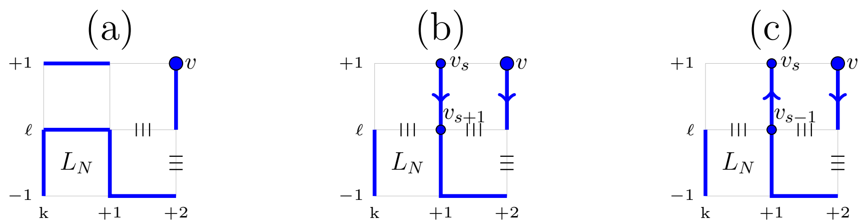

Fig. 4.93. Case 2.1 (b) with: (a) L_N closed; (b) L_N open and $v(k+1,l)=v_{s+1}$, and (c) L_N open and $v(k+1,l)=v_{s-1}$.

Note that if L_N is closed (Figure 4.93 (a)), since $v(k+1,l+1)$ is not an end-vertex of H, then we must have that $e(k,k+1;l+1)\in H$. But then $L_N+(0,1)\mapsto L_N$ is a short edge weakening, so we may assume that L_N is open. Then $e(k+1;l,l+1)\in H$.
Let $v(k+1,l+1)=v_s$. Then $v(k+1,l)=v_{s-1}$ (Figure 4.93 (b)) or $v(k+1,l)=v_{s+1}$ (Figure 4.93 (c)). If $v(k+1,l)=v_{s+1}$, then $L_N+(1,1)$ is parallel, and then $\mathrm{Sw}(L_N+(1,1))$, $L_N\mapsto L_N+(1,0)$ is a short edge weakening. Assume $v(k+1,l)=v_{s-1}$. Then after $bb_v(west)$, v relocates to $v(k+1,l)$, and we are back to Case 1.2(a). End of Case 2.1(b).

CASE 2.1 (c): $v=v(k+3,l)$. By NAA $v(k+2,l)$ and $v(k+3,l-1)$ are not end-vertices of H. Then we must have that $e(k+2,k+3;l)\in H$, and $e(k+3;l-2,l-1)\in H$. It follows that $L_E=L_N+(2,-2)$, and that L_E is open. Then $L_E\mapsto L_E+(0,1)$ is a short edge weakening. See Figure 4.94. End of Case 2.1(c). End of Case 2.1.

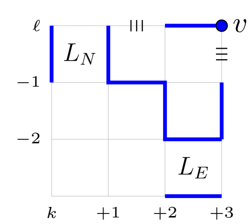

Fig. 4.94. Case 2.1(c).

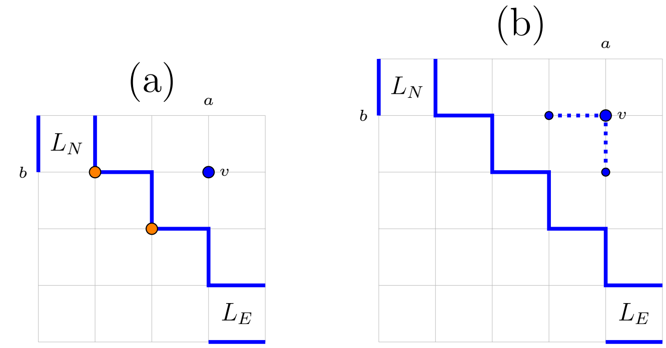

Fig. 4.95. (a) Case 2.2, an illustration with $d(T) = 4$. Possible locations of the end- vertex after $bb_v(west)$ in orange. (b) Case 2.3, an illustration with $d(T) = 5$.

CASE 2.2: $v \in V_2$. We may assume that $V_2 \neq \emptyset$. Note that this implies that $d(T) \geq 4$. Let $v = v(a, b)$. Now, after $bb_v(west)$, the end-vertex in the resulting Hamiltonian path is at $v(a-2, b)$ or on $v(a-1, b-1)$. See Figure 4.95 (a). Since $\{v(a-2, b), v(a-1, b-1)\} \subset$ ew-Terminal(T), $bb_v(west)$ is a short end-vertex weakening. End of Case 2.2.

CASE 2.3: $v \in V_3$. We may assume that $V_3 \neq \emptyset$. Note that this implies that $d(T) \geq 5$. Let $v = v(a, b)$. By NAA $v(a-1, b)$ and $v(a, b-1)$ are not end-vertices of H. But then we must have that $e(a; b-1, b) \in H$ and $e(a-1, a; b) \in H$, which conflicts with our assumption that $v(a, b)$ is an end-vertex of H. See Figure 4.95 (b). End of Case 2.3. End of Case 2. End of proof for part (a).

Proof of Lemma 4.30(b). Let T' be a lengthening of T and assume that $\mu_1, \ldots, \mu_s$ is an end-vertex weakening of T'. By Observation 4.20(b) ew-Terminal $(T') \subset$ Flank(T), so after μ_s, the resulting Hamiltonian path has an end-vertex u' in Flank(T). Then, by part (a), T has a weakening. The argument that $s + 1 \leq \min(m, n)$ is the same as the one used in Case 1.3(a) of Lemma 4.23, so we omit it. End of proof for part (b). □

Corollary 4.31. Let G be an $m \times n$ grid graph and let H be a Hamiltonian path of G with end-vertices u and v not adjacent. Let T be a turn of H with $d(T) \geq 3$. Then:

(i) If $v \notin$ Sector(T) then T has a weakening contained in Sector(T) that fixes v.

(ii) The last move in every edge weakening of T is a non-backbite move involving one of its leaves.

Proof. Both (i) and (ii) follow by inspecting all the weakenings used in Lemmas 4.30, 4.26, and 4.29, and Propositions 4.27 and 4.28. Note that part (i) makes implicit use of Observation 4.20 (d). □

Lemma 4.32. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian path or cycle of G. Let e_F be an edge in H followed by a j-stack of A_0's, which is then followed by an A_1-type with switchable middle-box W. Let X and Y be the boxes adjacent to W that are not its H-neighbours. Assume that $F = G\langle N[P(X,Y)]\rangle$ is a standard looping fat path. Let T be an admissible turn of F, and let $\mu_1, \ldots, \mu_s$ be an edge-weakening of T that is contained in Sector(T), avoids the j-stack of A_0's and the boxes incident on e_F, and is such that after the application of μ_{s-1}, F remains a standard looping fat path in the resulting Hamiltonian path or cycle. Then we can extend $\mu_1, \ldots, \mu_s$ by a non-backbite cascade of length at most two, after which W is switched.

Proof. For definiteness, assume that T is northeastern with northern leaf $L_N = R(a,b)$ and (by Corollary 4.31 (ii)) that μ_s is the move $L_N \mapsto L'_N$, $L'_N \mapsto L_N$, or Sw(L_N). By Corollary 4.19(b), $L_N \in N[P] \setminus P$ and $L_N + (0,-1) \in P$. Note that we must have $S_{\rightarrow}(a, b-1; a+1, b-2) \in H$, $e(a, a+1; b) \notin H$ and $e(a+1; b-1, b) \notin H$. Now, either $L_N + (0,-1)$ is an end-box of P (Figure 4.96) or it is not (Figure 4.97).

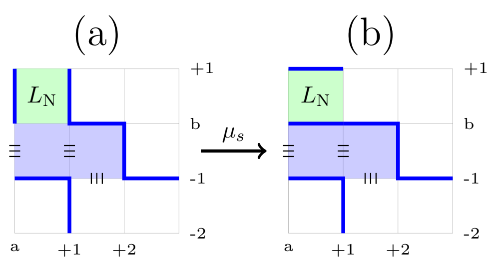

Fig. 4.96. L_N+$(0,-1)$ is not an end-box. (a) Before μ_s. (b) After μ_s.

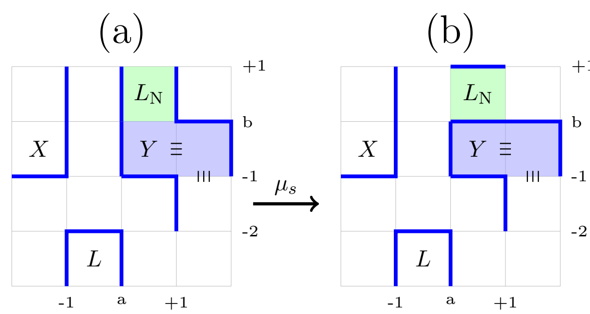

Fig. 4.97. L_N+$(0,-1)$ is an end-box. (a) Before μ_s. (b) After μ_s.

CASE 1: $L_N + (0,-1)$ is not an end-box of P. Then $e(a; b-1, b) \notin H$. Then, after μ_s, $L_N + (0,-1) \in P(X,Y)$ is switchable. Then, by Proposition 4.5 (a), there is a non-backbite cascade of length at most two, after which W is switched. End of Case 1.

CASE 2: $L_N + (0,-1)$ is an end-box of P. Assume that $L_N + (0,-1) = Y$. Then we have that $e(a; b-1, b) \in H$, that F is northern, that $e_F = e(a-1, a; b-2)$, and that $W = Y + (-1, 0)$. Then μ_s can be followed by $W \mapsto Y$. □

Now we are ready to give a proof of Lemma 3.13.

Lemma 3.13. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian cycle of G. Let C be a small cookie of G. Assume that G has only one large cookie, and that there is a j-stack of A_0 starting at the A_0-type containing C. Let L be the leaf in the top (j^{th}) A_0 of the stack, and assume that L is followed by an A_1-type. Let X and Y be the boxes adjacent to the middle-box of the A_1-type that are not its H-neighbours. If $P(X, Y)$ has no switchable boxes, then either:

(i) there is a cascade of length at most $\min(m, n)$, which avoids the stack of A_0's, and after which $P(X, Y)$ gains a switchable box, or

(ii) there is a cascade of length at most $\min(m, n) + 1$, that collects L and avoids the stack of A_0's.

Proof. Suppose that $P(X, Y)$ has no switchable boxes. Note that X and Y belong to the same H-component. Since H is a cycle, it follows that $P(X, Y)$ is contained in a standard looping fat path $F = G\langle N[P(X, Y)]\rangle$. By Lemma 4.21, F has an admissible turn T such that Sector(T) and the j-stack of A_0's are disjoint. Then, by Corollary 4.19(a), $d(T) \geq 3$. By Proposition 4.24, T has a weakening $\mu_1, \ldots, \mu_s$ of length at most $\min(m, n)$. By Observation 4.25, $\mu_1, \ldots, \mu_s$ is contained in Sector(T), and thus it avoids the j-stack of A_0's. By Lemma 4.32, and the fact that H is a cycle, after μ_s, either $P(X, Y)$ gains a switchable box Z, or $W \mapsto W'$ is a valid move, where $W' = X$ or $W' = Y$. If the former, then (i) holds; and if the latter, apply $W \mapsto W'$, and then (ii) holds. □.

Remark. As per the proof of Lemma 3.13, the first valid move that F will admit, either results in a switchable box for $P(X, Y)$, or makes available one of the moves $W \mapsto X$ and $W \mapsto Y$. Let $\min(m, n) = s$. It can be shown that there is no $j_0 \in \{1, \ldots, s\}$ such that after applying the cascade $\mu_1, \ldots, \mu_{j_0}$, $P(X, Y)$ gains a switchable box, or one of $W \mapsto X$ and $W \mapsto Y$ becomes available. However, it is easier to note that if such a j_0 were to exist, then we can just as well use this shorter cascade.

4.7 Summary

In this chapter we examined the structure of fat paths and used it to prove Lemma 3.13, which guarantees the existence of cascades required for the 1LC algorithm. Key definitions include fat paths, turns, weakenings, and the fat path conditions (FPC-1 through FPC-4).

In Section 4.1 we proved some structural properties of looping fat paths. Let F be such a path. Then F has no colinear edges other than those in its A_1-type (Proposition 4.14), and its boundary $B(F)$ is a Hamiltonian cycle of F (Lemma 4.13). Proposition 4.5 generalized the characterization of double-switch moves for Hamiltonian cycles of $m \times n$ grid graphs (Proposition 3.3) to Hamiltonian paths in polyominoes.

The two main structural results are that every fat path contains an admissible turn

(Proposition 4.18, Section 4.2) and that every admissible turn has a weakening (Propositions 4.24, 4.27, and 4.28 in Sections 4.4-4.6). We identified two types of weakenings: edge weakenings and end-vertex weakenings. Both are useful, with end-vertex weakenings required specifically when H is a Hamiltonian path and certain end-vertex configurations arise Section 4.6.

5 Reconfiguration algorithm for Hamiltonian paths

In this chapter, we prove that any two Hamiltonian paths on an $m \times n$ grid graph can be reconfigured into one another (Theorem 5.9). The proof relies on Theorem 2.1, which handles reconfiguration of Hamiltonian cycles. Our strategy is to reconfigure an arbitrary path into one whose end-vertices lie in corners of R_0. Once both end-vertices are positioned in corners, Proposition 5.4 shows that we can effectively treat the resulting Hamiltonian path as a Hamiltonian cycle, and so Theorem 2.1 applies.

We move the end-vertices to corners in two stages. First, the End-vertex-to-Boundary (EtB) algorithm moves an end-vertex to a side of R_0 (Proposition 5.6). Second, the End-vertex-to-Corner (EtC) algorithm moves the end-vertex from a side to a corner (Proposition 5.7). Both algorithms draw on Sections 1.3–1.5 of Chapter 1 and the fat path-turn-weakening pipeline from Chapter 4.

Definitions. Let G be an $m \times n$ grid graph. Throughout this paper, we position grid graphs in the first quadrant of the plane, so that $v(0,0)$ is the bottom-left corner and $v(m-1, n-1)$ is the top-right corner. We call the vertices $v(0,0)$, $v(0, n-1)$, $v(m-1, n-1)$, and $v(m-1, 0)$ the *corners* of R_0. A vertex $v(a, b)$ is called *even* if $a+b$ is even, and *odd* if $a+b$ is odd. We write $p(v)$ to denote the parity of a vertex v.

Let H be a Hamiltonian path of G with end-vertices u and v, with parities $p(u)$ and $p(v)$, respectively. We write $e(v)$ to denote the edge of H incident on v, and $e(u)$ to denote the edge of H incident on u.

Let s be a side of G with corners a and b. If $p(u) \in \{p(a), p(b)\}$ and $p(v) \in \{p(a), p(b)\}$, we call that side a *parity-compatible* side of G.

The existence of Hamiltonian paths on grid graphs was studied by Itai et al. in [14]. The following is a consequence of the main result (Theorem 3.2) in that paper.

Corollary 5.1. Let H be a Hamiltonian path in an $m \times n$ grid graph G with end-vertices u and v. Then:

(i) If m and n are both odd, then both u and v have even parity.

(ii) If at least one of m and n is even, then u and v have different parities. □

We record below Observation 5.2, which follows immediately from Corollary 5.1.

Observation 5.2. Let H be a Hamiltonian path in an $m \times n$ grid graph G, such that

$m, n \geq 3$. Then G has at least one pair of opposite sides that are parity-compatible. $\square$

Definitions. Let G^+_{north} be the $m \times (n+1)$ grid graph with vertex set $\{v(i,j) : 0 \leq i \leq m-1,\ 0 \leq j \leq n\}$ and edge set consisting of all edges between pairs of vertices at Euclidean distance 1. We refer to G^+_{north} as the *northern extension* of G. Define eastern extensions of G analogously.

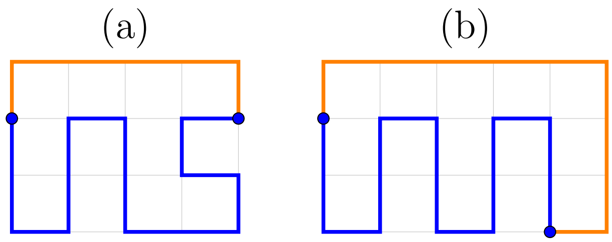

Fig. 5.1. (a) A northern path with a northern handle. (b) A northwest-southeast path with a northeastern handle.

A *northern Hamiltonian path* in G is a Hamiltonian path H_{north} with end-vertices $v(0, n-1)$ and $v(m-1, n-1)$, the top-left and top-right corners of G (Figure 5.1 (a)). Define eastern, southern, and western Hamiltonian paths of G analogously. A *northwest-southeast Hamiltonian path* in G is a Hamiltonian path $H_{\text{NW-SE}}$ with end-vertices $v(0, n-1)$ and $v(m-1, 0)$, the top-left and bottom-right corners of G (Figure 5.1 (b)). Define a southwest-northeast Hamiltonian path analogously.

We call the set of edges $\{e(0; n-1, n), e(0,1; n), \ldots, e(m-2, m-1; n), e(m-1; n-1, n)\}$ the *northern handle* of H_{north} and denote it by $\text{Handle}_{\text{north}}(H)$. Define the eastern, southern, and western handle of a Hamiltonian path analogously. We call the set of edges $\{e(0; n-1, n), e(0,1; n), \ldots, e(m-2, m-1; n),\ e(m-1, m; n), e(m; n-1, n), \ldots,\ e(m; 0, 1), e(m-1, m; 0)\}$ the *northeastern handle* of $H_{\text{NW-SE}}$ and denote it by $\text{Handle}_{\text{NE}}(H)$. Define the *southwestern handle* of $H_{\text{SW-NE}}$ analogously

We call the Hamiltonian cycle of G^+_{north} obtained by adding $\text{Handle}_{\text{north}}(H)$ to a northern Hamiltonian path H the *handle-path cycle* of H, and we denote it by $\overline{H}$. Note that $\overline{H}$ lies in G^+_{north} and is defined only when H is a northern Hamiltonian path. Analogous statements hold for eastern, southern, and western handle-path cycles.

Lemma 5.3. Let G be an $m \times n$ grid graph with $m, n \geq 3$, let H be a northern Hamiltonian path in G and let X and Y be boxes of G. Assume that $X \mapsto Y$ is a valid move for $\overline{H}$, after which we obtain $\overline{H}'$. Then $X \mapsto Y$ is valid for H, yielding a northern Hamiltonian path H', where $H' = \overline{H}' \setminus \text{Handle}_{\text{north}}(H)$. An analogous statement holds for eastern, southern, and western Hamiltonian paths of G.

Proof. Let $H = v_1, \ldots, v_r$ be a northern Hamiltonian path of G, and let $v_{r+1} \ldots, v_1$ be its northern handle. Let $\overline{H} = v_1, .., v_r, v_{r+1} \ldots, v_1$, and let $\overline{H}'$ be the Hamiltonian cycle obtained after applying $X \mapsto Y$ to $\overline{H}$. Let (v_s, v_{s+1}) and (v_t, v_{t+1}) be the edges of X in $\overline{H}$. By Corollary 1.5.2, X is anti-parallel. Then v_s is adjacent to v_{t+1} and v_{s+1} is adjacent to v_t. After switching X we obtain two cycles $H_1 = P(v_{s+1}, v_t), \{v_{s+1}, v_t\}$ and $H_2 = P(v_1, v_s), \{v_s, v_{t+1}\}, P(v_{t+1}, v_r),\ P(v_r, v_1)$. Since the (H_1, H_2)-port Y is contained in G, neither of its edges may belong to $P(v_r, v_1)$. Then Y must be a port between the cycle H_1 and the sub-path $P(v_1, v_s), \{v_s, v_{t+1}\}, P(v_{t+1}, v_r)$ of H_2. Since $X \mapsto Y$ does not interact with $\text{Handle}_{\text{north}}(H)$, it follows that $X \mapsto Y$ is a valid move for H and that $H' = \overline{H}' \setminus \text{Handle}_{\text{north}}(H)$, where H' is the northern Hamiltonian path obtained after applying $X \mapsto Y$ to H. □

Proposition 5.4. Let H and K be northern Hamiltonian paths in an $m \times n$ grid graph G with $n \geq m \geq 3$. Then there is a sequence of at most n^2m double-switch moves that reconfigures H into K. An analogous statement holds for eastern, southern, and western Hamiltonian paths of G.

Proof. By Theorem 2.1, there is a sequence $\mu_1, \ldots, \mu_s$ of double-switch moves that reconfigures $\overline{H}$ into $\overline{K}$. Recall that no new cookies are created by the RtCF algorithm, and that by definition, cascades create no new cookies. Because there are no cookie necks on $\text{Boxes}(G^+_{\text{north}}) \setminus \text{Boxes}(G)$, and because even when the neck of a large cookie is relocated, it remains on the same side of the boundary,[8] the entire sequence $\mu_1, \ldots, \mu_s$ is contained within G. Then, by Lemma 5.3 and induction, after applying $\mu_1, \ldots, \mu_s$ to H, we obtain K. The proof for the case where the Hamiltonian paths are eastern, southern, or western is similar, so we omit it. □

Corollary 5.4.1. Let H and K be either both northwest–southeast Hamiltonian paths or both southwest–northeast Hamiltonian paths of an $m \times n$ grid graph G, with $n \geq m \geq 3$. Then there is a sequence of at most n^2m double-switch moves that reconfigures H into K.

Proof. For definiteness, assume that H is northwest-southeast. We can use a northeastern handle to obtain a handle-path cycle of H. Then the result follows by using essentially the same arguments as in Lemma 5.3 and Proposition 5.4. □

---

[8] These moves appear in the proof of Lemma 4.23 in Case 1.2(b), the first paragraph of Case 2, and Case 2.1(b_1).

Proposition 5.5. Let G be an $m \times 2$ grid graph, and let H and K be any two Hamiltonian paths of G. Then there is a sequence of at most m backbite moves that reconfigures H into K.

Proof. Let H be a Hamiltonian path of the $m \times n$ grid graph G with end-vertices u and v. By Corollary 5.1, since $n = 2$, u and v have distinct parities. This means that G does have Hamiltonian e-cycles. We define the canonical forms of G to be the set of e-cycles of G. We show below (Lemma 5.10) that any two e-cycles can be reconfigured into one another by at most m backbite moves. It remains to show that a Hamiltonian path of G can be reconfigured into an e-cycle by a sequence of at most m backbite moves.

If H is already an e-cycle we are done, so assume that u and v are not adjacent. For definiteness, assume that u is east of v, and that $v = v(k, 1)$ lies on the northern side of R_0. It is enough to check that there is a backbite move that relocates v to a vertex with x-coordinate one unit greater than v. Let $e(v)$ be the edge of H incident on v. Either $k = 0$, or $k > 0$.

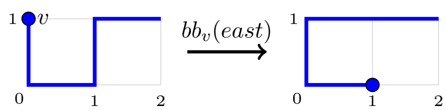

Fig. 5.2. Case 1.

CASE 1: $k = 0$. Note that if $e(v)$ is eastern, then $u = v(0, 0)$, and H is an e-cycle, so we may assume that $e(v)$ is not eastern. Then $e(v)$ must be southern. Then $e(0, 1; 0) \in H$ and $S_{\uparrow}(1, 0; 2, 1) \in H$. Then $bb_v(east)$ relocates v east, to $v(1, 0)$. See Figure 5.2. End of Case 1.

CASE 2: $k > 0$. Then $e(v)$ is eastern, southern, or western. We will see that the first two possibilities lead to invalid configurations.

CASE 2.1: $e(v)$ is eastern. Then we must have that $e(k-1, k; 0) \in H$ and $e(k, k+1; 0) \in H$. Let $v_s = v(k, 0)$. Either $e(k+1; 0, 1) \in H$, or $e(k+1; 0, 1) \notin H$.

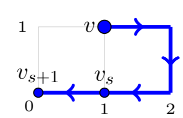

Fig. 5.3. Case 2.1(a).

CASE 2.1(a): $e(k+1; 0, 1) \in H$. Then $v_{s+1} = v(k-1, 0)$. Since the x-coordinate of u is greater than k, and the x-coordinate of v_{s+1} is $k-1$, the subpath $P(v_{s+1}, u)$ must contain a vertex of H with x-coordinate k. But the only two vertices of G with x-coordinate k are v_1 and v_s, and neither belongs to $P(v_{s+1}, u)$. See Figure 5.3. End of Case 2.1(a).

CASE 2.1(b): $e(k+1; 0, 1) \notin H$. Then, by Corollary 1.5.2, $R(k, 0)$ is anti-parallel. Then $v_{s+1} = v(k-1, 0)$, and now, we can use the same argument as in Case 2.1(a) to check that this configuration is invalid. End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $e(v)$ is southern. The arguments for this case are similar to the one used in Case 2.1(a), so we omit the proof. End of Case 2.2.

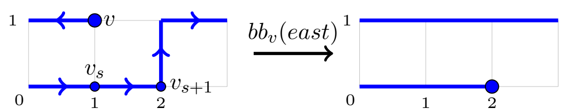

Fig. 5.4. Case 2.3.

CASE 2.3: $e(v)$ is western. Then we must have $S_{\uparrow}(k+1,0;k+2,2) \in H$, $e(k-1,k;0) \in H$, and $e(k,k+1;0) \in H$. Let $v_s = v(k,0)$. Note that, using an argument similar to the one in Case 2.1, we can show that, whether $e(k-1;0,1) \in H$ or $e(k-1;0,1) \notin H$, $v_{s+1} = v(k+1,0)$. Then, $bb_v(east)$ relocates v east, to $v(k+1,0)$. See Figure 5.4. End of Case 2.3. End of Case 2. □

Proposition 5.6. *(End-vertex-to-Boundary (EtB) Algorithm)* Let H be a Hamiltonian path of an $m \times n$ grid graph G, with end-vertices u and v, and let s and s' be a pair of opposite parity-compatible sides of R_0[9]. Then:

(i) If neither of u and v is on s or s', then there is a cascade after which one of u and v is on s or s'.

(ii) If u is on s and v is not on s', then there is a cascade that fixes u, after which v is on s'.

Either cascade has length at most $n^2 + mn + 4n$.

Proposition 5.7. *(End-vertex-to-Corner (EtC) Algorithm)* Let G be an $m \times n$ grid graph, and let H be a Hamiltonian path of G with end-vertices u and v. Let s_u and s_v be a pair of opposite sides of R_0, and assume that u is on s_u, and v is on s_v. Assume further that c_v is a corner of s_v that has the same parity as v, and that c_u is a corner of s_u that has the same parity as u. Then:

(i) If u is not at c_u and v is not at c_v, then there is a cascade after which v is at c_v and u is on s_u.

(ii) If u is at c_u and v is not at c_v, then there is a cascade after which v is at c_v and u is at c_u.

Either cascade has length at most $\frac{n^2}{2} + mn + \frac{5n}{2}$.

The proofs of Propositions 5.6 and 5.7 are given in Sections 5.1 and 5.2, respectively.

Proposition 5.8. Let H be a Hamiltonian path in an $m \times n$ grid graph G with $n \geq m \geq 3$.

[9] By Observation 5.2, there is such a pair.

Then there is a cascade of length at most $\frac{7n^2}{2}+5mn+\frac{31n}{2}$ that reconfigures H into a northern or eastern Hamiltonian path.

Proof. Since m and n must be both odd, both even, or differ in parity, we consider three cases.

CASE 1: m and n are odd. Then all four corners of R_0 have even parity, and by Corollary 5.1(i), so do u and v. Using EtB, we send u and v to the western and eastern sides of R_0, respectively. Using EtC we send u and v to the northwestern and northeastern corners of R_0, respectively. Then the resulting Hamiltonian path is northern. End of Case 1.

CASE 2: m and n are even. By Corollary 5.1(ii), u and v have distinct parities. WLOG assume that u is odd and v is even. Note that the corners $v(0,0)$ and $v(m-1,n-1)$ are even, and the corners $v(0,n-1)$ and $v(m-1,0)$ are odd. This means that each side of R_0 is incident on an even corner and an odd one. Using EtB, we send u and v to opposite sides of R_0. WLOG assume that we send u to the western side and v to the eastern side. Using EtC, we send u to $v(0,n-1)$ and v to $v(m-1,n-1)$. Then the resulting Hamiltonian path is northern. End of Case 2.

CASE 3: m and n have different parities. By Corollary 5.1(ii), u and v have distinct parities. WLOG assume that u is even and v is odd, and that m is even and n is odd. Then $v(0,0)$ and $v(0,n-1)$ are even, and $v(m-1,0)$ and $v(m-1,n-1)$ are odd. Using EtB, we send one of u and v to the northern side of R_0, and the other to the southern side. Either u has been sent to the northern side or to the southern side.

CASE 3.1: u has been sent to the northern side. Using EtC, we send u to $v(0,n-1)$ and v to $v(m-1,0)$. Using EtC again, we send v to $v(m-1,n-1)$. Then the resulting Hamiltonian path is northern. End of Case 3.1.

CASE 3.2: u has been sent to the southern side. Using EtC, we send u to $v(0,0)$ and v to $v(m-1,n-1)$. Using EtC again, we send u to $v(0,n-1)$. Then the resulting Hamiltonian path is northern. End of Case 3.2. End of Case 3.

Note that it takes at most two applications of EtB and three applications of EtC to reconfigure a Hamiltonian path into an eastern or a northern one. So we need at most $2(n^2+mn+4n)+3(\frac{n^2}{2}+mn+\frac{5n}{2})=\frac{7n^2}{2}+5mn+\frac{31n}{2}$ moves to complete the reconfiguration. □

Theorem 5.9. Let G be an $m \times n$ grid graph with $n \geq m$, and let H and K be any two Hamiltonian paths of G. Then there is a cascade with at most $n^2m + 7n^2 + 10mn + 31n$ moves that reconfigures H into K.

Proof. By Proposition 5.5, we may assume that $m, n \geq 3$. By Proposition 5.8, we may assume WLOG that H and K can be reconfigured into the northern Hamiltonian paths H' and K' by the cascades $\mu_1, \ldots, \mu_s$ and $\nu'_1, \ldots, \nu'_t$. By Proposition 5.4, H' can be reconfigured into K' by a cascade $\eta_1, \ldots, \eta_q$. For $j \in \{1, \ldots, t\}$ let ν_j be the inverse of the move ν'_j. Then $\mu_1, \ldots, \mu_s$, $\eta_1, \ldots, \eta_q$, $\nu_t, .., \nu_1$ is a cascade that reconfigures H into K.

The bound is the sum of the bound of Theorem 2.1 with twice that of Proposition 5.8. □

The rest of the chapter is structured as follows. Section 5.1 handles certain configurations that appear frequently in the proofs of both Propositions 5.6 and 5.7. We isolate these scenarios as separate lemmas, which we then invoke as needed in the main proofs. The proof of Proposition 5.6 is given in Section 5.2, and the proof of Proposition 5.7 is given in Section 5.3. Both proofs draw on Sections 1.3–1.5 of Chapter 1, the fat path-turn-weakening pipeline from Chapter 4, and the lemmas from Section 5.1.

5.1 Recurring Scenarios

Lemma 5.10. Let G be an $m \times n$ grid graph, and let H be Hamiltonian path with adjacent end-vertices. Then there is a cascade of backbite moves that reconfigures H into an e-cycle.

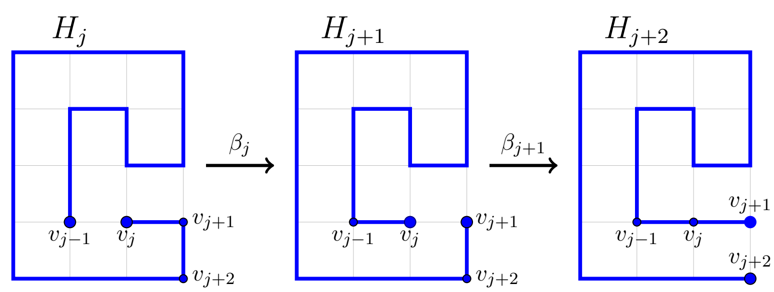

Fig. 5.5. An illustration of a cascade that reconfigures a Hamiltonian path H_j into an e-cycle H_{j+2}.

