On 3-terminal positions in Hex

Eric Demer, UCLA
Peter Selinger, Dalhousie University

Abstract

This paper is about 3-terminal regions in Hex. A 3-terminal region is a region of the Hex board that is completely surrounded by black and white stones, in such a way that the black boundary stones form 3 connected components. We characterize Hex as the universal planar Shannon game of degree 3. This ensures that every Hex position can be decomposed into 3-terminal regions. We then investigate the combinatorial game theory of 3-terminal regions. We show that there are infinitely many distinct Hex-realizable values for such regions. We introduce an infinite family of 3-terminal positions called superswitches and investigate their properties. We also present a database of Hex-realizable 3-terminal values, and illustrate its utility as a problem-solving tool by giving various applications. The applications include the automated verification of connects-both templates and pivoting templates, a new handicap strategy for 11×11 Hex, and a method for constructing witnesses for the non-inferiority of probes in many Hex templates. These methods allow us to disprove a conjecture by Henderson and Hayward.

1 Introduction

Hex is a perfect information game for two players. It was invented by Piet Hein in 1942 [7], and is usually played on an $n \times n$ rhombic grid of hexagonal cells like this:

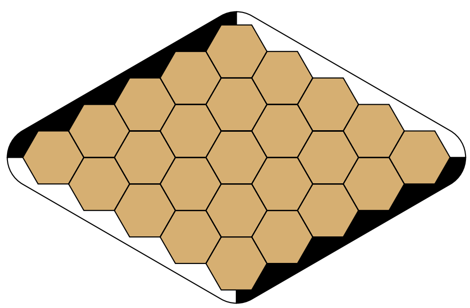

(1)

Two opposing board edges are colored black, and the other two are colored white. The players, called Black and White, alternately place a stone of their color on an empty cell, with Black going first. The winner is the player who connects their two board edges. It is an interesting property of Hex that draws are not possible: there is always exactly one winner [7]. Another interesting property is that the first player has a theoretical winning strategy on all boards of size $n \times n$; however, the proof is non-constructive, and no concrete winning strategy is known except for some very small board sizes [13].

In principle, the game ends as soon as one player connects their edges. In practice, most actual games end long before this happens, with the losing player resigning. However, for theoretical purposes, it is often simpler to assume that the game continues until the board is completely filled; since the surplus moves cannot change the winner, this assumption is without loss of generality. Also, since Black has a theoretical winning strategy and a significant practical advantage, in competitive play the *swap rule* is used: right after Black makes the first move, White has the option to switch colors. The swap rule makes the game more fair, but since it only affects the first two moves of the game, it is not relevant to most of the results of this paper. Except for one application in Section 6.5, we do not consider the swap rule here.

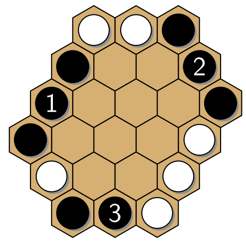

Figure 1: A 3-terminal region

Combinatorial game theory is a general theory of two-player sequential games with perfect information. It was introduced by Conway [2] and Berlekamp, Conway, and Guy [1]. Combinatorial game theory was initially developed for *normal play* games, in which a player loses when they cannot make a move. Because Hex is not a normal play game, it requires an adaptation of combinatorial game theory; such an adaptation was given in [15].

This paper is about 3-terminal regions in Hex. A *region* of the Hex board is just a subset of its cells. A region that is completely surrounded by black and white stones is called an *n-terminal region* when its boundary has n black components and n white components. We call these boundary components the black and white *terminals*. Figure 1 shows an example of a 3-terminal region.

When Hex is played on the whole board, there are only two possible outcomes at the end of a game: either Black wins or White wins. However, when we consider play inside a 3-terminal region, the situation is more complicated. There are five possible outcomes within the region: either Black connects all three black terminals, or White connects all three white terminals, or Black and White each connect exactly two of their terminals, which can happen in three different ways.

Outline of the paper

In Section 2, we explain why 3-terminal regions (and not, say, 4-terminal regions or 5-terminal regions) are fundamental for Hex: we show that every Hex position can be decomposed into 3-terminal regions. We do this by showing that Hex is equivalent to a class of games called planar Shannon games of degree 3. Section 2 does not require combinatorial game theory, but the rest of the paper does.

In Section 3, we provide background material on some basic definitions and results from the combinatorial game theory of Hex.

In Section 4, we introduce some interesting families of 3-terminal positions called *superswitches*, *simpleswitches*, and *tripleswitches*. We use this to prove that there are infinitely many non-equivalent 3-terminal Hex positions. As an application, we show how these switches can be used, in conjunction with a Hex solver, to verify a certain kind of Hex template called a *connects-both template*.

In Section 5, we present a database of Hex-realizable 3-terminal values. We explain how we generated the database and how it can be used.

In Section 6, we apply our theory of 3-terminal positions to the verification of another kind of Hex template, called a *pivoting template*. We characterize both sente and gote pivoting templates in terms of combinatorial game theory, and we use the database from Section 5 to find a particular 3-terminal context that can be used, in conjunction with a Hex solver, to verify pivoting templates. As another application, we describe a new handicap winning strategy for 11×11 Hex.

Finally, in Section 7, we use our theory to solve an open problem. A move in a Hex region is called *inferior* if it can never be the unique winning move. Henderson and Hayward [8] conjectured that 5 of the 8 possible intrusions into the region called the *ziggurat* or *4-3-2 edge template* are inferior. We disprove the conjecture by exhibiting specific Hex positions in which each of the 8 intrusions is the unique winning move. We describe a general method for computing such witnessing positions, and also settle the (non-)inferiority of intrusions into many other Hex templates.

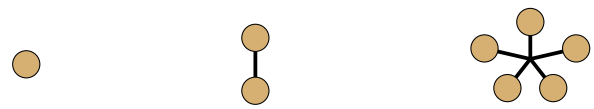

Figure 2: A vertex, an edge, and a hyperedge.

2 Hex is the universal 3-planar Shannon game

In this section, we prove that Hex has a nice mathematical characterization as the universal 3-planar Shannon game. We start by explaining what this means.

2.1 Set coloring games and Shannon games

We first define a general class of games called *set coloring games*. When X and Y are sets, we write Y^X for the set of all functions from X to Y.

Definition 2.1. Let $\mathbb{B} = \{\text{black}, \text{white}\}$. A *set coloring game* over $\mathbb{B}$ is given by a finite set X, whose elements are called *cells*, and a function $\pi : \mathbb{B}^X \to \mathbb{B}$, called the *payoff function*. It is played as follows: there is a game board, initially empty, whose cells are in one-to-one correspondence with the elements of X. The players, Black and White, take turns choosing an empty cell and placing a stone of their color on it. The game ends when the board is completely filled, at which point the coloring of the cells determines a function $f : X \to \mathbb{B}$. The winner is determined by the payoff function: Black wins if $\pi(f) = \text{black}$, and White wins otherwise.

We say that two set coloring games (X, π) and (X', π') are *isomorphic* if there is a bijection $\phi : X' \to X$ respecting the payoff function, i.e., such that for all $f : X \to \mathbb{B}$, we have $\pi(f) = \pi'(f \circ \phi)$. A *position* in a set coloring game is an assignment of black and/or white stones to some subset of the cells. Each such position can itself be regarded as a set coloring game on the remaining empty cells.

Hex is evidently a set coloring game. It will also be useful to consider another class of set coloring games, the so-called *Shannon games*. The Shannon games we consider in this paper are vertex Shannon games on hypergraphs. Recall that a *hypergraph* is a pair (V, E), where V is a set whose elements are called *vertices*, and E is a set of subsets of V, called *hyperedges*. A hypergraph where each hyperedge contains exactly two vertices is also called a *graph*, and in this case, the hyperedges are usually just called *edges*. In pictures, we will represent vertices as circles, edges as lines connecting two vertices, and hyperedges as clusters of lines joining any number of vertices, as in Figure 2.

The Shannon game is played on a finite hypergraph with two distinguished vertices called *terminals*. The cells are the non-terminal vertices. As usual, the players take turns coloring the cells, and at the end of the game, Black wins if and only if the two terminals are connected by an uninterrupted path of black vertices, i.e., if there exists a sequence of vertices $v_1, \ldots, v_n$ such that v_1 and v_n are the terminals, each v_i is connected to v_{i+1} by a hyperedge, and $v_2, \ldots, v_{n-1}$ are colored black. If Black does not have such a connection, then White wins. We note in passing that the class of Shannon games is not self-dual, i.e., exchanging the roles of Black and White in a Shannon game does not necessarily result in a Shannon game.

To see why Hex is a special case of the Shannon game, consider the following graph, which corresponds to 5×5 Hex:

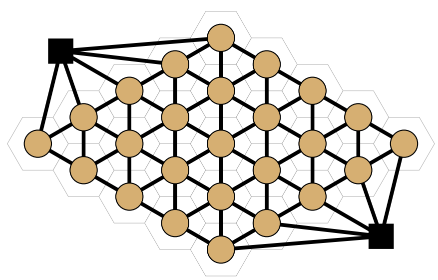

(2)

Here, the brown circles represent the non-terminal vertices and correspond to the cells of the Hex board, which are faintly shown in the background. The two black squares are the terminals and correspond to Black's board edges. Note that two vertices are adjacent in the graph if and only if the corresponding cells or board edges are adjacent on the Hex board. Therefore, the games in (1) and (2) are isomorphic.

A Shannon game is *planar* if its underlying hypergraph is planar, with terminals on the outside. Recall that a hypergraph is *planar* if it can be drawn in the 2-dimensional plane without any hyperedges crossing. The spaces between the hyperedges are called *faces*, and every finite planar hypergraph has an *outside face*, i.e., a face that stretches to infinity. When we require the terminals to be *on the outside*, we mean that they must be adjacent to this outside face. For example, the Shannon game in (2) is planar.

The *degree* of a vertex in a hypergraph is the number of hyperedges that contain it. We say that a Shannon game has degree n if every vertex has degree at most n. For example, the game in (2) has degree 6. If a game is planar and of degree n, we also say that the game is n-*planar*.

Remark 2.2. In the literature, vertex Shannon games are usually played on graphs, rather than on hypergraphs. The reason we consider hypergraphs instead of graphs is that it yields a more convenient definition of planarity and degree. However, for games of degree at least 3, the two definitions are equivalent if we permit games where some of the cells have already been colored at the start. Here is an example:

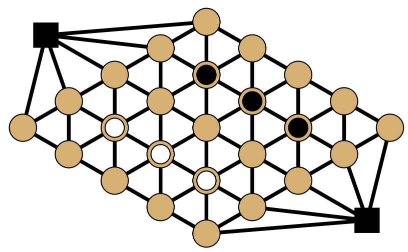

(3)

In this game, some of the cells are already occupied by black or white stones. Such a game can be simplified by deleting white vertices and contracting black (non-terminal) vertices. When deleting a white vertex, we also remove the vertex from any hyperedges that contained it. To contract a black vertex, we remove the vertex, and merge all of its incident hyperedges into a single hyperedge. We do this repeatedly until no more black or white cells are left. This results in an isomorphic game; for example, the position in (3) is isomorphic to the following:

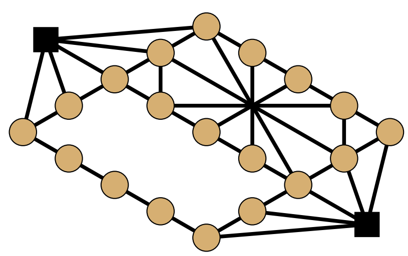

(4)

Note that the operations of deleting a white vertex or contracting a black one preserve the planarity of the hypergraph, and do not increase the degree. (Remember that the degree measures the number of hyperedges at a given vertex, not the number of vertices at a given hyperedge). Therefore, any position in a (planar) graph Shannon game is isomorphic to a (planar) hypergraph Shannon game of the same degree.

