

An update on the coin-moving game on the square grid

Florian GALLIOT^{1,3}, Sylvain GRAVIER^{1,3}, and Isabelle SIVIGNON^{2,3}

¹Univ. Grenoble Alpes, CNRS, Institut Fourier, 38000 Grenoble, France

²Univ. Grenoble Alpes, CNRS, Grenoble INP, GIPSA-lab, 38000 Grenoble,
France

³Univ. Grenoble Alpes, Maths à Modeler, 38000 Grenoble, France

Abstract

This paper extends the work started in 2002 by Demaine, Demaine and Verill (DDV) on coin-moving puzzles. These puzzles have a long history in the recreational literature, but were first systematically analyzed by DDV, who gave a full characterization of the solvable puzzles on the triangular grid and a partial characterization of the solvable puzzles on the square grid. This article specifically extends the study of the game on the square grid. Notably, DDV's result on puzzles with two "extra coins" is shown to be overly broad: this paper provides counterexamples as well as a revised version of this theorem. A new method for solving puzzles with two extra coins is then presented, which covers some cases where the aforementioned theorem does not apply.

Puzzles with just one extra coin seem even more complicated, and are only touched upon by DDV. This paper delves deeper, studying a class of such puzzles that may be considered equivalent to a game of "poking" coins. Within this class, some cases are considered that are amenable to analysis.

Introduction

We study a one-player game that is played on an undirected simple graph $G = (V, E)$, referred to as the *game graph* or *board*. Throughout the game, there will be coins sitting on some of the vertices of G (at most one per vertex). The coins are indistinguishable and define a *configuration*, i.e. a finite subset $C \subseteq V$ where we see each element of C as a coin sitting on the corresponding vertex. A legal *move* consists in moving a single coin to a free vertex so that, after the move, that coin has at least two other coins adjacent to it. This is called the *2-adjacency* restriction. Given two configurations A and B , we want to know whether the *puzzle* $A \xrightarrow{?} B$ is solvable: starting from A , is it possible to reach B using only legal moves? In the positive case, we would like an explicit winning sequence of moves. This game falls into the category of reconfiguration problems on graphs.

Instances of this game, or rather a variation with tightly packed coins that can only be slid in the plane without collision, appear in the literature as early as the 1950s in [Lan51] and [Lan53].

Figure 1 features a couple of classic puzzles on the triangular grid as well as a more rare puzzle on the square grid. Such examples also appear in [Gar75] and [BCG82] among others, but it is not until 2002 that general puzzles with these rules have been studied, in [DDV02] which serves as foundation for the present paper. The authors give a full characterization for solvable puzzles on the triangular grid (up to a minor omission which is easily settled: see [Gal19]). Furthermore, they address a large family of puzzles on the square grid, providing polynomial time solving algorithms. Another coin-moving game on the square grid, with different rules but similar methods and also polynomial time algorithms in many cases, is studied in [DP06].




Figure 1. The game graph is the triangular grid in the left and middle puzzles, and the square grid in the right puzzle. The left (resp. middle, resp. right) puzzle is solvable in 2 (resp. 3, resp. 4) moves.

In this paper, we exclusively study the case where G is the square grid (each vertex has four neighbors: left, right, top, bottom). A first natural question is : what can we reach starting from a configuration A ? A central observation is that, during the moves, all coins remain inside of the *span* of A , which is a finite set (union of rectangles) obtained from A by including all vertices that have at least two neighbors in A and iterating this process until no more vertex can be included. Any configuration B that we wish to obtain from A must therefore satisfy $B \subseteq \text{span}(A)$, which implies $\text{span}(B) \subseteq \text{span}(A)$. In the study that is made in [DDV02], a key information is the number of coins that can be removed from A while maintaining a span containing that of B : we call them *extra coins in A relatively to B* . The more extra coins at our disposal, the more flexibility with respect to the span constraint, hence, the easier a puzzle. Section 1 goes over fundamental observations and notions introduced in [DDV02], including that of span and extra coins, with the addition of some basic results of our own that will be used in this paper. Section 2 then recalls the method that is employed to solve puzzles in [DDV02], which we will also use ourselves.

In [DDV02], the authors claim that, up to an additional condition on B , two extra coins in A relatively to B are enough to solve any puzzle. Unfortunately, the proof only works if A and B have the same span. In Section 3, we extend this result to the case where each connected component of $\text{span}(A)$ contains at most one connected component of $\text{span}(B)$.

If this latter condition is not satisfied however, then Section 4 provides, for any n , an example of an unsolvable puzzle where the span of A is an $n \times n$ square and A has about $\frac{n}{2}$ extra coins relatively to B (with B satisfying the additional condition from [DDV02], and some more). It is worth noticing that n , in that case, is also the minimum cardinality of a configuration having an $n \times n$ span. We then give a new sufficient condition for a puzzle to be solvable, which shows in particular that this number $\frac{n}{2}$ is somehow tight.

As mentioned in [DDV02], the case of a single extra coin is even more complicated. The first non-trivial results for this case are presented in Section 5, under the restriction that A and B have the same span and contain exactly one coin (the extra coin) more than the minimum possible for their span. We show that our game then reduces to a *poking game* where, instead of the 2-adjacency rule, a coin can be slid onto a neighboring vertex under certain conditions. We give necessary and sufficient conditions for poking puzzles involving what is called a chain of coins.

1 Preliminaries

1.1 Notations and first observations

We use the (self-)dual grid for graphical representations of the game: each vertex, or *position*, is seen as a square and coins are placed at the center of squares.




Figure 2. A configuration C . From C , a possible move would be $c \mapsto p_1$ or $c \mapsto p_2$ for example. However, c cannot be moved to p_3 , because that position only has one neighboring coin other than c .

Notation 1.1. We introduce the following notations:

- Moving a coin c to a valid destination p is denoted by $c \mapsto p$. Note that a coin c is simply an element of V representing its location (since the coins are indistinguishable, a coin c is nothing but an occupied position p). Typically, the notation c is used to designate an occupied position, and the notation p is used for a position that is either unoccupied or not necessarily occupied.
- If there is a single move $c \mapsto p$ from A to B , we write $A \xrightarrow{c \mapsto p} B$ or simply $A \mapsto B$.
- If there exists a sequence of moves from A to B , we write $A \rightarrow B$, otherwise we write $A \not\rightarrow B$. By convention, an empty sequence of moves is allowed so that $A \rightarrow A$.
- If $A \rightarrow B$ and $B \rightarrow A$ then we may write $A \leftrightarrow B$.

Proposition 1.2. *Let $A \neq B$ be configurations. If $A \rightarrow B$ then there exists a coin in B that has at least two neighboring coins in B .*

Proof. The last moved coin has at least two neighboring coins by the 2-adjacency rule. □

Proposition 1.3. *Let $A \neq B$ be configurations. If there exists a coin $b \in B$ such that $B \setminus \{b\}$ consists of all isolated coins, then $A \rightarrow B$ if and only if $A \mapsto B$.*

Proof. Suppose that $A = A_0 \xrightarrow{c_1 \mapsto p_1} A_1 \xrightarrow{c_2 \mapsto p_2} \dots \xrightarrow{c_T \mapsto p_T} A_T = B$ and that $B \setminus \{b\}$ consists of all isolated coins. In particular all coins in $B \setminus \{b\}$ have at most one neighboring coin in B , so $p_T = b$ by the 2-adjacency rule. Since $A_{T-1} \setminus \{c_T\} = B \setminus \{p_T\}$, this means $A_{T-1} \setminus \{c_T\}$ consists of all isolated coins. In particular all coins in $A_{T-1} \setminus \{c_T\}$ have at most one neighboring coin in A_{T-1} , so $p_{T-1} = c_T$ by the 2-adjacency rule. Continuing so, we get $p_{T-1} = c_T, p_{T-2} = c_{T-1}, \dots, p_1 = c_2$. In conclusion, we have moved the same coin each time, so we could have moved it just once instead and got $A \xrightarrow{c_1 \mapsto p_T} B$. □

1.2 Picking up and dropping coins

Definition 1.4. [DDV02] Consider the following actions:

- *Pick up* a coin: remove a coin from the board, without any restriction on its position.
- *Drop* a coin: put a previously picked up coin back on the board, with the 2-adjacency restriction.

Picked up coins that have not yet been dropped may be referred to as *coins in hand*.

It is shown in [DDV02] that the game is unchanged if the player is allowed to pick up and drop coins additionally to moving them:

Proposition 1.5. [DDV02] *A puzzle is solvable (by moving coins) if and only if it is solvable by moving, picking up and dropping coins.*

Notation 1.6. In this variation, the state of the game at any given moment is described by the configuration A on board and the number k of coins in hand: we denote this information by A^{+k} . For example, the notation $A^{+k} \rightarrow B^{+k'}$ means that, from the configuration A with k coins in hand at the start, it is possible (via moving, picking up and dropping coins) to reach the configuration B with $k' = k + |A| - |B|$ coins in hand at the end.

1.3 Span of a configuration

As mentioned in the introduction, the notion of span is central to the study of the game on the square grid.

Notation 1.7. Let C be a configuration, we denote by $\text{Adj}(C) \subseteq V \setminus C$ the set of all positions outside C that have at least two neighbors in C .

Definition 1.8. [DDV02] The *span* of a configuration C , denoted by $\text{span}(C)$, is the limit of the non-decreasing sequence of configurations $(C_i)_{i \geq 0}$ defined recursively by $C_0 = C$ and $C_{i+1} = C_i \cup \text{Adj}(C_i)$. In other words, $\text{span}(C)$ is the set of all positions that could be reached from C if we had unlimited coins to add to the board at successive positions satisfying the 2-adjacency rule.

Proposition 1.9. [DDV02] *The span never increases during moves: if $A \rightarrow B$ then $\text{span}(A) \supseteq \text{span}(B)$.*

Definition 1.10. Identifying the square grid as $\mathbb{Z}^2$, an $m \times n$ *rectangle* R is a set of positions of the form $I \times J$ where I and J are intervals of cardinality m and n respectively, so that each *row* of R contains m positions and each *column* of R contains n positions. We say R is *even* (resp. *odd*) if its *half-perimeter* $m + n$ is even (resp. odd).

Proposition 1.11. [DDV02] *The span of any configuration C is a union of rectangles at distance at least 3 from each other (these rectangles are called the components of $\text{span}(C)$).*

An example can be seen in Figure 3. We add the following elementary property of the span which will be useful later.