Proof. Orient $H = H_1$ as $v_1, \ldots, v_r$. For each $j \in \{1, \ldots, r-1\}$, let $H_j = v_j, v_{j+1}, \ldots, v_r, v_1, \ldots, v_{j-1}$, which is a Hamiltonian path. For each j, let $(v_j, v_{j+1}) \mapsto \{v_{j-1}, v_j\}$ be the backbite move β_j. Then that β_j reconfigures H_j into H_{i+1}. Let j_0 be the smallest index in $\{1, \ldots, r\}$ such that the edge $\{v_{j_0}, v_{j_0+1}\}$ of $G \setminus H_{j_0+1}$ is in R_0. Note that H_{j_0+1} is an e-cycle. So, using j_0 backbite moves, we have reconfigured H_1 into the e-cycle H_{j_0+1}. See Figure 5.5.

Lemma 5.11. Let H be a Hamiltonian path of an $m \times n$ grid graph G. Let $e(k, k+1; l)$ be an edge of H followed southward by a j-stack of A_0's, which is not followed by an A_1-type. Let $R(k, b)$ be the leaf in the j^{th} A_0 of the stack, and assume that the rectangle $R(k-1, k+2; b-2, b-1)$ does not contain an end-vertex of H. Then there is a cascade of flips, after which $e(k, k+1; l-1)$ is in the resulting Hamiltonian path of G. Analogous results apply for the cases where the j-stack follows an edge in another direction.

Proof. Let $R(k, b)$ be the southern leaf of the last (j^{th}) A_0-type in the j-stack of A_0's that follows $e(k, k+1; l)$. Now either $e(k, k+1; b-1) \in H$, or $e(k, k+1; b-1) \notin H$.

CASE 1: $e(k, k+1; b-1) \in H$. Then $R(k, b-1) \mapsto R(k, b)$, $R(k, b+1) \mapsto R(k, b+2)$, $\ldots$, $R(k, l-3) \mapsto R(k, l-2)$, is a cascade, after which $e(k, k+1; l-1)$ is in the resulting Hamiltonian path. See Figure 5.6 (a). End of Case 1.

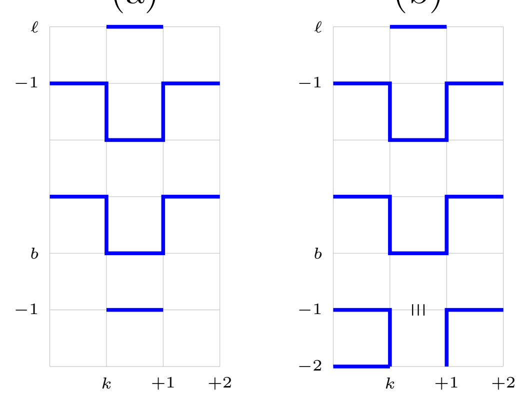

Fig. 5.6. (a) Case 1. (b) Case 2.

CASE 2: $e(k, k+1; b-1) \notin H$. Since the rectangle $R(k-1, k+2; b-2, b-1)$ does not contain an end-vertex of H, we must have $S_{\rightarrow}(k-1, b-1; k, b-2) \in H$ and $S_{\uparrow}(k+1, b-2; k+2, b-1) \in H$. Since $R(k, b)$ is not followed by an A_1-type, at least one of $e(k-1, k; b-2)$ and $e(k+1, k+2; b-2)$ is in H. By symmetry, we may assume WLOG that $e(k-1, k; b-2) \in H$. Then, after $R(k, b-2) \mapsto R(k-1, b-2)$, $e(k, k+1; b-1)$ is in the resulting Hamiltonian cycle, and we are back to Case 1. See Figure 5.6 (b). End of Case 2. □

Lemma 5.12. Let H be a Hamiltonian path of an $m \times n$ grid graph G, and let $n-1$ be even. Fix $k \in \{0, .., m-1\}$, and assume that the end-vertices of H have x-coordinate at most k. A configuration in which for each $i \in \{0, 2, \ldots, n-1\}$ we have that $e(k, k+1; i) \in H$ and that $e(k, k+1; i+1) \notin H$ is not a possible configuration of H. An analogous statement applies for the case where these edges are vertical.

Proof. Let $EE(q)$ be the statement "For every $i \in \{0, 2, \ldots, n-1\}$, $e(q, q+1; i) \in H$ and $e(q, q+1; i+1) \notin H$". Let $(*_1)$ be the statement "$EE(q)$ implies $EE(q+2)$, whenever $q \geq k$", and assume for now that $(*_1)$ is true. For a contradiction, assume that $EE(k)$ is true. Then $m-1-k$ is odd or $m-1-k$ is even.

CASE I. $m-1-k$ is odd. By $(*_1)$ and induction, $EE(m-2)$ is true. Then we must have $e(m-1;0,1) \in H$ and $e(m-1;1,2) \in H$. But now, since v_1 and v_r are west of $x=k$, H misses $v(m-1,3)$. See Figure 5.7(a). End of proof for Case I.

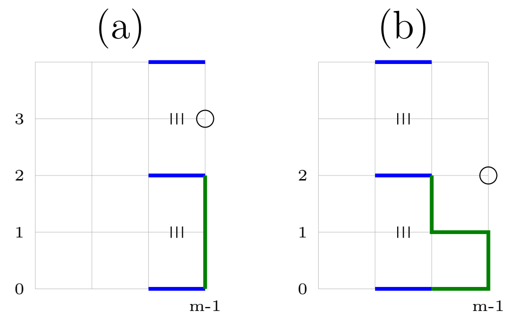

Fig. 5.7. (a) Case I. (b) Case II.

CASE II. $m-1-k$ is even. By $(*_1)$ and induction, $EE(m-3)$ is true. Then we must have $S_{\rightarrow}(m-2,0;m-1,1) \in H$. It follows that $S_{\downarrow}(m-2,2;m-1,1) \in H$. But now, since v_1 and v_r are west of $x=k$, H misses $v(m-1,2)$. See Figure 5.7 (b). End of proof for Case II. It remains to prove $(*_1)$.

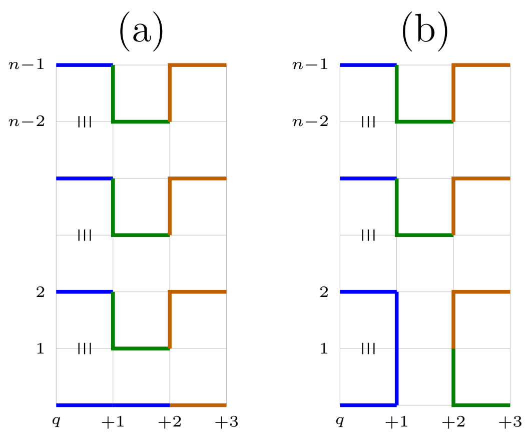

Fig. 5.8. (a) Case 1. (b) Case 2.

Proof of $(*_1)$. Assume that $EE(q)$ is true. Then, either $e(q+1,q+2;0) \in H$ or $e(q+1;0,1) \in H$. See Figure 5.8. Initial configurations in blue.

CASE 1: $e(q+1,q+2;0) \in H$. Then we must have that $S_{\downarrow}(q+1,2;q+2,1) \in H$, which implies $S_{\downarrow}(q+1,4;q+2,3) \in H$, which implies $\ldots$, which implies $S_{\downarrow}(q+1,n-1;q+2,n-2) \in H$. It follows that $S_{\uparrow}(q+2,n-2;q+3,n-1) \in H$, which implies $\ldots$, which implies $S_{\uparrow}(q+2,1;q+3,2) \in H$. Then $e(q+2,q+3;0)$ must be in H as well, and so, $EE(q+2)$ is true. End of proof for Case 1.

CASE 2: $e(q+1;0,1) \in H$. Either $e(q+1;1,2) \in H$ or $e(q+1,q+2;1) \in H$.

CASE 2.1: $e(q+1;1,2) \in H$. As in Case 1, we must have $S_{\downarrow}(q+1,4;q+2,3) \in H, \ldots, S_{\downarrow}(q+1,n-1;q+2,n-2) \in H$. It follows that $S_{\uparrow}(q+2,n-2;q+3,n-1) \in H, \ldots, S_{\uparrow}(q+2,1;q+3,2) \in H$. Then we must have $S_{\downarrow}(q+2,1;q+3,0)$ in H as well, and so, $EE(q+2)$ is true. End of proof for Case 2.1.

CASE 2.2: $e(q+1,q+2;1) \in H$. Let $Q(j)$ be the statement "$e(q+1;j,j+1) \in H$ and $e(q+1,q+2;j+1) \in H$". Note that $Q(0)$ is true. Then either $Q(j)$ is true for each $j \in \{0,2,\ldots,n-3\}$ or there is some $j_0 \in \{0,2,\ldots,n-3\}$ such that $Q(j)$ is true for all $j \le j_0$ but $Q(j_0+2)$ is not true.

CASE 2.2(a): $Q(j)$ is true for each $j \in \{0,2,\ldots,n-3\}$. This implies that $S_{\downarrow}(q+2,1;q+3,0),\ldots,S_{\downarrow}(q+1,n-2;q+2,n-3)$ belong to H. Then we must have $e(q,q+1;n-1) \in H$ and $e(q+1,q+2;n-1) \in H$ as well, and so, $EE(q+2)$ is true. See Figure 5.9. End of proof for Case 2.2(a).

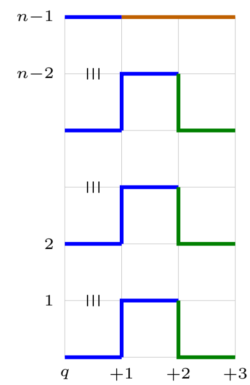

Fig. 5.9. Case 2.2(a).

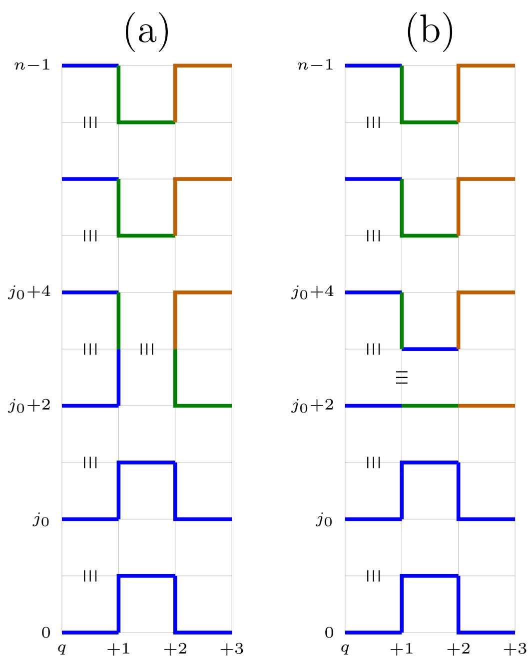

Fig. 5.10. (a) Case 2.2(b_1). (b) Case 2.2(b_2).

CASE 2.2(b): There is some $j_0 \in \{0,2,\ldots,n-3\}$ such that $Q(j)$ is true for all $j \le j_0$ but $Q(j_0+2)$ is not true. Then $S_{\downarrow}(q+2,1;q+3,0),\ldots,S_{\downarrow}(q+2,j_0+1;q+3,j_0)$ belong to H. Note that if $e(q+1;j_0+2,j_0+3) \notin H$ and $e(q+1,q+2;j_0+3) \notin H$ then H misses $v(q+1,j_0+3)$. Therefore we only need to check two cases: Case b_1: $e(q+1;j_0+2,j_0+3) \in H$ and $e(q+1,q+2;j_0+3) \notin H$ and b_2: $e(q+1;j_0+2,j_0+3) \notin H$ and $e(q+1,q+2;j_0+3) \in H$.

CASE 2.2(b_1): $e(q+1;j_0+2,j_0+3) \in H$ and $e(q+1,q+2;j_0+3) \notin H$. Then $e(q+1;j_0+3,j_0+4) \in H$, $S_{\downarrow}(q+2,j_0+3;q+3,j_0+2) \in H$ and $S_{\downarrow}(q+1,j_0+6;q+2,j_0+5),\ldots,S_{\downarrow}(q+1,n-1;q+2,n-2)$ all belong to H. It follows that $S_{\uparrow}(q+2,n-2;q+3,n-1),\ldots,S_{\uparrow}(q+2,j_0+3;q+3,j_0+4)$ all belong to H as well, and so $EE(q+2)$ is true. See Figure 5.10 (a). End of proof for Case 2.2(b_1).

CASE 2.2(b_2): $e(q+1;j_0+2,j_0+3) \notin H$ and $e(q+1,q+2;j_0+3) \in H$. We must have $e(q+1,q+2;j_0+2) \in H$ and $e(q+1;j_0+3,j_0+4) \in H$. It follows that $S_{\downarrow}(q+1,j_0+6;q+2,j_0+5),\ldots,S_{\downarrow}(q+1,n-1;q+2,n-2)$ all belong to H. Then we must have

$S_{\uparrow}(q+2, n-2; q+3, n-1), \ldots, S_{\uparrow}(q+2, j_0+3; q+3, j_0+4)$ in H as well. Then we must have $e(q+2, q+3; j_0+2) \in H$, and so $EE(q+2)$ is satisfied. See Figure 5.10 (b). End of proof for Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2. End of proof for $(*_1)$. □

Lemma 5.13. Let H be a Hamiltonian path of an $m \times n$ grid graph G. Fix $k \in \{0, .., m-1\}$, and assume that the end-vertices of H have x-coordinate at most k. Assume that $n-1$ is odd, and let $R(k+1, 0)$ be a northern leaf on the southern boundary. The configuration in which there is a full stack of A_0's starting at $R(k+1, 0)$ is not a possible configuration of H.

Analogous statements apply for the other compass directions.

Proof. Note that if $k+3 = m-1$, then $e(k+3; 0, 1) \in H$ and $e(k+3; 1, 2) \in H$. Then $e(k+3; 3, 4) \in H$, or $e(k+3; 4, 5) \in H$. If the former, then $v(k+3, 3)$ is an end-vertex of H, contradicting the assumption that the end-vertices have x-coordinate at most k; and if the latter, then H misses $v(k+3, 3)$. Then we must have that, $m-1 \geq k+4$. See Figure 5.11(a).

If $k+4 = m-1$, then we must have that $S_{\rightarrow}(k+3, 0; k+4, 1) \in H$. This implies that $S_{\downarrow}(k+3, 2; k+4, 1) \in H$. But then again, as in the previous paragraph, either H misses $v(k+4, 2)$ or $v(k+4, 2)$ is an end-vertex. So, we may assume that $m-1 \geq k+5$. See Figure 5.11(b).

Let $SS(q)$ be the statement "There is a full stack of A_0s starting at $R(q, 0)$". Let $(*_2)$ be the statement "$SS(q)$ implies $SS(q+2)$, whenever $q > k$", and assume for now that $(*_2)$ is true. For a contradiction, assume that $SS(q+1)$ is true. Then, $m-1-k$ is odd or it is even. In each case, the arguments are similar to those in Cases I and II of Lemma 5.12, so we omit them.

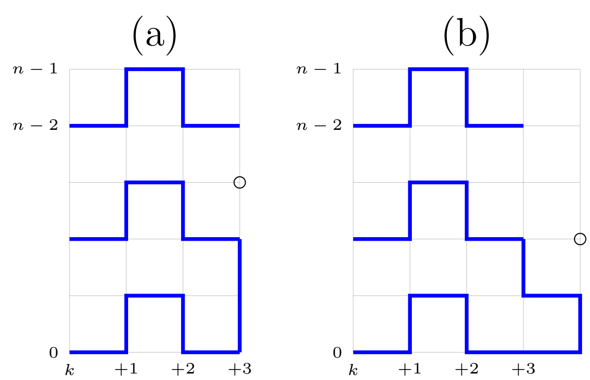

Fig. 5.11.

It remains to prove $(*_2)$.

Proof of $(*_2)$. Assume $SS(q)$. Then we must have that $S_{\uparrow}(q+2, n-2; q+3, n-1) \in H$, which implies $S_{\uparrow}(q+2, n-4; q+3, n-3) \in H$, which implies $\ldots$, which implies $S_{\uparrow}(q+2, 0; q+3, 1) \in H$. It follows that $S_{\downarrow}(q+3, 1; q+4, 0) \in H$, which implies $\ldots$, which implies $S_{\uparrow}(q+3, n-1; q+4, n-2) \in H$. So, $SS(q+2)$ is true. □

Observation 5.14. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian path of G with end-vertices u and v not adjacent (NAA). Let J be an H-component with more than one box, and let N_J be the neck of J. If N_J is not switchable, then:

(a) exactly one of u and v (say, u) is incident on the neck-edge of N_J,

(b) u is not at a corner of R_0, and

(c) $e(u) \in R_0$.

Lemma 5.15. Let G be an $m \times n$ grid graph, and let H be a Hamiltonian path of G with end-vertices u and v not adjacent. Let W be a switchable box in H, and let X and Y be the boxes adjacent to W that are not its H-neighbours. Assume that X and Y belong to distinct H-components J_X and J_Y. Then:

(a) If at least one of u and v does not lie on R_0, then there exists a cascade $\mu_1, \dots, \mu_s$, such that μ_s switches W and fixes both u and v, and $s \in \{1, 2\}$.

(b) If at least one of u and v lies on a corner of R_0, then there exists a cascade $\mu_1, \dots, \mu_s$ such that μ_s switches W and fixes both u and v, and $s \in \{1, 2\}$.

(c) If both end-vertices u and v lie on R_0, then there exists a cascade that keeps v on the same side of R_0 and that switches W and fixes u, and there exists another cascade that keeps u on the same side of R_0 and that switches W and fixes v. Each cascade has length at most three.

Proof. Orient H as $v = v_1, v_2, \dots, v_r = u$. Note that if W is parallel then we are done after $\mathrm{Sw}(W)$, so we may assume that W is anti-parallel. Let e_X and e_Y be the edges of W in H, let $X = \Phi(e_X, \mathrm{right})$, let $\Phi(e_Y, \mathrm{right}) = Y$, and WLOG assume that e_X precedes e_Y in H. Let N_{J_X} and N_{J_Y} be the necks of J_X and J_Y, respectively, and let $\{v_s, v_t\}$ be the neck-edge of J_X. As in Corollary 1.3.8, the neck-edge of J_X splits H into three subpaths $P_1 = P(v_1, v_s)$, $P_2 = P(v_s, v_t)$ and $P_3 = P(v_t, v_r)$. Now, either at least one of N_{J_X} and N_{J_Y} is switchable, or neither is.

CASE 1. At least one of N_{J_X} and N_{J_Y} is switchable. It follows from Observation 5.14 that if at least one of u and v does not lie on R_0 or at least one of u and v lies on a corner of R_0, then at least one of N_{J_X} and N_{J_Y} is switchable. Therefore, parts (a) and (b) will follow from this case alone.

WLOG, assume that N_{J_X} is switchable. Then N_{J_X} is parallel or N_{J_X} is anti-parallel.

CASE 1.1: N_{J_X} is anti-parallel. Then, J_X contains neither end-vertex, or it contains exactly one end-vertex, or it contains both.

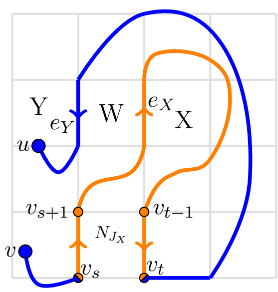

Fig. 5.12. Case 1.1(a).

CASE 1.1(a): J_X contains neither end-vertex. We will first check that $e_X \in P_2$ and $e_Y \notin P_2$. Since $X = \Phi(e_X, \text{right})$, by Corollary 1.3.10, $e_X \in \overrightarrow{K}_{J_X}$. By Corollary 1.3.11, since J_X contains no end-vertices of H, $\overrightarrow{K}_{J_X} = \overrightarrow{K}_{P_2}$. By construction, the directed subpath P_2 of H is identical to $\overrightarrow{K}_{P_2}$. Thus, $e_X \in P_2$. By Corollary 1.3.11, $\Phi(P_2, \text{right}) = J_X$. For contradiction, assume that $e_Y \in P_2$. By Observation 1.2.4 (a), $\Phi(P_2, \text{right}) \supset \Phi(\overrightarrow{K}(e_X, e_Y), \text{right})$. But then there is an H-path $P(X, Y) \subseteq \Phi(\overrightarrow{K}(e_X, e_Y), \text{right}) \subset \Phi(P_2, \text{right}) = J_X$, contradicting that X and Y belong to distinct components. Thus, $e_Y \notin P_2$.

Next we check that the edges of N_{J_X} in H are (v_s, v_{s+1}) and (v_{t-1}, v_t). Since $\Phi(P_2, \text{right}) = J_X$, N_{J_X} must be added by an edge in P_2. Since the neck-edge of N_{J_X} is $\{v_s, v_t\}$, the edge of H incident on v_s adding N_{J_X} is either (v_{s-1}, v_s) or (v_s, v_{s+1}). But $v_{s-1} \notin P_2$. Thus N_{J_X} must be added by (v_s, v_{s+1}). Using the same argument, we see that the other edge of N_{J_X} in H must be (v_{t-1}, v_t). See Figure 5.12.

Now, after switching N_{J_X}, $e_X \in H_c = P(v_{s+1}, v_{t-1}), \{v_{t-1}, v_{s+1}\}$ and $e_Y \in H_p = P(v_1, v_s), \{v_s, v_t\}, P(v_t, u)$. Then W is an (H_c, H_p)-port and so $N_{J_X} \mapsto W$ is a valid move that switches W and fixes both end-vertices of H. End of Case 1.1(a).

CASE 1.1(b): J_X contains exactly one end-vertex. WLOG assume that J_X contains v. Corollary 1.5.5 implies that $v \notin J_X \setminus B(J_X)$. Then $v \in B(J_X)$. By Corollary 1.3.11(d), $B(J_X) = P_2, \{v_s, v_t\}$. It follows that $v = v_s$, so $P_1 = v$ is trivial. Then, by Corollary 1.3.11(b), $\overrightarrow{K}_{J_X} = \overrightarrow{K}_{P_1}, \overrightarrow{K}_{P_2} = \overrightarrow{K}_{P_2}$. The rest of the argument is the same as in Case 1.1(a), so we omit it. End of Case 1.1(b)

CASE 1.1(c): J_X contains both end-vertices. By Corollary 1.5.5, either u and v are both in $J_X \setminus B(J_X)$, or neither is.

CASE 1.1(c_1): Neither u nor v is in $J_X \setminus B(J_X)$. Then both u and v must belong to $B(J_X)$. Using Corollary 1.1.13(d) again, we see that we must have $v = v_s$ and $u = v_t$. As in Case 1.1(b), this means that $\overrightarrow{K}_{J_X} = \overrightarrow{K}_{P_2}$. The rest of the argument is the same as in Case 1.1(a), so we omit it. End of Case 1.1(c_1).

CASE 1.1(c_2): both u and v belong to $J_X \setminus B(J_X)$. Then J_Y contains neither end-vertex. By Observation 5.14, N_{J_Y} is switchable. Corollary 1.5.3 implies that N_{J_Y} is anti-parallel. Then, by Case 1.1(a), $N_{J_Y} \mapsto W$ is a valid move, that switches W and fixes both end-vertices of H. End of Case 1.1(c_2). End of Case 1.1(c). End of Case 1.1.

CASE 1.2: N_{J_X} is parallel. Then there are two subcases to consider:

(a) $e_X \in P_i$, $e_Y \in P_j$, with $(i, j) \in \{(1, 1), (2, 2), (3, 3), (1, 3), (3, 1)\}$.

(b) $e_X \in P_i$, $e_Y \in P_j$, with $(i, j) \in \{(1, 2), (2, 1), (2, 3), (3, 2)\}$, and

CASE 1.2(a). We will check that each subcase in (a) is impossible.

CASE 1.2(a_1): $(i, j) = (1, 1)$ or $(i, j) = (3, 3)$. By symmetry, assume WLOG that $(i, j) = (1, 1)$. By Corollary 1.3.8 (b), $\Phi(\overrightarrow{K}_{P_1}, \text{right}) \subset J_X$ or $\Phi(\overrightarrow{K}_{P_1}, \text{right}) \subset G_{-1} \setminus J_X$. If the former, then there is an H-path $P(X, Y)$ contained in $\Phi(\overrightarrow{K}(e_X, e_Y), \text{right}) \subseteq \Phi(\overrightarrow{K}_{P_1}\text{right})$, contradicting that X and Y belong to distinct components. If the latter, then $X = \Phi(e_X, \text{right}) \subset \Phi(\overrightarrow{K}_{P_1}, \text{right})$ implies that $X \in G_{-1} \setminus J_X$, contradicting $X \in J_X$. End of Case 1.2(a_1).

CASE 1.2(a_2): $(i, j) = (2, 2)$. Since $X = \Phi(e_X, \text{right})$, Corollary 1.3.8 (a) $\Phi(P_2, \text{right}) \subseteq J_X$. For contradiction, assume the $e_Y \in P_2$ as well. Then there is an H-path $P(X, Y)$ contained in $\Phi(\overrightarrow{K}(e_X, e_Y), \text{right}) \subset \Phi(P_2, \text{right})$, contradicting that X and Y belong to distinct components. End of Case 1.2(a_2).

CASE 1.2(a_3): $(i, j) = (1, 3)$ or $(i, j) = (3, 1)$. By symmetry, assume WLOG that $(i, j) = (1, 3)$. Let $e_X = (v_x, v_{x+1})$ and $e_Y = (v_{y-1}, v_y)$. By definition of $\overrightarrow{K}_{P_1}$ (recall definition in Section 1.3), the edge (v_{x+1}, v_x) of H^* (recall definition in Section 1.2) belongs to $\overrightarrow{K}_{P_1}$. Since $X = \Phi((v_x, v_{x+1}), \text{right})$, we have that $W = \Phi((v_{x+1}, v_x), \text{right})$. Similarly, (v_y, v_{y-1}) belongs to $\overrightarrow{K}_{P_3}$, and $W = \Phi((v_y, v_{y-1}), \text{right})$. So we have that $W \in \Phi(\overrightarrow{K}_{P_1}, \text{right}) \cap \Phi(\overrightarrow{K}_{P_3}, \text{right})$. Since $X \in J_X \cap \Phi(\overrightarrow{K}_{P_1}, \text{right})$, by Corollary 1.3.8 (b), $\Phi(\overrightarrow{K}_{P_1}, \text{right}) \subset J_X$. Since $W \in \Phi(\overrightarrow{K}_{P_1}, \text{right})$, $W \in J_X$. Since $W \in J_X \cap \Phi(\overrightarrow{K}_{P_3}, \text{right})$, by Corollary 1.3.8 (b), $\Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset J_X$. But then $Y = \Phi(e_Y, \text{right})$ belongs to $\Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset J_X$, contradicting that X and Y belong to distinct H-components. End of Case 1.2(a_3) End of Case 1.2 (a).

CASE 1.2(b). By Corollary 1.4.2 (ii), if we apply $\text{Sw}(N_{J_X})$, W becomes parallel. Then, $\text{Sw}(N_{J_X})$, $\text{Sw}(W)$, is a cascade that switches W and fixes both end-vertices of H. End of Case 1.2(b). End of Case 1.2. End of Case 1. End of proof for parts (a) and (b), as well as for the scenario in part (c) where at least one of N_{J_X} and N_{J_Y} is switchable.

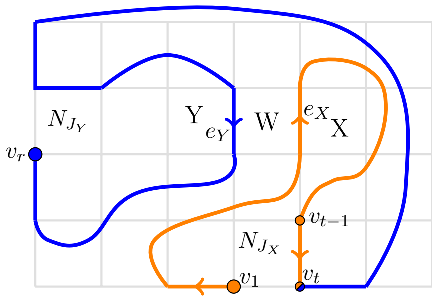

Fig. 5.13. Case 2.

CASE 2. Neither N_{J_X} nor N_{J_Y} is switchable. By Observation 5.14, the neck-edges of J_Y and J_X are incident on distinct end-vertices of H. In keeping with the notation of Corollary 1.3.8, we note that in this case, $v_1 = v_s$, so P_1 is the trivial path consisting of just the vertex v_1. Thus the neck-edge of J_X splits H into the subpaths $P_2 = P(v_1, v_t)$ and $P_3 = P(v_t, v_r)$, where v_r is incident on the neck edge of N_{J_Y}. By Observation 5.14, (v_1, v_2) must be in R_0 and neither u nor v may be at a corner of R_0. See Figure 5.13. There are two cases to consider:

1) $e_X \in P_i, e_Y \in P_j$, with $i = j$
2) $e_X \in P_i, e_Y \in P_j$, with $i \neq j$

CASE 2.1: $e_X \in P_i, e_Y \in P_j$, with $i = j$. Since P_1 is the trivial path, there are two possibilities: $e_X \in P_2, e_Y \in P_2$, and $e_X \in P_3, e_Y \in P_3$. The case where $e_X \in P_2, e_Y \in P_2$ was shown to be impossible in Case 1.2(a_2), so we omit it. We will check that the case where $e_X \in P_3, e_Y \in P_3$ is impossible as well.

Assume that $e_X \in P_3$ and $e_Y \in P_3$. By Observation 1.2.4 (a), $\Phi(\overrightarrow{K}_{P_3}, \text{right}) \supset \Phi(P_3, \text{right}) \supset \Phi(\overrightarrow{K}(e_X, e_Y), \text{right})$. By Corollary 1.3.8 (b), $\Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset J(X)$ or $\Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset G_{-1} \setminus J_X$. If the former, there is an H-path $P(X, Y) \subseteq \Phi(\overrightarrow{K}(e_X, e_Y), \text{right}) \subseteq \Phi(P_3, \text{right}) \subseteq \Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset J_X$, contradicting that X and Y belong to distinct components. And if the latter then $X = \Phi(e_X, \text{right}) \in \Phi(P_3, \text{right}) \subseteq \Phi(\overrightarrow{K}_{P_3}, \text{right}) \subset G_{-1} \setminus J_X$, contradicting that $X \in J_X$. End of Case 2.1.

CASE 2.2: $e_X \in P_i, e_Y \in P_j$ with $i \neq j$. Since both end-vertices of H are in R_0, J_X contains neither end-vertex. By NAA, $v_t \neq u$. Then, by Case 1.1(a), we have that $e_X \in P_2$, $e_Y \notin P_2$, and that (v_{t-1}, v_t) is the edge of N_{J_X} in H that is incident on v_t. It follows that $e_Y \in P_3$, and that $(v_t, v_{t+1}) \in R_0$.

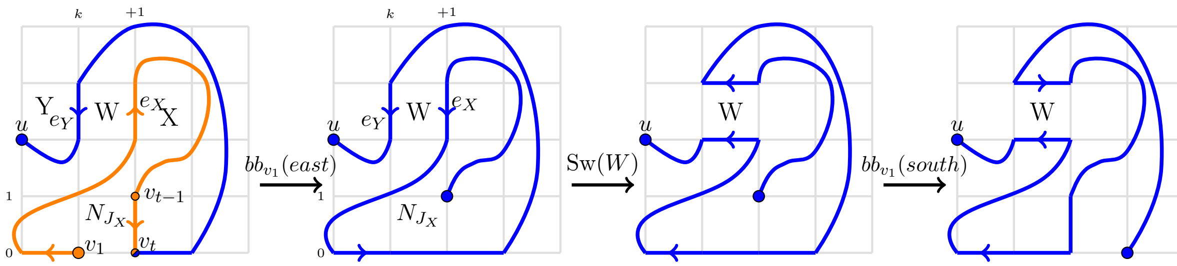

Fig. 5.14. Case 2.2.