Conversely, any (planar) Shannon hypergraph game of degree 3 or higher is isomorphic to a position in a (planar) Shannon graph game of the same degree in which some cells have possibly already been colored black. Namely, each hyperedge with k vertices, as shown on the left, can be replaced by a cluster of $k-2$ additional black-colored cells, as shown on the right:

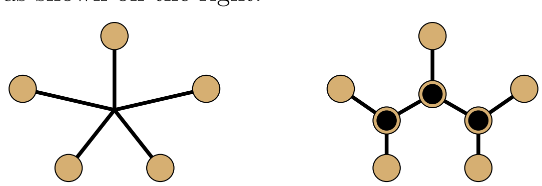

2.2 Hex is the universal 3-planar Shannon game

We are now ready to state the main theorem of this section:

Theorem 2.3. *Every 3-planar Shannon game is isomorphic to a Hex position. Conversely, every Hex position is isomorphic to a 3-planar Shannon game.*

The reader is invited to try to prove this theorem now, before looking at the proof. Note that the second part is not obvious, since Hex, as shown in (2), appears to be of degree 6, rather than 3. Although the proof is extremely simple, to our knowledge, this result was not previously known.

Proof of Theorem 2.3. For the first claim, consider any 3-planar Shannon game, such as the one shown in Figure 3(a). It is obvious that this game can be embedded in a sufficiently large Hex board in such a way that every non-terminal vertex corresponds to an empty cell, every terminal vertex corresponds to a black board edge, every hyperedge corresponds to a connected group of black cells, and everything else (i.e., the space between the hyperedges) is filled with white cells. An example of such an embedding is shown in Figure 3(b). Specifically, since each vertex has degree at most 3, there are at most 3 hyperedges that must be attached to it, and this is possible because the corresponding cell on the Hex board is a hexagon. Also, since the terminals are on the outside of the hypergraph, they can be connected to opposite black board edges.

For the converse, refer to Figure 4. Consider a Hex board, such as the one in (a). Insert black and white polygons between the cells, as shown in (b). Observe that for any pair of adjacent cells of (a), the corresponding cells in (b) are connected by both a black and a white polygon. Therefore, for any position in which Black's cells are connected in (a), Black's cells are also connected in (b), and dually for White. It follows that the game in (a) is isomorphic to the one in (b). Moreover, in (b), it is obvious that each cell has degree at most 3. In fact, the game in (b) is isomorphic to the 3-planar Shannon game shown in (c). This works for boards of any size. Moreover, if we start from a Hex position, rather than an empty board, we can color the corresponding vertices in (c), and then delete the white ones and contract the black ones as in Remark 2.2, which still yields a 3-planar Shannon game. □

Remark 2.4. Due to Theorem 2.3, every problem about Hex positions is equivalent to a problem about 3-planar Shannon games, and vice versa. But since Hex is usually played starting from an empty board, and not from some arbitrary initial position, there are many elements of Hex strategy that are specific to Hex, such as templates, ladders, etc. In other words, a good Hex player is not necessarily a good player of arbitrary 3-planar Shannon games. Naturally, positions such as the one in Figure 3 do not normally arise in actual Hex games.

2.3 Application: Embedding Hex in itself

The proof of Theorem 2.3 immediately implies that Hex can be isomorphically embedded in a larger version of itself. For example, Figure 5 shows two different ways of embedding 4×4 Hex in 10×10 Hex. In fact, the second of these embeddings is an iterated version of the first.

2.4 Bridg-It is the universal 2-planar Shannon game

Bridg-It is a connection game that was invented by David Gale. It was described in Martin Gardner's 1958 Scientific American column [5] under the name "Gale". It is played on a square board such as the one shown in Figure 6(a).

On Black's turn, Black connects two adjacent black dots by a straight line. On White's turn, White similarly connects two adjacent white dots. Crossing lines are not permitted. Black's goal is to connect the black dots on the left with the ones on the right, and White's goal is to connect the white dots at the top with the ones at the bottom. Figure 6(b) shows what the game may look like after a few moves.

Like Hex, Bridg-It has the property that one player always wins: draws are not possible. In fact, Bridg-It is a 2-planar Shannon game. Its hypergraph is shown in Figure 6(c). Note that the black dots of

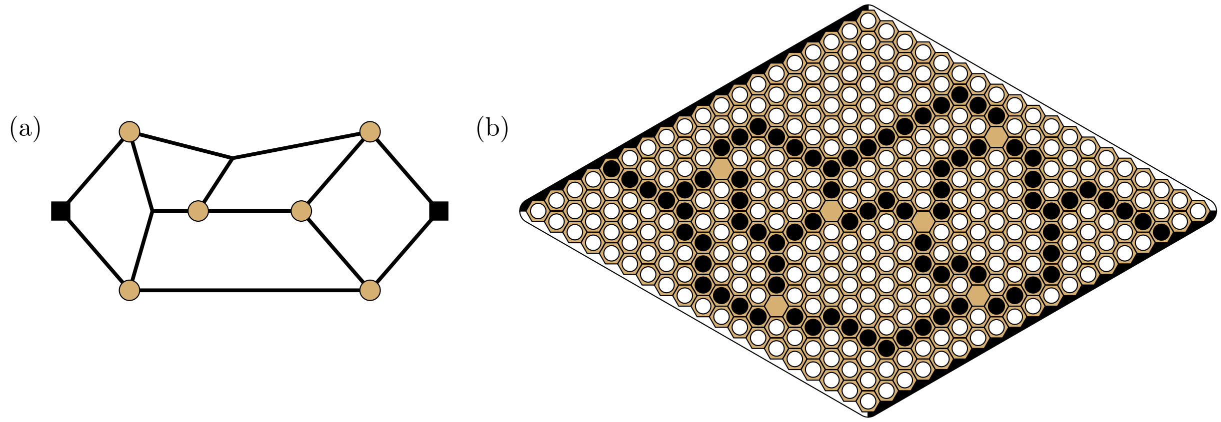

Figure 3: Converting a 3-planar Shannon game to a Hex position.

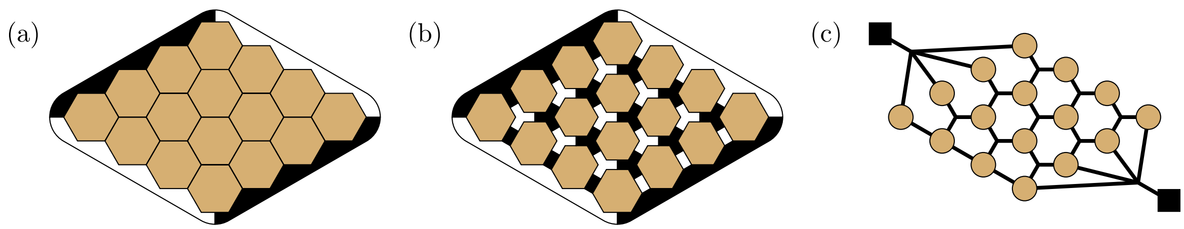

Figure 4: Converting a Hex position to a 3-planar Shannon game.

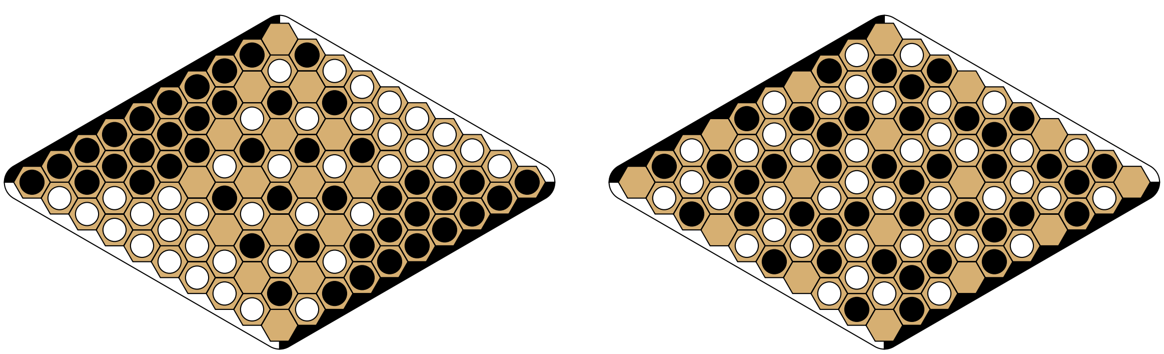

Figure 5: Two different embeddings of 4×4 Hex in 10×10 Hex.

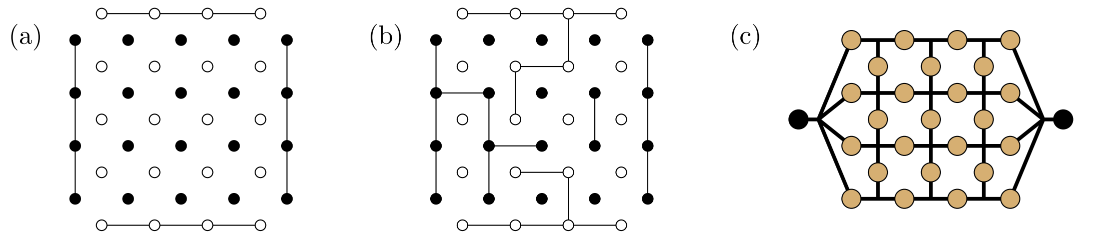

Figure 6: (a) The game of Bridg-It. (b) A partially completed game. (c) An isomorphic Shannon game.

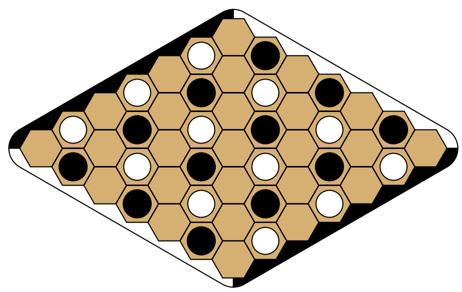

Figure 7: A Hex realization of Bridg-It

the Bridg-It board become hyperedges of the Shannon game, and the potential black edges of the Bridg-It board become vertices. Moreover, Bridg-It is the universal 2-planar Shannon game. On the one hand, the game in Figure 6(c) is obviously of degree 2. On the other hand, any 2-planar Shannon game can be embedded on a sufficiently large Bridg-It board, in a way analogous to how it was done for Hex in Figure 3.

Since every planar Shannon game of degree 2 is also of degree 3, it follows that Bridg-It can be isomorphically embedded in Hex. Indeed, Figure 7 shows a Hex realization of the game from Figure 6(a). This realization previously appeared in [6, Fig. 6.5].

The converse is not true: Hex cannot be isomorphically embedded in Bridg-It. In fact, there is a qualitative difference between the two games. In Bridg-It, there exists a computationally efficient algorithm for determining the winner of any position [12]. The corresponding problem in Hex is PSPACE-complete [14]. Therefore, the passage from degree 2 to degree 3 is non-trivial.