Figure 3. The span of a configuration. In this example, there are four components.

Proposition 1.12. *Let C be a configuration and let R be a component of $\text{span}(C)$. Then C contains at least one coin in each of the following: the top row of R , the bottom row of R , the leftmost column of R , the rightmost column of R , any union of two consecutive rows in R , any union of two consecutive columns in R .*

Proof. Using the symmetries, we only address the case of the top row and the union of two consecutive rows. Let $C = C_0, C_1, \dots, C_s = \text{span}(C)$ as in Definition 1.8 (indeed, the sequence $(C_i)_{i \geq 0}$ is stationary by finiteness of the span).

- Suppose for a contradiction that the top row R' contains no coin in C i.e. $C \cap R' = \emptyset$. Let $i \in \{0, \dots, s\}$ be smallest such that $C_i \cap R' \neq \emptyset$: we have $i \geq 1$ and $C_{i-1} \cap R' = \emptyset$. Obviously, any position in R' has at most one neighbor in $R \setminus R'$, so $\text{Adj}(C_{i-1}) \cap R' = \emptyset$. Since $C_i = C_{i-1} \cup \text{Adj}(C_{i-1})$, we get $C_i \cap R' = \emptyset$ which is a contradiction.
- The previous proof still works if we replace R' by any union of two consecutive rows, since the key argument that any position in R' has at most one neighbor in $R \setminus R'$ still holds. $\square$

1.4 Minimal/minimum configurations

Definition 1.13. Let C be a configuration.

- We say C is *minimal*, if the removal of any coin in C decreases the span.
- We say C is *minimum* if there is no configuration C' with same span as C such that $|C'| < |C|$ (in particular, C is then minimal).

Proposition 1.14. [DDV02] *Any move played from a minimal configuration decreases the span.*

As noticed in [DDV02], the cardinality of minimum configurations (as well as a lot of information on their structure) is well known thanks to the following classic problem from folklore. In a rectangular parcel R consisting of small squares arranged in a grid, some squares are initially invaded by weeds. Time passes, and at each time step, any square that is adjacent to at least two weeds-covered squares gets invaded in turn. How many squares need to be covered initially for the entire parcel to be invaded in the end? Since the rule for the propagation of the weeds is

exactly the same as for the construction of the span, the answer coincides with the cardinality of a minimum configuration with span R . This problem was first published in [Kva86] for a 10×10 parcel. An elegant solution is obtained via an invariant which is the perimeter of the invaded area:

Proposition 1.15. *Let R be an $m \times n$ rectangle and let M be a minimum configuration with span R . We have $|M| = \left\lceil \frac{m+n}{2} \right\rceil$. Moreover:*

- *If R is even, then all coins in M are isolated.*
- *If R is odd, then all coins in M are isolated except possibly for a single pair of adjacent coins.*

Proof. Define $\text{Perim}(C)$ as the perimeter of the union of unit squares on board corresponding to all occupied positions in a configuration C . For any configuration C and any $p \notin C$, we have $\text{Perim}(C \cup \{p\}) = \text{Perim}(C) + 4 - 2k$ where k is the number of coins in C that are a neighbor of p : indeed, in terms of the perimeter, each side of p removes 1 if it neighbors a coin in C or adds 1 if it does not. Therefore, any configuration C satisfies $\text{Perim}(C) = 4|C| - 2l$ where l is the number of pairs of adjacent coins in C . Moreover, the perimeter cannot increase during the construction of the span, since each added coin has at least two neighbors amongst the coins that are already present: this implies that any configuration C with span R satisfies $2(m+n) = \text{Perim}(\text{span}(C)) \leq \text{Perim}(C) \leq 4|C|$ hence $|C| \geq \left\lceil \frac{m+n}{2} \right\rceil$. Conversely, it is easy to find configurations with span R that have exactly $\left\lceil \frac{m+n}{2} \right\rceil$ coins (think of a diagonal of coins if $m = n$, or more generally an ‘L’ shape as in Definition 2.4), so the cardinality of minimum configurations with span R is exactly $\left\lceil \frac{m+n}{2} \right\rceil$. Finally, let M be a minimum configuration with span R and let l be the number of pairs of adjacent coins in M , we have $2(m+n) \leq \text{Perim}(M) = 4 \left\lceil \frac{m+n}{2} \right\rceil - 2l$, so in conclusion:

- If R is even, then $2(m+n) \leq 4 \frac{m+n}{2} - 2l$ hence $l = 0$ i.e. all coins in M are isolated.
- If R is odd, then $2(m+n) \leq 4 \frac{m+n+1}{2} - 2l$ hence $l \in \{0, 1\}$ i.e. all coins in M are isolated except possibly for a single pair of adjacent coins. $\square$

1.5 Extra coins and redundant coins

We recall the notion of extra coins defined in the introduction:

Definition 1.16. Let A be a configuration and $k \in \mathbb{N}$.

- A *set of extra coins in A* is a subset $A' \subset A$ such that $\text{span}(A \setminus A') = \text{span}(A)$. We say A *has k extra coins* if it contains a set of k extra coins. For example, A has one extra coin if and only if A is not minimal.
- Let B be a configuration. A *set of extra coins in A relatively to B* is a subset $A' \subset A$ such that $\text{span}(A \setminus A') \supseteq \text{span}(B)$. We say A *has k extra coins relatively to B* if it contains a set of k extra coins relatively to B . If A and B have same span, then this definition coincides with the previous one.

For example, in the puzzle $A \xrightarrow{?} B$ from Figure 4, $\{a_1, a_2\}$ is a set of extra coins in A (in particular, it is also a set of extra coins in A relatively to B).

Proposition 1.17. *Let A and B be distinct configurations. If $A \rightarrow B$ then A has one extra*




Figure 4. An example of a puzzle with two extra coins a_1, a_2 (Definition 1.16) and two redundant coins b_1, b_2 (Definition 1.18).

coin relatively to B .

Proof. Suppose A does not have one extra coin relatively to B , and consider the first move $c \mapsto p$ made from A . This move can be decomposed as follows: first we remove c from the board, then we put it back at p . After removing c from the board, the span is $\text{span}(A \setminus \{c\})$. When we put the coin back, the span stays the same because of the 2-adjacency rule. Therefore the span after the first move is $\text{span}(A \setminus \{c\})$, which does not contain $\text{span}(B)$ because $\{c\}$ is not a set of extra coins relatively to B . By Proposition 1.9, the span cannot increase during the moves, hence $A \not\rightarrow B$. $\square$

We now introduce the notion of redundant coins, which is also present in [DDV02] (though unnamed). They are, to the target configuration B , the relevant analogue of what extra coins are to the starting configuration A .

Definition 1.18. Let B be a configuration. A *set of redundant coins in B* is a subset $B' \subset B$ of the form $B' = \{b_1, \dots, b_k\}$ where, for all $i \in \{1, \dots, k\}$, b_i has at least two neighbors in $B \setminus \{b_1, \dots, b_{i-1}\}$. We say B *has k redundant coins* if it contains a set of k redundant coins.

For example, in the puzzle $A \xrightarrow{?} B$ from Figure 4, $\{b_1, b_2\}$ is a set of redundant coins in B . Note that redundant coins are extra coins, but the converse is not true in general (think of a_1 and a_2 in Figure 4).

Proposition 1.19. *Let A and B be distinct configurations. If $A \rightarrow B$ then B has one redundant coin.*

Proof. This is exactly Proposition 1.2. $\square$

Proposition 1.20. *Let B be a configuration and let $B' = \{b_1, \dots, b_k\}$ be a set of redundant coins in B , ordered as in the definition. Then $B \leftrightarrow (B \setminus B')^{+k}$. In particular, for any configuration A , we have $A \rightarrow B$ if and only if $A \rightarrow (B \setminus B')^{+k}$.*

Proof. We perform $B \rightarrow (B \setminus B')^{+k}$ by picking up the coins in B' , and $(B \setminus B')^{+k} \rightarrow B$ by dropping the coins in hand at $b_k, b_{k-1}, \dots, b_1$ successively (which respects the 2-adjacency rule since b_i has at least two neighbors in $B \setminus \{b_1, \dots, b_{i-1}\}$). $\square$

Therefore, more redundant coins can only make a puzzle easier. One can think of extra coins in A (or extra coins in A relatively to B) as the first coins that we move, and of redundant coins in B as the last coins that we place.

2 Canonical configurations

In [DDV02], the authors define a reference minimum configuration for a given span, called the canonical configuration. Their method to solve a puzzle $A \xrightarrow{?} B$, if $\text{span}(A) = \text{span}(B)$ for instance, consists in going from A to B by routing through their common canonical configuration. We will reuse this principle in Sections 3, 4 and 5. This section provides a summary of the definitions and results about canonical configurations that are presented in [DDV02].

2.1 Definitions

Notation 2.1. We denote by dist the usual distance in the square grid.

Definition 2.2. [DDV02] A *chain* between some coins c and c' is the configuration denoted by $[c_1, \dots, c_N]$ which is formed by a sequence of coins ($c = c_1, c_2, \dots, c_N = c'$) such that $\text{dist}(c_i, c_{i+1}) \in \{1, 2\}$ for all $i \in \{1, \dots, N-1\}$.

Proposition 2.3. [DDV02] *The span of a chain coincides with its smallest enclosing rectangle.*

Definition 2.4. [DDV02] An ‘L’ of size $m \times n$ is a minimum chain L between two opposite corners of an $m \times n$ rectangle R and hugging two consecutive sides of R . We say L is *even* (resp. *odd*) if R is even (resp. odd) i.e. if $m+n$ is even (resp. odd). See Figure 5 (right).

In accordance with Proposition 1.15, an $m \times n$ ‘L’ has cardinality $\left\lceil \frac{m+n}{2} \right\rceil$, and consecutive coins in an even ‘L’ are at distance exactly 2 whereas consecutive coins in an odd ‘L’ are at distance exactly 2 except for a single pair of adjacent coins. An even ‘L’ is entirely defined by its span and orientation, whereas for an odd ‘L’ we also need the localization of the two adjacent coins. By Proposition 2.3, the span of an ‘L’ is its smallest enclosing rectangle i.e. R in Definition 2.4.