For definiteness, assume that $v = v_1 = v(k, 0)$. It follows that $v_2 = v(k-1, 0)$ and $v_t = v(k+1, 0)$. Note that $W \neq N_{J_X}$ and $W \neq N_{J_X} + (1, 0)$. Now, by Corollary 1.4.5 (ii), after $bb_{v_1}(east)$, W is parallel in the resulting Hamiltonian path H', the end-vertex v is relocated to $v(k+1, 1)$, and the other end-vertex u is fixed. Since W is neither N_{J_X}, nor adjacent to N_{J_X}, applying Sw(W) to H' to obtain H'' does not remove the edges $e(k, k+1; 0)$ or $e(k+1, k+2; 0)$. Then we can apply $bb_{v(k+1,1)}(south)$ to H'', after which the end-vertex on $v(k+1, 1)$ is relocated to R_0, and u is fixed. It follows that, starting at H, $bb_{v_1}(east)$, Sw(W), $bb_{v(k+1,1)}(south)$, is a cascade that switches W, keeps v on the same side of R_0, and fixes u. See Figure 5.14.

There is a similar cascade that starts with a backbite move on the end-vertex u incident on the neck-edge of N_{J_Y} that switches W, keeps u on the same side of R_0, and fixes v. End of Case 2.2 □

Lemma 5.16. Let G be an $m \times n$ grid graph and let H be a Hamiltonian path of G with an end-vertex v incident on R_0, and assume that the end-vertices of H are not adjacent. Let the edge e of H be followed by an A_1-type with switchable middle-box W, and let X and Y be the boxes adjacent to W that are not its H-neighbours. Assume that X and Y belong to the same H-component, and let $P(X, Y)$ be the looping H-path of W. If v is incident on $P(X, Y)$, then $P(X, Y)$ has a switchable box Z incident on R_0.

Proof. We will use NAA repeatedly and implicitly. First note that since the A_1-type follows e, W cannot be incident on R_0. Let v be an end-vertex of H incident on a box $Z \in P(X, Y) = P$. For definiteness, assume that $Z = R(k, 0)$ and $v = v(k, 0)$. Note that if $e(v)$ is northern, then we must have that $S_{\downarrow}(k+1, 1; k+2, 0) \in H$. This implies that Z is an end-box of P, which is not possible since Z does not fit the description of an end-box of P. See Figure 5.15. The observation that W is not incident on R_0 implies that $W \neq Z + (\pm 1, 0)$; and $W \neq Z + (0, 1)$, because that means that $P(X, Y) = Z$. See Figure 5.17. Then $e(v)$ must be either eastern, or western.

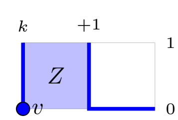

Fig. 5.15. $e(v)$ is northern.

CASE 1: $e(v)$ is eastern. By NAA, it is easy to observe that $k \geq 2$. Then $e(k-1;0,1) \in H$. Either $Z+(-1,0) \in P$, or $Z+(-1,0) \notin P$.

CASE 1.1: $Z+(-1,0) \notin P$. Note that the observation that W is not incident on R_0, together with the assumption that W is the switchable middle-box of an A_1-type, imply that Z cannot be an end-box of P. Then the H-neighbours of Z in P must be $Z+(0,1)$ and $Z+(1,0)$. Then $S_{\rightarrow}(k-1,1;k,2) \in H$, and $e(k+1,k+2;0) \in H$. There are three possibilities: $v(k+1,1)$ is the other end-vertex u of H and $e(u) = e(k+1,k+2;1)$, $v(k+1,1)$ is the other end-vertex u of H and $e(u) = e(k+1;1,2)$, and $v(k+1,1)$ is not the other end-vertex of H.

CASE 1.1(a): $u = v(k+1,1)$, $e(u) = e(k+1,k+2;1)$. Note that $Z+(1,0)$ does not fit the description of an end-box of P. Then $Z+(1,0) \in P$ must be switchable. See Figure 5.16 (a). End of Case 1.1(a).

CASE 1.1(b): $u = v(k+1,1)$, $e(u) = e(k+1;1,2)$. Note that $Z+(0,1)$ does not fit the description of an end-box of P. Then $Z+(0,1) \in P$ must be switchable. See Figure 5.16 (b). End of Case 1.1(b).

CASE 1.1(c): $v(k+1,1) \neq u$. Then $S_{\downarrow}(k+1,2;k+2,1) \in H$. Now, $W \neq Z+(2,0)$, and since $Z+(1,1)$ is not switchable, $W \neq Z+(1,1)$. It follows that $Z+(1,0)$ is not an end-box of P. Then $Z+(1,0)$ must be switchable. See Figure 5.16 (c). End of Case 1.1(c). End of Case 1.1.

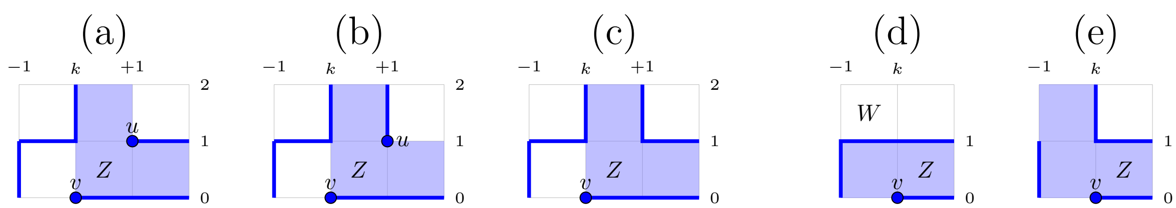

Fig. 5.16. (a). Case 1.1(a). (b). Case 1.1(b). (c). Case 1.1(c). (d). Case 1.2(a). (e). Case 1.2(b).

CASE 1.2: $Z+(-1,0) \in P$. Now, either $Z+(-1,0)$ is an end-box of P or it is not.

CASE 1.2(a): $Z+(-1,0)$ is an end-box of P. This implies that $W = Z+(-1,1)$, and that $(k-1,k;1) \in H$, $(k,k+1;1) \in H$, and that Z is not an end-box of P. Then Z must be switchable. See Figure 5.16 (d). End of Case 1.2(a).

CASE 1.2(b): $Z+(-1,0)$ is not an end-box of P. Then we must have that $e(k-1,k;1) \notin H$, and that $S_{\downarrow}(k,2;k+1,1) \in H$. Now, $W \neq Z+(1,0)$, and since $Z+(0,1)$ is not switchable, $W \neq Z+(0,1)$, so Z is not an end-box of P. It follows that Z must be switchable. See Figure 5.16 (e). End of Case 1.2(b). End of Case 1.2. End of Case 1.

CASE 2: $e(v)$ is western. Then $e(k+1;0,1) \in H$. Either Z is an end-box of P, or it is not.

CASE 2.1: Z is an end-box of P. Then we must have that $W = Z+(0,1)$, so $e(k-1,k;1) \in H$, and $e(k,k+1;1) \in H$. Then $Z+(-1,0) \in P$ is not an end-box of P and is switchable. See Figure 5.17 (a). End of Case 2.1.

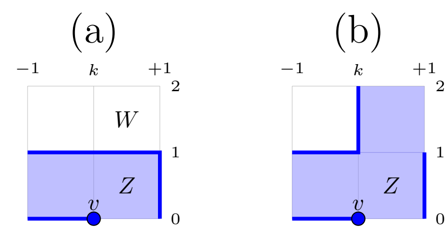

Fig. 5.17. (a) Case 2.1. (b) Case 2.2(b).

CASE 2.2: Z is not an end-box of P. Then the H-neighbours of Z in P must be $Z+(-1,0)$ and $Z+(0,1)$, and so we must have that $S_{\rightarrow}(k-1,1;k,2) \in H$. Once again, $W \neq Z+(-2,0)$, and since $Z+(-1,1)$ is not switchable, $W \neq Z+(-1,1)$. It follows that $Z+(-1,0) \in P$ is not an end-box and is switchable. See Figure 5.17 (b). End of Case 2.2. End of Case 2. □

5.2 Algorithm for relocating an endpoint to the boundary (EtB)

In this section, we prove Proposition 5.6. There is a lot of overlap between the proofs of (i) and (ii), so we will combine the proofs. Several times during the proof, we encounter subcases that require detailed analysis. We handle these separately as Claims 5.17–5.21, which can be read after the main proof if preferred. A flowchart (Flowchart 2) illustrating the proof structure appears at the end of this section.

Proof of Proposition 5.6. Orient H as $v = v_1, v_2, \ldots, v_r = u$, and recall that $e(v)$ is the edge of H incident on v. We may assume WLOG that the eastern and western sides of R_0 are parity-compatible, and that $v_1 = v(k,l)$ is an easternmost end-vertex of H.

Note that if one of v or u is incident on a side that is not parity-compatible, we can apply one backbite move that relocates that end-vertex to $G \setminus R_0$. So, for the proof of (i), we may assume that v and u are not incident on R_0, and neither is incident on the eastern or western side of R_0. In order to prove (i), it is enough to show that if neither of u and v is incident on R_0, then there is a cascade after which one of them lies closer to the eastern side of R_0 than either did prior to the cascade. We call this statement $(\dagger_1)$ for reference.

Similarly, for the proof of (ii), we may assume that u is on the western side of R_0, and that v is not on R_0. To prove (ii), it is enough to show that there is a cascade that fixes u, after which v has moved closer to the eastern side of R_0. We call this statement $(\dagger_2)$ for reference.

First we check that the case where the end-vertices are adjacent reduces to the case where (i) holds and end-vertices are not adjacent. Suppose the end-vertices of H are adjacent. By Lemma 5.10, we can reconfigure H into an e-cycle with one end-vertex u on the western side of R_0, so (i) holds. Now, it is easy to check that we can apply a backbite move that keeps u fixed and moves v to a vertex not adjacent to u, lying strictly closer than u to the eastern side of R_0. Therefore, from here on we assume that the end-vertices of H are not adjacent. Furthermore, if u and v are equally easternmost, we assume WLOG that v is north of u. Now, either $e(v)$ is eastern, or it is not.

CASE 1: $e(v)$ is not eastern. We check that $(\dagger_1)$ and $(\dagger_2)$ hold. Since $e(v)$ is not eastern, $e(k,k+1;l) \notin H$. Let $v+(1,0) = v_s$. Then $v_{s-1} = v_1 + (1,1)$, $v_{s-1} = v_1 + (2,0)$ or $v_{s-1} = v_1 + (1,1)$. See Figure 5.21. Then, after $bb_v(east)$, by Lemma 1.4.4, the end-vertex v' of the resulting Hamiltonian path is located on the vertex of G where v_{s-1} was located in H. Thus, v' is at least one unit closer to the eastern boundary, and so $(\dagger_1)$ and $(\dagger_2)$ are satisfied. See Figure 5.18 (a). End of Case 1.

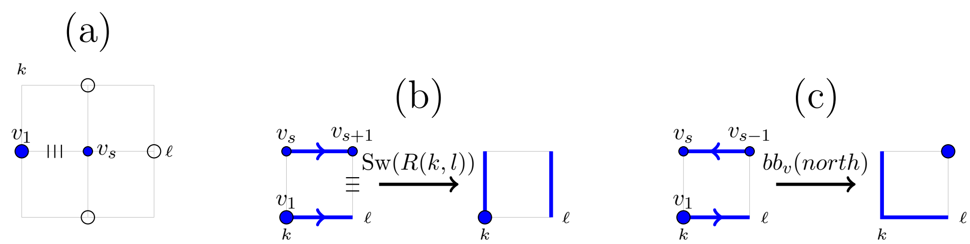

Fig. 5.18. (a) Case 1. (b) Case 2.1 with $v(k+1,l+1) = v_{s+1}$. (c) Case 2.1 with $v(k+1,l+1) = v_{s-1}$.

CASE 2: $e(v)$ is eastern. Then, either at least one of $e(v)+(0,1)$ and $e(v)+(0,-1)$ belongs to H, or neither does.

CASE 2.1: At least one of $e(v)+(0,1)$ and $e(v)+(0,-1)$ belongs to H. We check that $(\dagger_1)$ and $(\dagger_2)$ hold. By symmetry, we may assume WLOG that $e(v)+(0,1) \in H$. Orient H so that $v = v_1$ and let $v_s = v_1 + (0,1)$. Then $v(k+1,l+1) = v_{s-1}$ (Figure 5.18 (b)) or $v(k+1,l+1) = v_{s+1}$ (Figure 5.18 (c)) . If $v(k+1,l+1) = v_{s+1}$, then $e(k+1;l,l+1) \notin H$.

It follows that $R(k,l)$ is parallel. Apply $\mathrm{Sw}(R(k,l))$. Now the edge incident on v in the resulting Hamiltonian path is northern, and we are back to Case 1.

Assume that $v(k+1,l+1) = v_{s-1}$. Then, after $bb_v(north)$, v is relocated to $v(k+1,l+1)$, and so $(\dagger_1)$ and $(\dagger_2)$ are satisfied. End of Case 2.1

CASE 2.2: Neither $e(v)+(0,1)$ *nor* $e(v)+(0,-1)$ *belongs to* H. Note that if $l+1 = n-1$, then the other end-vertex of H must be on $v(k,l+1)$. Since this conflicts with NAA, we may assume that $n-1 \geq l+2$, and by symmetry, that $l-2 \geq 0$. See Figure 5.19 (a). Also, if $m-1 = k+1$, then the other end-vertex of H must be on $v(k+1,l+1)$ or on $v(k+1,l-1)$. See Figure 5.19 (b) for an illustration. Since both possibilities conflict with the assumption that v is an easternmost end-vertex, we must have that $m-1 \geq k+2$. Since $v \notin R_0$, we may assume that $k-1 \geq 0$. See Figure 5.19 (c). Then, by NAA, $S_{\rightarrow}(k-1,l-1;k,l-2) \in H$ and $S_{\rightarrow}(k-1,l+1;k,l+2) \in H$. By the assumption that v is an easternmost end-vertex of H, $v(k+1,l+1)$ and $v(k+1,l-1)$ are not end-vertices of H. Now, either $e(k+1;l,l+1) \in H$, or $e(k+1;l,l+1) \notin H$. If the former, then we must have that $S_{\uparrow}(k+1,l-2;k+2,l-1) \in H$; and if the latter, then we must have that $S_{\downarrow}(k+1,l+2;k+2,l+1) \in H$. Either way, $e(v)$ is followed by an A-type southward or northward. See Figures 5.23 (d) and (e).

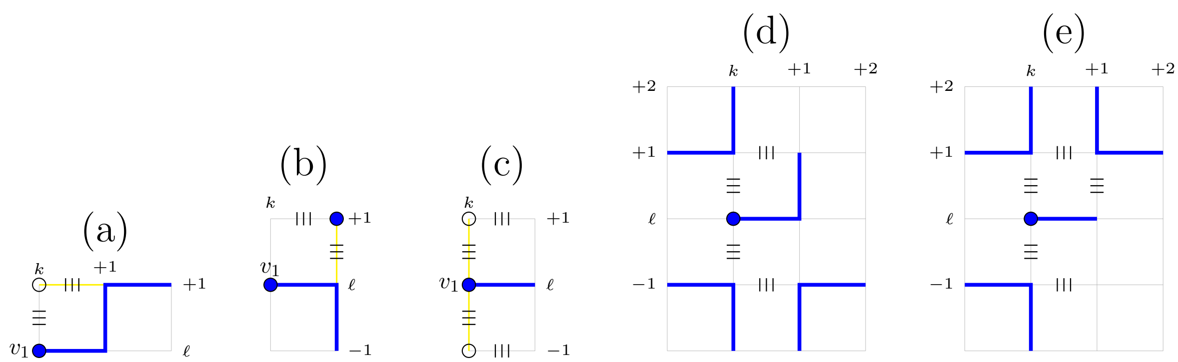

Fig. 5.19. (a) $n-1 = l+1$. (b) $m-1 = l+1$. (c) $k-1 \geq 0$. (d) $e(k+1;l,l+1) \in H$. (e) $e(k+1;l,l+1) \notin H$.

WLOG we may assume that $e(v)$ is followed by an A-type southward (as in Figure 5.19 (d)). Note that if $e(k-1,k;l-2) \in H$, then, after $R(k,l-2) \mapsto R(k-1,l-2)$, we are back to Case 2.1, so we may assume that $e(k-1,k;l-2) \notin H$. By symmetry, we may also assume that $e(k+1,k+2;l-2) \notin H$. Note that if $v(k,l-2) = u$ is the other end-vertex of H, then we are back to Case 1. Therefore we may assume that $v(k,l-2)$ is not the other end-vertex of H. Thus we have that $e(v)$ is followed southward by an A_0-type or by an A_1-type.

By Claim 5.17, we may assume that $m-1 \geq k+3$ and $0 \leq k-2$. Now, there are three possibilities [10]:

(a) $e(v)$ is followed southward by a j-stack of A_0s, not followed by an A_1-type $(e(v), js\text{-}A_0, \neg A_1, \text{South})$,

(b) $e(v)$ is followed southward by a j-stack of A_0s followed by an A_1-type $(e(v), js\text{-}A_0, A_1, \text{South})$,

(c) $e(v)$ is followed southward by a full stack of A_0s $(e(v), fs\text{-}A_0, \text{South})$.

CASE 2.2(a): $(e(v), js\text{-}A_0, \neg A_1, \mathit{South})$. Let $R(k,b)$ be the southern leaf of the last (j^{th}) A_0-type in the j-stack of A_0's that follows $e(v)$.

If the other end-vertex u of H is not incident on the rectangle $R(k-1, k+2; b-2, b-1)$, then, by Lemma 5.11, there is a cascade of flips, after which $e(k, k+1; l-1)$ is in the resulting Hamiltonian path, and we are back to Case 2.1. Therefore we may assume that u is incident on the rectangle $R(k-1, k+2; b-2, b-1)$.

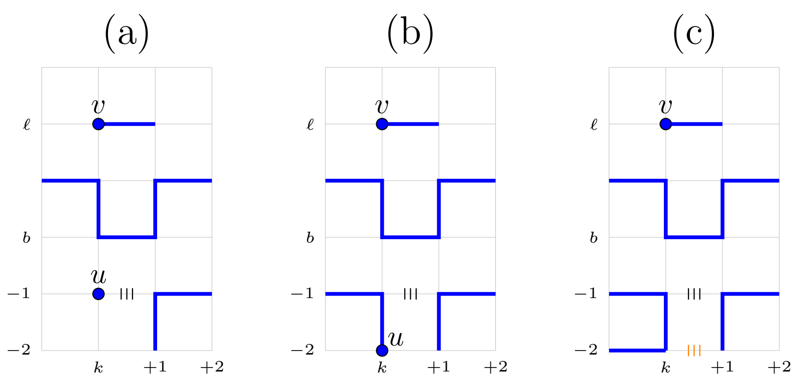

Fig. 5.20. Case 2.2(a) with: (a) $v(k, b-1) = u$, (b) $v(k, b-2) = u$. (c) $e(k-1, k; b-2) \in H$ (green).

If $e(k, k+1; b-1) \in H$, then we can use Case 1 of Lemma 5.11 to return to Case 2.1 again, so we may assume that $e(k, k+1; b-1) \notin H$. Now, the assumption that v is an eastmost end-vertex implies that $S_{\uparrow}(k+1, b-2; k+2, b-1) \in H$.

If $v(k, b-1)$ is the other end-vertex u of H, then we can apply $bb_u(east)$, after which $(\dagger_1)$ and $(\dagger_2)$ are satisfied (Figure 5.20 (a)). So, we may assume that $v(k, b-1)$ is not the other end-vertex of H. Then $S_{\rightarrow}(k-1, b-1; k, b-2) \in H$.

Now there are three possibilities: $v(k, b-2) = u$, $e(k, k+1; b-2) \in H$, or $e(k-1, k; b-2) \in H$. If the first, then we can apply $bb_u(east)$, after which $(\dagger_1)$ and $(\dagger_2)$ are satisfied. (Figure 5.20 (b)). If the second, then we can apply then $R(k, b-2) \mapsto R(k-1, b-2)$, and then we can use Lemma 5.11 again to return to Case 2.1 (Figure 5.20 (b)). The third is not possible, as it would mean that the length of the j-stack of A_0's is $j+1$, contradicting the assumption

[10] We added the abbreviations as this trio comes up several times in later cases.

that the j-stack of A_0's has length j (marked off in orange in Figure 5.20 (c)). End of Case 2.2(a).

CASE 2.2(b): $(e(v), js\text{-}A_0, A_1, \text{South})$. Let e_c be the edge in $G \setminus H$ between the two corners of the A_1-type, and let $W = R(k, l-2j-2)$. Note that W is the switchable box of the A_1-type following the j-stack of A_0's. If W is parallel, then, after Sw(W), we are back to Case 2.2(a), so we may assume that W is anti-parallel. Let X and Y be the boxes adjacent to W that are not its H-neighbours. Note that Claim 5.17, the fact that the A_1-type has edges south of W, and the fact that the A_1-type is south of $e(v)$, imply that X and Y are not incident on R_0. Now, either X and Y belong to different H-components, or they belong to the same H-component.

If X and Y belong to distinct H-components, then, by Lemma 5.15 (a), there is a non-backbite cascade, after which W is switched, and we are back to Case 2.2(a). Therefore we may assume that X and Y belong to the same H-component. In this case, $P(X,Y)$ is unique. Let $F = G\langle N[P(X,Y)]\rangle$. Then either F is a sturdy looping fat path or it is not.

If $P(X,Y)$ has a switchable box Z, then, by Proposition 4.5, either Sw(Z), Sw(W) is a cascade or $Z \mapsto W$ is a valid move, and we are back to Case 2.2(a). If F fails FPC-4, apply the available switch or flip or transpose move. If now, one of $W \mapsto X$, or $W \mapsto Y$ is valid, after applying that move, we are back to Case 2.2(a); and if $P(X,Y)$ has gained a switchable box Z, we can use Proposition 4.5 again to return to Case 2.2(a). Therefore we may assume that F satisfies FPC-2,3 and 4. Now, either F satisfies FPC-1 or it does not.

CASE 2.2(b_1): F does not satisfy FPC-1. Recall that $u \neq v$ is the other end-vertex of H. Then, at least one of u and v is incident on $P(X,Y)$. Then, either u is incident on $P(X,Y)$, or it is not.

CASE 2.2(b_1).(i): u is incident on $P(X,Y)$. Either u is on R_0 or it is not. If u is in R_0, then by Lemma 5.16, $P(X,Y)$ has a switchable box, contradicting the assumption that F satisfies FPC-3.

Assume that u is not in R_0. Then, by Claim 5.18, either there is a cascade after which ($\dagger_1$) and ($\dagger_2$) hold, or there is a cascade that fixes v, after which we are back to Case 2.2(a). End of Case 2.2(b_1).(i).

CASE 2.2(b_1).(ii): u is not incident on $P(X,Y)$ Then v must be incident on $P(X,Y)$. Then, by Claim 5.20, ($\dagger_1$) and ($\dagger_2$) hold. End of Case 2.2(b_1).(ii). End of Case 2.2(b_1)

CASE 2.2(b_2): F satisfies FPC-1. This means that F is a sturdy looping fat path. By Corollary 4.12(a), F is a standard looping fat path. By Corollary 4.22, F has an admissible turn T such that Sector(T) avoids the j-stack of A_0's and the boxes incident on $e(v)$, and is below the lines $y = x + l - 2j - k - 2$. Note that this implies that Sector(T) avoids v as well. By Corollary 4.19(a), $d(T) \geq 3$. Now, either $u \in R_0$ or $u \notin R_0$.

CASE 2.2(b_2).(i): $u \in R_0$. By Proposition 4.27, T has a double-switch weakening $\mu_1, \dots, \mu_s$. By Corollary 4.31(i), the weakening $\mu_1, \dots, \mu_s$ fixes v, and is contained in Sector(T). By Lemma 4.32, μ_s can be followed by a non-backbite cascade of length at most two, after which W is switched. End of Case 2.2(b_2).(i).

CASE 2.2(b_2).(ii): $u \notin R_0$. By Proposition 4.28, T has a weakening $\mu_1, \dots, \mu_s$. By Corollary 4.31(i), $\mu_1, \dots, \mu_s$ is contained in Sector(T).

If there is a q in $\{1, \dots, s-1\}$ such that after μ_q, X and Y belong to distinct H-components, then we can use Lemma 5.15(a) to extend $\mu_1, \dots, \mu_q$ into a cascade of length at most $q + 2$ that switches W and fixes v. So, we may assume that no such q exists. If there is a p in $\{1, \dots, s-1\}$ such that after μ_p, u is incident on $P(X, Y)$, then we can use Claim 5.18 to extend $\mu_1, \dots, \mu_p$ into a cascade of length at most $p + 2$, after which we are back to Case 2.2(a), or into a cascade of length at most $p + 3$, after which ($\dagger_1$) and ($\dagger_2$) hold. Thus, we may assume that no such p exists either.

Let H' be the resulting Hamiltonian path after μ_s. Since v and Sector(T) are disjoint, the cascade $\mu_1, \dots, \mu_s$ fixes v. Let u' be the other end-vertex of H'. Since $\mu_1, \dots, \mu_s$ is a weakening of T, after applying μ_s, either one of the terminal edges of T is in H', or u' is in ew-Terminal(T). If the former, then by Lemma 4.32, μ_s can be followed by a non-backbite cascade of length at most two, after which W is switched; and if the latter, by Observation 4.20(a), u' is incident on a box of $P(X, Y)$. Then, we can use Claim 5.18 again to extend $\mu_1, \dots, \mu_s$ into a cascade of length at most $s + 2$ after which we are back to Case 2.2(a), or a cascade of length at most $s + 3$ after which ($\dagger_1$) and ($\dagger_2$) hold. End of Case 2.2(b_2).(ii). End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2.

CASE 2.2(c): ($e(v), fs$-A_0, South). Note that if l is odd, by NAA, $e(k, k+1; 0) \in H$, and we are back to Case 2.2(a), so we may assume that l is even. We have that $S_{\downarrow}(k+2, 1; k+3, 0) \in H$. Note that $l = 2j$, and $j \geq 1$. If $j \geq 2$, then for each $i \in \{0, 2, \dots, l-4\}$, $S_{\downarrow}(k+2, i+1; k+3, i) \in H$ implies $S_{\downarrow}(k+2, i+3; k+3, i+2) \in H$. Thus we have that for each $i \in \{0, 2, \dots, l-2\}$, $S_{\downarrow}(k+2, i+1; k+3, i) \in H$. See Figure 5.26 (a). Now,

$e(k+1,k+2;l) \in H$ or $e(k+1;l,l+1) \in H$.

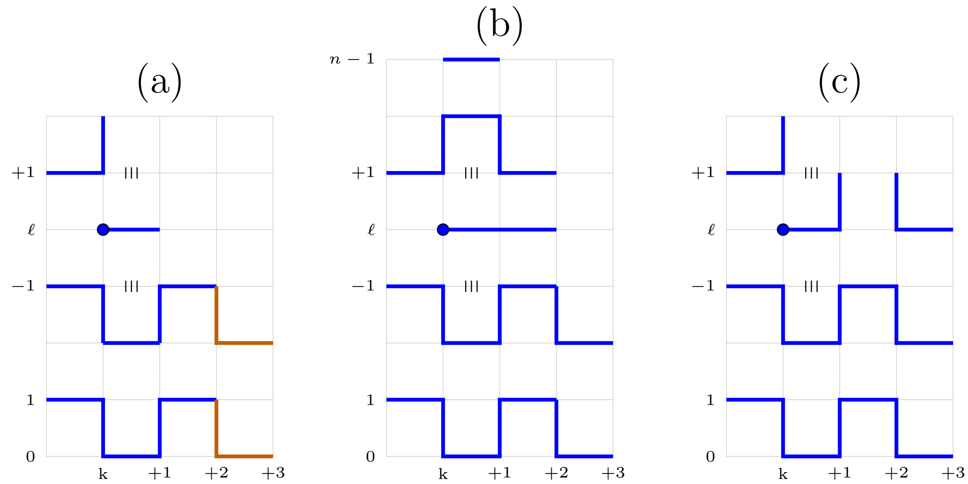

Fig. 5.21. (a) Case 2.2(c). (b) Case 2.2(c_1) with $(e(v), fs\text{-}A_0, \text{North})$. (c) Case 2.2($c_2$).

CASE 2.2(c_1): $e(k+1,k+2;l) \in H$. Then $S_{\downarrow}(k+1,l+2;k+2,l+1) \in H$. There are three possibilities:

1. $(e(v), js\text{-}A_0, \neg A_1, \text{North})$,
2. $(e(v), js\text{-}A_0, A_1, \text{North})$,
3. $(e(v), fs\text{-}A_0, \text{North})$.

If the first, then, by symmetry, this is the same as Case 2.2(a). If the second, then, by symmetry, this is the same as Case 2.2(b). Assume the third. If $n-1$ is odd, we have that $e(k,k+1;n-1) \in H$. Then, after the cascade of flips, $R(k,n-2) \mapsto R(k,n-3),\ldots,R(k,l+2) \mapsto R(k,l+1)$, we are back to Case 2.1. And if $n-1$ is even, this is impossible by Lemma 5.12. See Figure 5.26 (b). End of Case 2.2(c_1).

CASE 2.2(c_2): $e(k+1;l,l+1) \in H$. Then we have that $S_{\downarrow}(k+2,l+1;k+3,l) \in H$. Again, there are three possibilities:

1. $(e(k+1,k+2;l-1), js\text{-}A_0, \neg A_1, \text{North})$,
2. $(e(k+1,k+2;l-1), js\text{-}A_0, A_1, \text{North})$,
3. $(e(k+1,k+2;l-1), fs\text{-}A_0, \text{North})$.

Assume the first. By the assumption that v is north of u, the rectangle $R(k,k+3;l,n-1)$ does not contain u. Then, by Lemma 5.11 and symmetry, there is a cascade of flips after which $e(k+1,k+2;l) \in H$, and then we are back to Case 2.2(c_1). If the third, then $n-1$ is even or odd. If $n-1$ is even, then we must have that $e(k+1,k+2;n-1) \in H$. Then, after the cascade of flips $R(k+1,n-2) \mapsto R(k+1,n-3),\ldots,R(k+1,l+1) \mapsto R(k+1,l)$, we are back to Case 2.2(c_1). By Lemma 5.13, the case where $n-1$ is odd is not a possible configuration.

Assume the second. Let $W = R(k+1, l+2j)$. Note that W is the switchable box of the A_1-type following the j-stack of A_0's. Let X and Y be the boxes adjacent to W that are not its H-neighbours.
The fact that the A_1-type has edges located south of W, and the fact that the A_1-type is north of $e(v)$, imply that X and Y are not incident on the southern side or the northern side R_0. Since $k > 0$, they are not incident on the eastern side of R_0. An argument like the one in Cases I and II of Lemma 5.12 can be used to show that $m-1 \geq k+4$, and thus X and Y cannot be incident on the western side of R_0.

Now, either X and Y belong to distinct H-components, or they belong to the same H-component. If X and Y belong to distinct H-components, then, by Lemma 5.15 (a) and (b), there is a non-backbite cascade, after which W is switched. Then, by Lemma 5.11, there is a cascade of flips after which $e(k+1, k+2; l)$ is in the resulting Hamiltonian path, and we are back to Case 2.2(c_1). Therefore we may assume that X and Y belong to the same H-component. In this case, $P(X,Y)$ is unique. Let $F = G\langle N[P(X,Y)]\rangle$.