We finish this subsection by remarking that there is a variant of Shannon games called *edge Shannon games*. These games are played on a graph, and the players alternately color *edges*, rather than vertices. As usual, Black wins if at the end of the game, two distinguished terminals are connected by black edges. Note that every hypergraph has a *dual*, where the vertices of the dual hypergraph are the hyperedges of the original hypergraph, and vice versa. A hypergraph of degree 2 is just the dual of an ordinary graph. In this way, an edge Shannon game is essentially the same thing as a hypergraph vertex Shannon game of degree 2, and Bridg-It is the universal planar edge Shannon game. In fact, like Bridg-It, the edge Shannon game is also solved: there exists an efficient algorithm for determining which player has a winning strategy [12]

2.5 Planar Shannon games of degree greater than 3

We just saw that Bridg-It is the universal 2-planar Shannon game, Hex is the universal 3-planar Shannon game, and the passage from degree 2 to degree 3 has a dramatic effect on the complexity of planar Shannon games. It is therefore natural to ask whether games of degree 4 or higher get even more complex or interesting. We may also ask what class of games is universal for n-planar Shannon games when $n \geqslant 4$.

The answer to the first question is no: Hex is already PSPACE-complete, and the problem of deciding a winning strategy for any Shannon game is in PSPACE. Therefore, planar Shannon games of degree 4 and

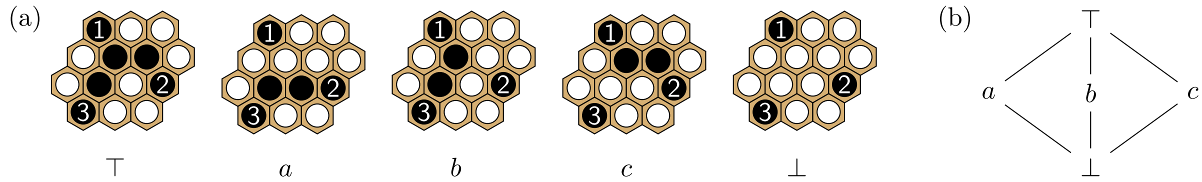

Figure 8: (a) The five possible outcomes of a 3-terminal region. (b) The outcome poset for 3-terminal regions.

higher are not more complex than Hex from a computational point of view.

The answer to the second question is also easy: a universal planar Shannon game of degree 4 is Hex with octagons. More precisely, it is not necessary for *all* cells to be octagons; there just need to be an unbounded number of them as the board size increases. One way to make a (topological) octagon on a Hex board is to merge two adjacent cells into a single cell. Having an unbounded number of octagons is then easily seen to be sufficient for embedding any planar Shannon game of degree 4, in a manner analogous to Figure 3. Similarly, a universal planar Shannon game of degree 5 is Hex with sufficiently many decagons, and so on.

3 A quick primer on the combinatorial game theory of Hex

The fundamentals of the combinatorial game theory of Hex were developed in [15]. Here, we briefly summarize some definitions and results that are useful in the rest of this paper. For a more complete and thorough introduction, see [15]. For combinatorial game theory in general, see [2, 1]. In this section, we will mostly be talking about Hex, but the same concepts also apply to more general Shannon games.

3.1 Regions and outcome posets

As mentioned in the introduction, a *region* of a Hex board is a subset of its cells; some or all of the cells may already be occupied by black and/or white stones. By a *completion* of the region, we mean any possible way of filling its empty cells with black and/or white stones. If a and b are completions of the same region, we say that $a \leqslant b$ if for every way of filling the rest of the board (i.e., the outside of the region) with stones, if a is a win for Black then so is b. We say that a and b are *equivalent* if $a \leqslant b$ and $b \leqslant a$. An *outcome* for the region is an equivalence class of its completions. The outcomes are partially ordered via $\leqslant$, and we refer to this partial order as the *outcome poset* for the region.

The outcome poset for an arbitrary Hex region can in general be very complicated. But in this paper, we are primarily interested in 3-terminal regions, and as mentioned in the introduction, there are only five possible outcomes for such a region. Specifically, if the region's black terminals are labelled 1, 2, and 3, we name the outcomes as follows:

- $\bot$: None of the black terminals are connected.
- a: Terminals 2 and 3 are connected, but terminal 1 is not.
- b: Terminals 1 and 3 are connected, but terminal 2 is not.
- c: Terminals 1 and 2 are connected, but terminal 3 is not.
- $\top$: All of the black terminals are connected.

An example of each outcome is shown in Figure 8(a). Let $P_3 = \{\bot, a, b, c, \top\}$ be this outcome poset, with the partial order shown in Figure 8(b).

G =

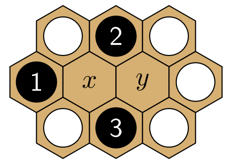

=

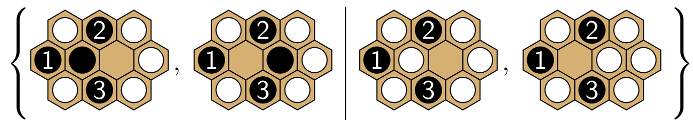

$$= \quad \{ \quad \{\top \mid \top\}, \qquad \{\top \mid a\} \quad \mid \quad \{a \mid \bot\}, \qquad \{\top \mid \bot\} \quad \}$$

Figure 9: The game form of a Hex position

3.2 Game forms

An outcome represents the state of a region when it is completely filled with stones, i.e., at the end of the game. To understand the value of a region that still has some empty cells in it, we need the concept of an abstract game form.

Definition 3.1 (Game forms over an outcome poset [15, Def. 4.1])**.** Fix an outcome poset A. The *game forms over* A are defined as follows:

- Whenever $a \in A$, then $[a]$ is a game form, called an *atomic* game.
- Whenever L and R are non-empty sets of game forms, $\{L \mid R\}$ is a game form, called a *composite* game.

Moreover, the class of game forms is the smallest class generated by the above two rules.

We often shorten “game form” to “game” when there is no potential for confusion. We use some customary notations from combinatorial game theory. For an atomic game, we often write a instead of $[a]$. If $G = \{L \mid R\}$ is a composite game, the elements of L and R are called the *left* and *right options* of G, respectively. The represent the positions that the left, respectively right, player can reach in the next move. Atomic games have no options. If G is any game, we write G^L for a typical left option and G^R for a typical right option of G. Thus, when we write “for all G^L”, we mean “for all left options G^L of G”; such a statement is vacuously true when G is atomic. Similarly, “there exists G^L” means “there exists a left option G^L of G”, and such a statement is trivially false when G is atomic. We also use the notation $G^{(L)}$ to mean G^L if G is composite, and G itself if G is atomic, and similarly for $G^{(R)}$.

The *followers* of a game G are G itself, all of its options, the options of the options, and so on. A game is *short* if it has finitely many followers. In this paper, we only consider short games, because all Hex positions are naturally short. The *depth* of a short game is defined in the obvious way: atomic games have depth 0, and the depth of a composite game is one more than the maximum of the depths of its options.

Every Hex region can be converted to a game form over its outcome poset. This may best be seen in an example. Consider the region in Figure 9. We identify Black with Left and White with Right. The game G has two left options, corresponding to Black making a move at x or y. The game G also has two right options, corresponding to White making a move at x or y. The resulting positions can be recursively converted to game forms, with atomic positions acting as the base case. For example, if Black occupies both x and y, all three black terminals are connected, so the outcome is $\top$. If Black occupies y and White occupies x, terminals 2 and 3 are connected but terminal 1 is not, so the outcome is a, etc. Thus, by recursively expanding G, we find that its game form is $\{\{\top \mid \top\}, \{\top \mid a\} \mid \{a \mid \bot\}, \{\top \mid \bot\}\}$.

Note that, as always in combinatorial game theory, we do not enforce alternating turns. This is because if Black makes a move in the region, White may choose to make a move elsewhere on the board (outside the region), and then Black may make another move in the region. Thus, although the players alternate in the global game, it is possible for a player to make consecutive moves within a region.

3.3 The order relations

To be able to compare game forms, we define two relations $\leqslant$ and $\lhd$. Intuitively, $G \leqslant H$ means that H is at least as good for Black as G, whereas $G \lhd H$ means that moving first in H is at least as good for Black as moving second in G.

Definition 3.2 (Order relations [15, Def. 4.3])**.** The relations $\leqslant$ and $\lhd$ on game forms are defined by mutual recursion as follows:

- $G \leqslant H$ if both of the following hold:
  - for all $G^{(L)}$, $G^{(L)} \lhd H$, and
  - for all $H^{(R)}$, $G \lhd H^{(R)}$.
- $G \lhd H$ if at least one of the following holds:
  - there exists G^R with $G^R \leqslant H$, or
  - there exists H^L with $G \leqslant H^L$, or
  - $G = [a]$ and $H = [b]$ are both atomic and $a \leqslant b$.

We say that two game forms G, H (over the same outcome poset) are *equivalent*, in symbols $G \simeq H$, if $G \leqslant H$ and $H \leqslant G$. Many useful properties of the order and of equivalence are proved in [15]. Among them are several transitivity properties [15, Lem. 4.6]:

- If $G \leqslant H$ and $H \leqslant K$, then $G \leqslant K$.
- If $G \lhd H$ and $H \leqslant K$, then $G \lhd K$.
- If $G \leqslant H$ and $H \lhd K$, then $G \lhd K$.

For short games, each equivalence class has a unique simplest member called the game's *canonical form* [15, Lem. 4.23]. Moreover, canonical forms can be efficiently computed. From now on, by a *value*, we mean an equivalence class of game forms. Values are usually given in canonical form.

3.4 Monotone and passable games

Among all of the game forms, those that are realizable as Hex positions have an important additional property. In Hex, an additional black stone can only improve Black's position, and dually, an additional white stone can only improve White's position. This is the property of *monotonicity*, which we can formulate for game forms as follows: for all G^L and G^R, we have $G^R \leqslant G \leqslant G^L$. Unfortunately, it turns out that the class of monotone game forms is not very robust; for example, the canonical form of a monotone game need not be monotone (see [15, Example 4.29]). This problem is solved by introducing the class of *passable* games in the following definition.

Definition 3.3 (Monotone and passable games [15, Defs. 4.24 and 6.2])**.** A game form G is called *monotone* if all left options satisfy $G \leqslant G^L$, all right options satisfy $G^R \leqslant G$, and recursively all options are monotone. A game form G is called *passable* if $G \lhd G$ and recursively all options are passable.

The concept of passable games is justified by the following theorem.

Theorem 3.4 (Fundamental theorem of monotone games [15, Cor. 6.6])**.** *Let A be an atom poset with top and bottom elements, and let G be a short game over A. Then G is equivalent to a monotone game if and only if G is equivalent to a passable game, if and only if the canonical form of G is passable.*

The following lemma is simple but useful.

Lemma 3.5. *If G is monotone and composite and all G^R satisfy $\top \lhd G^R$, then $\top \leqslant G$.*

Proof. To show $\top \leqslant G$, we only need to show two things: first, that all G^R satisfy $\top \lhd G^R$, which holds by assumption, and second that $\top \lhd G$. Since G is composite, it has some right option G^R, and since G is monotone, we have $G^R \leqslant G$. By assumption, $\top \lhd G^R \leqslant G$, which implies $\top \lhd G$ by transitivity. □

3.5 Disjunctive sums

We will also need the notion of disjunctive sum of games. Informally, the sum of two or more Hex positions is obtained by putting them side by side and combining them into a larger position in a prescribed way. For an example, consider equation 5 below, which illustrates one particular way of combining two 3-terminal positions into a single 3-terminal position.