Definition 2.5. [DDV02] Let R be an $m \times n$ rectangle: the *canonical ‘L’ with span R* is the ‘L’ with span R that is oriented like the letter L, with the additional property if R is odd that the two adjacent coins are in the top-left corner (if n is even) or bottom-right corner (if m is even). Let C be a configuration with span $\bigcup_{i=1}^s R_i$ where $R_1, \dots, R_s$ are rectangles at distance at least 3 from each other, as per Proposition 1.11: the *canonical configuration associated to C* is the configuration denoted by L_C with same span as C such that, for all $1 \leq i \leq s$, $L_C \cap R_i$ is the canonical ‘L’ with span R_i . See Figure 5.

2.2 Transformations of ‘L’s

Transformations of ‘L’s are the main subroutines used in the solving algorithms from [DDV02]. We have the following:

Proposition 2.6. [DDV02] *Let L_1 and L_2 be two ‘L’s with same span. Then $L_1^{+2} \leftrightarrow L_2^{+2}$.*

In particular, two coins in hand are enough to *flip* any ‘L’, which means turning it into the




Figure 5. Left: a configuration C . Right: the associated canonical configuration L_C (the top-left and bottom-left ‘L’s are even, the top-right and bottom-right ‘L’s are odd).

mirrored ‘L’ hugging the other two sides of the span. Let us introduce our own routines for this particular transformation, one reason being that our method to flip odd ‘L’s only uses one extra coin, which will be useful in Section 5. Figure 7 explains how to flip an even ‘L’ with two coins in hand, using the subroutines from Figure 6. Figure 10 explains how to flip an even ‘L’ with one coin in hand, using the subroutines from Figure 9 as well as the *leapfrog* technique described in Figure 8: a leapfrog means relocating the unique pair of adjacent coins inside of an odd ‘L’. In all figures throughout this paper, an encircled coin represents a coin that has just been dropped, while a crossed out coin represents a coin that we pick up.




Figure 6. Subroutines used to flip an even ‘L’. The bottom subroutine is only used if both sides are even.




Figure 7. Flipping an even ‘L’, with odd sides (top) or even sides (bottom).




Figure 8. A leapfrog. Intermediary states cover all possible locations of the pair of adjacent coins.




Figure 9. Subroutines used to flip an odd ‘L’.




Figure 10. Flipping an odd ‘L’.

2.3 Canonicalization process

The crucial result is that two coins in hand are enough to turn any configuration into its associated canonical configuration in a reversible manner:

Lemma 2.7. [DDV02] *For any configuration C , we have $C^{+2} \leftrightarrow L_C^{+2+|C|-|L_C|}$.*

From there, a method to solve a puzzle $A \xrightarrow{?} B$ would roughly be to: pick up two coins in A ; canonicalize; reverse into B minus two coins; drop two coins to finish B . To do this however:

- We need two coins in A that we can pick up without breaking the inclusion of spans at the start i.e. 2 extra coins in A relatively to B .
- We need two appropriate spots in B to drop our two coins in hand at the end i.e. 2 redundant coins in B .

Moreover, unless the spans are equal once the first two coins $\{a_1, a_2\}$ have been picked up, we need a way to go from $L_{A \setminus \{a_1, a_2\}}$ to L_B . This reasoning can be summed up as follows:

Corollary 2.8. *Let A and B be configurations such that $|A| = |B|$ and:*

- (i) *A has 2 extra coins relatively to B .*
- (ii) *B has 2 redundant coins.*

Let $A_0 := A \setminus \{a_1, a_2\}$ where $\{a_1, a_2\}$ is a set of extra coins in A relatively to B .

If $L_{A_0}^{+|A|-|L_{A_0}|} \rightarrow L_B^{+|B|-|L_B|}$, then $A \rightarrow B$.

Proof. Let $B_0 := B \setminus \{b_1, b_2\}$ where $\{b_1, b_2\}$ is a set of redundant coins in B . Note that $L_B = L_{B_0}$ since $\text{span}(B) = \text{span}(B_0)$. We go from A to B in five steps:

- We get $A \rightarrow A_0^{+2}$ by picking up a_1 and a_2 .
- Lemma 2.7 ensures that $A_0^{+2} \rightarrow L_{A_0}^{+2+|A_0|-|L_{A_0}|}$.
- We have $L_{A_0}^{+2+|A_0|-|L_{A_0}|} = L_{A_0}^{+|A|-|L_{A_0}|} \rightarrow L_B^{+|B|-|L_B|} = L_{B_0}^{+2+|B_0|-|L_{B_0}|}$ by assumption.
- Lemma 2.7 ensures that $L_{B_0}^{+2+|B_0|-|L_{B_0}|} \rightarrow B_0^{+2}$.
- Finally, Proposition 1.20 yields $B_0^{+2} \rightarrow B$. □

3 Two extra coins: with additional constraints

In [DDV02], the authors studied the case where A has 2 extra coins relatively to B , which seemed enough to solve most puzzles. In this section, we go back on their result and then give it a slight improvement.

3.1 Previous result

The main result in [DDV02] is the following:

Theorem 3.1. [DDV02] *Let A and B be configurations such that $|A| = |B|$, and suppose that:*

- (i) $\text{span}(A) = \text{span}(B)$.
- (ii) A has 2 extra coins relatively to B .
- (iii) B has 2 redundant coins.

Then $A \rightarrow B$.

As mentioned in Section 2, this theorem uses canonical configurations as an intermediary. Indeed, it follows immediately from Corollary 2.8 (with $L_{A_0} = L_A = L_B$ since $\text{span}(A) = \text{span}(B)$). An example of a puzzle that is solvable according to Theorem 3.1 is the one in Figure 4, where $\{a_1, a_2\}$ is a set of extra coins in A relatively to B and $\{b_1, b_2\}$ is a set of redundant coins in B .

In fact, a stronger version of this theorem is claimed in [DDV02], where condition (i) is not required. The authors reduce to the case where $\text{span}(A) = \text{span}(B)$ by picking up all coins in $A \setminus \text{span}(B)$. However, this does not work in general, because this might actually cause the span to become strictly smaller than that of B . It turns out that, without the added condition (i), some puzzles are solvable and some are not, as illustrated by Figure 11. The puzzle on the left is solvable in 12 moves. The puzzle on the right is unsolvable (as we will later prove) and therefore is a counterexample to the version of the theorem in [DDV02]. It is easy to check that no smaller counterexample exists, be it in terms of number of coins or half-perimeter of the starting span. A generalized family of counterexamples will be exhibited in Section 4: in all of them, the problem comes from the fact that $\text{span}(B)$ is split into two far apart components.




Figure 11. Two puzzles satisfying conditions (ii) and (iii) but not (i), since $\text{span}(A) \supsetneq \text{span}(B)$. The left one is solvable but the right one is not.

3.2 A slight improvement

First of all, we show that Theorem 3.1 can be slightly improved: instead of asking for the spans to be equal, it is sufficient that each component of the starting span contains at most one component of the target span. Indeed, while splitting an ‘L’ into two separate components can be difficult (as we have just seen), shrinking an ‘L’ with two coins in hand is not a problem.

Lemma 3.2. *If L_1 and L_2 are canonical ‘L’s with $\text{span}(L_1) \supseteq \text{span}(L_2)$, then $L_1^{+2} \rightarrow L_2^{+2+|L_1|-|L_2|}$.*

Proof. First of all, we *trim* L_1 to the right if needed, as follows (see Figure 12):

1. If L_1 is odd, we use a leapfrog to put the pair of adjacent coins to the far right.
2. We make sure there is a coin c at the rightmost position that we want to keep, by dropping one there if needed.
3. We finish by simply picking up all coins that are further right than c .




Figure 12. Trimming a 7×4 ‘L’ to its right, making it 4×4 (top) or 5×4 (bottom). The numbers above the arrows refer to the three steps.

We then trim our ‘L’ at the top, in analogous fashion. We now flip it, so it is now ready to be trimmed to the left and at the bottom. Once this is done, we flip it back and use a leapfrog if needed to make it canonical. $\square$

Theorem 3.3. *Let A and B be configurations such that $|A| = |B|$, and suppose that:*

- (i) *A has 2 extra coins relatively to B , and more precisely: there exist $a_1 \neq a_2$ in A such that $\text{span}(A \setminus \{a_1, a_2\}) \supseteq \text{span}(B)$ and each component of $\text{span}(A \setminus \{a_1, a_2\})$ contains at most one component of $\text{span}(B)$.*
- (ii) *B has 2 redundant coins.*

Then $A \rightarrow B$.

Proof. Let $A_0 := A \setminus \{a_1, a_2\}$. By Corollary 2.8, it suffices to show that $L_{A_0}^{+|A|-|L_{A_0}|} \rightarrow L_B^{+|B|-|L_B|}$. Since each component of $\text{span}(A_0)$ contains at most one component of $\text{span}(B)$, we can use

Lemma 3.2 to shrink each ‘L’ in L_{A_0} to the size of the corresponding ‘L’ in L_{B_0} . This is always possible, because we start off with $|A| - |L_{A_0}| \geq |A| - |A_0| = 2$ coins in hand and this number cannot decrease each time we shrink an ‘L’. $\square$

Note that puzzles satisfying condition (i) but not condition (ii) (while still satisfying the fact that B has 1 redundant coin, otherwise we would be in a trivially unsolvable case by Proposition 1.19) also may or may not be solvable as shown in Figure 13.



Figure 13. Two puzzles satisfying condition (i) but not (ii). The left puzzle is solvable in 4 moves, while the right puzzle is unsolvable.

4 Two extra coins: general case

What if some component of $\text{span}(A)$ contains two or more components of $\text{span}(B)$, so that Theorem 3.3 does not apply? A first natural guess would be that we then need more extra coins and/or more redundant coins than just two. Nevertheless, we now exhibit a family of unsolvable puzzles which proves that, even with the inclusion of spans, no constant number of extra coins in A (relatively to B or not) or redundant coins in B can guarantee that a puzzle is solvable in general. Next, we present a new sufficient condition, which shows in particular that the aforementioned family consists of just about worst-case puzzles.

4.1 Worst-case puzzles

Puzzles like those from Figure 11 require to split the span, at some point during the moves, in a way that we now prove impossible without a certain amount of total coins relative to the size of the rectangles involved.