If W is parallel (that is, if F fails FPC-2), then, after Sw(W), by Case 2.2(a) and symmetry, there is a cascade of flips after which we are back to Case 2.2(c_1), so we may assume that FPC-2 holds. If F fails FPC-3, then $P(X,Y)$ has a switchable box Z. By Proposition 4.5, either Sw(Z), Sw(W) is a cascade or $Z \mapsto W$ is a valid move. Then, by Lemma 5.11 and symmetry, there is a cascade of flips, after which we are back to Case 2.2(c_1). If F fails FPC-4, apply the available switch or flip or transpose move. If now one of $W \mapsto X$ or $W \mapsto Y$ is valid, we apply it. Then, by Lemma 5.11 and symmetry, there is a cascade of flips, after which we are back to Case 2.2(c_1). If otherwise $P(X,Y)$ has gained a switchable box Z, we can use Proposition 4.5 and Lemma 5.11 again to return to Case 2.2(c_1). Therefore we may assume that F satisfies FPC-2, 3 and 4. Now, either F satisfies FPC-1 or it does not.

CASE 2.2(c_2).(i): F does not satisfy FPC-1. Then there is an end-vertex of H incident on $P(X,Y)$. Recall that $u \neq v$ is the other end-vertex of H. Then, at least one of u and v is incident on $P(X,Y)$. If u is incident on $P(X,Y)$, then we can use Claim 5.19, to show that either there is a cascade of length three after which ($\dagger_1$) and ($\dagger_1$) hold, or there is a cascade of length two, after which, W is switched. If the latter, then we can use Lemma 5.11 and symmetry to return to Case 2.2(c_1). Therefore we may assume that v is incident on $P(X,Y)$ and that u is not. Then, by Claim 5.21, ($\dagger_1$) and ($\dagger_2$) hold. End of 2.2(c_2).(i)

CASE 2.2(c_2).(ii): F satisfies FPC-1. This means that F is a sturdy looping fat path. The case where $u \in R_0$ uses the same argument as Case 2.2(b_2).(i). The case where $u \notin R_0$ uses the same argument as Case 2.2(b_2).(ii). In each case, we either find a cascade after which we are back to Case 2.2(c_1), or we find a cascade after which ($\dagger_1$) and ($\dagger_2$) hold. End of Case 2.2(c_2).(ii). End of Case 2.2(c_2). End of Case 2.2(c). End of Case 2.2. End of Case 2. □

It remains to prove Claims 5.17–5.21. Claim 5.17 deals with boundary proximity checking. Claims 5.18–5.21 all handle similar scenarios: an end-vertex is incident on a looping H-path that otherwise satisfies FPC-2, FPC-3, and FPC-4. We postpone proving the bound on the length of the cascade required by the EtB algorithm until after we prove these claims.

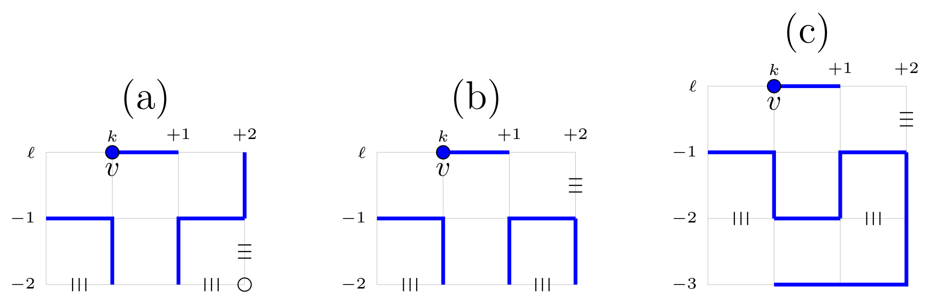

Fig. 5.22. Continuation of Case 2.2 of proof of Proposition 5.6 with: (a) $k+2 = m-1$ and $e(k+2; l-1, l) \in H$; (b) $k+2 = m-1$ and $e(k+2; l-1, l) \notin H$; (c) $k+2 = m-1$, $e(k+2; l-1, l) \notin H$, and $l-3 = 0$.

Claim 5.17. Recall that this is a continuation of Case 2.2 of the proof or Proposition 5.6, with $e(v)$ followed southward by an A_0 or an A_1. We will check that if $k+2 = m-1$, then ($\dagger_1$) and ($\dagger_2$) are satisfied. Assume that $k+2 = m-1$. Note that if $e(k+2; l-1, l) \in H$ (Figure 5.22 (a)), then $v(k+2; l-2)$ must be an end-vertex, which contradicts the assumption that v is an easternmost end-vertex. So, we may assume that $e(k+2; l-1, l) \notin H$ (Figure 5.22 (b)). It follows that $e(k+2; l-2, l-1) \in H$.

If $l-2 = 0$, then again, $v(k+2, l-2)$ must be an end-vertex, which contradicts the assumption that v is an easternmost end-vertex, so we may assume that $l-3 \geq 0$ (Figure 5.22 (c)). Then $e(k+2; l-3, l-2) \in H$.

Note that if $l-3 = 0$, then we must have $e(k+1, k+2; l-3) \in H$, $e(k, k+1; l-3) \in H$, and $e(k, k+1; l-2) \in H$. But then, after $R(k, l-3) \mapsto R(k, l-2)$, we are back to Case 2.1. So we may assume that $l-4 \geq 0$.

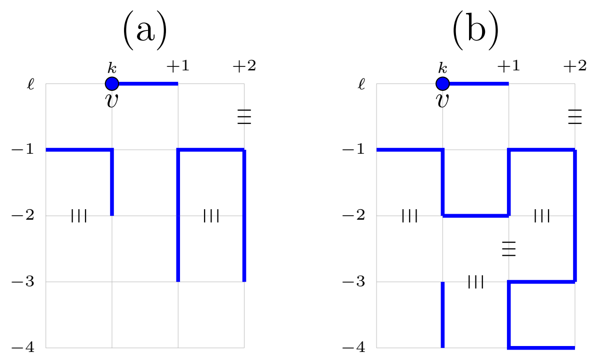

Fig. 5.23. Continuation of Case 2.2 of proof of Proposition 5.6 with $k+2 = m-1$, $e(k+2; l-1, l) \notin H$, $l-4 \geq 0$, and: (a) $e(k+1; l-3, l-2) \in H$; (b) $e(k+1; l-3, l-2) \notin H$

If $e(k+1; l-3, l-2) \in H$, then, after $R(k+1, l-3) \mapsto R(k, l-2)$, we are back to Case 2.1, so we may assume that $e(k+1; l-3, l-2) \notin H$. It follows that $e(k, k+1; l-2) \in H$. As before, if $e(k, k+1; l-3) \in H$, then, after $R(k, l-3) \mapsto R(k, l-2)$, we are back to Case 2.1, so we may assume that $e(k, k+1; l-3) \notin H$. Then $S_{\uparrow}(k+1, l-4; k+2, l-3) \in H$. Since v is easternmost, there are no end-vertices east of the line $x = k$. Now, either $e(k; l-4, l-3) \in H$, or $e(k; l-4, l-3) \notin H$. If the former (Figure 5.23(a)), then we must have $e(k+1, k+2; l-4) \in H$ as well. But then there is $R(k, l-4) \mapsto R(k+1, l-4)$, followed by $R(k, l-3) \mapsto R(k, l-2)$, we are back to Case 2.1 once more. If the latter (Figure 5.23(a)) then $u = v(k, l-3)$, and we can apply $bb_u(east)$, after which $(\dagger_1)$ and $(\dagger_2)$ hold.

Therefore, we may assume that $m-1 \geq k+3$, and by symmetry, that $0 \leq k-2$. $\square$

Claim 5.18. This is a continuation of Case 2.2$(b_1).(i)$. It will also be used by Case 2.2$(b_2).(ii)$.

We restate the setup here. $e(v)$ is followed by a j-stack of A_0s followed by an A_1-type southward. $W = R(k, l-2j-2)$ is the switchable middle-box of the A_1-type, X and Y are the boxes adjacent to W that are not its H-neighbours, and $P(X, Y)$ is contained in an H-component of G. $F = G\langle N[P(X, Y)]\rangle$ satisfies FPC-2,3, and 4, and fails FPC-1. We assume that u is incident on $P(X, Y)$, and that both u and v have x-coordinate at most k. We will show that there is a cascade of length at most three after which $(\dagger_1)$ and $(\dagger_2)$ hold, or that there is a cascade of length at most two, after which we return to Case 2.2(a).

Proof. Let $l' = l-2j-2$, let $v_y = v(k, l'+1)$ and $v_x = v(k+1, l'+1)$. Either u is incident on a box that contains v_x or v_y, or it is not.

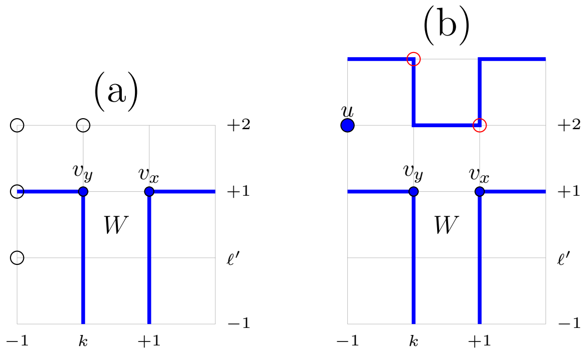

Fig. 5.24. Continuation of Case 2.2(b_1).(i) of proof of Proposition 5.6, with u incident on $P(X, Y)$. (a) Case 1. (b) Case 1.2 with $j > 0$

CASE 1: u is incident on a box that contains v_x or v_y. Since the x-coordinate of u is at most k, there are four possibilities: $u = v(k, l' + 2)$, $u = v(k - 1, l' + 2)$, $u = v(k - 1, l' + 1)$, or $u = v(k - 1, l')$. See Figure 4.24 (a).

CASE 1.1: $u = v(k, l' + 2)$. If $j = 0$, $u = v(k, l) = v$, which is impossible; and if $j > 0$, then we must have $S_{\downarrow}(k, l' + 3; k + 1; l' + 2) \in H$, which is again impossible. So $u \neq v(k, l' + 2)$. End of Case 1.1.

CASE 1.2: $u = v(k - 1, l' + 2)$. $j = 0$ contradicts NAA, so assume $j > 0$. Then after $bb_u(east)$, u is either relocated to $v(k + 1; l' + 2)$ or to $v(k; l' + 3)$. If the former, then $\dagger_1$ and $\dagger_2$ hold; and if the latter, then we can follow the first move with $bb_{v(k,l'+3)}(east)$, after which $(\dagger_1)$ and $(\dagger_2)$ hold. See Figure 4.24 (b). End of Case 1.2.

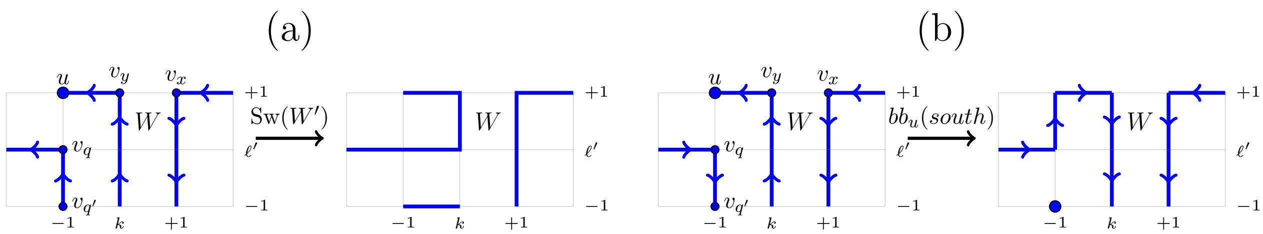

Fig. 5.25. Continuation of Case 2.2 of proof of Proposition 5.6. (a) Case 1.3(a) with $W' = W + (-1, -1)$; (b) Case 1.3(b_1).

CASE 1.3: $u = v(k - 1, l' + 1)$. Then we must have $S_{\rightarrow}(k - 2, l'; k - 1, l' - 1) \in H$. Let $v_q = v(k - 1, l')$ and $v_{q'} = v(k - 1, l' - 1)$. Either $q' < q$ or $q' > q$.

CASE 1.3(a): $q' < q$. It follows that $W + (-1, -1)$ is parallel. Then, after Sw$(W + (-1, -1))$, $W \mapsto W + (0 - 1, 0)$, we are back to Case 2.2(a) of Proposition 5.6. See Figure 5.25 (a) End of Case 1.3(a).

CASE 1.3(b): $q' > q$. Then, by Lemma 1.4.4, $bb_u(south)$ is the move $e(k-1; l'-1, l') \mapsto e(k-1; l', l'+1)$. By Proposition 4.4, we may partition H into the subpaths $P_0 = P(v_x, v_y)$, $P_1 = P(v_1, v_x)$ and $P_2 = P(v_y, v_r)$ such that every box of $P(X, Y)$ is incident on a vertex of P_0 and a vertex of P_1 or P_2. Now, either $v_{q'} \in P_0$, or $v_{q'} \notin P_0$.

CASE 1.3(b_1): $v_{q'} \in P_0$. This means that the order of vertices of H is $v_1, \ldots v_x, \ldots, v_q, \ldots, v_y$, u. Then, by Corollary 1.4.5, after $bb_u(south)$, W becomes parallel, so we can apply Sw(W), and return to Case 2.2(a) of Proposition 5.6. See Figure 5.25 (b). End of Case 1.3(b_1).

CASE 1.3(b_2): $v_{q'} \notin P_0$. Then $v_{q'}$ must belong to P_1. Note that this implies that $e(k-1, k; l'-1) \notin H$, and, by Lemma 1.4.4, that $bb_u(south)$ fixes P_0. Then, after $bb_u(south)$, $bb_{v(k-1,l'-1)}(east)$ is the move $e(k; l'-1, l') \mapsto e(k-1, k; l'-1)$, and we apply it. Now, after $bb_{v(k,l')}(east)$, the end-vertex is relocated to $v(k+1, l'-1)$, so ($\dagger_1$) and ($\dagger_2$) hold. See Figure 5.26. Note that we have used Lemma 1.4.4 for the previous two moves. End of Case 1.3(b_2). End of Case 1.3(b). End of Case 1.3.

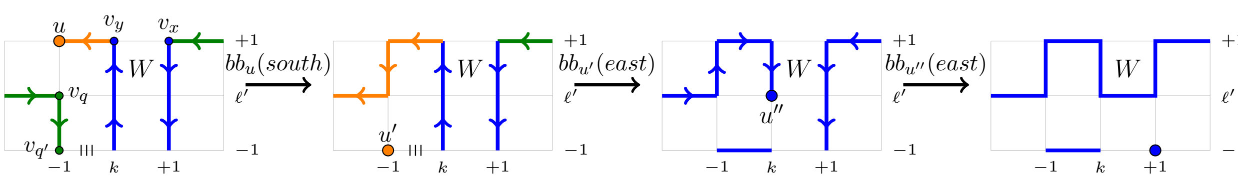

Fig. 5.26. Case 1.3(b_2), where $u' = v(k-1, l'-1)$ and $u'' = v(k, l')$.

CASE 1.4: $u = v(k-1, l')$. Let $v_q = v(k, l')$ and $v_{q'} = v(k, l'-1)$. Then $q < q'$ or $q > q'$.

CASE 1.4(a): $q < q'$. By Lemma 1.4.4, $bb_u(east)$ is the backbite move $e(k; l'-1, l') \mapsto e(k-1, k; l')$, and we apply it. Then, after $W \mapsto W + (-1, 0)$, we are back to Case 2.2(a) of Proposition 1.17. See Figure 5.27 (a). End of Case 1.4(a).

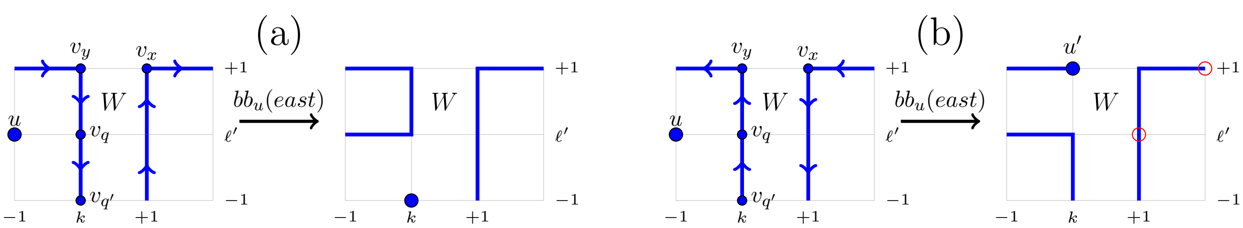

Fig. 5.27. (a) Case 1.4(a). (b) Case 1.4(b).

CASE 1.4(b): $q > q'$. By Lemma 1.4.4, $bb_u(east)$ is the backbite move $e(k; l', l'+1) \mapsto e(k-1, k; l')$, and we apply it. Now, after $bb_{v(k,l'+1)}(east)$, the end-vertex is relocated to

$v(k+1,l')$, or to $v(k+2;l'+1)$. Either way, ($\dagger_1$) and ($\dagger_2$) hold. See Figure 5.27 (b). End of Case 1.4(b). End of Case 1.4. End of Case 1.

CASE 2: u is not incident on a box that contains v_x or v_y. Let Z be the box of $P(X,Y)$ on which u is incident. By Proposition 4.4, we may partition H into the subpaths $P_0 = P(v_x, v_y)$, $P_1 = P(v_1, v_x)$ and $P_2 = P(v_y, v_r)$ such that every box of $P(X,Y)$ is incident on a vertex of P_0 and a vertex of P_1 or P_2. Let v_s be the vertex of Z that belongs to P_0. WLOG, assume that $Z = R(a,b)$, and $u = v(a+1,b)$. Either v_s is adjacent to u, or it is not.

CASE 2.1: v_s is adjacent to u. For definiteness, assume that $v_s = v(a,b)$. Note that the edge$\{u, v_s\}$ must belong to $G \backslash H$, and that the order of vertices of H is $v_1, \ldots v_x, \ldots, v_s, \ldots, v_y$, u. Then, by Corollary 1.4.5, after $bb_u(west)$, W becomes parallel, so we can apply Sw(W), and return to Case 2.2(a) of Proposition 5.6. End of Case 2.1.

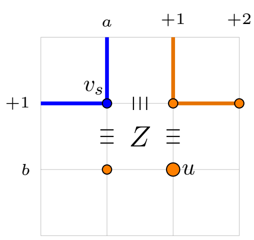

Fig. 5.28. Case 2.2.

CASE 2.2: v_s is not adjacent to u. For definiteness, assume that $v_s = v(a, b+1)$. If $v(a,b) \in P_0$ or $v(a+1,b+1) \in P_0$ then we are back to Case 1, so we may assume that $v(a,b)$ and $v(a+1,b+1)$ belong to P_1 or P_2, and by NAA, that neither vertex is v_1. By the premise of Case 2, the edges $e(a;b,b+1) \notin H$, $e(a,a+1;b+1) \notin H$. This implies that $S_{\rightarrow}(a-1,b+1;a,b+2) \in H$. Since u is an end-vertex, at least one of $e(a,a+1;b)$ and $e(a+1;b,b+1)$ does not belong to H. By symmetry, we may assume WLOG that $e(a+1;b,b+1) \notin H$. Then we must have $S_{\downarrow}(a+1,b+2;a+2,b+1) \in H$, so $v(a+2,b+1)$ is in P_1 or P_2. See Figure 5.28. Now, either $e(a,a+1;b) \in H$ or $e(a,a+1;b) \notin H$.

CASE 2.2(a): $e(a,a+1;b) \in H$. Either $Z+(1,0) \in P(X,Y)$ or $Z+(1,0) \notin P(X,Y)$.

CASE 2.2(a_1): $Z+(1,0) \in P(X,Y)$. Now, either $v(a+2;b) \in P_0$ (Figure 5.29 (a)), or $v(a+2;b) \notin P_0$ (Figure 5.29 (b)). If $v(a+2,b) \in P_0$, then we are back to Case 2.1. Assume that $v(a+2,b) \notin P_0$. Then, by Lemma 4.1, we must have that $v(a+2,b+1) \in P_0$. But this implies that $v(a+2,b+1) = v_x$ or $v(a+2,b+1) = v_y$, contradicting the premise of Case 2. End of Case 2.2(a_1).

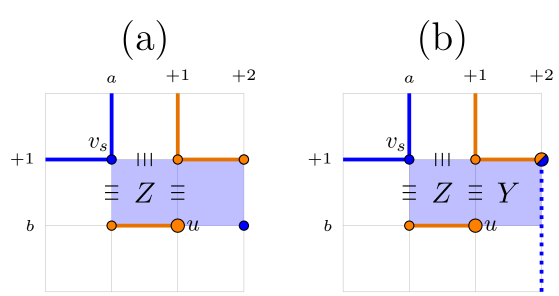

Fig. 5.29. Case 2.2(a_1) with: (a) $v(a+2,b) \in P_0$. (b) $v(a+2,b) \notin P_0$.

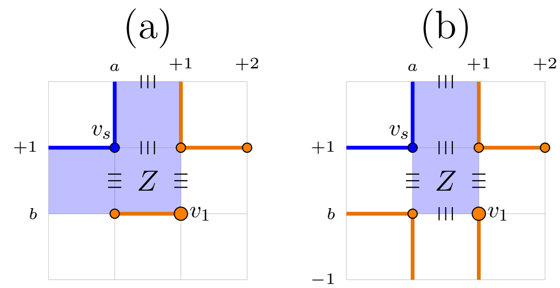

Fig. 5.30. (a) Case 2.2(a_2) (b) Case 2.2(b).

CASE 2.2(a_2): $Z + (1,0) \notin P(X,Y)$. Then Z must have exactly two H-neighbours in $P(X,Y)$. In particular, $Z+(0,1) \in P(X,Y)$. Note that $Z + (0,1)$ can not be X or Y. Therefore, it must also have exactly two H-neighbours in $P(X,Y)$. This implies that $e(a, a+1; b+2) \notin H$. But then $Z+(0,1) \in P$ is switchable, which contradicts that F satisfies FPC-3. See Figure 5.30 (a). End of Case 2.2(a_2). End of Case 2.2(a)

CASE 2.2(b): $e(a, a+1; b) \notin H$. This implies that $S_{\rightarrow}(a-1, b; a, b-1) \in H$. Now, exactly one of $e(a+1; b-1, b)$ and $e(a+1, a+2; b)$ belongs to H. By symmetry, we may assume without loss of generality that $e(a+1; b-1, b) \in H$. If $Z + (1,0) \in P$, then we're back to Case 2.2(a_1), so we may assume that $Z + (1,0) \notin P$. Note that Z cannot be X or Y, so it must have at least two H neighbours in P. Then, at least one of them must be $Z + (-1,0)$ or $Z + (0,1)$, and either possibility brings us back to Case 2.2(a_2). See Figure 5.30 (b). End of Case 2.2(b). End of Case 2.2. End of Case 2. □

Claim 5.19. This is the version of Claim 5.18 suitable for use in Case 2.2(c_2) of Proposition 5.6. The setup and cases here are very similar to the setup in Claim 5.18. We begin with the scenario $((k+1, k+2; l-1), js\text{-}A_0, A_1, \text{North})$, and the rest of the starting assumptions are the same. Case 2 is the same as in Claim 5.18.

Case 1 cannot occur: The assumption that the x-coordinates of the end-vertices are at most k, precludes Case 1.1. The assumption that v is located north of u precludes Cases 1.2, 1.3, and 1.4. □

Claim 5.20. Recall that this is a continuation of Case 2.2(b_1).(ii) of Proposition 5.6. We are in the scenario where $(e(v), js\text{-}A_0, A_1, \text{South})$, FPC-2,3, and 4 hold, u is not incident on $P(X,Y)$, but v is. We will show that in this case, $(\dagger_1)$ and $(\dagger_2)$ hold.

Proof. For definiteness, assume that $X = W + (1,0)$, $Y = W + (-1,0)$ and let v_x and v_y be the corners of the A_1-type incident on X and Y, respectively. Consider the j-stack of A_0's following $e(v)$. Either $j = 0$, or $j > 0$.

CASE 1: $j = 0$. Then $W = R(k, l-2)$. Let $v_x = v(k+1, l-1)$, and $v_y = v(k, l-1)$ be the corners of the A_1-type of $P(X,Y)$. Now, either $v_{y-1} = v(k, l-2)$ (Figure 5.31), or $v_{y-1} = v(k-1, l-1)$. If the former, then $bb_v(south)$ must be the move $(v_{y-1}, v_y) \mapsto \{v_y, v_1\}$, and the we are back to Case 1 of Proposition 5.6. We will show that the latter is not possible.

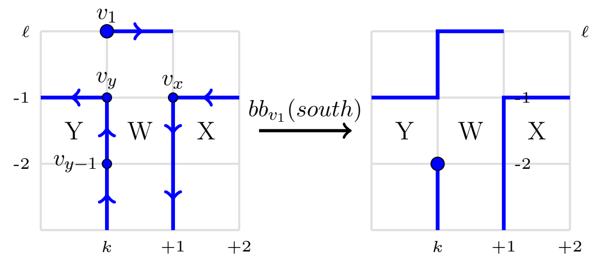

Fig. 5.31. Case 1 with $v_{y-1} = v(k, l-2)$

For contradiction, assume that $v_{y-1} = v(k-1, l-1)$. This means that $v_{x+1} = v(k+2, l-1)$. Then $x < y$, or $x > y$.

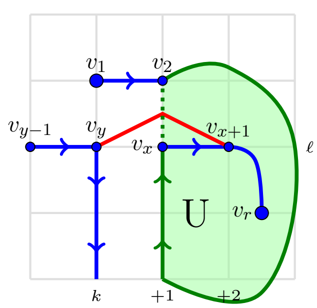

Fig. 5.32. Case 1.1.

CASE 1.1: $x < y$. Let U be the region bounded by the cycle Q of G consisting of the subpath $P(v_2, v_x)$ of H and the edge $\{v_x, v_2\}$ of $G \setminus H$. JCT and Corollary 1.1.5 imply that v_y and v_{x+1} are on distinct sides of Q. Either $v_{x+1} \in U$ or $v_{x+1} \in G \setminus U$.

Assume $v_{x+1} \in U$. By JCT, $P(v_{x+1}, v_r)$ is contained in U as well. But since $x < y$, $v_y \in P(v_{x+1}, v_r)$, contradicting v_y and v_{x+1} are on distinct sides of U. See Figure 5.32. The case where $v_{x+1} \in G \setminus U$ is similar so we omit it.

CASE 1.2: $x > y$. Since v is incident on $P(X,Y)$ at least one of $W + (0,2)$, $W + (-1,2)$, $W + (-1,1)$, and $W + (0,1)$ belongs to $P(X,Y)$. First we check that at least one of $W + (0,2)$, $W + (-1,2)$, and $W + (-1,1)$ belongs to $P(X,Y)$. If $W + (0,1) \in P(X,Y)$, since $W + (0,1)$ is not X or Y, both of its H-neighbours belong to $P(X,Y)$. In particular, $W + (-1,1)$ would belong to $P(X,Y)$. And if $W + (0,1) \notin P(X,Y)$, then at least one of $W + (0,2)$, $W + (1,2)$, and $W + (-1,1)$ must belong to $P(X,Y)$. Either way, at least one of $W + (0,2)$, $W + (1,2)$, and $W + (-1,1)$ belongs to $P(X,Y)$.

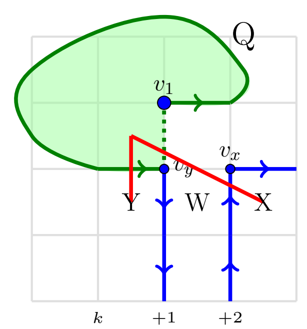

Fig. 5.33. Case 1.2.

Let U be the region bounded by the cycle Q of G consisting of the subpath $P(v_1, v_y)$ of H and the edge $\{v_y, v_1\}$ of $G \setminus H$. See Figure 5.41. JCT and Corollary 1.1.5 imply that Y and $W + (-1,1)$ belong to distinct sides of Q. Similarly, X and $W + (-1,1)$ belong to distinct sides of Q. Let X' be a box from the set $\{W + (0,2), W + (1,2), W + (-1,1)\}$ that belongs to $P(X,Y)$, and let $P(X,X')$ and $P(X',Y)$ be such that $P(X,X') \cap P(X',Y) = X'$. See Figure 5.33.

Let $c_1, \ldots, c_s$ be the centers of the boxes $X = X_1, \ldots, X' = X_s$ of $P(X,X')$. Note that for each $j \in \{1, \ldots, q-1\}$, the segment $[c_j, c_{j+1}]$ intersects the gluing edge of X_j and X_{j+1}, and no other edge of G. Since c_1 and c_s are on distinct sides of Q, the polygonal path $[c_1, c_2], \ldots, [c_{s-1}, c_s]$ must intersect Q. Since each segment intersects gluing edges of $P(X, Y)$, the polygonal path $[c_1, c_2], \ldots, [c_{s-1}, c_s]$ must intersect Q at the edge $\{v_y, v_1\}$. It follows that $W + (0, 1)$ and $W + (-1, 1)$ both belong to $P(X, X')$. A similar argument show that $W + (0, 1)$ and $W + (-1, 1)$ must belong to $P(X', Y)$. But this contradicts our assumption that $P(X, X') \cap P(X', Y) = X'$. End of Case 1.2. End of Case 1.

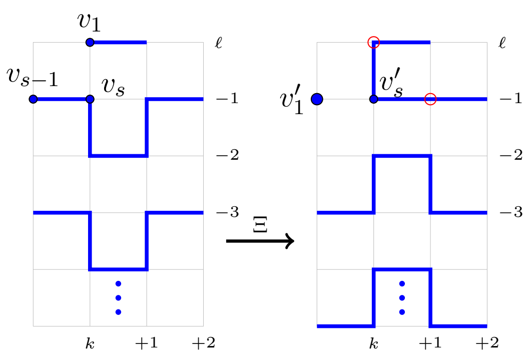

Fig. 5.34. Case 2.1(a).

CASE 2: The j-stack of A_0's has length strictly greater than zero. Let P_0, P_1, P_2 be the A_1 partitioning of H given in Lemma 4.1. That is, $P_1 = P(v_1, v_x)$, $P_0 = P(v_x, v_y)$, and $P_2 = P(v_y, v_r)$. Then $v(k, l-1) \in P_0$ or $v(k, l-1) \notin P_0$.

CASE 2.1: $v(k, l-1) \in P_0$. Let $v(k, l-1) = v_s$. Then $1 < x < s < y < r$, and one of the edges of W in H is in $P(v_1, v_s)$, while the other is in $P(v_s, v_r)$. Now, $v_{s-1} = v(k-1, l-1)$, or $v_{s-1} = v(k, l-2)$.