Because there can be multiple different ways of summing two games, our definition of the sum of abstract game forms is parameterized by a function f that determines how to combine atomic outcomes. For example, Figure 11 illustrates several cases of how the sum operation of equation 5 acts on atomic outcomes. We have the following definition for the sum of game forms:

Definition 3.6 (Sum)**.** Let A, B, C be posets and let $f : A \times B \to C$ be a monotone function, i.e., such that $a \leqslant a'$ and $b \leqslant b'$ implies $f(a, b) \leqslant f(a', b')$. Given a game G over A and a game H over B, their sum $G +_f H$ is a game over C, defined recursively as follows:

- $G +_f H = \{G^L +_f H, G +_f H^L \mid G^R +_f H, G +_f H^R\}$, when at least one of G or H is composite, and
- $[a] +_f [b] = [f(a, b)]$.

Note that the situation is different from normal play games [2, 1], where there is a single disjunctive sum operation. In our setting, there are many alternative ways of summing games, depending on the function f.

3.6 Special 3-terminal regions: corners and forks

As mentioned in Section 3.1, the outcome poset for a general 3-terminal region is $P_3 = \{\bot, a, b, c, \top\}$ with the partial order shown in Figure 8(b). However, some special kinds of 3-terminal regions have more specialized outcome posets. One way in which this happens is when two or more of the region's terminals are board edges. A 3-terminal region whose terminals include adjacent black and white board edges is called a *corner*, as in Figure 10(a). A 3-terminal region whose terminals include two opposite white board edges is called a *fork*, as in Figure 10(b). As before, outcome a means that terminals 2 and 3 are connected, outcome b means that terminals 1 and 3 are connected, and so on. But compared to a generic 3-terminal region, a corner has the additional property that $c \leqslant a$. Indeed, the rest of the board is a 4-terminal region, and the reader can verify that for each of the 14 possible outcomes for the rest of the board, if c is winning for Black in the corner, then so is a. For a fork, the situation is even more constrained: in this case, outcome c is equivalent to $\bot$, since in both cases, White's edges are connected within the region and White wins regardless of what happens on the rest of the board. The outcome posets for a corner and a fork are shown in Figure 10. Note that both of these posets are quotients of P_3.

4 Infinitely many non-equivalent 3-terminal positions

Because there are only five possible outcomes for a 3-terminal position, one may be tempted to think that there is not very much going on in such a position. Maybe there are only finitely many 3-terminal positions up to equivalence? In this section, we show that this is not the case. We will consider several infinite families of 3-terminal positions with interesting properties.

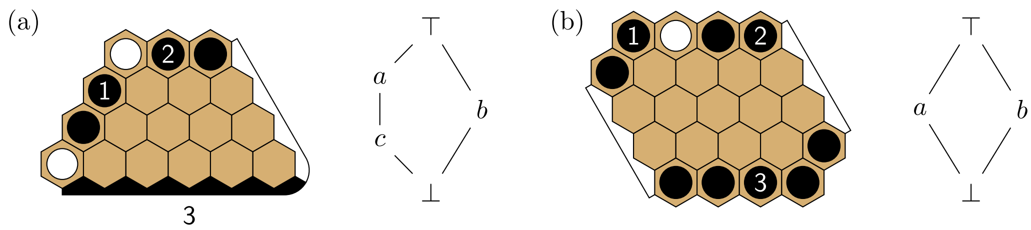

Figure 10: (a) A corner and its outcome poset. (b) A fork and its outcome poset.

4.1 Superswitches

Consider an outcome poset with two incomparable atoms a and b. It was shown in [15, Prop. 10.2] that there are infinitely many passable game values over these outcomes. Specifically, it was shown that the sequence of games defined by $G_0 = a$ and $G_{n+1} = \{a, b \mid G_n\}$ for all $n \geqslant 0$ is an infinite, strictly increasing sequence of passable game values. The first few values in the sequence are:

$$\begin{aligned}
&G_0 = a, \\
&G_1 = \{a, b \mid a\}, \\
&G_2 = \{a, b \mid \{a, b \mid a\}\}, \\
&G_3 = \{a, b \mid \{a, b \mid \{a, b \mid a\}\}\}, \\
&\dots
\end{aligned}$$

We call these games *superswitches*. The idea is that the default outcome is initially a, but the left player gets n chances to change ("switch") the outcome to b. In other words, even if the right player gets $n-1$ moves first, the left player can still choose between outcomes a and b.

The question was left open in [15] whether the superswitches are realizable as Hex positions, and if so, whether they are still distinct when regarded as Hex positions. We give positive answers to these questions below.

As before, we assume that the black terminals of all 3-terminal positions are numbered $1, 2, 3$. Recall that we write a for the outcome "Black connects terminals 2 and 3", b for the outcome "Black connects terminals 1 and 3", and c for the outcome "Black connects terminals 1 and 2". Our goal is to realize the superswitches as 3-terminal Hex positions.

We begin by considering the following operation on 3-terminal positions. If G and H are 3-terminal positions, their *concatenation*, written $G +_c H$, is the 3-terminal position shown schematically in the following diagram:

$$G +_c H \quad =$$

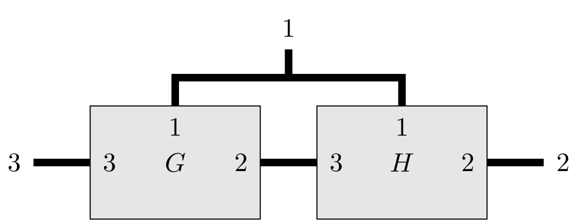

(5)

In words, $G +_c H$ is obtained by connecting terminal 1 of G to terminal 1 of H and terminal 2 of G to terminal 3 of H. Terminals 1, 2, and 3 of the combined position are terminals 1 of G and H, terminal 2 of H, and terminal 3 of G, respectively. As illustrated in Figure 11, the atoms a and b satisfy the following identities:

$$a +_c a = a, \quad a +_c b = b, \quad b +_c a = b, \quad b +_c b = b.$$

Lemma 4.1. *For all n, we have $G_n +_c G_1 \simeq G_{n+1}$.*

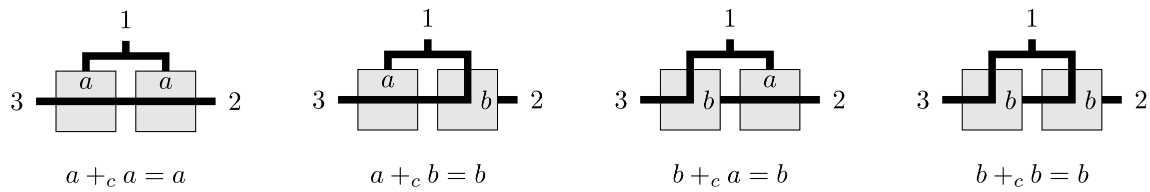

Figure 11: Some identities for the concatenation of atomic positions

Proof. First note that for all games G involving only the atoms a and b, we have $a +_c G \simeq G \simeq G +_c a$ and $b +_c G \simeq b \simeq G +_c b$. This is easily shown by induction on G.

We now prove the lemma by induction. The base case holds because $G_0 +_c G_1 = a +_c G_1 \simeq G_1$. Now consider some $n > 0$. Recall that, by definition, we have $G_n = \{a, b \mid G_{n-1}\}$. Also, by the induction hypothesis, we have $G_{n-1} +_c G_1 \simeq G_n$. Then

$$\begin{array}{rcl} G_n +_c G_1 & = & \{G_n^L +_c G_1, G_n +_c G_1^L \mid G_n^R +_c G_1, G_n +_c G_1^R\} \\ & \simeq & \{a +_c G_1, b +_c G_1, G_n +_c a, G_n +_c b \mid G_{n-1} +_c G_1, G_n +_c a\} \\ & \simeq & \{G_1, b, G_n, b \mid G_n, G_n\}. \end{array}$$

From [15, Prop. 10.2], we know that $G_1 \leqslant G_n$, and therefore the left option G_1 is dominated; therefore, the game $G_n +_c G_1$ simplifies to $\{G_n, b \mid G_n\}$. From $a \leqslant G_n$, we get $\{a, b \mid G_n\} \leqslant \{G_n, b \mid G_n\}$. On the other hand, from $G_n \leqslant G_{n+1}$, since the games are passable, we get $G_n \lhd G_{n+1} = \{a, b \mid G_n\}$, which implies $\{G_n, b \mid G_n\} \leqslant \{a, b \mid G_n\}$. It follows that $G_n +_c G_1 \simeq \{G_n, b \mid G_n\} \simeq \{a, b \mid G_n\} \simeq G_{n+1}$, as claimed. $\square$

4.2 Superswitches are Hex realizable

Lemma 4.2. *The following 3-terminal position has value $G_1 = \{a, b \mid a\}$.*

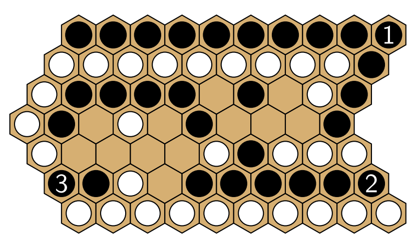

Proof. This can be checked by direct computation, by converting the position to a game form and then computing its canonical form. We will say more about how we actually found this position in Section 5. $\square$

Corollary 4.3. *Each value in the infinite sequence $G_0, G_1, \ldots$ is realizable as a 3-terminal Hex position.*

Proof. $G_0 = a$ is realizable since it is atomic, and G_1 is realizable by Lemma 4.2. It is also clear that if G and H are Hex realizable, then so is $G +_c H$, because all we have to do is combine G and H in the way shown in diagram (5). By repeated applications of Lemma 4.1, we have $G_n \simeq G_1 +_c G_1 +_c \ldots +_c G_1$ when $n \geqslant 1$. Therefore G_n is Hex realizable for all n. For example, the following is a Hex realization of

$G_3 \simeq G_1 +_c G_1 +_c G_1$. Notice that it is just 3 copies of G_1 concatenated together.

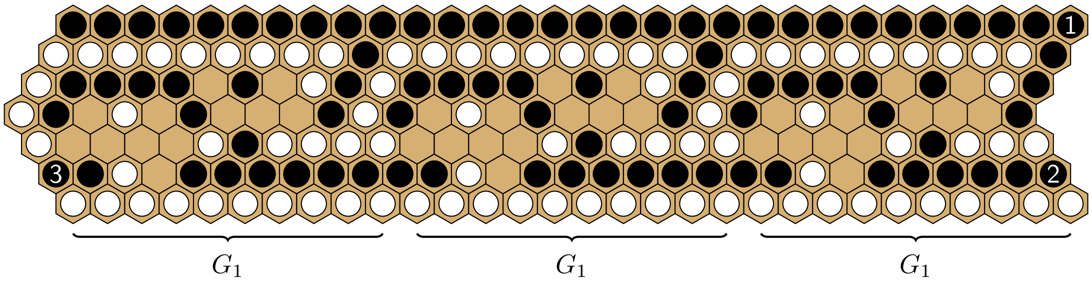

□

4.3 Superswitches are Hex distinguishable

Now that we have found Hex realizations of superswitches as 3-terminal positions, the question remains whether they are distinct as Hex positions. Although it was shown in [15, Prop. 10.2] that the superswitches $G_0, G_1, \ldots$ are distinct *as abstract game values*, it does not a priori follow that they are distinct in Hex, because to establish the latter, we need to find a *Hex realizable* context that distinguishes any pair of them. Fortunately, the superswitches themselves can be repurposed to provide such contexts.

First, let us consider the *dual superswitches*, defined by $G_0^{\mathrm{op}} = a$ and $G_{n+1}^{\mathrm{op}} = \{G_n^{\mathrm{op}} \mid a, b\}$. They work exactly like the superswitches, except that the roles of left and right are reversed. In other words, it is now the right player who gets n chances to switch the outcome from a to b. We can obtain Hex realizations of the dual superswitches by taking the Hex realizations of the corresponding superswitches, exchanging the roles of Black and White, and renumbering the terminals. For good measure, we also flip the position upside down. For example, the following is a realization of the dual superswitch G_1^{op}.