Definition 4.1. Let R_1 and R_2 be rectangles at distance at least 3 from each other, and let A be a configuration such that R_1 and R_2 are included in the same component of $\text{span}(A)$. An (R_1, R_2) -split of A is a sequence of moves $A = A_0 \mapsto A_1 \mapsto \dots \mapsto A_T$ ($T \geq 1$ necessarily) such that R_1 and R_2 are included in two separate components of $\text{span}(A_T)$.

Proposition 4.2. Let R_1 and R_2 be rectangles of size $m_1 \times n_1$ and $m_2 \times n_2$ respectively, whose projections on the x axis intersect, and whose projections on the y axis do not intersect with a gap of $h \geq 2$ rows separating them. Let A be a configuration such that R_1 and R_2 are included in the same component of $\text{span}(A)$. If there exists an (R_1, R_2) -split of A , then $|A| \geq \frac{m_1 + n_1 + m_2 + n_2 + h - 1}{2}$.

Proof. Let $A = A_0 \mapsto A_1 \mapsto \dots \mapsto A_T$ be an (R_1, R_2) -split of A with minimum number of moves, so that R_1 and R_2 are included in the same component R of $\text{span}(A_{T-1})$ but in two separate components R'_1 and R'_2 of $\text{span}(A_T)$ (see Figure 14). If R'_1 is of size $m'_1 \times n'_1$ and R'_2 is of size $m'_2 \times n'_2$ with a gap of h' rows separating them, then we have $m'_1 \geq m_1$, $m'_2 \geq m_2$ and

$n'_1 + n'_2 + h' \geq n_1 + n_2 + h$ hence $\frac{m'_1 + n'_1 + m'_2 + n'_2 + h' - 1}{2} \geq \frac{m_1 + n_1 + m_2 + n_2 + h - 1}{2}$. Therefore, the worst case for what we want to prove is if $R_1 = R'_1$ and $R_2 = R'_2$, which is what we assume from now.




Figure 14. Illustration of R , R_1 , R_2 , R'_1 , R'_2 .

We use A_{T-1} to count the coins and carry out the proof. We have:

$$|A| = |A_{T-1}| = |A_{T-1} \cap R_1| + |A_{T-1} \cap R_2| + |A_{T-1} \setminus (R_1 \cup R_2)|. \quad (1)$$

Moreover:

- The only coin in $A_T \cap R_1$ that might not be in $A_{T-1} \cap R_1$ is the coin that has been moved to go from A_{T-1} to A_T , however that coin does not contribute to the span since it has at least two coins adjacent to it in $A_T \cap R_1$. Since $\text{span}(A_T \cap R_1) = R_1$, we thus get $\text{span}(A_{T-1} \cap R_1) = R_1$. In particular, Proposition 1.15 yields:

$$|A_{T-1} \cap R_1| \geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil. \quad (2)$$

- Similarly, $\text{span}(A_{T-1} \cap R_2) = R_2$ and:

$$|A_{T-1} \cap R_2| \geq \left\lceil \frac{m_2 + n_2}{2} \right\rceil. \quad (3)$$

- By Proposition 1.12, there cannot be two consecutive rows of R without a coin in A_{T-1} . Since there is a gap of h rows between R_1 and R_2 , this yields:

$$|A_{T-1} \setminus (R_1 \cup R_2)| \geq \left\lfloor \frac{h}{2} \right\rfloor. \quad (4)$$

Combining (1), (2), (3) and (4), we get:

$$|A| \geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil + \left\lceil \frac{m_2 + n_2}{2} \right\rceil + \left\lfloor \frac{h}{2} \right\rfloor \geq \frac{m_1 + n_1 + m_2 + n_2 + h - 1}{2}. \quad \square$$

Corollary 4.3. *Let $n \geq 9$. We define configurations A_n and B_n as in Figure 15:*

- A_n consists of an $n \times n$ ‘L’ with $\left\lfloor \frac{n-2}{2} \right\rfloor$ coins added to the bottom row.
- B_n has the same smallest enclosing rectangle as A_n and contains an $n \times 1$ ‘L’ in both the top and bottom row with $\left\lfloor \frac{n-5}{2} \right\rfloor$ coins added to the bottom row.




Figure 15. Definition of the puzzle $A_n \xrightarrow{?} B_n$ (here $n = 9$). The shaded coins represent the extra/redundant coins.

Then $A_n \not\rightarrow B_n$ even though: A_n has $\lfloor \frac{n-2}{2} \rfloor$ extra coins, B_n has $\lfloor \frac{n-5}{2} \rfloor$ redundant coins, and $\text{span}(A_n) \supset \text{span}(B_n)$.

Proof. Let R_1 (resp. R_2) be the top row (resp. the bottom row) of $\text{span}(A_n)$: R_1 and R_2 are of size $n \times 1$ and separated by a gap of $n - 2$ rows. Solving this puzzle would mean performing an (R_1, R_2) -split of A_n , which is impossible by Proposition 4.2 because:

- If n is odd then $n + \frac{n-3}{2} = |A_n| = |B_n| = 2\frac{n+1}{2} + \frac{n-5}{2} = \frac{3n-3}{2} < \frac{n+1+n+1+(n-2)-1}{2}$.
- If n is even then $n + \frac{n-2}{2} = |A_n| = |B_n| = 2\frac{n+2}{2} + \frac{n-6}{2} = \frac{3n-2}{2} < \frac{n+1+n+1+(n-2)-1}{2}$. $\square$

Corollary 4.4. For any $k \in \mathbb{N}$, there exist configurations A and B with $|A| = |B|$ such that:

- $\text{span}(A) \supset \text{span}(B)$.
- A has k extra coins.
- B has k redundant coins.
- $A \not\rightarrow B$.

It is actually possible to improve Proposition 4.2: we now give a refined bound that even applies to the puzzle on the right of Figure 11 which, as previously mentioned, is the smallest counterexample to the version of Theorem 3.1 in [DDV02].

Proposition 4.5. Let R_1 , R_2 and A be as in Proposition 4.2. If there exists an (R_1, R_2) -split of A but none in two moves or less, then $|A| \geq \frac{m_1+n_1+m_2+n_2+h+2}{2}$.

Proof. Let us pick up where the proof of Proposition 4.2 ended. Since there exists no (R_1, R_2) -split of A in two moves or less, we have $T \geq 3$. We use the fact that A_{T-1} then has the following properties:

- (i) A_{T-1} contains a coin that is adjacent to at least two other coins.
- (ii) For all $c \in A_{T-1}$, $A_{T-1} \setminus \{c\}$ does not consist of all isolated coins.

Property (i) comes from Proposition 1.2 since $T-1 \geq 1$, and property (ii) comes from Proposition 1.3 since $T-1 \geq 2$ (indeed $A \not\rightarrow A_{T-1}$ because our sequence of moves has been chosen shortest). We distinguish between four cases:

- 1) Case 1: R_1 and R_2 are both odd; h is even.

This is the easiest case:

$$\begin{aligned}
|A| &\geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil + \left\lceil \frac{m_2 + n_2}{2} \right\rceil + \left\lfloor \frac{h}{2} \right\rfloor = \frac{m_1 + n_1 + 1}{2} + \frac{m_2 + n_2 + 1}{2} + \frac{h}{2} \\
&= \frac{m_1 + n_1 + m_2 + n_2 + h + 2}{2}.
\end{aligned}$$

- 2) Case 2: At least one of R_1 or R_2 is even (say R_2 is even); h is even.

We just need to find one coin more than what (2), (3) and (4) give us combined, because if we do then we can conclude that:

$$\begin{aligned}
|A| &\geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil + \left\lceil \frac{m_2 + n_2}{2} \right\rceil + \left\lfloor \frac{h}{2} \right\rfloor + 1 \\
&\geq \frac{m_1 + n_1}{2} + \frac{m_2 + n_2}{2} + \frac{h}{2} + 1 = \frac{m_1 + n_1 + m_2 + n_2 + h + 2}{2}.
\end{aligned}$$

Therefore, suppose for a contradiction that (2), (3) and (4) are all tight. For (4), this means that $A_{T-1} \setminus (R_1 \cup R_2)$ consists exactly of one coin every two rows in the gap between R_1 and R_2 . For (2) and (3), this means $A_{T-1} \cap R_1$ and $A_{T-1} \cap R_2$ are both minimum: by Proposition 1.15, all coins in $A_{T-1} \cap R_2$ are isolated and all coins in $A_{T-1} \cap R_1$ are isolated except possibly for a single pair of adjacent coins. We can see that the only way to satisfy property (i) is if $A_{T-1} \cap R_1$ contains a pair $\{c_1, c_2\}$ of adjacent coins such that c_1 is adjacent to one of the coins in $A_{T-1} \setminus (R_1 \cup R_2)$, as in Figure 16 (left). However, c_1 then violates property (ii), a contradiction.

- 3) Case 3: At least one of R_1 or R_2 is odd (say R_2 is odd); h is odd.

Again, we just need to find one coin more than what (2), (3) and (4) give us combined, because if we do then we can conclude that:

$$\begin{aligned}
|A| &\geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil + \left\lceil \frac{m_2 + n_2}{2} \right\rceil + \left\lfloor \frac{h}{2} \right\rfloor + 1 \\
&\geq \frac{m_1 + n_1}{2} + \frac{m_2 + n_2 + 1}{2} + \frac{h - 1}{2} + 1 = \frac{m_1 + n_1 + m_2 + n_2 + h + 2}{2}.
\end{aligned}$$

Therefore, suppose for a contradiction that (2), (3) and (4) are all tight. For (2) and (3), this means that $A_{T-1} \cap R_1$ and $A_{T-1} \cap R_2$ are both minimum: in particular, neither contains three coins such that one is adjacent to the other two. For (4), this means that $A_{T-1} \setminus (R_1 \cup R_2)$ consists exactly of one coin every two rows in the gap between R_1 and R_2 , none of which is adjacent to R_1 or R_2 since h is odd. See Figure 16 (middle). All in all, property (i) is violated, a contradiction.

- 4) Case 4: R_1 and R_2 are both even; h is odd.