CASE 2.1(a): $v_{s-1} = v(k-1, l-1)$. Then, by Lemma 1.4.5, after $bb_{v_1}(south)$, W becomes parallel. For each $i \in \{1, \ldots, j\}$, let μ_i be the move $W + (0, 2j-1) \mapsto W + (0, 2j)$. Let Ξ be the sequence of moves $bb_{v_1}(south)$, Sw(W), $\mu_1, \ldots, \mu_j$. Note that Ξ is a cascade and we apply it. Let H' be the resulting Hamiltonian path. Then $v_1' = v(k-1, l-1)$ and $v_s' = v(k, l-1)$ are the first and s^{th} vertices in H', respectively. Then $v_{s-1}' = v(k+1, l-1)$ or $v_{s-1}' = v(k, l)$. If the former, then after $bb_{v_1'}(east)$, the end-vertex in the resulting Hamiltonian path is at $v(k+1, l-1)$, so $(\dagger_1)$ and $(\dagger_2)$ are satisfied. And if the latter, then after $bb_{v_1'}(east)$, we are back to Case 2.1 of Proposition 5.6. See Figure 5.34. End of Case 2.1(a).

CASE 2.1(b): $v_{s-1} = v(k, l-2)$. As in Case 2.1(a), after $bb_{v_1}(south)$, W is switchable. If $j = 1$, then after $bb_{v_1}(south)$, Sw(W), we are back to Case 2.1 of Proposition 5.6, so assume that $j > 1$. For each $i \in \{1, \ldots, j-1\}$, let μ_i be the move $W + (0, 2j-1) \mapsto W + (0, 2j)$. Let Ξ be the sequence of moves $bb_{v_1}(west)$, Sw(W), $\mu_1, \ldots, \mu_{j-1}$. Note that Ξ is a cascade and we

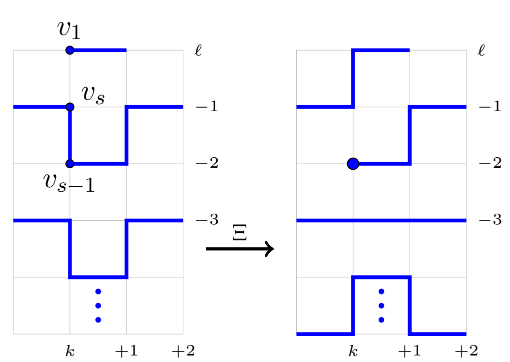

Fig. 5.35. Case 2.1(b).

apply it. Again, we are back to Case 2.1 of Proposition 5.6. See Figure 5.35. End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $v(k, l-1) \notin P_0$. Then the A_0-type following $e(v)$ southward must also belong to P_1 or P_2. Let $W' = R(k, l)$. We have that $V(W' + (0, -1)) \subset V(P_1) \cup V(P_2)$. Then, by Lemma 4.1, $W' + (0, -1) \notin P(X, Y)$. Now we check that $W' \in P(X, Y)$.

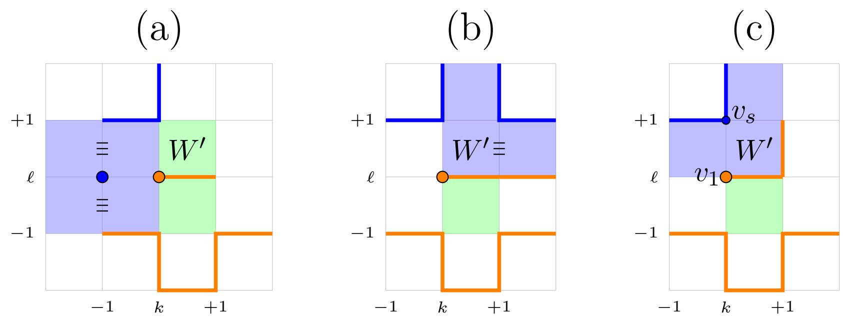

Fig. 5.36. (a) Case 2.2 with $W' \notin P$. (b) Case 2.2. with $e(k+1; l, l+1) \notin H$. (c) Case 2.2 with $W' \in P$, $e(k+1; l, l+1) \in H$, and $v(k, l+1) \in P_0$.

For a contradiction, assume that $W' \notin P(X, Y)$. Note that this implies that both $W'+(-1, 0)$ and $W'+(-1, -1)$ belong to $P(X, Y)$. Since neither $W'+(-1, 0)$ nor $W'+(-1, -1)$ is an end-box, their H-neighbours in $P(X, Y)$ must be $W' + (-2, 0)$ and $W' + (-2, -1)$, respectively. But this implies that $v(k, l-1)$ is an end-vertex of H, contradicting the NAA. Thus we must have that $W' \in P(X, Y)$. See Figure 5.36 (a).

Next, we check that $e(k+1; l, l+1) \in H$. For a contradiction, assume that $e(k+1; l, l+1) \notin H$. Then we have that $S_{\downarrow}(k+1, l+2; k+2, l+1) \in H$, and that $e(k+1, k+2; l) \in H$. Then, at least one of $W'+(0, 1)$ and $W'+(1, 0)$ must be an H-neighbour of W' in $P(X, Y)$. Suppose that $W' + (1, 0)$ is an H-neighbour of W' in $P(X, Y)$. Then $W' + (1, 0)$ is either switchable, or an end-box of $P(X, Y)$, neither of which is possible. The same goes for $W' + (0, 1)$. It follows that we must have $e(k+1; l, l+1) \in H$. See Figure 5.36 (b). Then the H-neighbours of W' in $P(X, Y)$ must be $W' + (-1, 0)$ and $W' + (0, 1)$.

Next we check that $v(k, l+1)$ must belong to P_0. Note that the vertices v_1, v_2, and v_3 belong to P_1, while $W' \in P(X, Y)$. Then, by Lemma 4.1, $v(k, l+1) \in P_0$. See Figure 5.36 (c).

Let $v(k, l+1) = v_s$. Then $1 < x < s < y < r$, and one of the edges of W in H is in $P(v_1, v_s)$, while the other is in $P(v_s, v_r)$. Now, $v_{s-1} = v(k, l+2)$ or $v_{s-1} = v(k-1, l+1)$.

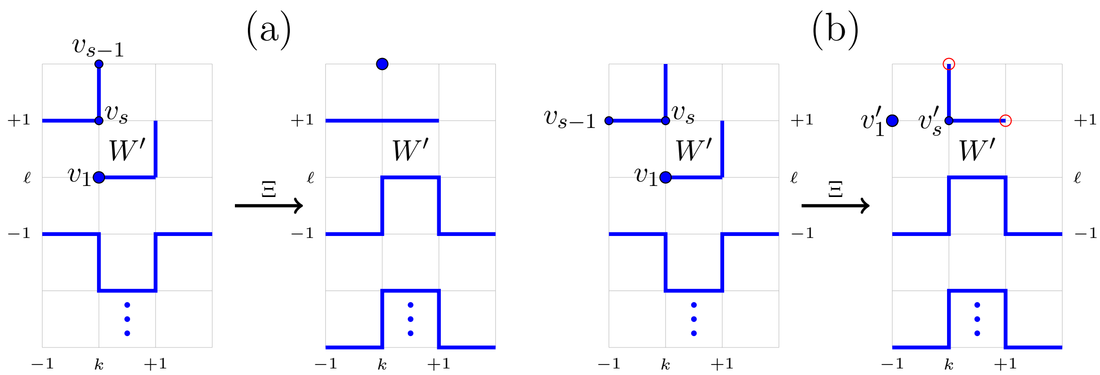

Fig. 5.37. (a) Case 2.2(a). (b) Case 2.2(b).

CASE 2.2(a): $v_{s-1} = v(k, l+2)$. Then, by Corollary 1.4.5, after $bb_{v_1}(north)$, W becomes parallel. For each $i \in \{1, \dots, j+1\}$, let μ_i be the move $W + (0, 2j-1) \mapsto W + (0, 2j)$. Let Ξ be the sequence of moves $bb_{v_1}(north)$, $\mathrm{Sw}(W)$, $\mu_1, \dots, \mu_{j+1}$. Note that Ξ is a cascade, and we apply it. Now, if $e(k, k+1; l+2) \notin H$, we are back to Case 1 of Proposition 5.6; and if $e(k, k+1; l+2) \in H$, we are back to Case 2.1 of Proposition 5.6. See Figure 5.37 (a). End of Case 2.2(a). End of Case 2.1.

CASE 2.2(b): $v_{s-1} = v(k-1, l+1)$. By Corollary 1.4.5, after $bb_{v_1}(north)$, W becomes parallel. For each $i \in \{1, \dots, j+1\}$, let μ_i be the move $W + (0, 2j-1) \mapsto W + (0, 2j)$. Let Ξ be the sequence of moves $bb_{v_1}(north)$, $\mathrm{Sw}(W)$, $\mu_1, \dots, \mu_{j+1}$. Note that Ξ is a cascade, and we apply it. Let H' be the resulting Hamiltonian path. Then $v'_1 = v(k-1, l+1)$ and $v'_s = v(k, l+1)$ are the first and s^{th} vertices in H', respectively. Then $v'_{s-1} = v(k+1, l+1)$ or $v'_{s-1} = v(k, l+2)$. See Figure 5.37 (b). If the former, then after $bb_{v'_1}(east)$, the end-vertex in the resulting Hamiltonian path is on $v(k+1, l+1)$, so $(\dagger_1)$ and $(\dagger_2)$ are satisfied.

Suppose then that $v'_{s-1} = v(k, l+2)$. Apply $bb_{v_1}(east)$. Now, if $e(k, k+1; l+2) \notin H$, we are back to Case 1 of Proposition 5.6; and if $e(k, k+1; l+2) \in H$, we are back to Case 2.1 of Proposition 5.6. End of Case 2.2(b). End of Case 2.2. End of Case 2. □

Claim 5.21. Recall that this is a continuation of Case 2.2(c_2).(i) of Proposition 5.6. We are in the scenario where $(e(k+1, k+2; l-1), js\text{-}A_0, A_1, \text{North})$, FPC-2,3, and 4 hold, u is not incident on $P(X, Y)$, but v is. We will show that in this case, $(\dagger_1)$ and $(\dagger_2)$ hold.

Let v_x and v_y be the corners of the A_1-type of $P(X, Y)$. Let P_0, P_1, P_2 be the A_1 partitioning of H given in Lemma 4.1. That is, $P_1 = P(v_1, v_x)$, $P_0 = P(v_x, v_y)$, and $P_2 = P(v_y, v_r)$. Then $v(k, l-1) \in P_0$ or $v(k, l-1) \notin P_0$.

CASE 1: $v(k, l-1) \in P_0$. Let $v_s = v(k, l-1)$. Then $1 < x < s < y < r$, and one of the edges of W in H is in $P(v_1, v_s)$, while the other is in $P(v_s, v_r)$. Now, $v_{s-1} = v(k, l-2)$ or $v_{s-1} = v(k-1, l-1)$.

CASE 1.1: $v_{s-1} = v(k, l-2)$. Then we apply $bb_v(south)$. By Corollary 1.4.5, W is now parallel. Apply Sw(W). Apply the sequence of flips $W + (0, -1) \mapsto W + (0, -2)$, …, $R(k+1, l-1) \mapsto R(k+1, l-2)$[11]. See Figure 5.38. Now we are back to a translation by $(0, -2)$ of Case 2.2(c_1) of Proposition 5.6. End of Case 1.1

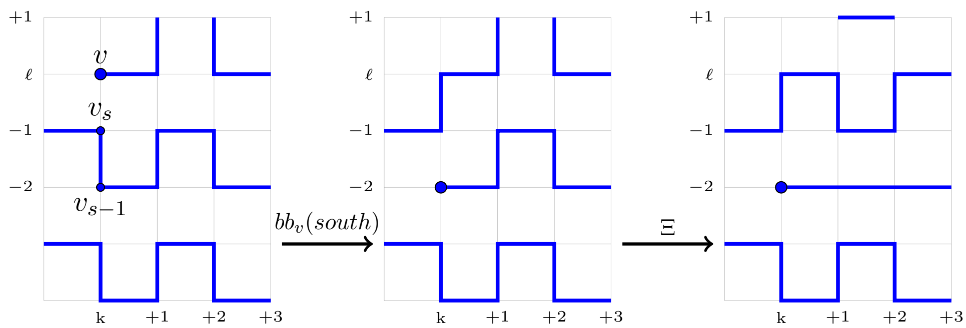

Fig. 5.38. Case 1.1, where Ξ is the cascade Sw(W), $W + (0, -1) \mapsto W + (0, -2), \ldots, R(k+1, l-1) \mapsto R(k+1, l-2)$.

CASE 1.2: $v_{s-1} = v(k-1, l-1)$. We apply $bb_v(south)$. By Corollary 1.4.5, W is now parallel. Apply Sw(W). Let Ξ be the sequence of flips $W + (0, -1) \mapsto W + (0, -2)$, …, $R(k+1, l+1) \mapsto R(k+1, l)$. Apply Ξ. See Figure 5.39. In the resulting Hamiltonian

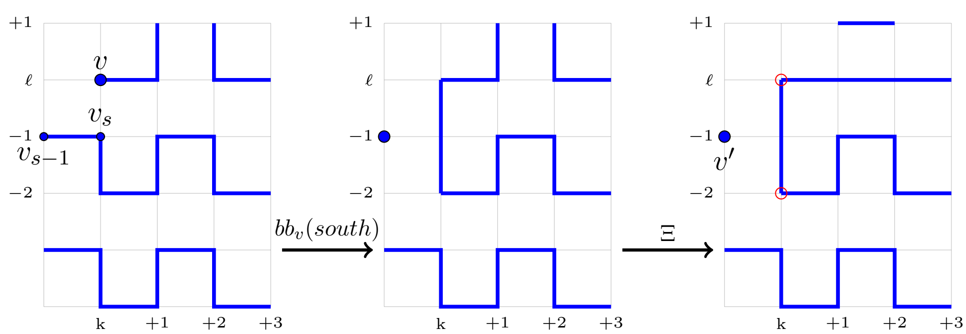

Fig. 5.39. Case 1.2, where Ξ is the cascade Sw(W), $W + (0, -1) \mapsto W + (0, -2), \ldots, R(k+1, l+1) \mapsto R(k+1, l)$.

path H', let $v(k-1, l-1) = v'$ and $v(k, l-1) = v_{s'}$. Then $v_{s'-1} = v(k, l)$ or $v_{s'-1} = v(k, l-2)$. If $v_{s'-1} = v(k, l)$ (Figure 5.40 (a)), then after applying $bb_{v(k-1,l-1)}(east)$, we are back to Case 2.2(c_1) of Proposition 5.6; and if $v_{s'-1} = v(k, l-2)$ (Figure 5.40 (b)), then after applying $bb_{v(k-1,l-1)}(east)$, $R(k+1, l-1) \mapsto R(k+1, l-2)$, we are back to a translation by $(0, -2)$ of Case 2.2(c_1) of Proposition 5.6. End of Case 1.2. End of Case 1.

[11] Note that if the j-stack of A_0s is empty, this sequence consists of only the move $W + (0, -1) \mapsto W + (0, -2)$.

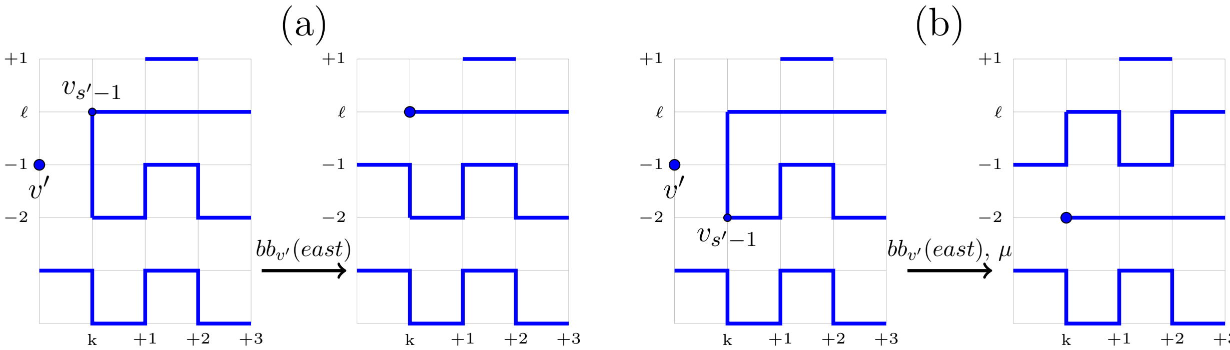

Fig. 5.40. (a) Case 1.2, after $bb_v(south)$, Ξ, with $v_{s'-1} = v(k,l)$. (b) Case 1.2, after $bb_v(south)$, Ξ, with $v_{s'-1} = v(k,l-2)$, and where μ is the move $R(k+1,l-1) \mapsto R(k+1,l-2)$.

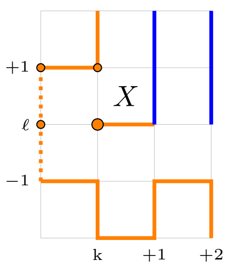

Fig. 5.41. Case 2.1, with $v(k,l+1) \notin P_0$.

CASE 2: $v(k,l-1) \notin P_0$. Then either $j = 0$ or $j > 0$.

CASE 2.1: $j = 0$. Note that if $X + (0,1) \in P(X,Y)$, then after $X + (0,1) \mapsto W$, we are back to Case 2.1 of Proposition 5.6, so we may assume that $X + (0,1) \notin P(X,Y)$. It follows that $X + (-1,0) \in P(X,Y)$. First we will show that $v(k,l+1)$ must belong to P_0. For a contradiction, assume that $v(k,l+1) \notin P_0$. Consider the box $X + (-1,0)$ of $P(X,Y)$. By Lemma 4.1, at least one vertex of $X + (-1,0)$ must belong to P_0. We have that $v(k,l) \in P_1$, and since $v(k,l+1) \notin P_0$, $v(k-1,l+1)$ is not in P_0 either. Now, either $e(k-1;l,l+1) \in H$, or $e(k-1;l,l+1) \notin H$. If $e(k-1;l,l+1) \in H$, then $v(k-1,l)$ is not in P_0 either, contradicting Lemma 4.1. And if $e(k-1;l,l+1) \notin H$, then we must have $(k-1;l-1,l) \in H$, and so again, $v(k-1,l)$ cannot be in P_0, contradicting Lemma 4.1. It follows that we must have $v(k,l+1) \in P_0$. See Figure 5.41.

Let $v_s = v(k,l+1)$. Then $1 < x < s < y < r$, and one of the edges of W in H is in $P(v_1,v_s)$, while the other is in $P(v_s,v_r)$. Now, $v(k,l+2) = v_{s+1}$ or $v(k,l+2) = v_{s-1}$. The former can be shown to be impossible; however it is shorter to note that if it were true (Figure 5.42), after $\mathrm{Sw}(X+(0,1))$, we are back to Case 2.1 of Proposition 5.6, so we may assume the latter. Now, either $e(k,k+1;l+2) \in H$, or $e(k,k+1;l+2) \notin H$.

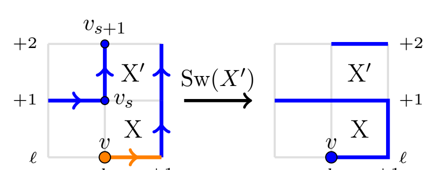

Fig. 5.42. Case 2.1, with $v(k,l+2) = v_{s+1}$; $X' = X+(0,1)$.

CASE 2.1(a): $e(k,k+1;l+2) \in H$. Then, after $bb_v(north)$, $W + (0,1)$ is switchable; and after Sw$(W+(0,1))$, we are back to a translation by $(0,2)$ of Case 2.2(c_1) of Proposition 5.6. See Figure 5.43 (a). End of Case 2.1(a).

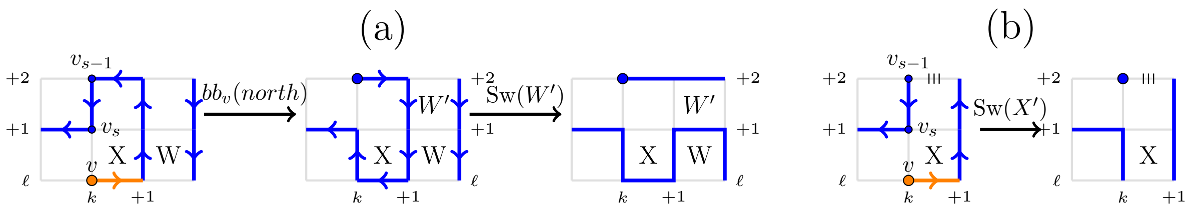

Fig. 5.43. (a) Case 2.1(a), with $W' = W + (0,1)$. (b) Case 2.1(b).

CASE 2.1(b): $e(k,k+1;l+2) \notin H$. Then, after $bb_v(north)$, we are back to a translation by $(0,2)$ of Case 1 of Proposition 5.6. See Figure 4.43(b). End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $j > 0$. Note that all vertices of $R(k,l)$, except perhaps for $v(k,l+1)$ belong to P_1. Then, by Lemma 4.1, $v(k,l+1) \in P_0$. Let $v(k,l+1) = v_s$. Then $v_{s-1} = v(k,l+2)$ or $v_{s-1} = v(k-1,l+1)$.

CASE 2.2(a): $v_{s-1} = v(k,l+2)$. Then we apply $bb_v(north)$. By Corollary 1.4.5, W is now parallel. Apply Sw(W). Apply the sequence of flips $W+(0,-1) \mapsto W+(0,-2), \ldots, R(k+1,l+3) \mapsto R(k+1,l+2)$[12]. See Figure 5.44. Now we are back to a translation by $(0,2)$ of Case 2.2(c_1) of Proposition 5.6. End of Case 2.2(a)

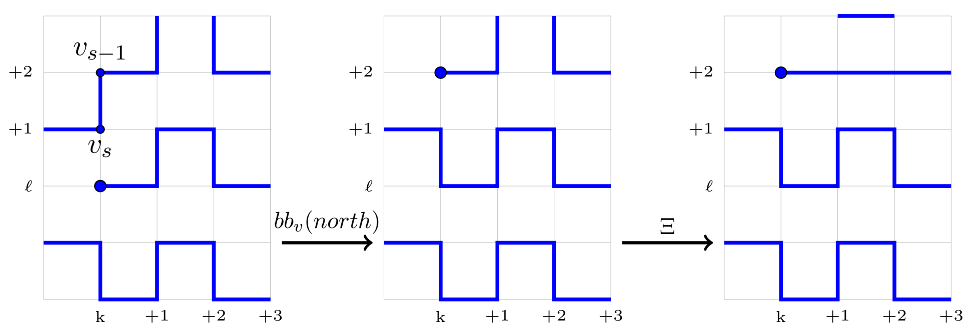

Fig. 5.44. Case 2.2(a), where Ξ is the cascade Sw(W), $W + (0,-1) \mapsto W+(0,-2), \ldots, R(k+1,l+3) \mapsto R(k+1,l+2)$.

CASE 2.2(b): $v_{s-1} = v(k-1,l+1)$. Then we apply $bb_v(north)$. By Corollary 1.4.5, W is now parallel. Apply Sw(W). Let Ξ be the sequence of flips $W+(0,-1) \mapsto W+(0,-2), \ldots, R(k+1,l+3) \mapsto R(k+1,l+2)$[12]. See Figure 5.45. Apply Ξ. In the resulting Hamiltonian path H', let $v(k-1,l+1) = v'$ and $v(k,l+1) = v_{s'}$. Then $v_{s'-1} = v(k,l+2)$ or $v_{s'-1} = v(k,l)$. If $v_{s'-1} = v(k,l+2)$, then after applying $bb_{v'}(east)$, we are back to a translation by (0,2) of

[12] Note that if $j = 1$, this sequence is empty.

Case 2.2(c_1) of Proposition 5.6 (Figure 5.46(a)); and if $v_{s'-1} = v(k,l)$, then after applying $bb_{v'}(east)$, $R(k+1,l+1) \mapsto R(k+1,l)$, we are back to Case 2.2(c_1) of Proposition 5.6 (Figure 5.46(b)). End of Case 2.2(b). End of Case 2.2. End of Case 2. □.

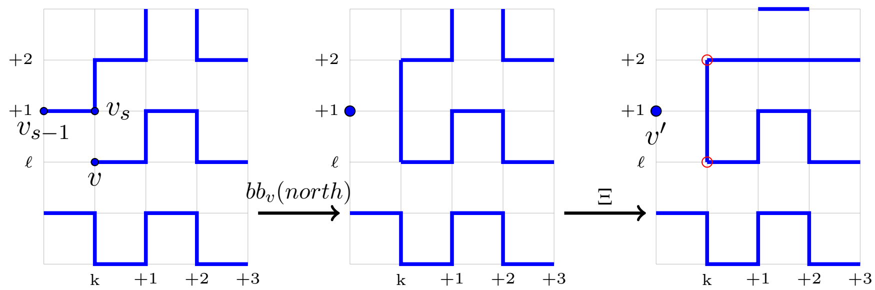

Fig. 5.45. Case 2.2(b), where Ξ is the cascade Sw(W), $W + (0,-1) \mapsto W + (0,-2), \ldots, R(k+1,l+3) \mapsto R(k+1,l+2)$.

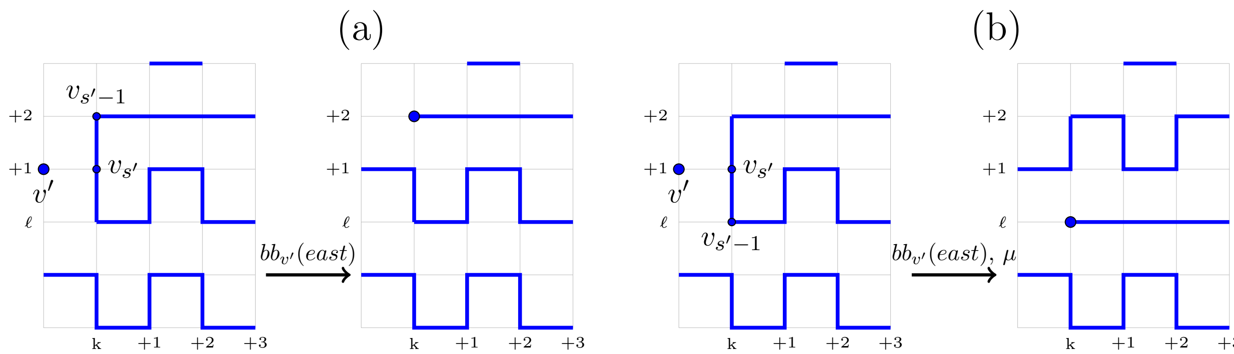

Fig. 5.46. (a) Case 2.2(b), after $bb_v(north)$, Ξ, with $v_{s'-1} = v(k, l+2)$. (b) Case 2.2(b), after $bb_v(north)$, Ξ, with $v_{s'-1} = v(k,l)$, and where μ is the move $R(k+1,l+1) \mapsto R(k+1,l)$.

Bound for Proposition 5.6 (EtB). The argument is similar to that of Observation 4.33. Let $n \geq m$. The longest cascade required to relocated an end-vertex closer to a desired side of R_0 arises when proceeding through Case 2.2(c_2).(ii) of Proposition 5.6. That is, $(e(k+1,k+2;l-1), js\text{-}A_0, A_1)$, where $P(X,Y)$ is contained in a standard looping fat path that has a turn T with a set of lengthenings $\mathcal{T}(T)$, with $|\mathcal{T}(T)|$ and j as large as possible, and where T has an edge weakening. Note that if T has an end-vertex weakening instead, by Claims 5.18, 5.19, 5.20, and 5.21, we can find shorter cascades after which ($\dagger_1$) and ($\dagger_2$) hold. As in Observation 4.33, j can be at most $\frac{n}{2}$. By Proposition 4.28, it takes at most m moves for an edge weakening of T. After that, by Lemma 4.32, it takes at most two moves to switch W, and then at most $j - 1 \leq \frac{n}{2} - 1$ moves to return to Case 2.2(c_1). Thus, it takes at most $\frac{n}{2} - 1 + m + 2 = \frac{n}{2} + m + 1$ to return to Case 2.2(c_1). By the same argument, we need at most another $\frac{n}{2} + m + 1$ moves to return to Case 2.1, and then at most another two moves for ($\dagger_1$) and ($\dagger_2$) to hold. Therefore, in the worst case scenario, at most $2(\frac{n}{2} + m + 1) + 2 = n + 2m + 4$ moves are required to relocate an end-vertex closer to the desired side of R_0. We wish to relocate the end-vertices to opposite sides of R_0, so EtB requires at most $n(n + 2m + 4) = n^2 + 2mn + 4n$ moves.