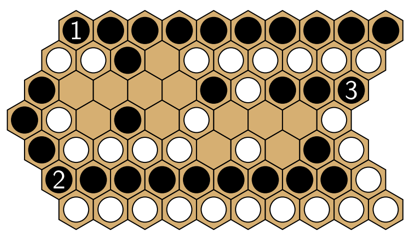

As usual, the other dual superswitches G_n^{op} are obtained by concatenation.

Next, we consider the following operation: If G, H are 3-terminal positions, their *juxtaposition*, written $G +_j H$, is the 2-terminal position shown schematically in the following diagram:

$G +_j H \quad =$

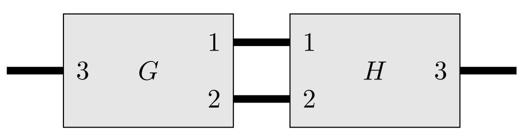

(6)

Since terminal 1 of G is connected to terminal 1 of H and terminal 2 of G is connected to terminal 2 of H, it is immediately obvious that the atoms a and b satisfy the following identities:

$$a +_j a = \top, \quad a +_j b = \bot, \quad b +_j a = \bot, \quad b +_j b = \top.$$

The following proposition shows what happens when we juxtapose a superswitch with a dual superswitch. Recall that $* = \{\top \mid \bot\}$ is the value of a game that is a first-player win.

Proposition 4.4. *We have*

$$G_n +_j G_m^{op} \simeq \begin{cases} \top & \text{if } n > m-1, \\ * & \text{if } n = m-1, \\ \bot & \text{if } n < m-1. \end{cases}$$

Before we prove the theorem, we give an intuitive explanation: the left player owns the left switch and the right player owns the right switch. Neither player wants to commit to a setting for their switch, or else the other player will set their own switch accordingly and win the game. Therefore, Left plays in the right switch as long as possible and Right plays in the left switch as long as possible. Whoever first finishes this "race to the bottom" wins, except that the left player has a small advantage since the "default" setting for both switches is a, which gives a left player win.

Proof. To ease the notation in this proof, we define the *sign games* by

$$S_n = \begin{cases} \top & \text{if } n > 0, \\ * & \text{if } n = 0, \\ \bot & \text{if } n < 0. \end{cases}$$

The claim of the proposition can then be stated as $G_n +_j G_m^{\mathrm{op}} \simeq S_{n-m+1}$. We note that for all integers n, we have $S_n \simeq \{S_{n+1} \mid S_{n-1}\}$.

We start by computing $G_n +_j a$ and $G_n +_j b$. Recall that the sequence $G_0, G_1, \dots$ is increasing. We have $G_0 +_j a = a +_j a = \top$, and since $G_0 \leqslant G_n$ for all n, it follows that $G_n +_j a \simeq \top$ as well. We have $G_0 +_j b = a +_j b = \bot$, $G_1 +_j b = \{a, b \mid a\} +_j b = \{\bot, \top \mid \bot\} \simeq *$, and $G_2 +_j b = \{a, b \mid G_1\} +_j b \simeq \{\bot, \top \mid *\} \simeq \top$. Therefore $G_n +_j b \simeq \top$ holds for all $n \geqslant 2$. In other words, $G_n +_j b \simeq S_{n-1}$. By a dual computation, we find that $a +_j G_m^{\mathrm{op}} \simeq S_{1-m}$ and $b +_j G_m^{\mathrm{op}} \simeq \bot$ for all $m \geqslant 0$.

We prove the proposition by induction on $n + m$. The base case for $m = 0$ holds because we just calculated that $G_n +_j G_0^{\mathrm{op}} = G_n +_j a \simeq \top = S_{n+1}$ for all $n \geqslant 0$; similarly, the base case for $n = 0$ holds because $G_0 +_j G_m^{\mathrm{op}} = a +_j G_m^{\mathrm{op}} \simeq S_{1-m}$ for all $m \geqslant 0$. Now consider the case $n, m > 0$. We have:

$$\begin{aligned} G_n +_j G_m^{\mathrm{op}} &= \{G_n^L +_j G_m^{\mathrm{op}}, G_n +_j G_m^{\mathrm{op}L} \mid G_n^R +_j G_m^{\mathrm{op}}, G_n +_j G_m^{\mathrm{op}R}\} \\ &= \{a +_j G_m^{\mathrm{op}}, b +_j G_m^{\mathrm{op}}, G_n +_j G_{m-1}^{\mathrm{op}} \mid G_{n-1} +_j G_m^{\mathrm{op}}, G_n +_j a, G_n +_j b\} \\ &\simeq \{S_{1-m}, \bot, S_{n-m+2} \mid S_{n-m}, \top, S_{n-1}\}. \end{aligned}$$

In the last step, we have used the above calculations as well as the induction hypothesis. Since $1 - m < n - m + 2$, we have $S_{1-m} \leqslant S_{n-m+2}$, and therefore S_{n-m+2} dominates the other left options; similarly, S_{n-m} dominates all other right options. Therefore $G_n \simeq \{S_{n-m+2} \mid S_{n-m}\} \simeq S_{n-m+1}$ as claimed. ☐

Corollary 4.5. *There are infinitely many Hex-distinguishable 3-terminal positions.*

Proof. Consider Hex realizations of the superswitches G_n and G_m, where $n < m$. To show that they are Hex-distinct, it suffices to find a Hex context in which G_n is a first player win for Black and G_m is a first player loss for Black. By Proposition 4.4, $(-) +_j G_{n-1}^{\mathrm{op}}$ is such a context. ☐

For example, the following shows the 2-terminal position $G_2 +_j G_1^{\mathrm{op}}$, which has value $*$ by Proposition 4.4.

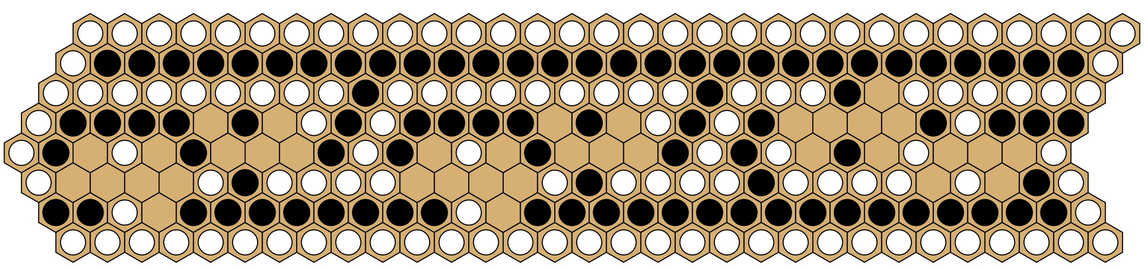

G_1 G_1 G_1^{op}

4.4 Cofinality

The increasing sequence of superswitches $G_0, G_1, \ldots$ has another useful property: it is cofinal among the set of all finite passable games H over $\{\bot, a, b, \top\}$ in which Left cannot achieve the outcome $\top$.

Proposition 4.6. *Let H be a finite passable game over $\{\bot, a, b, \top\}$. If $\top \ntrianglelefteq H$, then there exists some n such that $H \leqslant G_n$. If $\top \nleqslant H$, then there exists some n such that $H \lhd G_n$.*

Proof. By the fundamental theorem (Theorem 3.4), we can assume without loss of generality that H is monotone. We claim that the following hold for all monotone games:

(a) If $\top \nleqslant H$ and $d \geqslant \mathrm{depth}(H) + 2$, then $H \lhd G_d$.

(b) If $\top \ntrianglelefteq H$ and $d \geqslant \mathrm{depth}(H) + 2$, then $H \leqslant G_d$.

We prove these claims by induction on H. If H is atomic, then from either one of the assumptions $\top \nleqslant H$ or $\top \ntrianglelefteq H$, we know that $H \in \{\bot, a, b\}$. Then $H \leqslant G_2$ holds by immediate calculation, and therefore also $H \lhd G_d$ and $H \leqslant G_d$ for any $d \geqslant 2$. Now suppose that H is composite and let $d \geqslant \mathrm{depth}(H) + 2$. To prove (a), assume $\top \nleqslant H$. By the contrapositive of Lemma 3.5, there is some H^R satisfying $\top \ntrianglelefteq H^R$. By the induction hypothesis (b), $H^R \leqslant G_d$, which implies $H \lhd G_d$. To prove (b), assume $\top \ntrianglelefteq H$. To prove $H \leqslant G_d$, first consider any left option H^L. We must show $H^L \lhd G_d$. This follows from the induction hypothesis (a) because $\top \nleqslant H^L$. Next, consider the unique right option G_{d-1} of G_d. We must show $H \lhd G_{d-1}$. Let H^R be any right option of H. Since H^R is monotone, we have $H^R \leqslant H$; it follows that $\top \ntrianglelefteq H^R$. Also, H^R has smaller depth than H, so by the induction hypothesis (b), we have $H^R \leqslant G_{d-1}$. This implies $H \lhd G_{d-1}$ as claimed. □

4.5 Simpleswitches

We will see an application of cofinality in Section 4.6, but first we note that superswitches are relatively inefficient: each additional level of the superswitch (i.e., passing from G_n to G_{n+1}) requires a Hex position with 12 additional empty cells. It turns out that there is a more efficient cofinal sequence that requires only 6 empty cells per level. We call these the *simpleswitches*. They are defined by:

$$\begin{array}{lcl} H_0 & = & a, \\ H_1 & = & \{a, \{\top \mid b\} \mid a\}, \\ H_{n+1} & = & \{a, b \mid H_n\} \quad \text{for all } n \geqslant 1. \end{array}$$

Except for the variation in H_1, the simpleswitches are very similar to the superswitches. It is easy to see that $G_n \leqslant H_n \leqslant G_{n+1}$ for all $n \geqslant 0$, and therefore the simpleswitches form a strictly increasing sequence with the same cofinality as the superswitches.

By a proof very similar to that of Lemma 4.1, we find that we have $H_n +_c H_1 \simeq H_{n+1}$ for all n. Also, we can verify by direct evaluation that the following is a 3-terminal realization of H_1:

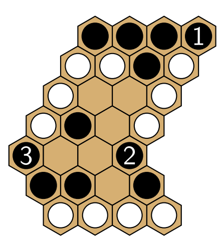

Thus, for example, we get the following realization for the simpleswitch $H_5 = H_1 +_c H_1 +_c H_1 +_c H_1 +_c H_1$:

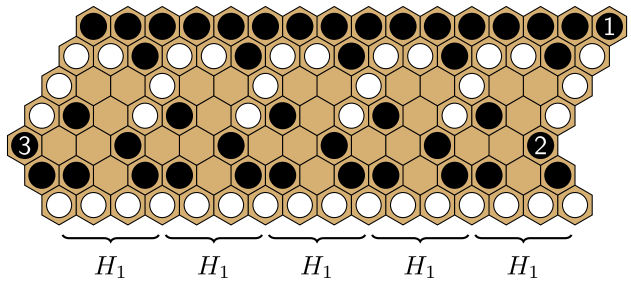

It is evident that the Hex realizations of simpleswitches are much more compact than those we gave of superswitches.

4.6 Application: Verifying connects-both templates

In Hex, an *edge template* is a region that includes a distinguished black stone and a black edge, such that the following two properties hold:

- *Validity:* Black can guarantee to connect the given stone to the edge within the region, even with White moving first.
- *Minimality:* Removing any empty cell or any black stone from the region makes the template invalid.