In this case, we need to find two coins more than what (2), (3) and (4) give us combined, because if we do then we can conclude that:

$$\begin{aligned}
|A| &\geq \left\lceil \frac{m_1 + n_1}{2} \right\rceil + \left\lceil \frac{m_2 + n_2}{2} \right\rceil + \left\lfloor \frac{h}{2} \right\rfloor + 2 \\
&\geq \frac{m_1 + n_1}{2} + \frac{m_2 + n_2}{2} + \frac{h - 1}{2} + 2 = \frac{m_1 + n_1 + m_2 + n_2 + h + 3}{2}.
\end{aligned}$$

First of all, the same proof as in Case 3 shows that (2), (3) and (4) cannot all be tight: there needs to be some coin c in A_{t-1} that has (at least) two coins adjacent to it. Now

suppose for a contradiction that all three inequalities become tight if we remove c i.e.: $|(A_{T-1} \setminus \{c\}) \cap R_1| = \lceil \frac{m_1+n_1}{2} \rceil$, $|(A_{T-1} \setminus \{c\}) \cap R_2| = \lceil \frac{m_2+n_2}{2} \rceil$, $|(A_{T-1} \setminus \{c\}) \setminus (R_1 \cup R_2)| = \lfloor \frac{h}{2} \rfloor$. See Figure 16 (right). If $c \in R_1$ then the two coins adjacent to c in A_{T-1} are inside R_1 as well (indeed, as we have seen in Case 3, the fact that h is odd means that none of the $\lfloor \frac{h}{2} \rfloor$ coins in $A_{T-1} \setminus (R_1 \cup R_2)$ is adjacent to R_1). Therefore $(A_{T-1} \setminus \{c\}) \cap R_1$ has span R_1 , and is minimum since $|(A_{T-1} \setminus \{c\}) \cap R_1| = \lceil \frac{m_1+n_1}{2} \rceil$. Similarly, $(A_{T-1} \setminus \{c\}) \cap R_2$ has span R_2 and is minimum. By Proposition 1.15, all coins in $(A_{T-1} \setminus \{c\}) \cap R_1$ and $(A_{T-1} \setminus \{c\}) \cap R_2$ are thus isolated. Since h is odd, the $\lfloor \frac{h}{2} \rfloor$ coins in $(A_{T-1} \setminus \{c\}) \setminus (R_1 \cup R_2)$ are also isolated. This contradicts property (ii). $\square$




Figure 16. The configuration A_{T-1} : a contradiction in Case 2 (left), Case 3 (middle) and Case 4 (right).

Corollary 4.6. *The puzzle on the right of Figure 11 is unsolvable.*

Proof. Let R_1 (resp. R_2) be the top row (resp. the bottom row) of the starting span: R_1 and R_2 are of size 3×1 and separated by a three row gap. Solving this puzzle would mean performing an (R_1, R_2) -split of the starting configuration, which is impossible by Proposition 4.5: indeed, we can easily check that it is impossible in two moves or less, and the puzzle contains $6 < \frac{3+1+3+1+3+2}{2}$ coins. $\square$

4.2 A new sufficient condition

We now present a result that holds even when some component of $\text{span}(A)$ contains two or more components of $\text{span}(B)$.

Notation 4.7. Let C be a configuration. We denote by $\min_C$ the cardinality of minimum configurations with same span as C . Note that, if this span is an $m \times n$ rectangle, then $\min_C = \lceil \frac{m+n}{2} \rceil$ by Proposition 1.15.

Theorem 4.8. *Let A and B be configurations such that $|A| = |B| =: N$ and:*

- (i) $\text{span}(A) \supseteq B$ is a single $m \times n$ rectangle.
- (ii) A has 2 extra coins.
- (iii) B has 2 redundant coins.

(iv) $N \geq \frac{3}{2} \max(\min_A, \min_B) + 2$.
Then $A \rightarrow B$.

We would have liked conditions (i) and (ii) to be replaced by the sole condition that A has 2 extra coins relatively to B , however we are not sure how the proof would work in that case. Apart from that, the additional assumption compared to Theorem 3.1 is condition (iv), which is not about the quality of the coins involved (extra/redundant) but purely about their quantity, as was suggested by the worst-case puzzles from Corollary 4.3. Moreover, these puzzles also show that the bound from condition (iv) is almost tight: indeed, if n is odd for instance, we have $\min_{A_n} = n$ and $\min_{B_n} = n + 1$ so that the puzzle $A_n \xrightarrow{?} B_n$ satisfies $N = \frac{3n-3}{2} = \frac{3}{2} \max(\min_{A_n}, \min_{B_n}) - 3$, just 5 coins away from this bound.

We now proceed with the proof of Theorem 4.8. As usual, we are going to route through the canonical configurations, which means the challenge is to go from L_A to L_B . The proof relies on an intuitive way to do so, which consists in forming a wave of coins (by flipping 'L's) to sweep across the board while dropping coins at all positions in L_B , as detailed in the proof of the following lemma. Note that this lemma is more general than we use, since the target configuration is not required to be canonical.

Lemma 4.9. *Let $m, n \geq 1$ and let L be the canonical 'L' of size $m \times n$. Let $k \in \mathbb{N}$. If a configuration $C \subseteq \text{span}(L)$ satisfies $|C| < \min\left(\left\lceil \frac{m+n}{2} \right\rceil - \left\lceil \frac{\min(m,n)}{2} \right\rceil + (k-1), 2(k-1)\right)$, then $L^{+k} \rightarrow C^{+k+|L|-|C|}$.*

Proof. We proceed by induction on the half-perimeter $m + n$. If $m = n = 1$, then $C = L$ or $C = \emptyset$ so the result is obvious. Suppose $m + n \geq 3$ and assume the result holds for any half-perimeter lesser than $m + n$. Up to a 90 degree rotation of the board, also assume $m \leq n$. We divide $R := \text{span}(L)$ into two rectangles: a bottom half R_1 and a top half R_2 (if n is even then both halves are equal, otherwise we choose one of them arbitrarily to be bigger than the other by one row). Define $C_1 := C \cap R_1$ and $C_2 := C \cap R_2$. Up to swapping the roles of R_1 and R_2 , assume $|C_1| \leq |C_2|$. Let $n' \in \left\{\left\lfloor \frac{n}{2} \right\rfloor, \left\lceil \frac{n}{2} \right\rceil\right\}$ be the number of rows of R_2 , so that R_2 is of size $m \times n'$. To reach C , we build the bottom half C_1 first, then the top half C_2 .

[1] We start by building the bottom half of C . We show that $L^{+k} \rightarrow C_0^{+k'}$ where $k' := k + |L| - |C_0|$ and C_0 is the configuration defined as follows: $C_0 \cap R_1 = C_1$, and $C_0 \cap R_2 = L_2$ is the canonical 'L' of size $m \times n'$. See Figure 17.

By our assumption on C , we have $k - 1 > \frac{|C|}{2} \geq |C_1|$ (recall that we have assumed without loss of generality that $|C_1| \leq |C_2|$ i.e. C_1 contains at most half of the coins in C) so $k \geq |C_1| + 2$. This allows us to view our k coins in hand as follows:

- We have 2 *supporting coins* that we will use to transform chains.
- We have $|C_1|$ *building coins* that we will drop at the right positions to build C_1 .
- If $k > |C_1| + 2$, the remaining $k - |C_1| - 2$ coins will be kept in hand.

Note that the supporting coins might not remain the same throughout the moves. For example, we might drop a supporting coin, perform some moves, and then pick up a coin: in that case, the picked up coin becomes a supporting coin even if it is not "physically" the same coin that we dropped initially. We proceed in four steps.

(a) Let p be the bottom-left corner of R_2 : we want to make sure there is a coin at p .




Figure 17. An example of a configuration C (left) with its associated configuration C_0 (right).

- Case (a1): even R , odd n' . There already is a coin at p .
- Case (a2): even R , even n' . We drop one of our supporting coins at p .
- Case (a3): odd R . If needed, we use a leapfrog to put one of the two adjacent coins at p . Recall that a leapfrog only uses one supporting coin.




Figure 18. The board after step (a). From left to right: case (a1), case (a2), case (a3).

- (b) The board now contains an ‘L’ whose extremal coins are at the bottom-left corner of R_2 and the bottom-right corner of R_1 , as highlighted in Figure 18. We now use our supporting coin(s) to flip this ‘L’, while building C_1 in the process (see Figure 20 for the desired result). To achieve this, we take advantage of the fact that the flip sweeps over the entirety of R_1 . In the subroutines from Figures 6 and 9, some positions are highlighted by a black outline: whenever one of these positions contains a coin in C_1 , we drop a building coin there at the appropriate moment during the subroutine (one example is detailed in Figure 19). Over the flip as a whole, these positions cover all of R_1 except its rightmost column, so that all of C_1 is correctly replicated at the end of this step apart from the coins in the rightmost column. Note that the ‘L’ that we flip is odd in case (a2), so that the lone remaining supporting coin is indeed enough to flip it. See Figure 20.
- (c) We now make sure there is a coin in the bottom-right corner of R_2 . If the newly flipped ‘L’ is odd, this is done with a leapfrog. If it is even, in particular we are not in case (a2), so we can afford to drop one of our two supporting coins at the bottom-right corner of R_2 if needed. In both cases, we still have at least one supporting coin at




Figure 19. How to drop a building coin at any desired position (here we drop four of them, but we can drop less). The building coins are shaded.




Figure 20. Left: a configuration C (the same example as in Figure 17). Middle: the board after step (a). Right: the board after step (b).

our disposal. We now correct the rightmost column of R_1 : we drop building coins where they are needed in the holes in-between the coins that are already on board, and then we pick up all coins in that column that are not in C_1 .

- (d) At this point, R_2 either contains an ‘L’ or a chain that is almost an ‘L’ apart from the fact it has two pairs of adjacent coins (this can happen if we have dropped a supporting coin at the bottom-right corner of R_2 in step (c)). In this latter case, we use a leapfrog to retrieve the coin in excess so that R_2 contains a true ‘L’. Finally, we leapfrog if needed to make this ‘L’ canonical, so that the full configuration on board is now exactly C_0 . These leapfrogs are always possible because we have at least one supporting coin at our disposal. See Figure 21.




Figure 21. Step (d) performed as follow-up to Figure 20.