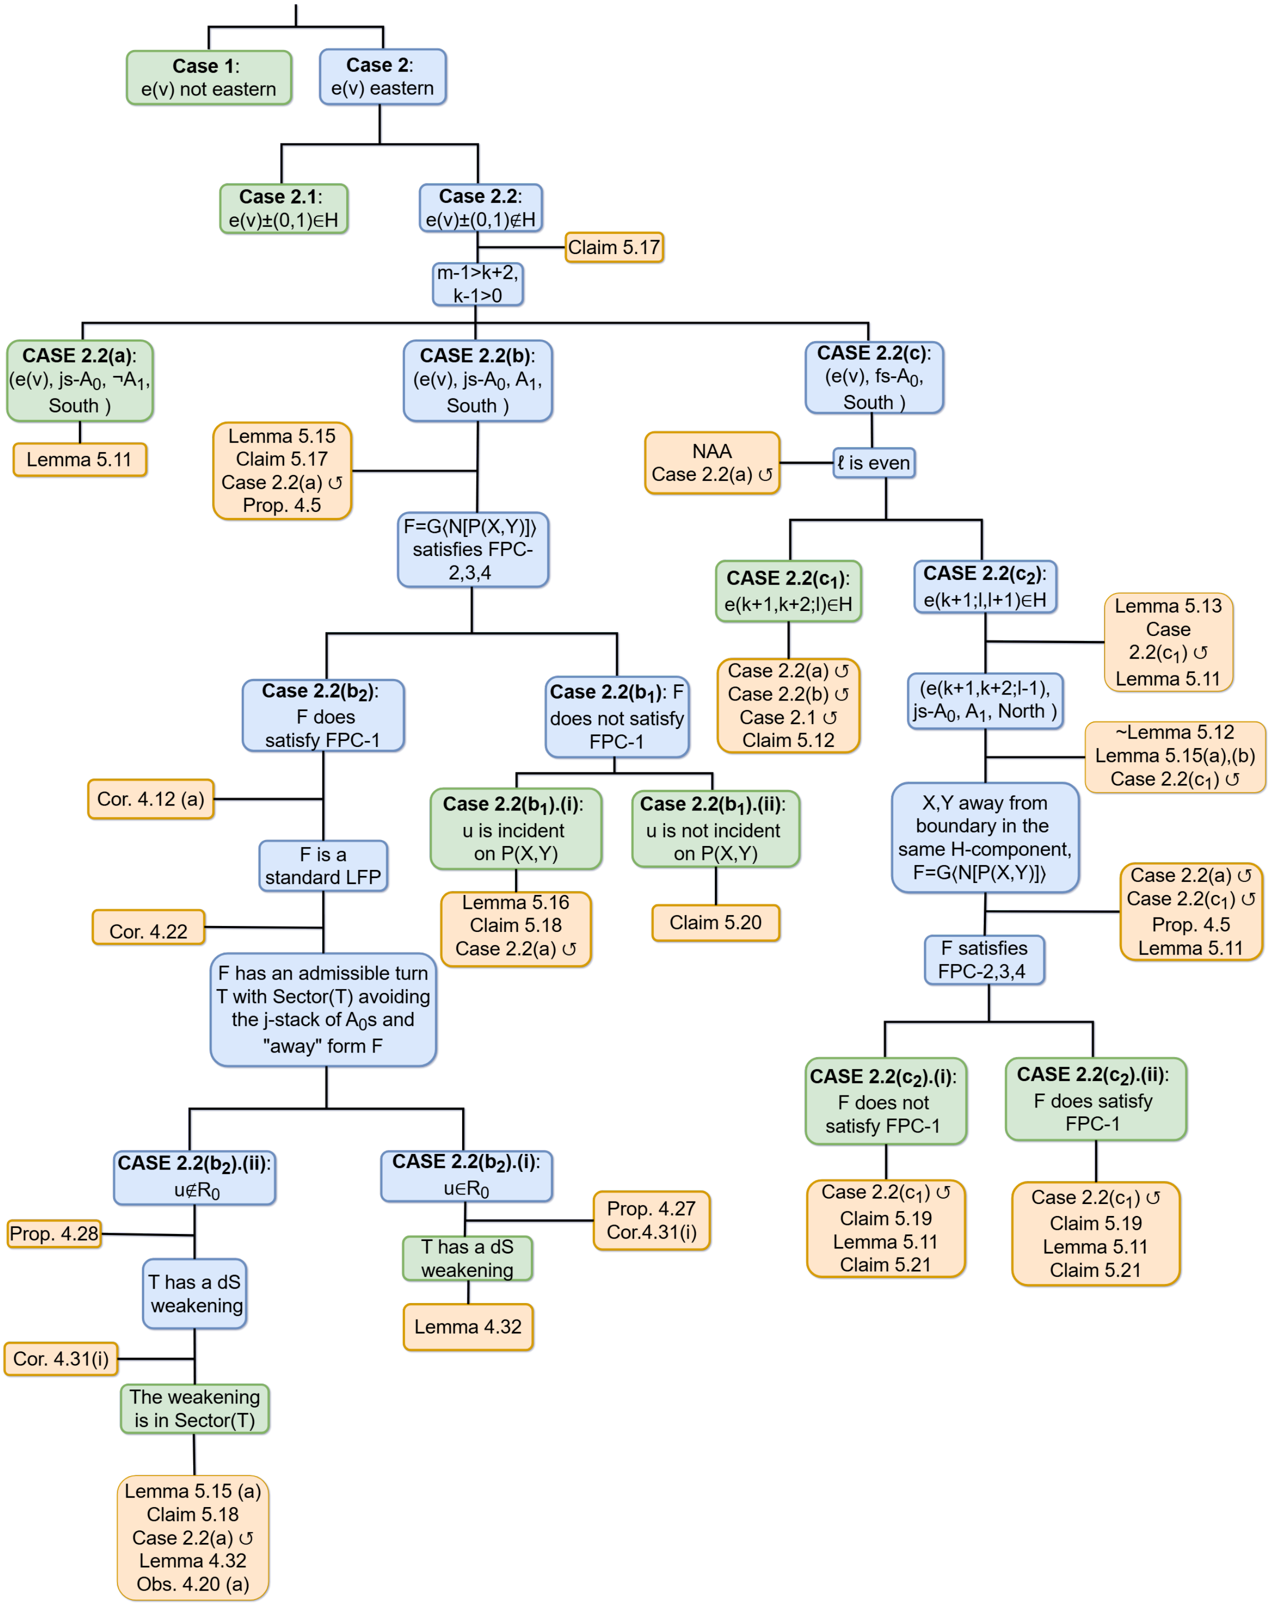

Flowchart 2. Proof structure for Proposition 5.6. Green boxes indicate terminal cases, blue

boxes indicate intermediate steps requiring further analysis. Orange boxes show the results applied at each step. When an orange box connects to a green box, that result completes the subcase. When an orange box appears on an edge between blue boxes, that result is used to transition between cases.

5.3 Algorithm for relocating an endpoint to a corner (EtC)

In this section, we prove Proposition 5.7. Four similar scenarios arise in the proof, each featuring an edge of H followed by an A-type. We consider those scenarios together in Claim 5.25. As with Proposition 5.6, the claim can be read after the main proof. For Claim 5.25, we will also need Lemma 5.22.

Lemma 5.22. Let G be an $m \times n$ grid graph, with $n \geq m \geq 3$. Let H be a Hamiltonian path of G with end-vertices u and v. Assume that v is on the eastern side of R_0, with y-coordinate at most b, and that u is on the western side of R_0. Let $a \in \{0, 1\}$ and let the edge $e = e(a; b, b+1)$ of H be followed by a j-stack of A_0's, which is then followed by an A_1-type with switchable middle-box W. Let X and Y be the boxes adjacent to W that are not its H-neighbours, and assume that $F = G\langle N[P(X, Y)]\rangle$ is a standard looping fat path. Then there is a double-switch cascade of length at most $m + 2$ that switches W, and that avoids the j-stack of A_0's and the boxes incident on e.

Analogous statements apply for the other compass directions.

Proof. By Corollary 4.22, F has an admissible turn T such that Sector(T) avoids the j-stack of A_0's and lies above the line $y = -x + a + b + 2j + 2$. See Figure 5.47. By Corollary 4.19(a), $d(T) \geq 3$. By Proposition 4.27, T has a double-switch weakening $\mu_1, \ldots, \mu_s$. By Corollary 4.31(i), $\mu_1, \ldots, \mu_s$ is contained in Sector(T), and thus it avoids the j-stack of A_0's. By Lemma 4.32 and the fact that there are no parallel switchable boxes, $\mu_1, \ldots, \mu_s$ can be extended by a double-switch move, after which W is switched. The bound follows from Proposition 4.27 and Lemma 4.32. □

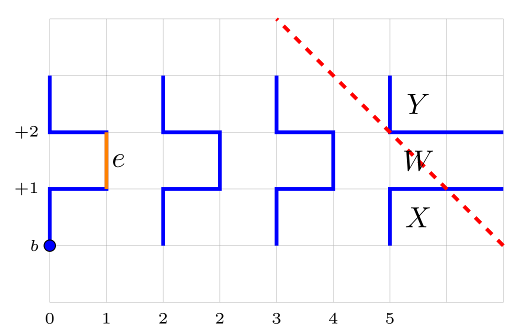

Fig. 5.47. An illustration with $a = 1$ and $j = 2$. The line $y = -x + a + b + 2j + 2$ in red.

Now we are ready to prove Proposition 5.7.

Proof of Proposition 5.7. Orient H as $v = v_1, \ldots, v_r = u$. By Corollary 1.5.2, all switchable boxes are anti-parallel. We will use this fact implicitly and repeatedly throughout the proof. We will also use Lemmas and Corollaries 1.4.1-1.4.5 repeatedly, and often implicitly. For definiteness, assume that $v = v(0, l)$ is on the western side of R_0, that u is on the eastern side, and that $c_v = v(0, n-1)$. To prove (i), it is enough to show that if u is not at c_u and v is not at c_v, then there is a cascade after which v is relocated to $v = v(0, l+2)$ or $v = v(0, l+4)$, and u remains on s_u. We call this statement ($\ddagger_1$) for reference. To prove (ii), it is enough to show that if u is at c_u and v is not at c_v, then there is a cascade after which v is relocated to $v = v(0, l+2)$ or $v = v(0, l+4)$, and u is at c_u. We call this statement ($\ddagger_2$) for reference.

There is a lot of overlap between the proofs of (i) and (ii), so we will combine the proofs. Either $e(0; l, l+1) \in H$ or $e(0; l, l+1) \notin H$.

CASE 1: $e(0; l, l+1) \notin H$. Then $S_{\downarrow}(0, l+2; 1, l+1) \in H$. Let $v(0, l+1) = v_q$. If $v(0, l+2) = v_{q-1}$, then $bb_v(north)$, v is relocated to $v(0, l+2)$ and ($\ddagger_1$) and ($\ddagger_2$) hold. So we may assume that $v(0, l+2) = v_{q+1}$. Let $Z = R(0, l+2)$. Either $e(1; l+1, l+2) \in H$ or $e(1; l+1, l+2) \notin H$.

CASE 1.1: $e(1; l+1, l+2) \notin H$. Either $e(0, 1; l+2) \in H$ or $e(0, 1; l+2) \notin H$.

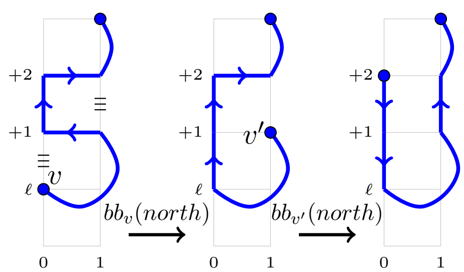

Fig. 5.48. Case 1.1(a). $v' = v(1, l+1)$.

CASE 1.1(a): $e(0, 1; l+2) \in H$. Then, after $bb_v(north)$, $bb_{v(1,l+1)}(north)$, ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 1.1(a). See Figure 5.57.

CASE 1.1(b): $e(0, 1; l+2) \notin H$. Since $u \neq v(0, n-1)$ and since v has the same parity as the corner $v(0, n-1)$, $n-1 \geq l+4$. Then $e(0; l+2, l+3) \in H$ and $e(1; l+2, l+3) \in H$. Note that if $e(0, 1; l+3) \in H$ (Figure 5.49 (a)), then we must have $e(0, 1; l+4) \in H$. Then, after $Z + (0, 1) \mapsto Z$, we are back to Case 1.1(a), so we may assume that $e(0, 1; l+3) \notin H$. Then we must have that $n-1 \geq l+4$, and that $e(0; l+3, l+4) \in H$. Let $v_s = v(1, l+1)$ and $v_t = v(1, l+2)$. Then, either $s < t$ or $s > t$.

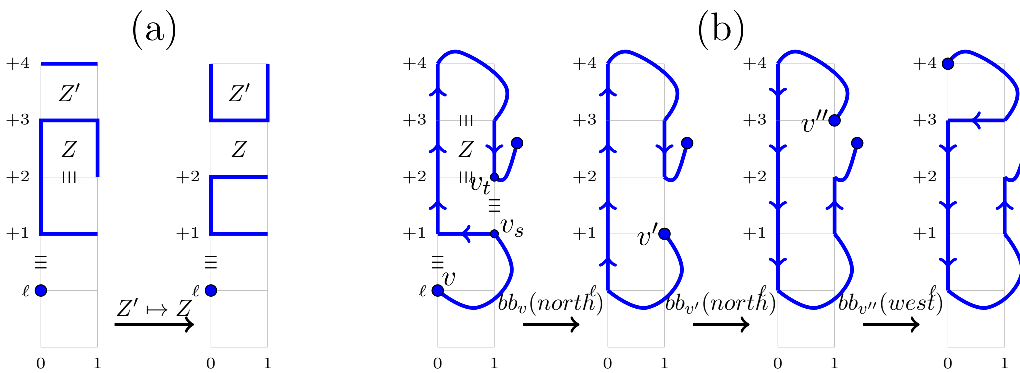

Fig. 5.49. (a) Case 1.1(b) with $e(0, 1; l + 3) \in H$. (b) Case 1.1(b_1). $v' = v(1, l + 1)$, $v'' = v(1, l + 3)$.

CASE 1.1(b_1): $s < t$. Then, after $bb_v(north)$, $bb_{v(1,l+1)}(north)$, $bb_{v(1,l+3)}(west)$, ($\ddagger_1$) and ($\ddagger_2$) hold. See Figure 5.49 (b).End of Case 1.1 (b_1).

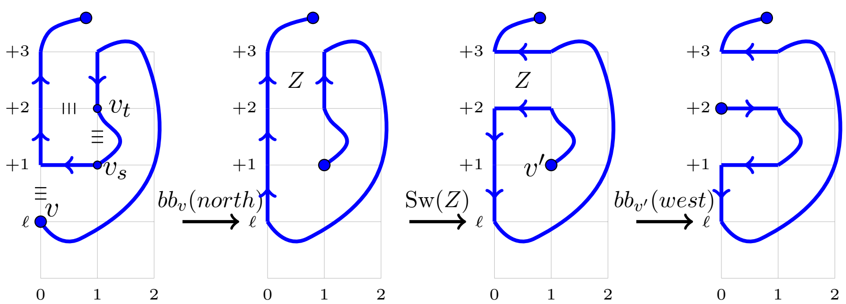

Fig. 5.50. Case 1.1(b_2). $Z = R(0, l + 2)$ and $v' = v(1, l + 1)$.

CASE 1.1(b_2): $s > t$. Then, after $bb_v(north)$, $\text{Sw}(Z)$, $bb_{v(1,l+1)}(west)$, ($\ddagger_1$) and ($\ddagger_2$) hold. See Figure 5.50. End of Case 1.1(b_2). End of Case 1.1(b). End of Case 1.1.

CASE 1.2: $e(1; l + 1, l + 2) \in H$. Then $e(0, 1; l + 2) \notin H$. Since $u \neq v(0, n - 1)$ and since v has the same parity as the corner $v(0, n - 1)$, $n - 1 \geq l + 4$. Then $e(0; l + 2, l + 3) \in H$. See Figure 5.51 (a). Now, either $e(1; l + 2, l + 3) \in H$, or $e(1; l + 2, l + 3) \notin H$.

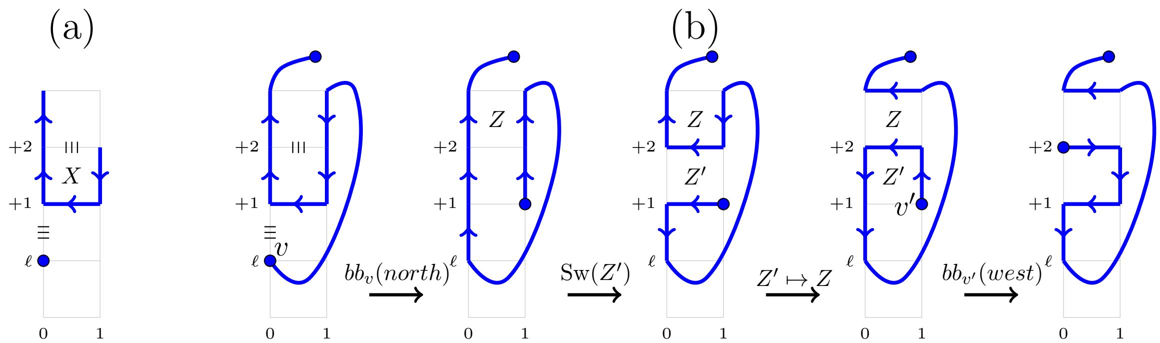

Fig. 5.51. (a) Case 1.2. (b) Case 1.2(a); $Z' = Z + (0, -1)$, $v' = v(1, l + 1)$.

CASE 1.2(a): $e(1; l+2, l+3) \in H$. Then there is a cascade $bb_v(north), \mathrm{Sw}(Z+(0,-1)), Z+(0,-1) \mapsto Z$, $bb_{v(1,l+1)}(west)$, after which ($\ddagger_1$) and ($\ddagger_2$) hold. See Figure 5.51 (b). End of Case 1.2(a).

CASE 1.2(b): $e(1; l+2, l+3) \notin H$. Then $e(1,2; l+2) \in H$. See Figure 5.52 (a). Now, either $e(0; l+3, l+4) \in H$ or $e(0; l+3, l+4) \notin H$.

CASE 1.2(b_1): $e(0; l+3, l+4) \in H$. Then $S_\downarrow(1, l+4; 2, l+3) \in H$. Note that if $e(2; l+3, l+4) \in H$ (Figure 5.52 (b)), then, after $Z+(1,0) \mapsto Z+(1,1)$, we are back to Case 1.2(a). And if $e(2; l+1, l+2) \in H$ (Figure 5.52 (c)), then, after $Z+(1,0) \mapsto Z+(1,-1)$, we are back to Case 1.1(b). So, we may assume that $e(2; l+3, l+4) \notin H$ and that $e(2; l+1, l+2) \notin H$. We call this configuration I_1 for reference (Figure 5.52 (d)). We will prove in Claim 5.23 below that in this case, ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 1.2(b_1).

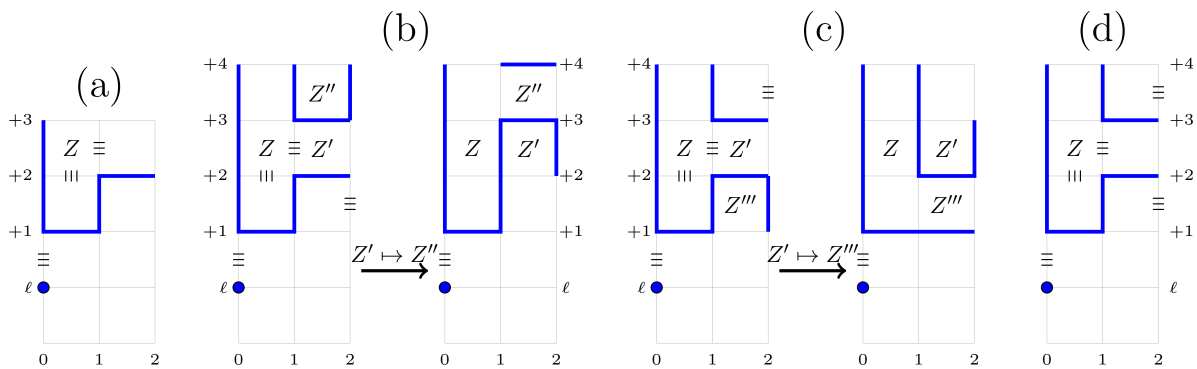

Fig. 5.52. (a) Case 1.2(b). (b) Case 1.2(b_1) with $e(2; l+3, l+4) \in H$. (c) Case 1.2(b_1) with $e(2; l+1, l+2) \in H$. (d) Configuration I_1.

CASE 1.2(b_2): $e(0; l+3, l+4) \notin H$. Then $e(0,1; l+3) \in H$. Since $u \neq v(0, n-1)$, $n-1 > l+4$. Then $S_\downarrow(0, l+5; 1, l+4) \in H$. Note that if $e(1,2; l+3) \in H$ (Figure 5.53(a)), then, after $Z+(1,0) \mapsto Z+(0,1)$, we are back to Case 1.2(a) so we may assume that $e(1,2; l+3) \notin H$. Then $e(1; l+3, l+4) \in H$.

If $e(2; l+3, l+4) \in H$ (Figure 5.53(b)), then, after $Z+(1,1) \mapsto Z+(0,1)$, we are back to Case 1.2(b_1), so we may assume that $e(2; l+3, l+4) \notin H$. Note that this implies that $m-1 \geq 3$. Now, if $e(1; l+3, l+4)$ is not followed by an A_0-type or an A_1-type, then, by Lemma 5.11, there is a cascade of flips after which $e(2; l+3, l+4) \in H$. Then we can apply $Z+(1,1) \mapsto Z$, and once again we are back to Case 1.2(b_1). Therefore we may assume that $e(1; l+3, l+4)$ is followed by an A_0-type or an A_1-type. We call this configuration I_2 for reference (Figure 5.53(c)). We will prove in Claim 5.23 below that in this case, ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 1.2(b_2). End of Case 1.2(b). End of Case 1.2. End of Case 1. □

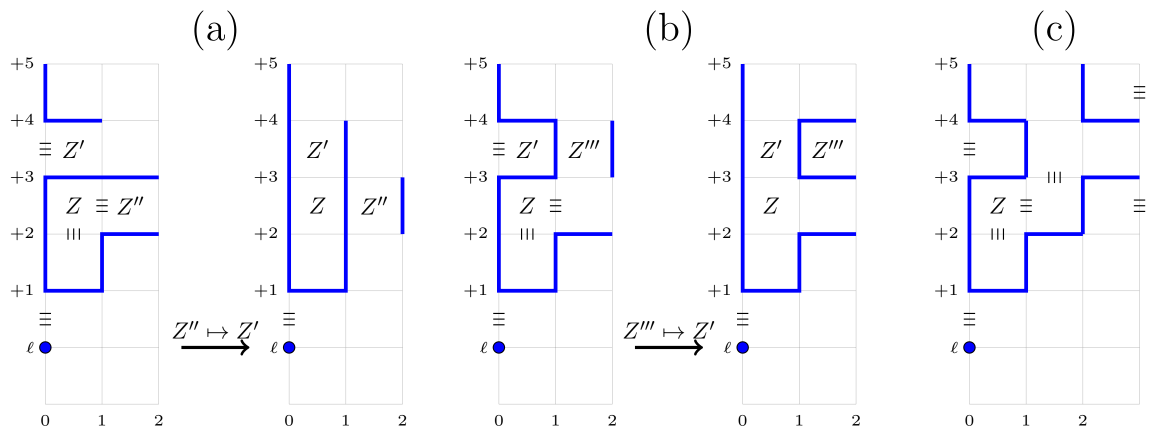

Fig. 5.53. (a) Case 1.2(b_2) with $e(1,2;l+3) \in H$. (b) Case 1.2(b_2) with $e(2;l+3,l+4) \in H$. (c) Configuration I_2.

CASE 2: $e(0;l,l+1) \in H$. Let $Z = R(0,l)$. Either $e(1;l,l+1) \in H$, or $e(1;l,l+1) \notin H$.

CASE 2.1: $e(1;l,l+1) \in H$. Now, either $e(0,1;l+1) \in H$ or $e(0,1;l+1) \notin H$.

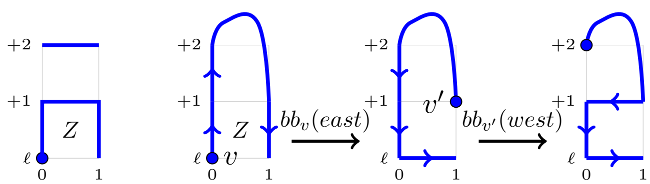

Fig. 5.54. (a) Case 2.1(a). (b) Case 2.1(b); $v' = (1,l+1)$.

CASE 2.1(a): $e(0,1;l+1) \in H$. Then we must have $e(0,1;l+2) \in H$. Then, after $Z+(0,1) \mapsto Z$, we are back to Case 1. See Figure 5.54 (a). End of Case 2.1(a).

CASE 2.1(b): $e(0,1;l+1) \notin H$. Then we must have $e(0;l+1,l+2) \in H$. Then we can apply $bb_v(east)$, $bb_{v(1,l+1)}(west)$, and then ($\ddagger_1$) and ($\ddagger_2$) hold. See Figure 5.54 (b). End of Case 2.1(b). End of Case 2.1.

CASE 2.2: $e(1;l,l+1) \notin H$. Since $m-1>1$, we must have $S_\uparrow(1,l-1;2,l) \in H$, and since u is on the eastern side of R_0, we must have $S_\uparrow(0,l-2;1,l-1) \in H$. Now, either $e(0;l+1,l+2) \in H$, or $e(0;l+1,l+2) \notin H$.
CASE 2.2(a): $e(0;l+1,l+2) \in H$. Then $S_\downarrow(1,l+2;2,l+1) \in H$. We call this configuration I_3 for reference. See Figure 5.55 (a). We will prove in Claim 5.23 below, that in this case ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 2.2(a).

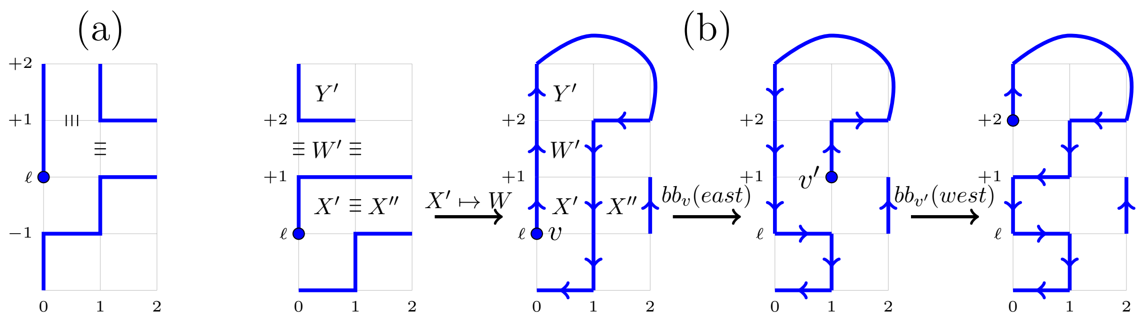

Fig. 5.55. (a) Case 2.2(a), Configuration I_3. (b) Case 2.2(b_1).(i); where $X'' = X' + (1,0)$ and $v' = v(1, l+1)$.

CASE 2.2(b): $e(0; l+1, l+2) \notin H$. Then $S_{\downarrow}(0, l+3; 1, l+2) \in H$. Let $W' = R(0, l+1)$, $X' = R(0, l)$ and $Y' = R(0, l+2)$. Now, either $e(1, 2; l+1) \in H$ or $e(1, 2; l+1) \notin H$.

CASE 2.2(b_1): $e(1, 2; l+1) \in H$. Then $e(1; l+1, l+2) \notin H$. Now, X' and Y' either belong to the same H-component, or they do not.

CASE 2.2(b_1).(i): X' and Y' belong to the same H-component. This implies that $e(2; l, l+1) \notin H$. Let J be the H component containing X' and Y'. Then $X' + (1, 0)$ is a switchable box in the H-path $P(X', Y')$ contained in J. Then, by Proposition 4.5, $X' + (1, 0) \mapsto W'$ is a valid move. Then we can apply $bb_v(east)$, $bb_{v(1,l+1)}(west)$, after which ($\ddagger_1$) and ($\ddagger_2$) hold. See Figure 5.55 (b). End of Case 2.2(b_1).(i).

CASE 2.2(b_1).(ii): X' and Y' belong to distinct H-components. By Lemma 5.15 (c), there is a cascade that keeps u on the eastern side of R_0, after which W' is switched. Then, we are back to Case 2.2(a). See Figure 5.56. End of Case 2.2(b_1).(ii). End of Case 2.2(b_1).

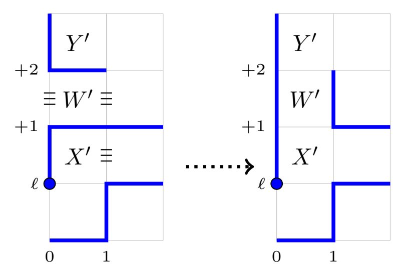

Fig. 5.56. Case 3.2(b_1).(ii).

CASE 2.2(b_2): $e(1, 2; l+1) \notin H$. Then $e(1; l+1, l+2) \in H$. By Lemma 5.11, if $e(1; l+1, l+2)$ is not followed by an A_0-type or an A_1-type, then there is a cascade after which $e(2; l+1, l+2) \in H$. Then we can apply $W' + (1, 0) \mapsto W'$, and then we are back to Case 2.2(a). Therefore we may assume that $e(1; l+1, l+2)$ is followed by an A_0-type or an A_1-type. We call this configuration I_4 for reference. See Figure 5.57 We will prove in Claim 5.23 below that in this case, ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 2.2(b_2). End of Case 2.2(b). End of Case 2.2. End of Case 2. □

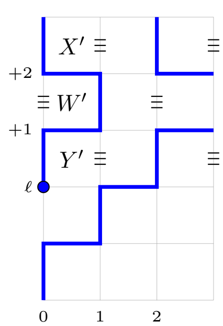

Fig. 5.57. Case 3.2(b_2). Configuration I_4.

It remains to prove Claim 5.23. We postpone proving the bound on the length of the cascade required by the EtC algorithm until after we prove Claim 5.23.

Claim 5.23. Suppose that we are in one of the configurations $I_1, \ldots, I_4$, encountered in the proof of Proposition 5.7. Then there is a cascade after which $(\ddagger_1)$ and $(\ddagger_2)$ are satisfied.

Proof. For $p \in \{1, \ldots, 4\}$, each configuration I_p (Figure 5.58) features an edge $e = e(a; b, b+1)$ with $a \in \{0, 1\}$ followed by an A_1-type or an A_0-type, and such that the end-vertex v has y-coordinate at most b. There are three possibilities:

1. $(e, js\text{-}A_0, \neg A_1, \text{East})$, 2. $(e, fs\text{-}A_0, \text{East})$, and 3. $(e, js\text{-}A_0, A_1, \text{East})$.

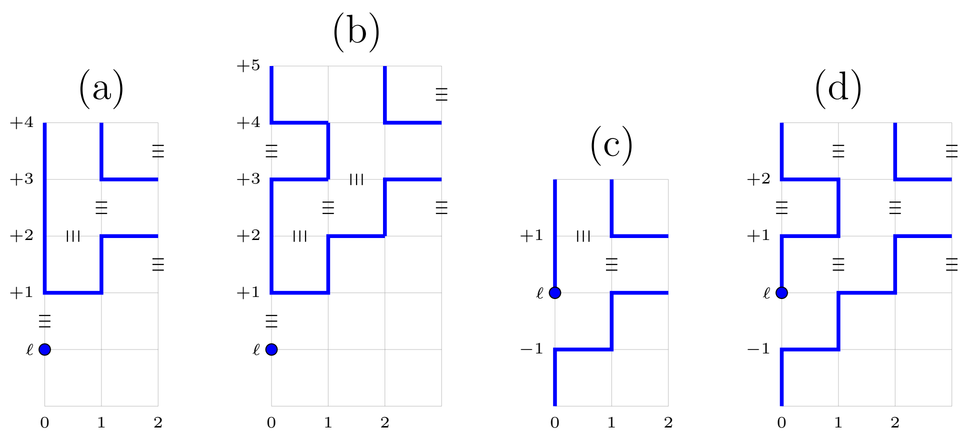

Fig. 5.58. (a) Configuration I_1. (b) Configuration I_2. (c) Configuration I_3. (d) Configuration I_4.

CASE 1: $(e, js\text{-}A_0, \neg A_1, \textit{East})$. Let $a' = 2j$. Since the j-stack of A_0s is not a full stack, $m - 1 \geq a' + 2$. We consider each $p \in \{1, 2, 3, 4\}$ separately.

CASE 1.1: $p = 1$. Note that if $e(a'+1; l+2, l+3) \in H$, then after $R(a', l+2) \mapsto R(a'-1, l+2), \ldots, R(2, l+2) \mapsto R(1, l+2)$, $R(1, l+2)$ is switched, and we are back to Case 1.2(a) of Proposition 5.7. Therefore we may assume that $e(a'+1; l+2, l+3) \notin H$. See Figure 5.59.

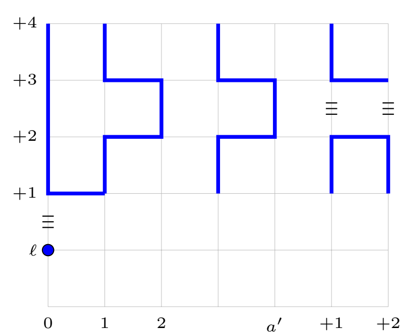

Fig. 5.59. Case 1.1.

If the rectangle $R(a'+1, a'+2; l+1, l+4)$ does not contain u, then, by Lemma 5.11, there is a cascade of flips after which $R(1, l+2)$ is switched, and we are back to Case 1.2(a) of Proposition 5.7. Therefore we may assume that $R(a'+1, a'+2; l+1, l+4)$ does contain u. This implies that $m - 1 = a' + 2$. Note that if

$e(m-1; l+2, l+3) \in H$, then stack of A_0s is a full one, contrary to the Case 1 assumption, so we may assume that $e(m-1; l+2, l+3) \notin H$. Then, at least one of $e(m-1; l+1, l+2)$ and $e(m-1; l+3, l+4)$ belongs to H. WLOG assume that $e(m-1; l+1, l+2) \in H$. Then, after $R(a'+1; l+2) \mapsto R(a'+1; l+1)$, $e(a'+1; l+2, l+3) \in H$, and we can reuse the cascade in the first paragraph. End of Case 1.1.