The following is an example of an edge template:

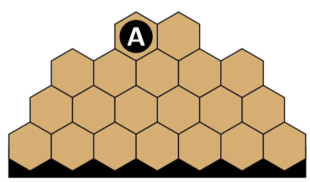

Here, the distinguished stone is marked "A", and it also happens to be the only stone in the template. For more information about edge templates, and a proof of the validity of the above template, see [16].

Large templates are difficult to verify by hand; for validity, one must consider every possible strategy by White, and minimality is even more difficult to verify manually. Fortunately, edge templates can easily be verified using Hex solver software such as Mohex [11]. Given a board position and a player to move, exactly one player must have a winning strategy, and the Hex solver computes whether that player is Black or White. For example, to verify the validity of the above template, it is sufficient to run a Hex solver on the following position, with White to move:

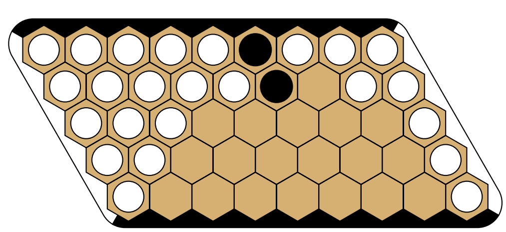

If the winner is Black, the template is valid. Moreover, minimality is easily checked by placing a white stone on one of the empty cells and re-running the solver. The winner should be White. If this can be repeated for every empty cell of the template, the template is minimal. (If the template contains additional black stones, one should also try removing these stones one at a time and check that the template is no longer valid).

Some edge templates have additional properties. For example, the following template has the property that Black can guarantee to connect *both* stones A and B to the edge, with White going first. In other words, Black does not have to choose which of A and B to connect, but can always connect both of them. We call such a template a *connects-both template*.

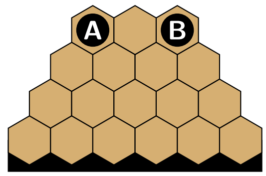

A priori, it is not clear how we can use a Hex solver to verify a connects-both template. A naive first attempt might be to connect both A and B to the opposite board edge, but in that case, Black would win if Black could connect A *or* B to the edge, rather than A *and* B. What we need is to put the would-be template into a context such that Black wins if and only if Black connects both stones. Unfortunately, as we will show in Section 4.8, there exists no single Hex context with this property. However, as we will now show, there is a *sequence* of contexts that will do the job in the limit.

We say that a 3-terminal position has the *connects-both* property if Black can guarantee to connect all three terminals to each other, with White to move first. In other words, this is the case if and only if the position has value $\top$.

Theorem 4.7. *A 3-terminal position G has the connects-both property if and only if $G +_j H_n^{op} = \top$ for all dual simpleswitches H_n^{op}.*

Proof. The left-to-right implication is easy, because an easy induction shows that $\top +_j H_n^{\mathrm{op}} \simeq \top$ holds for all n.

For the right-to-left direction, assume $G +_j H_n^{\mathrm{op}} = \top$ for all dual simpleswitches H_n^{op}. Assume, for the sake of obtaining a contradiction, that $G \not\simeq \top$. Then $G < \top$. Note that because G is a Hex position, G is necessarily finite and monotone. Let G' be the game obtained from G by replacing each occurrence of the atom c by $\perp$. Since the function mapping c to $\perp$ is monotone, G' is a monotone game over $\{\perp, a, b, \top\}$. We also have $G' \leqslant G$, thus $G' < \top$. By Proposition 4.6, there is some superswitch G_n such that $G' \lhd G_n$. It follows that $G' +_j H_{n+2}^{\mathrm{op}} \leqslant G' +_j G_{n+2}^{\mathrm{op}} \lhd G_n +_j G_{n+2}^{\mathrm{op}} \simeq \perp$, where the last equivalence holds by Proposition 4.4. On the other hand, juxtaposition satisfies $c +_j x = \perp +_j x$ for every atom x. Therefore, since G' only differs from G by replacing some atoms c by $\perp$, we have $G' +_j H_{n+2}^{\mathrm{op}} = G +_j H_{n+2}^{\mathrm{op}}$. Putting this together, we have $G +_j H_{n+2}^{\mathrm{op}} \lhd \perp$, contradicting the assumption that $G +_j H_{n+2}^{\mathrm{op}} = \top$. ☐

Theorem 4.7 provides a method for verifying connects-both templates using a Hex solver: we merely need to put the template in the context of a sufficiently high tower of dual simpleswitches. An example of this is shown in Figure 12. Here, the template itself appears in the lower right corner, and the rest of the board is taken up by a dual simpleswitch of depth 4. If the position is a second-player win for Black for dual simpleswitches of *all* depths, then the template is valid.

In practice, of course, one cannot check this for infinitely many dual simpleswitches. Theoretically, it is sufficient to limit the depth of the dual simpleswitch to slightly more than the depth of the candidate template, i.e., the number of its empty cells. But in most cases, this depth is too large to be feasibly checkable. However, with some manual work, it is possible to verify a connects-both template using only a single dual simpleswitch of some fixed small depth. If the solver determines that White can win going first, the template is not valid. Otherwise, for every possible white intrusion, the solver will suggest a potential black winning response. If the response is such that it connects terminals A and B to each other, then the depth of the simpleswitch no longer matters; otherwise, one must recursively consider all white follow-up moves. In this way, the validity of most templates can be verified in a relatively small number of steps.

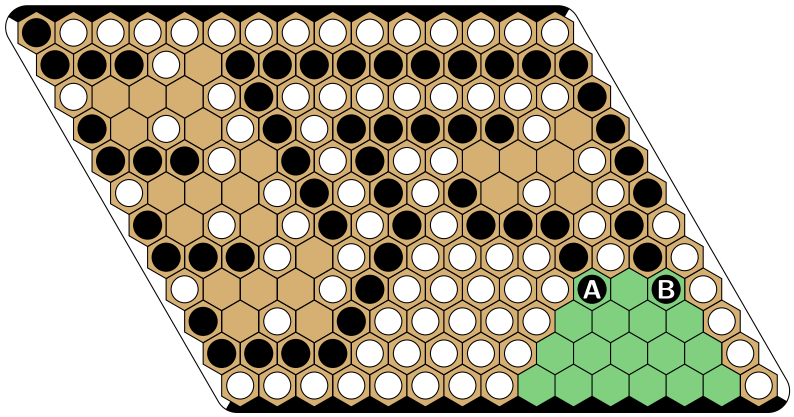

Figure 12: Verifying a connects-both template

For the proof of minimality, one can as usual consider every possible way of placing a single white stone in the template or removing a black stone, and uses the solver to show that the resulting region is a first-player win for White. As long as this procedure succeeds, the depth of the used simpleswitches does not matter, because by Theorem 4.7, a dual simpleswitch of any one depth is sufficient to invalidate a template.

4.7 Tripleswitches

In Proposition 4.6, we showed that the sequence of superswitches is cofinal for the class of games over $\{\bot, a, b, \top\}$ in which Left cannot achieve the outcome $\top$. There is an analogous cofinal sequence of games over $\{\bot, a, b, c, \top\}$. Define the following games:

$$\begin{array}{lcl} T_0 & = & \bot, \\ T_{n+1} & = & \{a, b, c \mid T_n\} \quad \text{for all } n \geqslant 0. \end{array}$$

For example, $T_3 = \{a, b, c \mid \{a, b, c \mid \{a, b, c \mid \bot\}\}\}$. We call these games *ideal tripleswitches*. It is easy to see that they form an increasing sequence of passable games. The following proposition establishes the cofinality of the sequence:

Proposition 4.8. *Let H be a finite passable game over $\{\bot, a, b, c, \top\}$. If $\top \ntrianglelefteq H$, then there exists some n such that $H \leqslant T_n$. If $\top \nleqslant H$, then there exists some n such that $H \lhd T_n$.*

Proof. The proof is exactly the same as that of Proposition 4.6, with one more atom. $\square$

With the exception of T_0, T_1, and T_2, we do not know whether the ideal tripleswitches are realizable as 3-terminal Hex positions. However, there is a sequence of 3-terminal Hex positions that has the same cofinality as the ideal tripleswitches. We call these the *Hex tripleswitches*.

Definition 4.9. If G_1, G_2, and G_3 are 3-terminal positions, their *pinwheel composition*, which we write as

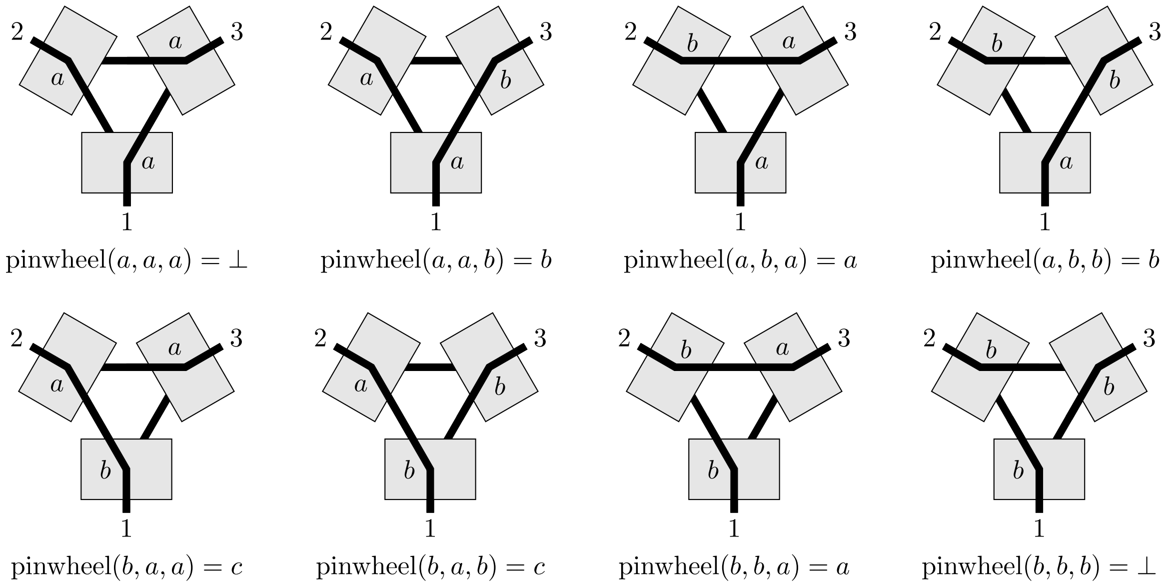

Figure 13: Some identities for pinwheels of atomic positions

pinwheel(G_1, G_2, G_3), is the 3-terminal position shown schematically in the following diagram:

$$\text{pinwheel}(G_1, G_2, G_3) \quad = \tag{7}$$

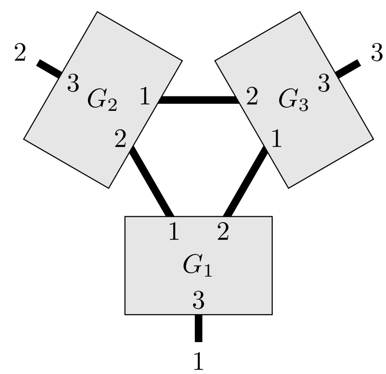

A number of identities for atomic pinwheels are illustrated in Figure 13.