- [2] We now build the top half of C . Since the bottom half is built already, we will not touch it, therefore we want to show that $L_2^{+k'} \rightarrow C_2^{+k'+|L_2|-|C_2|}$. By our induction hypothesis, since R_2 is of half-perimeter $m + n' < m + n$ (recall that $m + n \geq 3$, so $n \geq 2$ hence $n' < n$), it suffices to show that $|C_2| < \min\left(\left\lceil \frac{m+n'}{2} \right\rceil - \left\lceil \frac{\min(m,n')}{2} \right\rceil + (k' - 1), 2(k' - 1)\right)$. We now show both inequalities. Beforehand, since the total number of coins at our disposal is

$|L| + k$ and $|C_1| + |L_2|$ of them are on the board at this point, note that:

$$k' = (|L| + k) - (|C_1| + |L_2|) = \left(\left\lceil \frac{m+n}{2} \right\rceil + k \right) - \left(|C_1| + \left\lceil \frac{m+n'}{2} \right\rceil \right). \quad (5)$$

- By our assumption on C , we have $|C_1| = |C| - |C_2| < \left\lceil \frac{m+n}{2} \right\rceil - \left\lceil \frac{\min(m,n)}{2} \right\rceil + (k-1) - |C_2|$. Using equality (5), we get:

$$k' > |C_2| + 1 + \left\lceil \frac{\min(m,n)}{2} \right\rceil - \left\lceil \frac{m+n'}{2} \right\rceil \geq |C_2| + 1 + \left\lceil \frac{\min(m,n')}{2} \right\rceil - \left\lceil \frac{m+n'}{2} \right\rceil,$$

where the last inequality comes from the fact that $n > n'$. We thus get the first desired inequality: $|C_2| < \left\lceil \frac{m+n'}{2} \right\rceil - \left\lceil \frac{\min(m,n')}{2} \right\rceil + (k' - 1)$.

- By our assumption on C , we have $k-1 > \frac{|C|}{2} \geq |C_1|$ hence $k-1-|C_1| > 0$. Moreover $k-1-|C_1| = k-1-|C|+|C_2| > |C_2| - \left\lceil \frac{m+n}{2} \right\rceil + \left\lceil \frac{\min(m,n)}{2} \right\rceil = |C_2| - \left\lceil \frac{m+n}{2} \right\rceil + \left\lceil \frac{m}{2} \right\rceil$. Since all integers $x > y$ with $x > 0$ satisfy $x \geq \frac{y}{2} + 1$, we get $k-1-|C_1| \geq \frac{1}{2} \left(|C_2| - \left\lceil \frac{m+n}{2} \right\rceil + \left\lceil \frac{m}{2} \right\rceil \right) + 1$.

Recall that $k' - 1 = (k - 1 - |C_1|) + \left\lceil \frac{m+n}{2} \right\rceil - \left\lceil \frac{m+n'}{2} \right\rceil$ by equality (5). Therefore:

$$\begin{aligned} k' - 1 &\geq \frac{|C_2|}{2} + \frac{1}{2} \left\lceil \frac{m+n}{2} \right\rceil + \frac{1}{2} \left\lceil \frac{m}{2} \right\rceil - \left\lceil \frac{m+n'}{2} \right\rceil + 1 \\ &\geq \frac{|C_2|}{2} + \frac{m+n}{4} + \frac{m}{4} - \frac{m + \frac{n+1}{2} + 1}{2} + 1 = \frac{|C_2|}{2} + \frac{1}{4} \\ &> \frac{|C_2|}{2}, \end{aligned}$$

from which $|C_2| < 2(k' - 1)$ which concludes. $\square$

Proof of Theorem 4.8. In fact, we prove a more general result where condition (iv) is replaced by the following double inequality:

$$N > \min_A + \frac{\min_B}{2} + 1 \quad (6)$$

$$N > \min_B + \left\lceil \frac{\min(m,n)}{2} \right\rceil + 1 \quad (7)$$

Let us first check that this assumption is indeed weaker. Suppose that (iv) holds, then:

- $N \geq \min_A + \frac{\min_B}{2} + 2$. Therefore, (6) holds.
- $N \geq \min_B + \frac{\min_A}{2} + 2$. Since N is an integer, this yields $N \geq \min_B + \left\lceil \frac{\min_A}{2} \right\rceil + 2$. Moreover $\left\lceil \frac{\min_A}{2} \right\rceil = \left\lceil \frac{1}{2} \left\lceil \frac{m+n}{2} \right\rceil \right\rceil = \left\lceil \frac{m+n}{4} \right\rceil \geq \left\lceil \frac{\min(m,n)}{2} \right\rceil$, where we have used the fact that any real number x satisfies $\left\lceil \frac{x}{2} \right\rceil = \left\lceil \frac{\lceil x \rceil}{2} \right\rceil$. Therefore, we get $N \geq \min_B + \left\lceil \frac{\min(m,n)}{2} \right\rceil + 2$, so (7) holds.

Assume that conditions (i),(ii),(iii),(6),(7) all hold. As already mentioned, we want to use Corollary 2.8 (in this case $L_{A_0} = L_A$ because of our assumption that A has 2 extra coins) so we need to show that $L_A^{+k} \rightarrow L_B^{+k+|L_A|-|L_B|}$ where $k := |A| - |L_A| = N - \min_A = N - \left\lceil \frac{m+n}{2} \right\rceil$. By Lemma 4.9, it suffices to show that $\min_B = |L_B| < \min \left(\left\lceil \frac{m+n}{2} \right\rceil - \left\lceil \frac{\min(m,n)}{2} \right\rceil + (k-1), 2(k-1) \right)$.

- By (6), we have $\min_B < 2(N - \min_A - 1) = 2(k-1)$.

- By (7), we have $\min_B < N - \left\lceil \frac{\min(m,n)}{2} \right\rceil - 1 = \left\lceil \frac{m+n}{2} \right\rceil - \left\lceil \frac{\min(m,n)}{2} \right\rceil + (k-1)$. $\square$

Even though, as we have mentioned, the bound in Theorem 4.8 is tight up to an additive factor $O(1)$ for a square span ($n = m$, as in A_n and B_n), things might be different for a rectangular $m \times n$ span in general.

5 One extra coin: minimum+1 configurations

The case where A only has 1 extra coin relatively to B is even more complicated. In this section, we initiate its study with the following restrictions:

- $\text{span}(A) = \text{span}(B)$ is a single rectangle.
- A is *minimum+1*, which means that A consists of a minimum configuration with same span as A plus one coin added to it. In particular, A has 1 extra coin (and not more).

Note that under these restrictions, and further assuming that B has 1 redundant coin b which we know is necessary, B is also *minimum+1*: indeed, $B \setminus \{b\}$ has same span as A so it is minimum by cardinality.

5.1 Even span

The case of an even span is straightforward: if A is a *minimum+1* configuration with even span, then the only way to make moves from A without decreasing the span is to move the same coin over and over again.

Proposition 5.1. *Let $A \neq B$ be minimum+1 configurations such that $\text{span}(A) = \text{span}(B)$ is an even rectangle. Then $A \rightarrow B$ if and only if $A \mapsto B$.*

Proof. Write $B = M \cup \{b\}$ where M is a minimum configuration with same span as B . Since this span is an even rectangle, we know M has all isolated coins by Proposition 1.15, therefore $A \rightarrow B$ if and only if $A \mapsto B$ by Proposition 1.3. $\square$

5.2 Odd span

On the contrary, the slightly higher density of coins allows for some non-trivial movement when the span is odd. For instance, we have seen (and it was already noticed in [DDV02]) that one coin in hand is enough to flip an odd ‘L’, so we know that the puzzle on the left of Figure 22 is solvable. In this example, the big difference compared to the even case is the presence of a pair of adjacent coins, which helps the moves and makes this “+1” situation more of a “+1.5” situation in a way. When it comes to a 90 degree *rotation* of an odd ‘L’ however, even though it is easy to do with two coins in hand, it seems impossible with just one (as noticed again in [DDV02]), which we will confirm shortly so that the puzzle on the right of Figure 22 is not solvable. The challenge is to figure out where the limit is between what is achievable and what is not, and we now bring our contribution in that direction. In what follows, we fix an odd $m \times n$ rectangle R .




Figure 22. Two examples of puzzles involving minimum+1 configurations with even span. The one on the left is solvable, the one on the right is not.

5.2.1 Reduction to the poking game

First of all, we can make an interesting observation:

Lemma 5.2. *Let A be a minimum+1 configuration, and consider a sequence of moves $A = A_0 \xrightarrow{c_1 \mapsto p_1} A_1 \xrightarrow{c_2 \mapsto p_2} A_2 \mapsto \dots$ such that the span is preserved throughout and no coin is moved twice in a row. Then:*

- For all $t \geq 1$, A_t is minimum+1.
- For all $t \geq 1$, c_{t+1} is a neighbor of p_t . In other words, the coin that we move (apart from the very first one) is always a neighbor of the most recently moved coin.

Proof. Let $t \geq 1$.

- $\text{span}(A_t \setminus \{p_t\}) = \text{span}(A_t) = \text{span}(A)$, therefore $A_t \setminus \{p_t\}$ is minimum by cardinality, so A_t is minimum+1.
- Suppose for a contradiction that c_{t+1} is not a neighbor of p_t . Since no coin is moved twice in a row, we know $p_t \neq c_{t+1}$, so this means that p_t has at least two neighbors in $A_t \setminus \{c_{t+1}\}$ by the 2-adjacency rule, which yields $\text{span}(A_t \setminus \{c_{t+1}\}) = \text{span}(A_t \setminus \{c_{t+1}, p_t\})$. Finally, since $A_{t+1} \setminus \{p_{t+1}\} = A_t \setminus \{c_{t+1}\}$, we get:

$$\text{span}(A_t) = \text{span}(A_{t+1}) = \text{span}(A_{t+1} \setminus \{p_{t+1}\}) = \text{span}(A_t \setminus \{c_{t+1}\}) = \text{span}(A_t \setminus \{c_{t+1}, p_t\}),$$

therefore A_t has 2 extra coins which is a contradiction. $\square$

Conceptually, it is helpful to forget about the "+1" coin, by ignoring the last moved coin at all times. This allows us to focus on the underlying minimum configuration and how it evolves during the moves. Since each coin that we move is a neighbor of the most recently moved coin, as guaranteed by Lemma 5.2, each move causes the underlying minimum configuration to see one of its coins effectively slide by one position towards another coin. An example is given in Figure 23 (the first move is ignored as it will be in the reduction to come). This inspires us to consider a new game, which is played directly on minimum configurations instead of minimum+1 configurations.

Definition 5.3. Let M be a minimum configuration with a pair $\{c_1, c_2\}$ of adjacent coins (recall that, by Proposition 1.15, M contains at most one such pair). A *poke* consists, for some $i \in \{1, 2\}$, in sliding c_i to one of its free neighboring positions p with the restriction that p must have at least one neighboring coin c other than c_i (we might say that c_i is poked *towards* c). This is denoted by $c_i \xrightarrow{\mathcal{P}} p$.