CASE 1.2: $p = 2$. By an argument similar to the one in Case 1.1, we find a cascade of flips after which $R(2, l+3)$ is switched. Then we can apply $R(1, l+3) \mapsto R(0, l+3)$, and then we are back to Case 1.1. End of Case 1.2.

CASE 1.3: $p = 3$. By an argument similar to the one in Case 1.1, we find a cascade of flips after which $R(1, l)$ is switched, and then we are back to Case 2.1(b) of Proposition 5.7. End of Case 1.3.

CASE 1.4: $p = 4$. By an argument similar to the one in Case 1.1, we find a cascade of flips after which $R(2, l+1)$ is switched. Then we can apply $R(1, l+1) \mapsto R(0, l+1)$, and then we are back to Case 1.3. End of Case 1.4. End of Case 1.

CASE 2: $(e, fs\text{-}A_0, East)$. Recall that one of the assumptions of Proposition 5.7 stated that u and c_u, and v and c_v, have matching parities. We will refer to this assumption as the "end-vertex-corner parity" assumption and write "ECP" whenever we refer to it. We will consider each p separately.

CASE 2.1. $p = 1$. If $m-1$ is odd (Figure 5.60 (a)), then we must have $e(m-1; l+2, l+3) \in H$. Then, after $R(m-2, l+2) \mapsto R(m-3, l+2), \ldots, R(2, l+2) \mapsto R(1, l+2)$, we are back to Case 1.2(a) of Proposition 5.7. Therefore we may assume that $m-1$ is even.

If $m-1 = 2$ (Figure 5.60 (b)), then $u = v(2, l+1)$. Now, since $u = v + (2, 1)$, u and v have different parities, while the corners c_v and c_u on $v(0, n-1)$ and $v(2, n-1)$ must have the same parity. Then either c_u and u have different parities, or c_v and v do. Either way, this contradicts the ECP assumption.

If $m-1 = 4$ (Figure 5.60 (c)), then we must have $S_{\uparrow}(2, l; 3, l+1) \in H$. It follows that $u = v(4, l+1)$, and this also contradicts the ECP assumption.

Suppose that $m-1 \geq 6$. Then, for each $i \in \{2, 4, \ldots, m-5\}$, $S_{\uparrow}(i, l; i+1, l+1) \in H$ implies that $S_{\uparrow}(i+2, l; i+3, l+1) \in H$. It follows that for $i \in \{2, 4, \ldots, m-3\}$, $S_{\uparrow}(i, l; i+1, l+1) \in H$. Then we must have that $u = v(m-1, l+1)$. But then again, this contradicts the ECP assumption. End of Case 2.1.

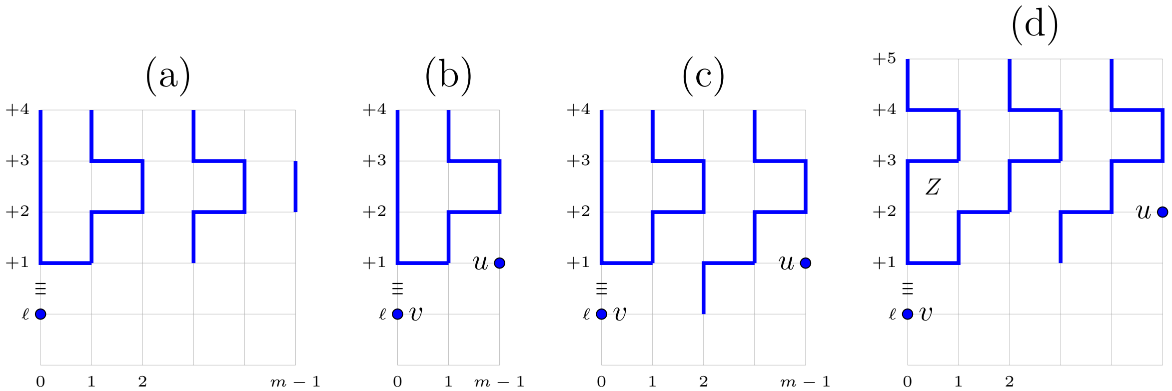

Fig. 5.60. (a) Case 2.1 with $m-1$ odd. (b) Case 2.1 with $m-1=2$. (c) Case 2.1 with $m-1=4$. (d) Case 2.2 with $m-1$ odd.

CASE 2.2. $p=2$. If $m-1$ is even, then we must have $e(m-1;l+3,l+4)\in H$. Then, after $R(m-2,l+3)\mapsto R(m-3,l+3),\dots,R(1,l+3)\mapsto R(0,l+3)$, we are back to Case 2.1.

Therefore we may assume that $m-1$ is odd. See Figure 5.60 (d). Then, using the same argument as in the last paragraph of Case 2.1, we can check that $u=v(m-1,l+2)$. See Figure 5.73(d). Now we see that both end-vertices have y-coordinate at most $l+2$, and there is a full stack of A_0's starting at the small eastern leaf $R(0,l+3)$. Then, by Lemma 5.13, this is an impossible configuration. End of Case 2.2.

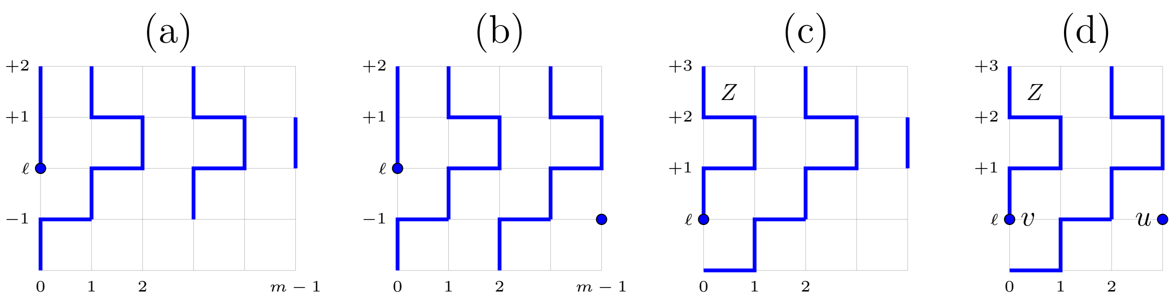

Fig. 5.61. (a) Case 2.3 with $m-1$ odd. (b) Case 2.3 with $m-1$ even. (c) Case 2.4 with $m-1$ odd. (d) Case 2.4 with $m-1$ even.

CASE 2.3. $p=3$. If $m-1$ is odd (Figure 5.61 (a)), then we must have $e(m-1;l,l+1)\in H$. Then, after $R(m-2,l)\mapsto R(m-3,l),\dots,R(2,l)\mapsto R(1,l)$, we are back to Case 2.1(b) of Proposition 5.7. And if $m-1$ is even (Figure 5.61 (b)), then, the same argument as in Case 2.1, shows that $u=v(m-1,l-1)$, which again, contradicts the ECP assumption. End of Case 2.3.

CASE 2.4. $p=4$. If $m-1$ is odd (Figure 5.61 (c)), then we must have $e(m-1;l+1,l+2)\in H$. Then, after $R(m-2,l+1)\mapsto R(m-3,l+1),\dots,R(2,l+1)\mapsto R(1,l+1)$, we are back to Case 2.3. Therefore, we may assume that $m-1$ is even. See figure 5.61 (d). Then, using the same argument as in Case 2.1, we can check that $u = v(m-1, l)$. But then, as in Case 2.2, we can use Lemma 5.13 again to show that this is an impossible configuration. See Figure 5.75(b). End of Case 2.4. End of Case 2.

CASE 3: $(e, js\text{-}A_0, A_1, East)$. Let W be the switchable box of the A_1-type, and let X and Y be the boxes adjacent to W that are not its H-neighbours. Either X and Y belong to distinct H-components, or they do not.

CASE 3.1: X and Y belong to distinct H-components. By Lemma 5.15 (c), there is a cascade after which W is switched, and which otherwise avoids I_p, for $p \in \{1, 2, 3, 4\}$. Then we are back to Case 1. End of Case 3.1.

CASE 3.2: X and Y belong to the same H-component. Let $F = G\langle N[P(X, Y)]\rangle$. Either F is a looping fat path or it is not.

CASE 3.2(a): F is a looping fat path. Let $J(F)$ be the H-component of G that contains F. Since both end-vertices of H are in R_0, no edge of a main trail of $J(F)$ can have an end-vertex incident on it in $G \setminus R_0$. Then, by Lemma 1.3.12, $J(F)$ is non-self-adjacent. Since F is an H-subtree of $J(F)$, F is also non-self-adjacent. By Proposition 4.11(i), F has no polyking junctions. It follows that F is a standard looping fat path. Then, by Lemma 5.22, there is a double-switch cascade $\mu_1, \ldots, \mu_q$ that switches W, that avoids the j-stack of A_0's, and that is above the line $y = -x + a + b + 2j + 2$. See Figure 5.47. Then, after μ_q, we are back to Case 1.

CASE 3.2(b): F is not looping fat path. By Corollary 1.5.2, W cannot be parallel, so FPC-2 must hold. If the end-vertex u is incident on $P(X, Y)$, by Lemma 5.16, $P(X, Y)$ has a switchable Z box incident on the eastern side of R_0. Then, after $Z \mapsto W$, we are back to Case 1. Now, either v is incident on $P(X, Y)$, or it is not.

CASE 3.2(b_1): v is incident on $P(X, Y)$. First we check that for each $p \in \{1, 2, 3, 4\}$, none of $R(a, b+1)$, $R(a, b)$, and $R(a+1, b)$ can belong to $P(X, Y)$.
Suppose that $p = 1$. In this case, $a = 0$ and $b = l + 2$. Either $e(2; l+2, l+3) \in H$ or $e(2; l+2, l+3) \notin H$.

Assume that $e(2; l+2, l+3) \in H$ (Figure 5.62 (a)). If $R(1, l+2) \in P$, then it must be an end-box of P, which it is not, so $R(1, l+2) \notin P$. If $R(0, l+2) \in P$, it is not an end-box, so $R(0, l+1)$ must belong to P. But then $R(0, l+1)$ must be an end-box of P, which it is not, so $R(0, l + 2)$ cannot be in P. A similar argument shows that $R(0, l + 3) \notin P$.

Assume that $e(2; l+2, l+3) \notin H$ (Figure 5.62 (b)). Then $W = R(1, l + 2)$. By Corollary 4.2(b), $R(1, l + 2) \notin P$. The arguments showing that $R(0, l + 2) \notin P$ and $R(0, l + 3) \notin P$ are similar to the ones used in the previous paragraph, so we omit them.

Hence, we have shown that if $p = 1$, then none of $R(a, b + 1)$, $R(a, b)$, and $R(a + 1, b)$ can belong to $P(X, Y)$. The arguments for the cases where $p = 2$, $p = 3$, and $p = 4$ are similar, so we omit them. Now, the assumption on v of Case 3.2(b_1) implies that $p \neq 3$ and $p \neq 4$. Then $p = 1$ or $p = 2$.

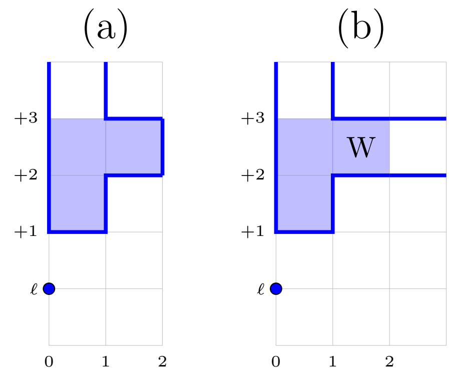

Fig. 5.62. (a) $p = 1$, $e(2; l+2, l+3) \in H$. (b) $p = 1$, $e(2; l + 2, l + 3) \notin H$.

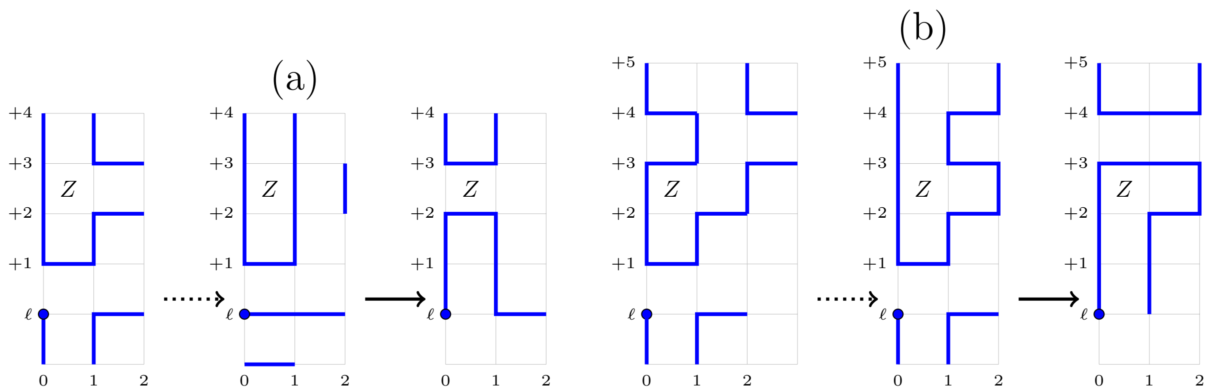

Fig. 5.63. Case 3.2(b_1).(i). (a) Configuration I_1, with $e(v)$ southern. Case 3.2(b_1).(ii). Configuration I_2, with $e(v)$ southern.

CASE 3.2(b_1).(i): $p = 1$. First we check that if $e(v)$ is eastern, then $R(0, l)$ and $R(0, l - 1)$ cannot belong to P. Suppose that $e(v)$ is eastern. Then, if $R(0, l) \in P$, it would have to be an end-box, which it is not, so $R(0, l) \notin P$. Since u is on the eastern side, we must have $S_{\uparrow}(0, l - 2; 1, l - 1) \in H$. Then again, $R(0, l - 1) \in P$, it would have to be an end-box, which it is not, so $R(0, l - 1) \notin P$.

Thus we only need to check the case where $e(v)$ is southern. See Figure 5.63 (a). Then $R(0, l - 1)$ must belong to P and be switchable. Then, after $R(0, l - 1) \mapsto W$, we apply the cascade of flips $W \mapsto W + (-1, 0), \ldots, R(2, l+2) \mapsto R(1, l+2)$, followed by $R(0, l) \mapsto R(2, l)$. See Figure 5.63 (a). Now we apply $bb_v(east)$, $bb_{v(1,l+1)}(west)$, (not shown in Figure 5.63 (a)) after which ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 3.2(b_1).(i).

CASE 3.2(b_1).(ii): $p = 2$. As in the previous case, we must have that $e(v)$ is southern, and that $R(0, l - 1)$ must belong to $P(X, Y)$ and be switchable (Figure 5.63 (b). Then, after

$R(0,l-1) \mapsto W$, we apply the cascade of flips $W \mapsto W+(-1,0), \ldots, R(1,l+3) \mapsto R(0,l+3)$, followed by $R(0,l) \mapsto R(0,l+3)$. See Figure 5.63 (b). Now we apply $bb_v(east)$, $bb_{v(1,l+1)}(west)$, (not shown in Figure 5.63 (b)) after which ($\ddagger_1$) and ($\ddagger_2$) hold. End of Case 3.2(b_1).(ii). End of Case 3.2(b_1).

CASE 3.2(b_2): v is not incident on $P(X,Y)$. Then $P(X,Y)$ must have a switchable box Z. We consider each $p \in \{1,2,3,4\}$ separately.

CASE 3.2(b_2).(i): $p=1$. Since v is not incident on $P(X,Y)$, $Z \neq R(0,l)$ and $Z \neq R(0,l-1)$. Then, after $Z \mapsto W$, we are back to Case 1.2(a) of Proposition 5.7. End of Case 3.2(b_2).(i).

CASE 3.2(b_2).(ii): $p=2$. Since v is not incident on $P(X,Y)$, $Z \neq R(0,l)$ and $Z \neq R(0,l-1)$. Then, after $Z \mapsto W$, we are back to Case 1.2. End of Case 3.2(b_2).(ii).

CASE 3.2(b_2).(iii): $p=3$. Then, after $Z \mapsto W$, we are back to Case 1.3. End of Case 3.2(b_2).(iii).

CASE 3.2(b_2).(iv): $p=4$. Then, after $Z \mapsto W$, we are back to Case 1.4. End of Case 3.2(b_2).(iv). End of Case 3.2(b_2). End of Case 3.2(b). End of Case 3.2. End of Case 3. □

Bound for Proposition 5.7 (EtC). Let $n \geq m$. The longest cascade required to relocate an end-vertex closer to a desired corner while keeping it on the same side of R_0 arises when proceeding through Case 3.2(a) of Claim 5.23, and considering configuration I_2. As before, $j \leq \frac{n}{2}$. By Lemma 5.22, switching W requires at most $m+2$ moves. Then, after $m+2+\frac{n}{2}-1$ moves, we are back to configuration I_1. Using the same argument, we see that at most $m+1+\frac{n}{2}$ moves are required to return to Case 1.1(b_1) of Proposition 5.7, and then three additional moves for ($\ddagger_1$) or ($\ddagger_2$) to be satisfied. Thus we need $n+2m+5$ moves to relocate an end-vertex closer to a desired corner while keeping it on the same side of R_0. Note that it takes at most $\frac{n}{2}$ such relocations to take an end-vertex from a vertex on a side of R_0 to a desired corner incident on that side. Therefore, EtC requires at most $\frac{n^2}{2}+mn+\frac{5n}{2}$ moves to complete (i) or (ii).

5.4 Summary

In this chapter, we proved Theorem 5.9: any two Hamiltonian paths on an $m \times n$ grid graph can be reconfigured into one another. The proof uses Theorem 2.1, which handles reconfiguration of Hamiltonian cycles. We showed that to apply Theorem 2.1 to paths, it suffices to move the end-vertices to corners of R_0. The End-vertex-to-Boundary algorithm (Section 5.2) moves an end-vertex to a side of R_0, and the End-vertex-to-Corner algorithm (Section 5.3) moves it from a side to a corner. Together, these algorithms extend the reconfiguration result from cycles to paths.

6 Conclusion

This chapter summarizes our main results (Section 6.1), discusses their computational complexity and implications for sampling algorithms (Sections 6.2 and 6.3), characterizes extremal reconfiguration cases (Section 6.4), explores the role of boundary structure (Section 6.5), and proposes several directions for future research (Section 6.6). Sections 6.4–6.6 present several conjectures for future investigation.

6.1 Summary of Results

In this dissertation we gave algorithms for reconfiguring Hamiltonian cycles and Hamiltonian paths on rectangular grid graphs. The main results are Theorem 2.1 and Theorem 5.9. Theorem 2.1 proves that any two Hamiltonian cycles on the same $m \times n$ rectangular grid graph with $m \leq n$ can be reconfigured into one another using only double-switch moves, with at most $O(n^2m)$ moves required. Theorem 5.9 proves the analogous result for Hamiltonian paths using double-switch, single-switch, and backbite moves, again with an $O(n^2m)$ move bound. Both algorithms have $O(m^2n^2)$ time complexity with details given in Section 6.2.

We first analyzed the structure that a Hamiltonian path H imposes on a polyomino, decomposing it into H-components. These H-components generalize the notion of cookies defined by Nishat and Whitesides in [17], originally defined for Hamiltonian cycles, so that the concept applies to both Hamiltonian cycles and Hamiltonian paths. Then we established properties of H-components, which we then used to prove the existence of required moves for the reconfiguration algorithms.

We introduce canonical forms for Hamiltonian cycles, designed to be simple yet sufficient as terminal configurations: once a cycle is reconfigured into a canonical form, the reconfiguration problem is essentially solved.

The double-switch move used throughout this dissertation is a generalization of the flip and transpose moves used in [17]. To extend the Hamiltonian cycle reconfiguration result to Hamiltonian paths, we included the single-switch and backbite moves.

The proof of Theorem 2.1 centers on the Reconfiguration to Canonical Form (RtCF) algorithm, supported by the OneLargeCookie (1LC) and ManyLargeCookies (MLC) algorithms. RtCF works iteratively from the outside in on a nested sequence of rectangles $R_0, R_1, \ldots, R_t$, reconfiguring the cycle so that all but two edges of each R_i belong to the cycle before proceeding inward. The 1LC and MLC algorithms establish the existence of required moves at each step.

Theorem 5.9 reduces path reconfiguration to cycle reconfiguration through two new algorithms: End-to-Boundary (EtB) and End-to-Corner (EtC). These algorithms transform any Hamiltonian path into a form where Theorem 2.1 applies. The most involved cases in all four algorithms – 1LC, MLC, EtB, and EtC – are resolved through the fat-path-turn-weakening machinery developed in Chapter 4.

6.2 Computational complexity

In this section we give bounds for the time and space complexity of the reconfiguration algorithms presented in this dissertation. We assume that we are reconfiguring Hamiltonian paths or cycles on an $m \times n$ grid graph, with $n \geq m$.

Space complexity: The grid graph can be represented by an adjacency list using $O(mn)$ space. An adjacency matrix representation would require $O(m^2n^2)$ space, but this is not necessary since the underlying graph is fixed. The current Hamiltonian cycle or path configuration requires $O(mn)$ space to store which edges are currently in use. Therefore, the total space complexity for the reconfiguration task is $O(mn)$. If one wishes to record the complete reconfiguration sequence, this requires additional space proportional to the number of moves. Since the reconfiguration requires at most $O(mn^2)$ moves, storing the complete sequence requires $O(mn^2)$ space. This can be done by storing an initial configuration and a sequence of length mn^2, where each term describes a move and costs $O(1)$.

Time complexity. Consider the following scenario: an edge e followed by a j-stack of A_0-types (with $j \leq n/2$), which is then followed by an A_1-type with switchable middle-box W. The middle-box W determines an H-path $P(X, Y)$. A key computational bottleneck is determining which boxes belong to $P(X, Y)$. Using breadth-first search, this requires $O(mn)$ time in the worst case, as $P(X, Y)$ may traverse up to $O(mn)$ boxes. We now show that the remaining tasks – checking FPC conditions, locating admissible turns, and finding turn weakenings – each require at most $O(mn)$ time.

Checking FPC conditions. Once $P(X, Y)$ is identified, we check the FPC conditions as follows. FPC-2 (checking whether W is parallel) requires checking which two edges of the Hamiltonian path are incident on W, which takes $O(mn)$ time, as well as one final check to see if the edges are parallel. So FPC-2 can be checked in $O(mn)$ time. FPC-1 (checking whether an end-vertex of the Hamiltonian path is incident on a box of $P(X,Y)$) requires checking each box of $P(X,Y)$ against the two end-vertices, taking $O(mn)$ time. FPC-3 (checking whether $P(X,Y)$ contains a switchable box) requires checking whether each box in $P(X,Y)$ is switchable, also $O(mn)$ time.

For FPC-4, we must check whether there exists a move (switch, flip, or transpose) after which W is switched, $P(X,Y)$ gains a switchable box, or $W \to X$ or $W \to Y$ becomes valid. Note that only moves in the vicinity of boxes in $N[P(X,Y)] \setminus P(X,Y)$ can affect $P(X,Y)$ in the required manner. More precisely, it suffices to check moves near leaves or switchable boxes in $N[P(X,Y)] \setminus P(X,Y)$, which can be identified during the initial $O(mn)$ scan of $P(X,Y)$. For each such location, we check a 3×3 region in $O(1)$ time. Since there are at most $O(mn)$ such locations, the total time for FPC-4 is $O(mn)$.

Weakening of an admissible turn in a standard looping fat path. Let$F = G\langle N[P(X,Y)]\rangle$ be a standard looping fat path and assume that we want to locate an admissible turn T of the boundary $B(F)$. Recall that we store the Hamiltonian path as a sequence of vertices with each term identified by its coordinates, and that we begin our search for T at an edge of the A_1-type of F. For definiteness, assume that T is northeastern. This can be determined by checking the direction of the first two edges e_1 and e_2 of T in $O(1)$ time. Next, we need to find the first western edge e_W encountered after e_1, which can be no farther than n units east of the first edge, requiring at most $2n$ direction checks. If the turn is admissible, we are done. If not, we repeat the procedure once more, but this time we need to find the first northern edge, which can be no farther than m units south of the first western edge, requiring at most $2m$ direction checks. Thus, finding an admissible turn requires $O(m+n)$ time in the worst case. Given this admissible turn $T = T_1$, we can find Sector(T) with cost $O(1)$ and check whether an end-vertex is incident on Sector(T) with cost $O(1)$ as well. From this information, we can determine exactly which potential lengthening of T may contain that end-vertex in its flank, avoiding the need to check every potential lengthening individually.

For each turn T_i in the sequence of lengthenings of T (where $i = 1, 2, \ldots, q$ and $q \le m$), we check for a short weakening (a cascade of at most three moves) in $O(1)$ time by examining two 3×3 rectangles centered at the leaves of T_i. If no short weakening exists, there must be a lengthening T_{i+1}, whose exact structure (whether each leaf is open or closed) can be determined during the same check in $O(1)$ time. This process continues for at most $q \le m$ iterations, requiring $O(m)$ time total. Once the final turn T_q is reached (guaranteed to have a short weakening), we apply the weakenings in reverse order from T_q back to T_1, requiring one move per turn for a total of $O(m)$ moves, each taking $O(1)$ time to identify. Thus the time complexity is $O(m+n)$.

The fat-path-turn-weakening procedure. Consider once more the scenario described above: an A_1-type with switchable middle-box W determining an H-path $P(X, Y)$, and the subgraph $F = G\langle N[P(X, Y)]\rangle$. We define the *fat-path-turn-weakening procedure* as the complete process consisting of: (i) identifying which boxes belong to $P(X, Y)$, (ii) checking the FPC conditions for F, (iii) locating an admissible turn T in F whenever F is a standard looping fat path, and (iv) finding a weakening of T. As shown above, this procedure requires $O(mn)$ time. We will reference this procedure in the complexity bounds that follow.

Time complexity of the 1LC algorithm. The 1LC algorithm handles Hamiltonian cycles with exactly one large cookie and at least one small cookie.

Identifying all small cookies requires $O(m + n)$ time. For each outermost small cookie C, the algorithm first checks whether C can be collected by a single move, which takes $O(1)$ time. If not, by Lemma 3.9, C is followed by a j-stack of A_0-types (with $j \leq n/2$) and an A_1-type with switchable middle-box W. Checking for this costs $O(n)$ operations.

In this case, collecting C requires the fat-path-turn-weakening procedure, which as shown above takes $O(mn)$ time. Following this, j flips are executed to collect C, requiring $O(n)$ additional time. Thus, collecting an outermost small cookie requires at most $O(mn)$ time.

Time complexity of Theorem 2.1. The time complexity for MLC is $O(m+n)$ (see pseudocode at the end of Section 3.1), so the overall time complexity is dominated by the 1LC operations.

The algorithm for Theorem 2.1 applies RtCF to both input Hamiltonian cycles to reduce them to canonical forms, then reconfigures between the canonical forms. Thus the total cost is twice the cost of RtCF plus the cost of canonical reconfiguration. The reconfiguration between canonical forms takes $O(m)$ time if we retain the cookie neck information from the RtCF executions, or $O(mn)$ time if computed from scratch. Either way, this is dominated by the RtCF cost.

Recall from the move bound for Theorem 2.1 at the end of Chapter 2 that at iteration j of RtCF, there are at most $(n + m - 4j)$ cookies on rectangle R_j of size $(m - 2j) \times (n - 2j)$. Since the time complexity of 1LC on R_0 is $O(mn)$ per cookie, collecting a cookie via 1LC on R_j requires $O((m - 2j)(n - 2j))$ time. Therefore, the total time complexity is:

$$T = 2 \cdot \sum_{j=0}^{\lfloor m/2 \rfloor} (n + m - 4j) \cdot O((m - 2j)(n - 2j)) + O(m) = O(m^2 n^2).$$

Time complexity of the EtB algorithm. The EtB algorithm moves the two endpoints of a Hamiltonian path to opposite boundaries of the grid. For each endpoint relocation step, either the endpoint can be moved directly in $O(1)$ time, or we are in the setting of Case 2.2 of the proof of Proposition 5.6. In Case 2.2, the costliest scenario requires at most two applications of the fat-path-turn-weakening procedure, along with at most $\frac{n}{2}$ flips. Since each application of the fat-path-turn-weakening procedure takes $O(mn)$ time and the flips take $O(n)$ time, the cost per relocation is $O(mn)$. Moving both endpoints to opposite boundaries may require at most $O(n)$ such relocation steps. Therefore, the total time complexity of the EtB algorithm is $O(n) \cdot O(mn) = O(mn^2)$.

Time complexity of the EtC algorithm. The EtC algorithm moves the two endpoints of a Hamiltonian path to specified corners of the grid. Each relocation step either completes in $O(1)$ time or uses Claim 5.23. Similarly, in Claim 5.23, either we complete relocation in $O(1)$, or we are required to use Lemma 5.22, which invokes the fat-path-turn-weakening procedure. Its cost of $O(mn)$ dominates all other operations in the relocation.

To move a single endpoint from a boundary to a corner requires at most $\frac{n}{2}$ such relocation steps. Therefore, the time complexity to move each endpoint to its target corner is $O\left(\frac{n}{2}\right) \cdot O(mn) = O(mn^2)$.

Time complexity of Theorem 5.9. Theorem 5.9 reconfigures an arbitrary Hamiltonian path into another arbitrary Hamiltonian path on an $m \times n$ grid graph. The strategy, as described in the proof, is to apply Proposition 5.8 twice to transform both Hamiltonian paths into northern or eastern Hamiltonian paths, which can be viewed as Hamiltonian cycles. We then may apply Theorem 2.1 to reconfigure between them.

From the proof of Proposition 5.8, we know that transforming a path into northern form requires at most two applications of EtB and three applications of EtC. Since EtB has time complexity $O(mn^2)$ and EtC has time complexity $O(mn^2)$, one application of Proposition 5.8 requires $2 \cdot O(mn^2) + 3 \cdot O(mn^2) = O(mn^2)$. So the total time complexity is $2O(mn^2) + O(m^2n^2) = O(m^2n^2)$, accounting for one application of Proposition 5.8 per path, and one application of Theorem 2.1 to reconfigure the resulting northern or eastern paths.

Summary. We have analyzed the computational complexity of the reconfiguration algorithms presented in this dissertation. The fat-path-turn-weakening procedure, which is invoked in the most time-intensive scenarios, requires $O(mn)$ time. The endpoint relocation algorithms (EtB and EtC) each require $O(mn^2)$ time. The complete reconfiguration of Hamiltonian cycles (Theorem 2.1) and Hamiltonian paths (Theorem 5.9) both require $O(m^2n^2)$ time. Space complexity is $O(mn)$ for basic reconfiguration or $O(mn^2)$ if storing the complete move sequence. Notably, the time complexities for Theorems 2.1 and 5.9 are within a factor of n of the move bounds for these theorems.

Algorithm	**Move Bound**	**Time Complexity**
Identifying $P(X, Y)$	N/A	$O(mn)$
MLC	2	$O(m+n)$
1LC	$n^2/2 + m + 3$	$O(mn)$
EtB	$n^2 + mn + 4n$	$O(mn^2)$
EtC	$n^2/2 + mn + 5n/2$	$O(mn^2)$
Theorem 2.1	n^2m	$O(m^2n^2)$
Theorem 5.9	$n^2m + 7n^2 + 10mn + 31n$	$O(m^2n^2)$

Table 1: Move bounds and time complexities for the main algorithms.