We define the Hex position $T'_n = \text{pinwheel}(G_{n+2}, G_{n+2}, G_{n+2})$, where G_{n+2} is a Hex realization of the $(n+2)$nd superswitch. Note that, since the only atoms occurring in the canonical value of G_{n+2} are a and b, by the identities in Figure 13, the only atoms that can possibly occur in the canonical value of T'_n are $\{\bot, a, b, c\}$. Therefore, outcome $\top$ is not achievable, and in particular, $\top \not\triangleleft T'_n$. It follows by Proposition 4.8 that $T'_n \leqslant T_m$ for some m. Conversely, we claim that $T_n \leqslant T'_n$, and therefore the sequence $(T'_n)_{n\in\mathbb{N}}$ has the same cofinality as the sequence $(T_n)_{n\in\mathbb{N}}$.

Lemma 4.10. *For all n, $T_n \leqslant T'_n$.*

Proof. Informally, the proof can be summarized by saying that even after Right gets n free moves in T'_n, Left can still set all 3 of the pinwheel's superswitches to atomic values of Left's choice, thereby achieving any desired outcome a, b, or c.

More formally, we first note that $a, b \leqslant \{a, b \mid \{a, b \mid a\}\} = G_2$, and therefore pinwheel$(x, y, z) \leqslant$ pinwheel$(G_2, G_2, G_2) = T'_2$ for all $x, y, z \in \{a, b\}$. By the identities in Figure 13, it follows that $a, b, c \leqslant T'_0$, and therefore also $a, b, c \triangleleft T'_0$.

We prove the lemma by induction. The base case is clear since $T_0 = \bot$. For the induction step, consider $n > 0$. To show $T_n \leqslant T'_n$, first consider any left option x of T_n. Then $x \in \{a, b, c\}$, and therefore $x \lhd T'_0 \leqslant T'_n$ as desired. Next, consider any right option H of $T'_n = \text{pinwheel}(G_{n+2}, G_{n+2}, G_{n+2})$. Using the canonical form of G_{n+2}, there are exactly three such right options, namely $\text{pinwheel}(G_{n+1}, G_{n+2}, G_{n+2})$, $\text{pinwheel}(G_{n+2}, G_{n+1}, G_{n+2})$, and $\text{pinwheel}(G_{n+2}, G_{n+2}, G_{n+1})$. In all three cases, it is the case that $\text{pinwheel}(G_{n+1}, G_{n+1}, G_{n+1}) \leqslant H$, or in other words, $T'_{n-1} \leqslant H$. By the induction hypothesis, we know $T_{n-1} \leqslant T'_{n-1}$, and since T_{n-1} is a right option of T_n, it follows that $T_n \lhd H$ as desired. □

4.8 Application: No single context verifies connects-both templates

We can use the Hex tripleswitches to prove something that was claimed in Section 4.6, namely, that there is no single Hex context for a 3-terminal position such that Black wins if and only if the position has the connects-both property. Here, by a "Hex context for a 3-terminal position", we mean any position C (atomic or not) on a Hex board with a "hole" into which a 3-terminal position G fits, as suggested by the following illustration:

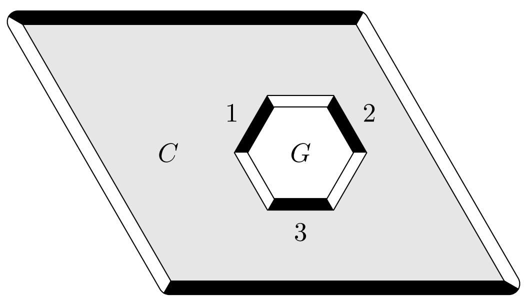

If C is such a context and G is a 3-terminal position, we write $G +_{ctx} C$ for the operation of plugging G into the hole in the context.

Proposition 4.11. *There is no Hex context C with the property that for all 3-terminal Hex positions G, we have $G +_{ctx} C \simeq \top$ if and only if $G \simeq \top$.*

Proof. We first observe that if C is an *atomic* context (i.e., if C has no empty cells) and if $\top +_{ctx} C$ is winning for Black, then one is of $a +_{ctx} C$, $b +_{ctx} C$, or $c +_{ctx} C$ is already winning. To see this, consider any Black winning path in $\top +_{ctx} C$. If the path lies completely outside the region G, the claim is trivial and $\bot +_{ctx} C$ is already winning in this case. Otherwise, since the path must both enter and exit the region G, one of G's terminals must be connected to the top board edge and one of G's terminals must be connected to the bottom board edge. It follows that one of $a +_{ctx} C$, $b +_{ctx} C$, or $c +_{ctx} C$ is winning.

Now suppose, for the purpose of obtaining a contradiction, that there is a Hex context C with the claimed property. By assumption, Black has a second-player winning strategy in $\top +_{ctx} C$. Let d be the depth of C, i.e., the number of empty cells in it. Let T'_{d+2} be the Hex tripleswitch defined in Section 4.7. Recall that $T'_{d+2} < \top$. We claim that Black has a second-player winning strategy in $T'_{d+2} +_{ctx} C$, contradicting our assumption about C.

By Lemma 4.10, we have $T_{d+2} \leqslant T'_{d+2}$, and therefore it is sufficient to show that Black has a second-player winning strategy in $T_{d+2} +_{ctx} C$. Black's strategy is as follows. If White plays in C, Black responds in C by following the second-player winning strategy for $\top +_{ctx} C$. If White plays in T_{d+2}, Black ignores White's move and makes an arbitrary move in C. Due to monotonicity, this move can only help Black in C, i.e., Black continues to have a second-player winning strategy in $\top +_{ctx} C$. In this way, after at most d moves by each player, the context C reaches an atomic position, say C_0. Meanwhile, Black has not yet made any moves in the T_d component of the game, and White has made at most d moves there, so that component's value is T_k for some $k \geqslant 2$. In summary, at the end of this phase, the game's value is $T_k +_{ctx} C_0$.

By assumption, $\top +_{ctx} C_0$ is a win for Black. By our first observation above, there is an atom $x \in \{a, b, c\}$ such that $x +_{ctx} C_0$ is a win for Black, i.e., $x +_{ctx} C_0 \simeq \top$. On the other hand, by properties of tripleswitches, we have $x \leqslant T_2 \leqslant T_k$, and therefore $T_k +_{ctx} C_0 \simeq \top$. It follows that Black has a winning strategy from the current position, which finishes the proof. □

5 A database of Hex realizable 3-terminal values

In Section 4.2, we used the fact that the abstract game value $\{a, b \mid a\}$ is realizable as a specific 3-terminal Hex position. Similarly, in Section 4.5, we used a specific Hex realization of the value $\{a, \{\top \mid b\} \mid a\}$. Both of these Hex positions are non-obvious, and the reader may wonder how we found them. More generally: given a passable game value G over the 3-terminal outcome poset, what is a practical method for finding a Hex realization of G?

While we do not currently know how to solve this problem in general, we have created a tool that has proven useful in practice. Namely, we created a database of more than a million Hex realizable 3-terminal values. In Section 5.1, we will briefly explain how the database was created, and in Section 5.2, we will illustrate how to use it.

5.1 How the database was created

When trying to create a list of Hex realizable values, the first idea is to enumerate many small 3-terminal Hex positions and record their values. However, this idea does not work well. The problem is that even if one restricts attention to positions that are small enough to be efficiently evaluated (in our case, positions with up to approximately 15 empty cells), the number of such positions is astronomical, and most of them yield values that are either repetitive or much too complicated to be of interest.

Instead, we started from a small initial set of realizable 3-terminal values, and then repeatedly closed the set under symmetries, duals, and the pinwheel construction of Definition 4.9. After each iteration, we reduced all values to canonical form, weeded out duplicates, and removed values that were unreasonably complicated (we arbitrarily chose to eliminate values whose canonical forms contained more than 50 atoms). We do not claim the method to be exhaustive (it is not clear whether all realizable 3-terminal values can be obtained by repeated applications of pinwheel composition). Nevertheless, and perhaps surprisingly, the method generated a very large number of interesting Hex-realizable values with a reasonable computational effort.

5.2 How to use the database

The database is available from [4]. It consists of two files in a format that is readable by both humans and machines. The file primaries.txt contains a list of 122 so-called *primary positions*. The first 32 of these are shown in Figure 14. We write P_n for the nth primary position, and P_n^{op} for its dual, which is obtained by switching the roles of black and white. To facilitate taking duals, we have labelled both the black terminals and the white terminals of the primary positions. For example:

$P_5(1, 2, 3) =$ 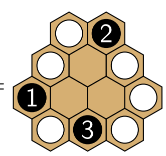

$P_5^{\mathrm{op}}(1, 2, 3) =$ 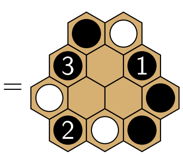

We write $P_n(x, y, z)$ for the analogous position whose terminals are numbered x, y, z instead of $1, 2, 3$. For example:

$P_5(x, y, z) =$ 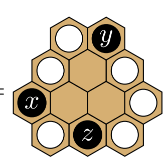

$P_5^{\mathrm{op}}(x, y, z) =$ 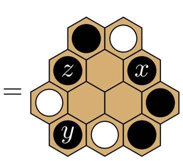

We also use the superscript $\times$ as a *chirality* annotation, indicating that the position should be mirrored, so that its terminals are arranged counterclockwise instead of clockwise. Finally, the notation $\mathrm{Eq}(x, y)$ denotes