Figure 23. Top: a sequence of moves (the last moved coin is shaded).
Bottom: evolution of the underlying minimum configuration obtained by deleting the last moved coin.

Notation 5.4. Let M and M' be minimum configurations with span R , each containing a pair of adjacent coins.

- If there is a single poke $c \xrightarrow{\mathcal{P}} p$ from M to M' , we write $M \xrightarrow{c \xrightarrow{\mathcal{P}} p} M'$ or simply $M \xrightarrow{\mathcal{P}} M'$.
- If there exists a sequence of pokes from M to M' , we write $M \xrightarrow{\mathcal{P}} M'$.

Proposition 5.5. *A poke is always reversible: if $M_0 \xrightarrow{c \xrightarrow{\mathcal{P}} p} M_1$ then $M_1 \xrightarrow{p \xrightarrow{\mathcal{P}} c} M_0$. In particular, for any two minimum configurations M and M' with span R each containing a pair of adjacent coins, we have $M \xrightarrow{\mathcal{P}} M'$ if and only if $M' \xrightarrow{\mathcal{P}} M$.*

Proof. If $\{c, c_2\}$ denotes the pair of adjacent coins in M_0 , then p is part of the pair of adjacent coins in M_1 and c_2 is a neighboring coin of the free position c in M_1 which makes $p \xrightarrow{\mathcal{P}} c$ a valid poke from M_1 . $\square$

Proposition 5.6. *Any poke preserves the span, and the configuration obtained after a poke is still a minimum configuration with a pair of adjacent coins.*

Proof. Let M be a minimum configuration with a pair $\{c_1, c_2\}$ of adjacent coins. Let M' be the configuration obtained from M after some poke $c_1 \xrightarrow{\mathcal{P}} p$ (up to swapping the roles of c_1 and c_2). By definition of a poke, p has a neighboring coin $c \in M \setminus \{c_1\}$. Obviously M' contains a pair of adjacent coins, namely $\{p, c\}$. To conclude, it suffices to show that M' has same span as M , because this implies that M' is minimum by cardinality. In M , the position p is free but has two distinct occupied neighbors c_1 and c , therefore $\text{span}(M) = \text{span}(M \cup \{p\})$. In M' , the position c_1 is free but has two distinct occupied neighbors c_2 and p , therefore $\text{span}(M') = \text{span}(M' \cup \{c_1\})$. Since $M \cup \{p\} = M' \cup \{c_1\}$, we get $\text{span}(M) = \text{span}(M')$. $\square$

We have thus introduced a *poking game* where, given two minimum configurations M and M' whose span is the same odd rectangle and both containing a pair of adjacent coins, we ask the question whether $M \xrightarrow{\mathcal{P}} M'$. The next result is a formal proof of the reduction illustrated in Figure 23 and thus confirms that we can now focus entirely on the poking game.

Proposition 5.7. *Let A and B be minimum+1 configurations with span R . Suppose $A \neq B$*

and $A \not\rightarrow B$. Then $A \rightarrow B$ if and only if there exists a move $A \xrightarrow{c_1 \mapsto p_1} A_1$ and $a \in A_1$, $b \in B$ such that:

- a is a neighbor of p_1 .
- b is a redundant coin in B .
- $A_1 \setminus \{a\} \xrightarrow{\mathcal{P}} B \setminus \{b\}$.

As a consequence, the case where A and B are minimum+1 configurations whose span is the same odd rectangle reduces to the poking game, up to a factor $O(N^2)$ where $N := |A| = |B|$.

Proof. First suppose $A \rightarrow B$. We write $A = A_0 \xrightarrow{c_1 \mapsto p_1} A_1 \xrightarrow{c_2 \mapsto p_2} \dots \xrightarrow{c_T \mapsto p_T} A_T = B$ where no coin is moved twice in a row. We know $T \geq 2$ because $A \neq B$ and $A \not\rightarrow B$. Setting $a = c_2$ and $b = p_T$, since c_2 is a neighbor of p_1 by Lemma 5.2 and p_T is a redundant coin in B by the 2-adjacency rule, it suffices to show that $A_1 \setminus \{c_2\} \xrightarrow{\mathcal{P}} B \setminus \{p_T\}$. For $2 \leq t \leq T$, we define $M_t := A_t \setminus \{p_t\}$, the intermediary configuration obtained when the moved coin from A_{t-1} to A_t is "in the air": M_t is minimum by cardinality because $\text{span}(M_t) = \text{span}(A_t)$. Note that $M_2 = A_1 \setminus \{c_2\}$ and $M_T = B \setminus \{p_T\}$, so our goal is to get $M_2 \xrightarrow{\mathcal{P}} M_T$. See Figure 24, which completes Figure 23 and illustrates the full reduction to the poking game.

- First of all, we show that $M_2, \dots, M_T$ all contain a pair of adjacent coins. Let $2 \leq t \leq T$, we have $M_t = A_{t-1} \setminus \{c_t\}$. Since $p_{t-1} \in A_{t-1}$, we know $p_{t-1} \in M_t$: indeed, no coin is moved twice in a row hence $p_{t-1} \neq c_t$. Moreover p_{t-1} has two neighbors in A_{t-1} by the 2-adjacency rule, so it has one neighbor in M_t .
- Let $2 \leq t \leq T-1$, we check that $c_{t+1} \xrightarrow{\mathcal{P}} p_t$ is a valid poke from M_t to M_{t+1} .
 - We have $(M_t \setminus \{c_{t+1}\}) \cup \{p_t\} = A_t \setminus \{c_{t+1}\} = A_{t+1} \setminus \{p_{t+1}\} = M_{t+1}$ so the poke $c_{t+1} \xrightarrow{\mathcal{P}} p_t$ (if valid) does transform M_t into M_{t+1} .
 - We know $c_{t+1} \in A_t$, moreover $c_{t+1} \neq p_t$ since no coin is moved twice in a row, so $c_{t+1} \in M_t$.
 - Obviously $p_t \notin M_t$ by definition of M_t .
 - By Lemma 5.2, c_{t+1} is a neighbor of p_t .
 - It only remains to check that c_{t+1} is part of the pair of adjacent coins in M_t . But if it were not the case, then the (illegal) poke would create a second pair of adjacent coins i.e. M_{t+1} would contradict Proposition 1.15.

In conclusion, we have $M_2 \xrightarrow{c_3 \mapsto p_2} M_3 \xrightarrow{c_4 \mapsto p_3} \dots \xrightarrow{c_T \mapsto p_{T-1}} M_T$. The proof of the converse is similar: given the first move $A \xrightarrow{c_1 \mapsto p_1} A_1$ and a sequence of pokes $(c_t \xrightarrow{\mathcal{P}} p_t)_{2 \leq t \leq T}$ from $A_1 \setminus \{a\}$ to $B \setminus \{b\}$ for some suitable a and b , we check that the sequence of moves $A \xrightarrow{c_1 \mapsto p_1} A_1 \xrightarrow{a \mapsto p_2} A_2 \xrightarrow{c_2 \mapsto p_3} A_3 \xrightarrow{c_3 \mapsto p_4} \dots \xrightarrow{c_{T-1} \mapsto p_T} A_T \xrightarrow{c_T \mapsto b} B$ is valid. As for the last statement of this proposition, it comes from the fact there are: $O(N^2)$ possibilities for the first move, at most four possibilities for a (since p_1 has at most four neighbors), and at most three possibilities for b (indeed, given a redundant coin b_0 in B , the only coins other than b_0 that might be individually redundant in B are those from the pair of adjacent coins in the minimum configuration $B \setminus \{b_0\}$, because the others are isolated in $B \setminus \{b_0\}$ and thus have at most one neighbor in B). $\square$

5.2.2 The poking game on minimum chains

It is natural to study the simplest minimum configurations first, namely minimum chains. We start by detailing their structure: a minimum chain with span R containing two adjacent coins goes from one corner of R to its opposite without taking any detour (i.e. in a "straight line").




Figure 24. Top: a sequence of moves. Bottom: the corresponding reduction to the poking game, obtained by making the first move, deleting a , making a sequence of pokes, and then adding b .

Proposition 5.8. *Let $M = [c_1, \dots, c_N]$ be a minimum chain with span R containing a pair of adjacent coins $\{c_{i_0}, c_{i_0+1}\}$ for some $1 \leq i_0 \leq N - 1$. Then:*

- *For all $i \in \{1, \dots, N - 1\}$: $\text{dist}(c_i, c_{i+1}) = 2$ if $i \neq i_0$ and $\text{dist}(c_i, c_{i+1}) = 1$ if $i = i_0$.*
- *c_1 and c_N are opposite corners of R .*
- *If c_1 is the top-left corner of R , then for all $i \in \{1, \dots, N - 1\}$, c_{i+1} is either to the right, to the bottom, or to the bottom-right of c_i . The analogous assertion holds if c_1 is any other corner of R .*

Proof. Following the chain from c_1 to c_N , we consider the number of jumps that are made in each of the four directions. For instance, if for some $i \in \{1, \dots, N - 1\}$ the coin c_{i+1} is to the bottom-right of c_i , then the transition from c_i to c_{i+1} counts as one right jump plus one bottom jump, whereas if c_{i+1} is two positions under c_i then this transition counts as two bottom jumps. The first assertion of this proposition follows directly from the definition of a chain and Proposition 1.15, and further ensures that there are $2(N - 2) + 1 = m + n - 2$ jumps in total (indeed, recall that $N = \frac{m+n+1}{2}$ by Proposition 1.15). By Proposition 1.12, there exist $t, b, l, r \in \{1, \dots, N\}$ such that c_t (resp. c_b , resp. c_l , resp. c_r) is in the top row (resp. bottom row, resp. leftmost column, resp. rightmost column) of R . Without loss of generality, assume $t \leq b$ and $l \leq r$, then at least $n - 1$ bottom jumps are made from c_t to c_b and at least $m - 1$ right jumps are made from c_l to c_r . Since there are $m + n - 2$ jumps in total, we conclude that: exactly $n - 1$ bottom jumps are made from c_t to c_b , exactly $m - 1$ right jumps are made from c_l to c_r , and those cover all jumps made, hence the third assertion. Since no top jump or left jump is ever made, in particular there is no top jump between c_1 and c_t and no left jump between c_1 and c_l , therefore c_1 is the top-left corner of R . For the same reason, c_N is the bottom-right corner of R , which proves the second assertion. $\square$

Recall that one coin in hand is enough to flip an odd ‘L’. The following theorem, which generalizes this observation, is the central result of this section. In particular, it gives a formal proof that rotating an odd ‘L’ is impossible with a single coin in hand.