6.3 Implications for sampling algorithms

Consider the state space of all Hamiltonian cycles (or paths) on an $m \times n$ rectangular grid graph. Recall from the Introduction chapter that in Markov chain theory, a chain is called *irreducible* if it is possible to reach any state from any other state through a sequence of transitions. For chains on finite state spaces with symmetric transition probabilities (as in our case), irreducibility – together with the easily satisfied condition of aperiodicity – implies *ergodicity*, the property that the chain eventually explores all possible configurations and converges to a unique equilibrium distribution. An *ergodicity class* is defined as a maximal subset of the state space in which any two configurations can reach each other through a sequence of moves. A chain is irreducible if and only if the entire state space forms a single ergodicity class. Otherwise, the state space partitions into multiple ergodicity classes, and a Markov chain started in one class can never reach configurations in another class. This issue is particularly relevant when using Markov chain Monte Carlo to sample self-avoiding walks (including high-density polymer models, which correspond to Hamiltonian paths on lattices). Madras and Slade observe in [21] that for certain Monte Carlo algorithms, the state space partitions into many ergodicity classes, with the largest class being exponentially smaller than the total number of configurations as the chain length increases. Oberdorf et al. [22] used a backbite-based Monte Carlo algorithm for sampling Hamiltonian walks on the two-dimensional square lattice. Their simulations suggested that the underlying Markov chain was ergodic, but they could not provide a proof. Theorem 2.1 proves that any Hamiltonian cycle can be reconfigured into any other using double-switch moves in rectangular grid graphs, thereby establishing irreducibility of the associated Markov chain. Theorem 5.9 extends this result to Hamiltonian paths. Together, these results make progress toward proving the ergodicity of the Markov chain studied in [22] and provide an alternative, rigorous algorithm for generating the Hamiltonian cycles considered there. Analogous Monte Carlo algorithms for three-dimensional cubic lattices [20, 23, 24] similarly assume ergodicity without proof. Proving analogous reconfiguration results in three dimensions remains open.

6.4 Resistant Hamiltonian cycles

In this section we consider Hamiltonian cycles on an $m \times n$ grid graph G with $n \geq m$ for which RtCF requires $\Theta(mn^2)$ moves to complete. We call such cycles *resistant* and denote their class by $\mathcal{H}_{\text{res}}(m, n)$. We give a heuristic characterization of resistant cycles (Conjecture 6.3) and argue that they comprise a vanishingly small fraction of all Hamiltonian cycles.

Recall how the upper bound of mn^2 moves arises for Theorem 2.1: essentially, the factor m comes from the iterations of RtCF over the nested rectangles; one factor n comes from the number of small cookies we may need to collect within a rectangle; and the other factor n comes from the cascade length required for the 1LC algorithm (Proposition 3.10) to collect a small cookie C. Part of this cascade's move cost is due to the sequence of flips needed to collect C, if C is followed by a j-stack of A_0s, and the other part due to a sequence of turn weakenings that may be required. (we refer to j as the *height* of the stack). We make the following assumption.

Assumption 6.1. Let H be a Hamiltonian cycle on an $m \times n$ grid graph G. If the cascade required to collect a small cookie C uses a sequence of j moves for turn weakenings, then C is followed by a stack of A_0s of height at least j, and so it also requires an additional sequence of at least j flips.

We will make two additional assumptions in this section and discuss them briefly at the end of the section.

It follows from Assumption 6.1 that a j-stack of A_0s following a small cookie C contributes at least as many moves to the cost of collecting C as a sequence of turn weakenings that may also be needed. Thus, up to a constant multiple, we may view the height of a j-stack of A_0s as the main contributor to the move cost of collecting C. So, if two Hamiltonian cycles each contain many stacks of A_0-types of large height, and these stacks do not coincide (their locations in G are pairwise disjoint), then RtCF will require $\Theta(mn^2)$ moves to terminate. The total move cost from large cookies alone is only $O(mn)$, so the only way to reach $\Theta(mn^2)$ is to have enough small cookies that each demand many moves. We make this more concrete below.

By 1LC, collecting a cookie followed by a j-stack of A_0s costs at least j moves. Moreover, after the collection, in the next iteration of RtCF the stack of A_0s persists, with one fewer A_0 (see Figure 3.14(d) in Chapter 3). Consider a single small cookie followed by a stack of A_0s with height $h = h(m, n)$. We make the following assumption:

Assumption 6.2. Let H be a Hamiltonian cycle on an $m \times n$ grid graph G. A stack of A_0s with height h contributes at least h small cookies to the total number of cookies that need to be collected by RtCF, one for each A_0 in the stack.

That is, if the first A_0 of the stack contains the small cookie C_i on rectangle R_i, after C_i has been collected via the flip move $X \mapsto C_i$, then on the iteration $i + 1$, there will be at least one small cookie C_{i+1} on R_{i+1} adjacent to X. See Figure 6.1 (the smallest grid we could construct containing a 3-stack). For $j \in \{i, i + 1, \ldots, i + h\}$, the cookie C_j on R_j requires $h - j$ moves to be collected, so the total move contribution to the RtCF move counter is $\frac{(h)(h+1)}{2} = \Theta(h^2)$. Let $c = c(m, n)$ be the number of such stacks of A_0s with height h. Then, to reach the bound of $\Theta(mn^2)$ moves, we need h and c to satisfy $ch^2 = \Theta(mn^2)$. Thus, there are many possibilities for h and c.

We record the above as Conjecture 6.3.

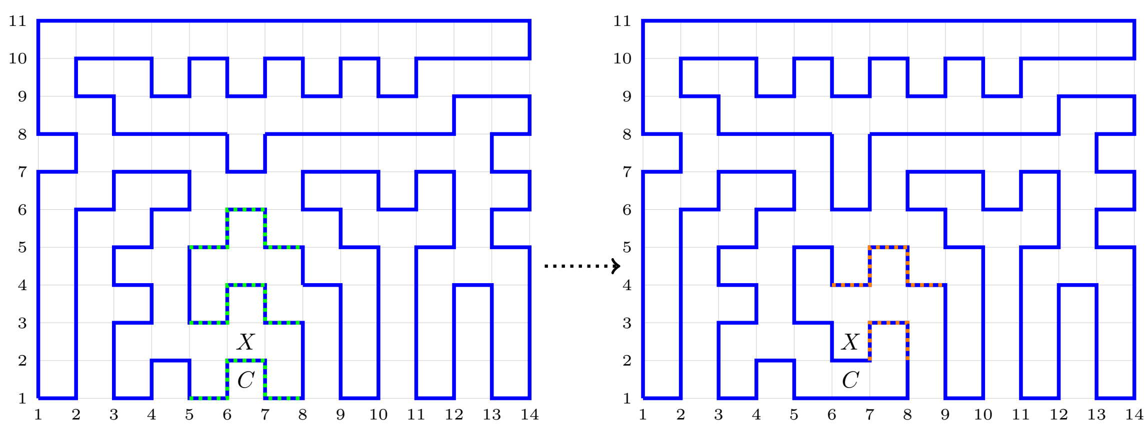

Fig. 6.1. A Hamiltonian cycle containing a 3-stack of A_0s before and after the application of a cascade of flips to collect the small cookie C. Note that after the cascade, $X + (1, 0)$ is a small cookie in the next iteration of RtCF, contained in a 2-stack of A_0s.

Conjecture 6.3. Let G be an $m \times n$ grid graph and let H be a Hamiltonian cycle of G. Then H is resistant if and only if there exist functions $h(m, n)$ and $c(m, n)$ such that H contains c stacks of A_0-types, each with height h, where $ch^2 = \Theta(mn^2)$.

A proof of Conjecture 6.3 would require proving Assumptions 6.1 and 6.2, and a detailed consideration of the converse. Note that the conjecture assumes the existence of Hamiltonian cycles with the required stacks, but we do not prove such cycles exist for all (m, n). As the heights of the A_0 stacks rise, verifying that a candidate graph is indeed a Hamiltonian cycle becomes increasingly difficult. While we have not produced examples with arbitrarily long stacks for sufficiently large grid dimensions, the existence of such resistant cycles seems plausible since the presence of A_0-stacks does not create obvious obstructions to Hamiltonicity. Moreover, by constructing Hamiltonian cycles with j-stacks of A_0s for small values of j, we noticed that accommodating the stacks seemed to require both $m \gtrsim j^2$ and $n \gtrsim j^2$. While we have only observed this for $j \leq 6$ (with the 6-stack requiring a grid of over 1000 vertices), the pattern suggests that both grid dimensions must grow substantially with stack height.

Conjecture 6.4. The n^2m bound on the number of moves is tight. That is, there exists a constant $\alpha \in (0, 1)$ such that for all sufficiently large m and n, there are Hamiltonian cycles H, K of an $m \times n$ grid graph G for which any reconfiguration sequence from H to K requires at least $\alpha n^2 m$ moves.

Conjecture 6.5. Resistant cycles are exponentially rare: $\lim_{m,n\to\infty} \frac{|\mathcal{H}_{\text{res}}(m,n)|}{|\mathcal{H}(m,n)|} = 0$.

We give a plausibility argument in support of Conjecture 6.5, using a heuristic but unproven assumption (Assumption 6.6). Since Hamiltonian cycles are self-avoiding walks, $|\mathcal{H}(m, n)| \leq 3^{mn}$. In Note 6.7 at the end of the section, we show that $(2^{1/6})^{mn} \leq |\mathcal{H}(m, n)|$, so the per-vertex growth rate τ of $|\mathcal{H}(m, n)|$ satisfies $2^{1/6} \leq \tau \leq 3$. The authors in [31] give better bounds for the related case of Hamiltonian paths on square grid graphs.

What is the smallest possible space occupied by stacks of A_0s in a resistant Hamiltonian cycle? Each stack of height h occupies $O(h)$ vertices, so c stacks occupy ch vertices total. From the characterization $ch^2 = \Theta(mn^2)$, we have $ch = \Theta(mn^2/h)$. Since m and n are fixed, to minimize the space occupied we must maximize h. The maximum possible stack height is $h = \Theta(n)$, in which case $c = \Theta(m)$. Thus, stacks must occupy at least $\Theta(mn^2/n) = \Theta(mn)$. That is, there exists a constant $a \in (0, 1)$ such that the stacks occupy at least amn vertices – a fraction a of the total grid. We now make the following assumption about how this constraint affects the count of Hamiltonian cycles.

Assumption 6.6. Constraining amn vertices with A_0-stacks reduces the possible number of Hamiltonian cycles as if these vertices were removed from the grid. That is, for each fixed placement of the stacks there are $\tau^{mn(1-a)}$ possible resistant Hamiltonian cycles.

In how many ways can we position the stacks? We need to choose $c = \Theta(m)$ positions from the mn vertices in the grid. For an upper bound, we use $(mn)^c = (mn)^m$. Thus:

$$|\mathcal{H}_{\text{res}}(m,n)| \lesssim (mn)^m \cdot \tau^{mn(1-a)}.$$

The proportion of resistant cycles to all cycles is then at most

$$\begin{aligned}\frac{|\mathcal{H}_{\text{res}}(m,n)|}{|\mathcal{H}(m,n)|} &\lesssim \frac{(mn)^m \cdot \tau^{mn(1-a)}}{\tau^{mn}} \\ &= (mn)^m \cdot \tau^{-amn} \\ &= (mn)^m \cdot \left(\frac{1}{\tau}\right)^{amn}.\end{aligned}$$

Since $(mn)^m = e^{m\ln(mn)}$ grows sub-exponentially compared to $\tau^{amn} = e^{amn\ln\tau}$, and since $\tau \geq 2^{1/6} > 1$ so $1/\tau < 1$, we have

$$(mn)^m \cdot \left(\frac{1}{\tau}\right)^{amn} \to 0 \quad \text{as } m,n \to \infty.$$

This heuristic suggests that the class of resistant Hamiltonian cycles is vanishingly small relative to all Hamiltonian cycles.

Before showing that $(2^{1/6})^{mn} \leq |\mathcal{H}(m,n)|$, we need a definition. Let G be a graph. A *disjoint cycle-cover* of G is a spanning subgraph that is a disjoint union of cycles. This is also known as a *2-factor* of G.

Note 6.7. To see that the number of Hamiltonian cycles on an $m \times n$ rectangular grid grows exponentially in the area, assume m is even and that n is divisible by 3 (recall that m is the horizontal dimension and n is the vertical dimension, with $m \leq n$). Partition G into $n/3$ horizontal strips, each of size $3 \times m$ (that is, 3 rows and m columns). By Lemma 3.4(d) in [18], each $3 \times m$ strip contains exactly $2^{(m-2)/2}$ Hamiltonian cycles. Fix a disjoint cycle cover $\{H_1, \ldots, H_{n/3}\}$ of G consisting of one Hamiltonian cycle per strip. Observe that there are at least two switchable boxes between every adjacent pair (H_i, H_{i+1}). Switching one of those boxes merges the two cycles into one without affecting the others, and repeating this process merges all strips into a single Hamiltonian cycle of G. See Figure 6.2. There are $(2^{(m-2)/2})^{n/3} = 2^{n(m-2)/6}$ such Hamiltonian cycles, so

$$|\mathcal{H}(m,n)|^{1/(mn)} \geq \left(2^{n(m-2)/6}\right)^{1/(mn)} = 2^{(m-2)/(6m)} \to 2^{1/6} \quad \text{as } m \to \infty.$$

Hence $\tau \geq 2^{1/6} \approx 1.122 > 1$. When $n \not\equiv 0 \pmod 3$, the construction can be extended by including the remaining 1 or 2 rows without affecting the exponential rate.
We conclude with remarks on the two key assumptions in this section. Proofs for Assumptions 6.1 and 6.2 seem within reach of our current tools, without requiring new techniques. We believe Assumption 6.1 would require substantial case analysis related to the structure of sturdy looping fat paths, and that Assumption 6.2 would be more straightforward. Assumption 6.6 on the other hand, seems more difficult and we have no idea how to prove it. We believe the insights about resistant cycles are valuable even with these assumptions remaining unproven.

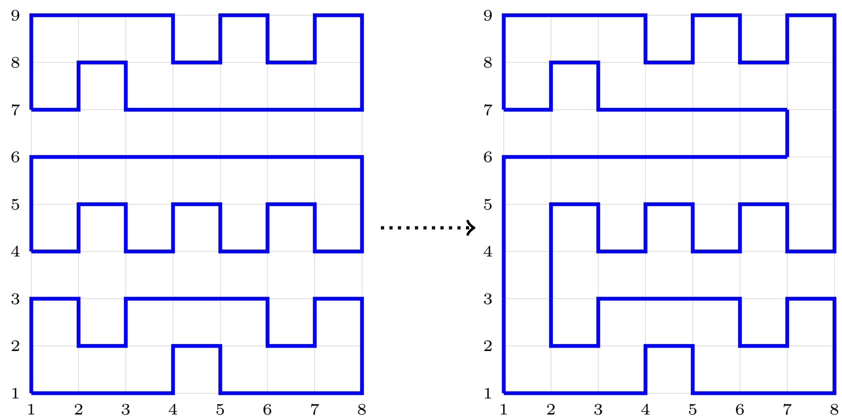

Fig. 6.2. Note that for $j \equiv 0 \mod 3$ the boxes $R(1, j)$ and $R(7, j)$ are always switchable.

6.5 Role of the boundary

Throughout this dissertation, many results are stated for rectangular grid graphs. In the proofs, we do not explicitly invoke this assumption again, but the rectangular structure is used implicitly at nearly every step: in the design of canonical forms and the RtCF algorithm (Chapter 2); in the MLC algorithm and Lemma 3.5 (Chapter 3); in the turn-weakening lemmas near boundaries (Lemmas 4.23, 4.26, 4.29, and 4.30 in Chapter 4); in the description of some impossible configurations (Lemmas 5.11 and 5.12 in Chapter 5); and in the EtB and EtC algorithms (Chapter 5). In this section we consider what happens when we relax the rectangular boundary condition or embed rectangular grids on a cylinder or torus. We will see that reconfiguration can fail: Hamiltonian cycles may become frozen, or the space of cycles may split into disconnected components.

Example 6.8. Hamiltonian cycles of a simply connected polyomino that are frozen under the double-switch move. See Figure 6.3 (a).

Example 6.9. Hamiltonian paths of a simply connected polyomino that belong to distinct ergodicity classes under the backbite move. See Figure 6.3 (b).

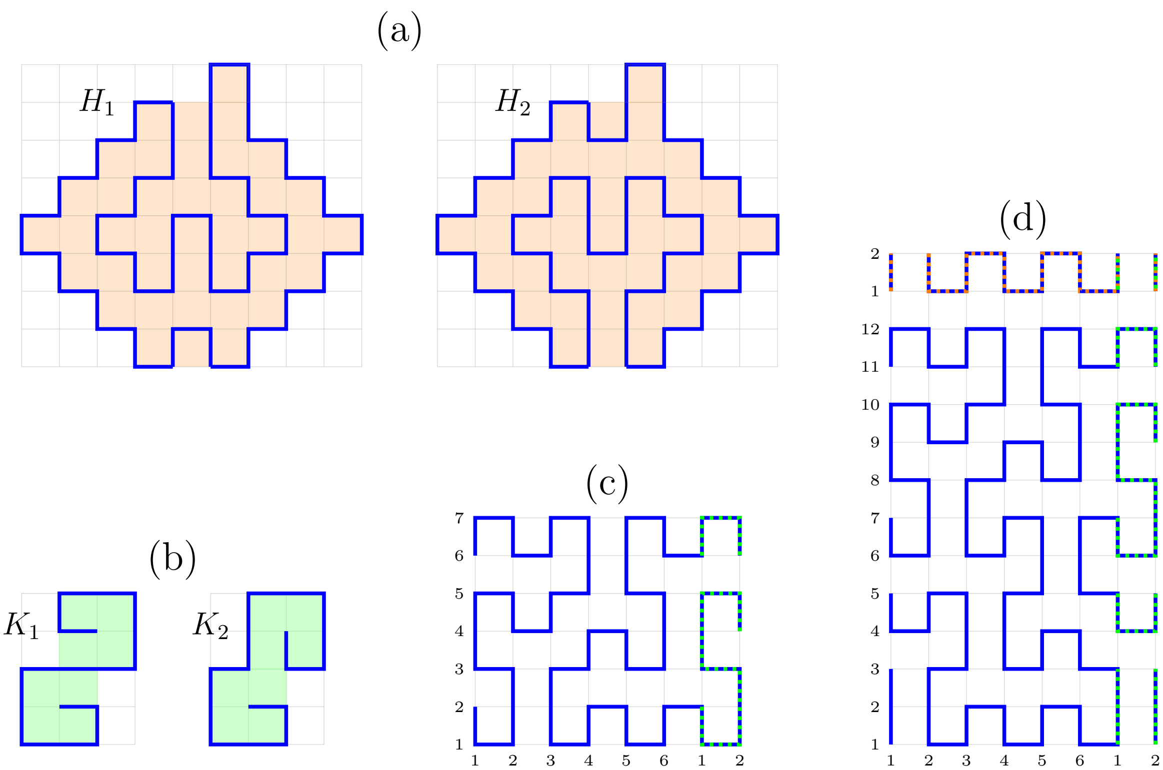

Fig. 6.3. (a) Two distinct Hamiltonian cycles H_1 and H_2 of the simply connected polyomino G (in orange) with non-rectangular boundary. (b) Two distinct Hamiltonian paths K_1 and K_2 of the simply connected polyomino G' (in green) with non-rectangular boundary. (c) A Hamiltonian cycle on a 6×7 grid graph embedded on a cylinder; horizontal seam dotted green. (d) A Hamiltonian cycle on a 6×12 grid graph embedded on a torus; horizontal seam dotted green, vertical seam dotted orange.

Example 6.10. A Hamiltonian cycle of a 6×7 grid graph embedded on a cylinder that is frozen under the double-switch move. This construction appears to generalize $k \times (k+1)$ grids, for each even $k \geq 6$. See Figure 6.3 (c).

Example 6.11. A Hamiltonian cycle of a 6×12 grid graph embedded on a torus that is frozen under the double-switch move. This construction appears to generalize to $k \times 2k$ grids for each even $k \geq 6$. See Figure 6.3 (d).

The counterexamples illustrated above raise the following question: can we reconfigure any Hamiltonian cycle on a polyomino to any other by expanding our move set? Can we do the same for rectangular grid graphs embedded on a cylinder or torus? We conclude this section with two conjectures. We first need a definition. Let $H = H_0$ be a Hamiltonian cycle of a polyomino G. A *quadruple-switch move* on H is defined as follows: for $i \in \{1, \ldots 4\}$, perform a switch σ_i on a switchable box of H_{i-1} to obtain the subgraph H_i. The move is *valid* if H_4 is a Hamiltonian cycle of G.

Conjecture 6.12. Let G be a polyomino and let H and K be two distinct Hamiltonian cycles of G. Then there exists a sequence of valid quadruple-switch moves that reconfigures H into K.

Conjecture 6.13. Let G be an $m \times n$ grid graph embedded on a cylinder or torus, and let H and K be two distinct Hamiltonian cycles of G. Then there exists a sequence of valid quadruple-switch moves that reconfigures H into K.

Why might a quadruple-switch move be enough for these more general reconfiguration questions? It seems that these moves might be sufficient to break down the stubborn structures arising from combinations of A-types and stairs subgraphs (recall definition at the beginning of Chapter 4), and compensate for the loss of control incurred by relaxing the conditions that the boundary be rectangular and that the graph be embedded in the plane.

6.6 Extensions and future work

In this section we give an overview of possible extensions of our results and state open problems for future research. These extensions include restricting the move set, reconfiguration of cycle covers (of which Hamiltonian cycles are a special case), reconfiguration on higher dimensional square lattices, and reconfiguration on the triangular lattice. We extend the usual $m \times n$ rectangular grid graph to higher dimensions as follows. For integers $n_1, \ldots, n_d \geq 2$, a *d-dimensional orthotope grid graph (or d-orthotope)* is the graph G with vertex set

$$V(G) = \{(x_1, \ldots, x_d) \in \mathbb{Z}^d : 0 \leq x_i < n_i \text{ for all } i\},$$

and with an edge between two vertices if they differ by 1 in exactly one coordinate and are equal in all others.

Nishat and Whitesides in [17] ask whether it is possible to reconfigure Hamiltonian cycles in rectangular grid graphs using just flip and transpose moves (recall the definitions from the Introduction chapter). Nishat in [18] also asks if given two disjoint cycle covers in a grid graph, we can reconfigure them into one another. We restate these questions here as conjectures.

Conjecture 6.14. Let G be an $m \times n$ grid graph and let H and K be two distinct Hamiltonian cycles of G. Then there exists a sequence of flips and transposes that reconfigures H into K.

We have frequently encountered the scenario consisting of a leaf L followed by an A_1-type with switchable middle-box W with looping H-path $P(X, Y)$. We have shown that we can always find a cascade ending with the move $Z \mapsto W$, where Z is a switchable box of $P(X, Y)$. We believe that Conjecture 6.14 hinges on showing that this last "long-range" double-switch move $Z \mapsto W$ can be replaced by a cascade of flips and transposes. The tools in Chapter 4 seem useful to address this conjecture, however we believe additional results would be needed regarding the structure of fat paths and turns, including a proof of Assumption 6.1.

Conjecture 6.15. Let G be an $m \times n$ grid graph and let H and K be two distinct disjoint cycle covers of G. Then there exists a sequence of double-switch moves that reconfigures H into K. Moreover, if G is a general polyomino, then quadruple-switch moves are enough for reconfiguration.

The rationale for the polyomino case is similar to that given in support of Conjectures 6.12 and 6.13: more flexible moves may compensate for the loss of the rectangular structure.

Our results are specific to square lattices. The reconfiguration problem remains open for triangular and hexagonal lattices. The backbite move may be particularly effective for reconfiguring Hamiltonian paths in the triangular lattice, since we have more edges to choose from: an interior end-vertex in a triangular lattice has five possible edges (vs three in the square lattice), available for initiating a backbite move. Moreover, for subgraphs with triangular boundaries, there are two possible edges toward each boundary side, suggesting that relocating an end-vertex to the boundary may be more straightforward. That is, an analogue of the EtB algorithm may be easier to prove for triangular-lattice graphs. This additional degree of freedom suggests that backbite-only reconfiguration may be more tractable in the triangular lattice.

Extending our results to higher dimensions presents new challenges. The tools developed in this dissertation relied on the 2-dimensional planar structure, in particular, the consequences of Jordan's curve theorem for polygons (Theorem 1.1.4). In three dimensions and beyond, there is no good analogue of this theorem, and much of our toolbox becomes unusable.

We have been researching Hamiltonian cycle reconfiguration in 3-dimensional grids with the help of GridLab, a tool with interactive visualization features that can be used to generate Hamiltonian paths and cycles in 3 dimensions, available online at [32]. This research suggests that reconfiguration may still be possible with appropriate move sets, leading us to the following conjecture.

Conjecture 6.16. Let G be a 3-dimensional cuboid grid graph (a 3-orthotope) and let H and K be two distinct Hamiltonian cycles of G. Then there exists a sequence of double-switch and switch moves that reconfigures H into K.

For Hamiltonian paths in higher dimensions, the backbite move may suffice on its own:

Conjecture 6.17. Let G be a d-dimensional orthotope grid graph (for $d \geq 3$) and let H and K be two distinct Hamiltonian paths of G. Then there exists a sequence of backbite moves that reconfigures H into K.

Conjecture 6.16 may hinge on resolving the following subproblem: Let $H = P(v_1, v_s)$, $P(v_s, v_t)$, $P(v_t, v_r)$ be a Hamiltonian path on a d-orthotope, where $P(v_s, v_t)$ is already in a "nice" form. Can we relocate an end-vertex, say v_1, to any vertex of $H \setminus P(v_s, v_t)$ that v_1 can reach using only backbite moves? We are currently researching this question for the 3-dimensional case, with a move set that includes single-switches and double-switches.

Index

References

[1] M. Celaya, K. Choo, G. MacGillivray, and K. Seyffarth. "Reconfiguring k-colourings of complete bipartite graphs". In: *Kyungpook Mathematical Journal* 56.3 (2016), pp. 647–655.

[2] P. Bonsma and L. Cereceda. "Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances". In: *Theoretical Computer Science* 410.50 (2009), pp. 5215–5226.

[3] M. Bonamy, M. Johnson, I. Lignos, V. Patel, and D. Paulusma. "Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs". In: *Journal of Combinatorial Optimization* 27.1 (2014), pp. 132–143.

[4] C. Feghali and J. Fiala. "Reconfiguration graph for vertex colourings of weakly chordal graphs". In: *Discrete Mathematics* 343.3 (2020), p. 111733.

[5] M. Johnson, D. Kratsch, S. Kratsch, V. Patel, and D. Paulusma. "Finding shortest paths between graph colourings". In: *Algorithmica* 75.2 (2016), pp. 295–321.

[6] R. Haas and K. Seyffarth. "The k-dominating graph". In: *Graphs and Combinatorics* 30.3 (2014), pp. 609–617.

[7] A. Suzuki, A. E. Mouawad, and N. Nishimura. "Reconfiguration of dominating sets". In: *Journal of Combinatorial Optimization* 32.4 (2016), pp. 1182–1195.

[8] E. D. Demaine et al. "Linear-time algorithm for sliding tokens on trees". In: *Theoretical Computer Science* 600 (2015), pp. 132–142.

[9] M. Bonamy and N. Bousquet. "Recoloring graphs via tree decompositions". In: *European Journal of Combinatorics* 69 (2018), pp. 200–213.

[10] J. van den Heuvel. "The complexity of change." In: *Surveys in Combinatorics* 409.2013 (2013), pp. 127–160.

[11] N. Nishimura. "Introduction to reconfiguration". In: *Algorithms* 11.4 (2018), p. 52.

[12] D. Barnette. "Conjecture 5". In: *Recent Progress in Combinatorics*. Ed. by W. T. Tutte. New York: Academic Press, 1969, p. 343.

[13] A. Takaoka. "Complexity of Hamiltonian cycle reconfiguration". In: *Algorithms* 11.9 (2018), p. 140.

[14] A. Itai, C. H. Papadimitriou, and J. L. Szwarcfiter. "Hamilton paths in grid graphs". In: *SIAM Journal on Computing* 11.4 (1982), pp. 676–686.

[15] C. Umans and W. Lenhart. "Hamiltonian cycles in solid grid graphs". In: *Proceedings 38th Annual Symposium on Foundations of Computer Science*. IEEE. 1997, pp. 496–505.

[16] S. D. Chen, H. Shen, and R. Topor. "An efficient algorithm for constructing Hamiltonian paths in meshes". In: *Parallel Computing* 28.9 (2002), pp. 1293–1305.

[17] R. I. Nishat and S. Whitesides. "Bend complexity and Hamiltonian cycles in grid graphs". In: *International Computing and Combinatorics Conference*. Springer. 2017, pp. 445–456.

[18] R. I. Nishat. "Reconfiguration of Hamiltonian cycles and paths in grid graphs". PhD thesis. University of Victoria, 2020.

[19] Albi Kazazi. *Reconfiguration of Hamiltonian cycles in grid graphs*. YouTube video. 2024. URL: https://www.youtube.com/watch?v=zPUfEAv_jjE&ab_channel=albik.

[20] M. L. Mansfield. "Monte Carlo studies of polymer chain dimensions in the melt". In: *The Journal of Chemical Physics* 77.3 (1982), pp. 1554–1559.

[21] N. Madras and G. Slade. *The Self-Avoiding Walk*. Springer Science & Business Media, 2013.

[22] R. Oberdorf, A. Ferguson, J. L. Jacobsen, and J. Kondev. "Secondary structures in long compact polymers". In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 74.5 (2006), p. 051801.

[23] J. L. Jacobsen. "Unbiased sampling of globular lattice proteins in three dimensions". In: *Physical Review Letters* 100.11 (2008), p. 118102.

[24] J. Deutsch. "Long range moves for high density polymer simulations". In: *The Journal of Chemical Physics* 106.21 (1997), pp. 8849–8854.

[25] R. I. Nishat, V. Srinivasan, and S. Whitesides. "The hamiltonian path graph is connected for simple s, t paths in rectangular grid graphs". In: *Journal of Combinatorial Optimization* 48.4 (2024), p. 31.

[26] R. I. Nishat, V. Srinivasan, and S. Whitesides. "1-Complex s, t Hamiltonian Paths: Structure and Reconfiguration in Rectangular Grids". In: *Journal of Graph Algorithms and Applications* 27.4 (2023), pp. 281–327.

[27] R. I. Nishat and S. Whitesides. "Reconfiguration of Hamiltonian Cycles in Rectangular Grid Graphs". In: *International Journal of Foundations of Computer Science* 34.07 (2023), pp. 773–793.

[28] R. I. Nishat and S. Whitesides. "Reconfiguring Hamiltonian cycles in L-shaped grid graphs". In: *International Workshop on Graph-Theoretic Concepts in Computer Science*. Springer. 2019, pp. 325–337.

[29] B. Mohar and C. Thomassen. *Graphs on Surfaces.* Johns Hopkins series in the mathematical sciences. Johns Hopkins University Press, 2001, pp. I–XI, 1–291.

[30] M. Axenovich. *Graph Theory Notes*. https://www.math.kit.edu/iag6/lehre/graphtheory2019w/media/main.pdf. Accessed: 7 August 2025. 2020.

[31] M. Bousquet-Mélou, A. J. Guttmann, and I. Jensen. "Self-avoiding walks crossing a square". In: *Journal of Physics A: Mathematical and General* 38.42 (2005), p. 9159.

[32] A. Kazazi. *3D GridLab*. 2025. URL: https://graphsofwrath.com/.