Theorem 5.9. *If M is a minimum chain with span R containing a pair of adjacent coins, then $M \xrightarrow{\mathcal{P}} M'$ if and only if M' is a minimum chain between the same two coins as M .*

Proof. Write $M = [c_1, \dots, c_N]$. We first show the necessity, and then we give a constructive

proof of the sufficiency.

- Let $i_0 \in \{1, \dots, N-1\}$ such that the pair of adjacent coins in M is $\{c_{i_0}, c_{i_0+1}\}$. We must show that, after a poke from M , we still obtain a chain between c_1 and c_N . Without loss of generality, assume that we poke c_{i_0} onto some position p : we want to show that $[c_1, \dots, c_{i_0-1}, p, c_{i_0+1}, \dots, c_N]$ is a proper chain.
 - We already know that, for all $i \in \{1, \dots, N-1\} \setminus \{i_0-1, i_0\}$: $\text{dist}(c_i, c_{i+1}) \in \{1, 2\}$.
 - Since p is a neighbor of c_{i_0} other than c_{i_0+1} , we have $\text{dist}(p, c_{i_0+1}) = 2$.
 - A consequence of the third assertion in Proposition 5.8 is that, for any $i \in \{1, \dots, N-1\} \setminus \{i_0\}$: $\text{dist}(c_{i_0}, c_i) = 2(i_0 - i)$ if $i < i_0$ and $\text{dist}(c_{i_0}, c_i) = 2(i - i_0) - 1$ if $i > i_0$. In particular, necessarily $i_0 \geq 2$ and c_{i_0} is poked towards c_{i_0-1} , because it is the only coin at distance exactly 2 from c_{i_0} . This means $\text{dist}(p, c_{i_0-1}) = 1$.

In conclusion, the configuration obtained after the poke is $[c_1, \dots, c_{i_0-1}, p, c_{i_0+1}, \dots, c_N]$, which is a chain between c_1 and c_N since $i_0 \notin \{1, N\}$. Moreover, it is minimum by cardinality, because the poke preserves the span according to Proposition 5.6.

- Suppose M' also forms a chain between c_1 and c_N , we want to show that $M \xrightarrow{\mathcal{P}} M'$. Without loss of generality, assume that $c = c_1$ is the top-left corner of $\text{span}(M)$. Let L be the ‘L’ with end-points c and c' that hugs the left and bottom sides of R and where the two adjacent coins are at the top-left. Since all pokes are reversible, it suffices to show that we can reach L : indeed, we would get $M \xleftrightarrow{\mathcal{P}} L \xleftrightarrow{\mathcal{P}} M'$. We prove this by induction on $|M|$. The case $|M| = 2$ is trivial. Now assume $|M| \geq 3$ and suppose the result holds for chains of cardinality lesser than $|M|$.
 1. Up to poking c_2 towards c_3 in the case where the pair of adjacent coins is $\{c_1, c_2\}$, we can assume that the chain $[c_2, \dots, c_N]$ contains two adjacent coins.
 2. We can now apply the induction hypothesis on the chain $[c_2, \dots, c_N]$.
 3. To finish, there are three possibilities depending on the direction of the jump from c_1 to c_2 . If c_2 is to the bottom of c_1 , then c_1 prolongs the ‘L’ we obtained by induction, so we just have to poke c_2 up towards c_1 to obtain the desired ‘L’. If c_2 is to the bottom-right (resp. right) of c_1 , then we proceed as in Figure 27 (resp. Figure 29). $\square$




Figure 25. The poking version of the leapfrog.

5.2.3 The poking game in general

We have just seen that, if the starting configuration M is a chain, then its end-points never move and all we can do is twist the string of coins that connects them. When playing the poking game on more general configurations, it quickly becomes apparent that certain coins are ever immobile and all we do is perform successive transformations of chains between some of them as in Theorem 5.9:




Figure 26. Subroutines used when c_2 is to the bottom-right of c_1 .




Figure 27. Case where c_2 is to the bottom-right of c_1 , using the subroutines from Figures 25 and 26. Top: even m , odd n . Bottom: odd m , even n .

- We start by transforming a certain chain inside a certain rectangle R_1 , until one of the two adjacent coins is at distance 2 from the first coin of a chain C coming out of R_1 .
- We can then poke towards this coin and get a new chain, inside a new rectangle R_2 , with C prolonging one half of the previous chain.
- We then transform this new chain, etc.

This is illustrated in Figure 30: in this example, all the action seems to happen inside three specific rectangles (R_1 , R_2 , R_3), with no possibility other than transitioning from one to the other via consecutive transformations of chains.

An incomplete study of the poking game in general is made in [Gal19], the results from which can be summarized as follows:

- First of all, Proposition 5.8 can be generalized so as to get a full understanding of the structure of minimum configurations with a pair of adjacent coins. In particular, further exploiting of the perimeter considerations glimpsed in the proof of Proposition 1.15 show that these configurations are forests in terms of the *2-connectivity* (connectivity at distance 1 or 2).
- A polynomial-time algorithm returns, given a configuration M , a specific set $F(M)$ of coins in M (in Figure 30, these are the shaded coins). These coins are forever immobile, and all that is possible is to perform successive transformations of chains between some coins in $F(M)$. The set $F(\cdot)$ is invariant throughout the pokes.
- (If M is a chain, recall that we already have Theorem 5.9.)
- If M is a tree in terms of the 2-connectivity, then $M \xrightarrow{\mathcal{P}} M'$ if and only if $F(M) = F(M')$.




Figure 28. Subroutines used when c_2 is to the right of c_1 .




Figure 29. Case where c_2 is to the right of c_1 , using the subroutines from Figures 25 and 28. Top: even m , odd n . Bottom: odd m , even n .

- If M is a forest in terms of the 2-connectivity (general case), then this does not work anymore. We conjecture that a refined algorithm provides us with an invariant that characterizes solvable puzzles, similarly to the case of trees.

Conclusion

We have made some contributions to the study of the coin-moving game initiated in [DDV02]. Our focus has been on configurations with either 2 extra coins or 1 extra coin, as suggested by the authors. Indeed, we use their canonicalization technique, which effectively nullifies the extra coins beyond the first two. Stepping outside of this framework however, more extra coins might be useful, so finding necessary and sufficient conditions in the case of 3 or more extra coins remains an open problem. For the same reason, we have not explored the case of 2 extra coins and just 1 redundant coin.

In the case where there are 2 extra coins in A and 2 redundant coins in B , Theorem 4.8 gives a sufficient condition that involves the total number of coins and the perimeters of the spans of A and B . The worst-case puzzles described in Corollary 4.3 show that the bound from the theorem is tight in the case of a square span, up to an additive factor $O(1)$.




Figure 30. Playing the poking game. The outlined rectangle shows the chain that is currently being transformed. The shaded coins seem unmovable.

However, for a rectangular $m \times n$ span in general, it looks like discrepancy between m and n tends to make puzzles easier. For instance, and provided we have the inclusion of spans, if $n = 1$ then it seems easy to show that 1 extra coin and 1 redundant coin are always sufficient, and if $n = 2$ then we suspect that 2 extra coins and 2 redundant coins are always sufficient. In general, we think that it could be possible to improve Theorem 4.8 by using only $\frac{\min(m,n)}{2} + O(1)$ additional coins.

We now know of two methods to solve general puzzles with 2 extra coins in A and 2 redundant coins in B : the one from Theorem 3.1 consists in going from A to ‘L’s then reverse into B , while the one from Theorem 4.8 consists in going from A to an ‘L’ then sweep across the board and drop coins to create B . It would be good to design other solving methods that would apply to some puzzles that do not meet the conditions of either theorem, such as the puzzle on the left of Figure 11.

As for the case of a lone extra coin, we have focused on the case where A and B are minimum+1 configurations with the same span, for which the game can be reduced to a poking game. The case of chains is completely solved. A deeper study is made in [Gal19]: a detailed description of the structure of minimum configurations is given, and allows for further results as well as leads towards a general solution. On the contrary, nothing is known about the structure of minimal configurations, so the general *minimal+1* case remains obscure. Nevertheless, it is interesting to note that the poking game can be extended to that case (see [Gal19] for the details).

Finally, it is easy to check that all proofs of our sufficiency results are constructive and provide polynomial-time algorithms to find a winning sequence of moves of polynomial length. However, the complexity of the game as a whole remains unknown.

Acknowledgments

We would like to thank the reviewers for providing us with helpful feedback and inspiring the abstract of this paper.

References

- [BCG82] Elwyn R. BERLEKAMP, John H. CONWAY and Richard K. GUY, *Winning Ways*, vol. 2, pp. 755-756. Academic Press, London, 1982.

- [DDV02] Erik D. DEMAINE, Martin L. DEMAINE and Helena A. VERILL, "Coin-Moving Puzzles", *More Games of No Chance*, pp. 405–431. Cambridge University Press, 2002. Collection of papers from the MSRI Combinatorial Game Theory Research Workshop, Berkeley, California, July 24–28, 2000.
- [DP06] Adrian DUMITRESCU and János PACH, "Pushing squares around", *Graphs and Combinatorics*, vol. 22, pp. 37–50, 2006.
- [Gal19] Florian GALLIOT, "A coin-moving game on graphs", Master's thesis, 2019. Available at <https://dumas.ccsd.cnrs.fr/dumas-03160998>.
- [Gar75] Martin GARDNER, *Mathematical Carnival*, chapter 2 "Penny Puzzles", pp. 12–26. Alfred A. Knopf, New York, 1975.
- [Kva86] "Problems from the 49th Moscow Mathematical Olympiad" (in Russian), *Kvant*, N°9, p. 57, 1986.
- [Lan51] Harry LANGMAN, "Curiosa 261: A disc puzzle", *Scripta Mathematica*, 17(1–2):144, March–June 1951.
- [Lan53] Harry LANGMAN, "Curiosa 342: Easy but not obvious", *Scripta Mathematica*, 19(4):242, December 1953.