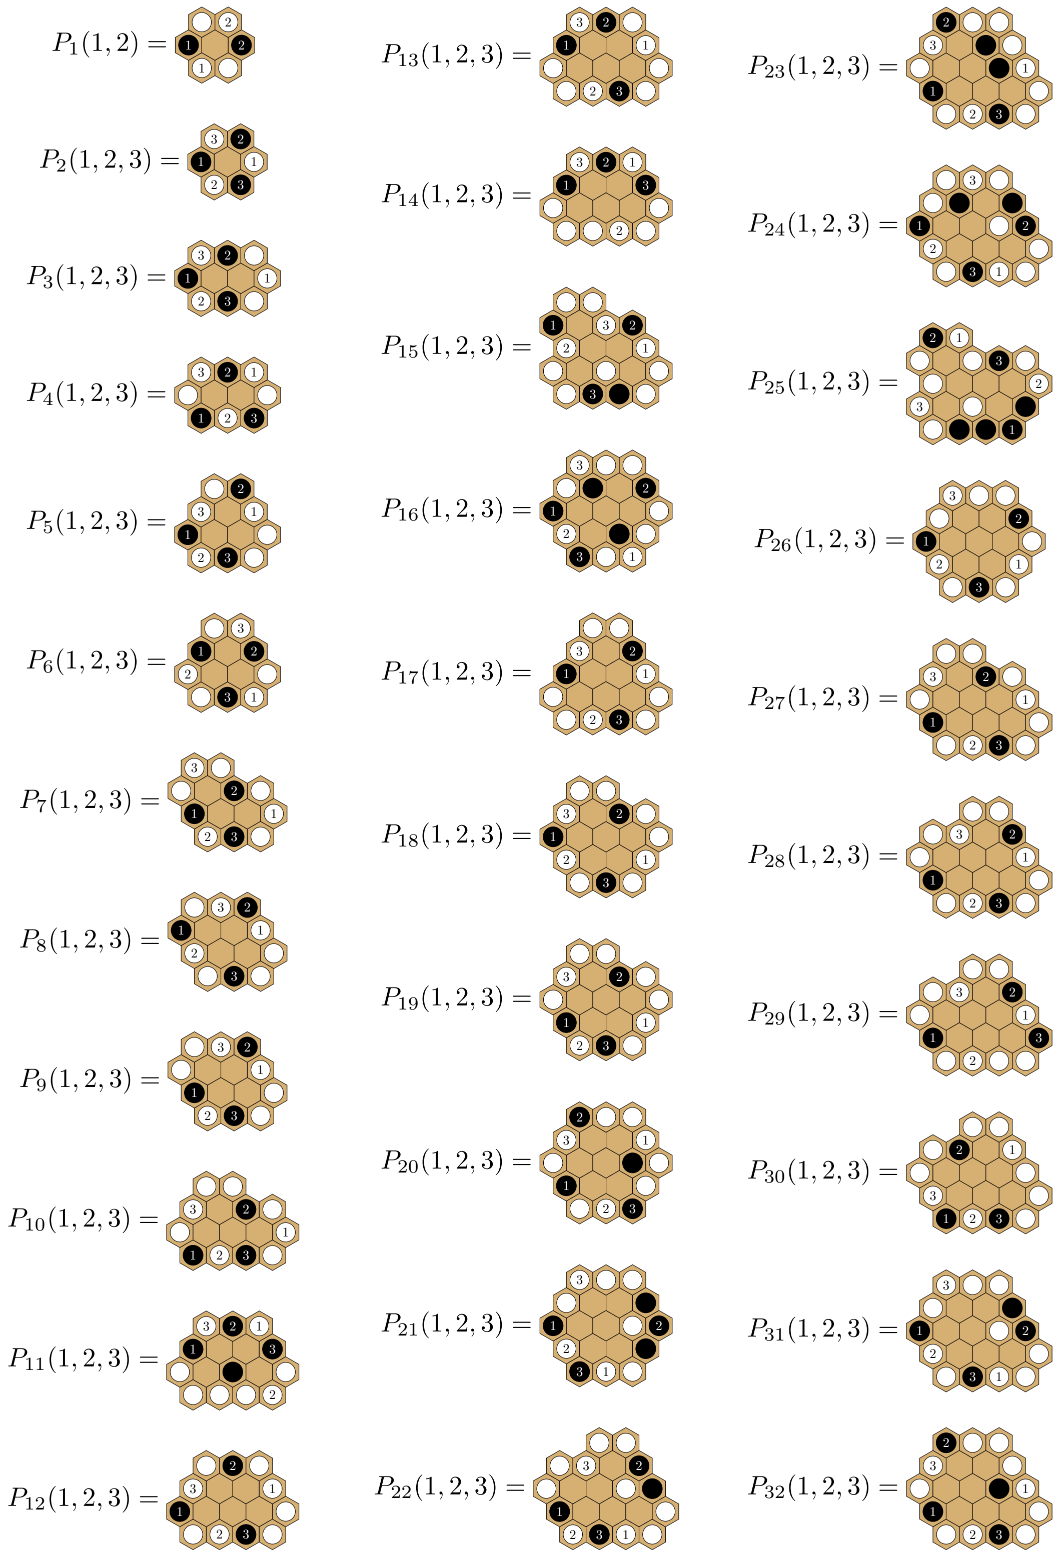

Figure 14: Some primary 3-terminal positions

0	$\bot$	$G(1,2,3) = \text{Empty}$
0	$\top$	$G(1,2,3) = \text{Eq}(1,2,3)$
0	a	$G(1,2,3) = \text{Eq}(2,3)$
0	b	$G(1,2,3) = \text{Eq}(1,3)$
0	c	$G(1,2,3) = \text{Eq}(1,2)$
1	$\{a \mid \bot\}$	$G(1,2,3) = P_1(2,3)$
1	$\{\top \mid a\}$	$G(1,2,3) = \text{Eq}(2,3), P_1(1,2)$
2	$\{\top \mid \{a \mid \bot\}\}$	$G(1,2,3) = P_3(1,2,3)$
2	$\{\{\top \mid b\} \mid \bot\}$	$G(1,2,3) = P_3^{\text{op}}(2,3,1)$
3	$\{a, \{\top \mid b\} \mid \bot\}$	$G(1,2,3) = P_5(1,2,3)$
3	$\{a, \{\top \mid c\} \mid \bot\}$	$G(1,2,3) = P_5(1,3,2)^\times$
5	$\{b, \{\top \mid a, \{c \mid \bot\}\} \mid \bot\}$	$G(1,2,3) = P_3^{\text{op}}(1,4,3), P_5(2,1,4)^\times$
6	$\{a, \{\top \mid b\} \mid a\}$	$G(1,2,3) = P_{25}^{\text{op}}(2,1,3)^\times$
12	$\{a, b \mid a\}$	$G(1,2,3) = P_1(3,4), P_3^{\text{op}}(4,5,3), P_3(2,4,6), P_5(2,1,6)^\times, P_9(2,5,4)$
14	$\{a, b \mid \{a, b \mid \{a \mid \bot\}\}\}$	$G(1,2,3) = P_5^{\text{op}}(4,5,1), P_{15}^{\text{op}}(4,6,2)^\times, P_{25}^{\text{op}}(6,5,3)^\times$

Figure 15: Excerpt from the 3-terminal database

two terminals that are connected to each other, and similarly $\text{Eq}(x, y, z)$. Thus:

$P_5(x, y, z)^\times =$ 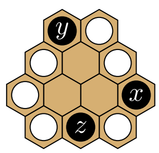

$\text{Eq}(x, y) =$ 

$\text{Eq}(x, y, z) =$ 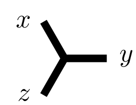

The file realizable-3terminal.txt describes one Hex position per line, in the format shown Figure 15. Each line contains three pieces of information about a position: its number of empty cells, its value in canonical form, and a description of the position itself. The latter expresses each position as a composition of primary positions.

For example, consider the second-to-last line of Figure 15. It states that the game value $\{a, b \mid a\}$ is realized by the position

$$G(1, 2, 3) = P_1(3, 4), P_3^{\text{op}}(4, 5, 3), P_3(2, 4, 6), P_5(2, 1, 6)^\times, P_9(2, 5, 4).$$

To decipher this notation, we first draw the five primary positions indicated on the right-hand side, dualizing and mirroring them as indicated by the "op" and "×" annotations:

$P_1(3, 4) =$ 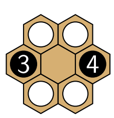

$P_3(2, 4, 6) =$ 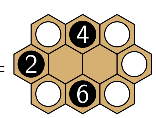

$P_3^{\text{op}}(4, 5, 3) =$ 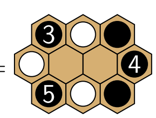

$P_5(2, 1, 6)^\times =$ 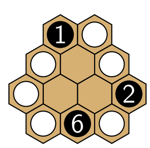

$P_9(2, 5, 4) =$ 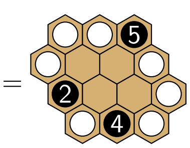

Next, we connect like-numbered terminals and use 1, 2, and 3 as the exterior terminals of the position

$G(1,2,3)$, like this:

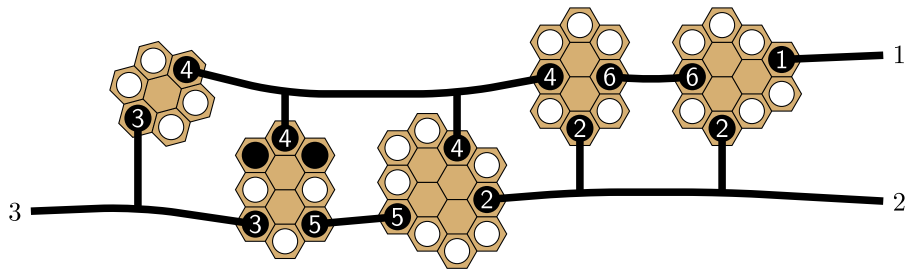

Finally, if desired, we can manually try to fit the position into a small amount of space. The following is a compact representation of the value $\{a, b \mid a\}$. We have numbered the internal terminals to make it easier to compare the position with the one above; note that terminal 5 is not a stone, but a direct connection between adjacent cells. This realization of the superswitch is in fact the one that we first saw in Section 4.2.

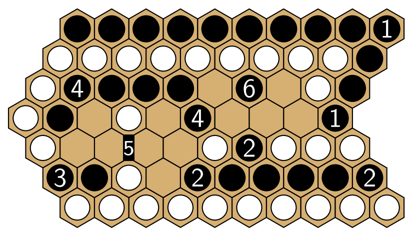

As another example, consider the realization of $\{a \mid \{\top \mid b\}|a\}$. Figure 15 states that this value is realized by the Hex position $G(1,2,3) = P_{25}^{\text{op}}(2,1,3)^{\times}$. Looking up the primary position P_{25} in Figure 14 and applying the indicated transformations, we obtain the realization

$$G(1,2,3) = P_{25}^{\text{op}}(2,1,3)^{\times} =$$ 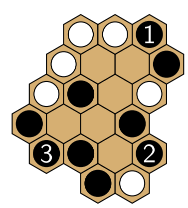

Modulo some padding, this is exactly the simpleswitch of Section 4.5.

Remark 5.1. Because of Theorem 2.3, we know that every Hex position can be decomposed into the single-cell primaries P_1 and P_2. Therefore, the remaining 120 primary positions are not strictly speaking necessary. However, using these additional primaries greatly decreases the size of the database, while also making it more useable.

Remark 5.2. From Section 3.6, we know that the outcome posets for some kinds of 3-terminal regions, such as corners and forks, are proper quotients of the outcome poset of a generic 3-terminal region. This means that distinct 3-terminal positions may become equivalent when regarded as corners or forks. We can therefore take our database of Hex-realizable 3-terminal values and extract smaller databases of Hex-realizable corners and forks. The 1382388 distinct Hex-realizable 3-terminal values in our database reduce to 369291 Hex-realizable corners and 40310 Hex-realizable forks.

6 Application: Verifying pivoting templates

In addition to edge templates, which guarantee that one or more stones can be connected to the relevant edge, one can also consider regions with other related properties. An example of this is a *pivoting template*.

Informally, a pivoting template is a region that includes a distinguished black stone A and an empty cell B, such that Black can continuously threaten to connect A to the edge until the point where Black either connects A to the edge or occupies B and connects B to the edge. Moreover, to be considered a template, the region should also be minimal with this property.

The following is an example of a pivoting template:

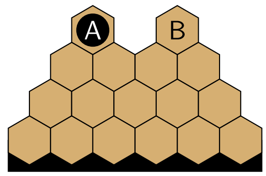

(8)

Depending on White's moves, Black may play as follows:

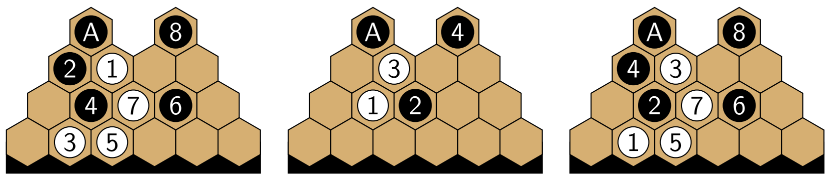

In all cases, Black is threatening to connect A to the edge (meaning that if White passes, Black could connect A to the edge in a single move), until Black occupies B.

6.1 Pivoting forks

We can use combinatorial game theory to formalize the concept of pivoting. Consider the outcome poset $P_2 = \{\bot, a, b, \top\}$, where a and b are incomparable. Slightly generalizing the terminology from Section 3.6, we will use the term *fork* for any game over this outcome poset (even for regions that are not connected to white board edges). Informally, we want to say that a fork is *pivoting* if Left can continuously threaten a until the point where either a or b is achieved. To make this more formal, we need to interpret the phrases "achieve", "threaten", and "continuously until" in combinatorial game theory terms. We say that a game G *achieves* b if $b \leqslant G$, and that G *threatens* a if $a \lhd G$. The former means that Left can guarantee outcome at least b, going second. The latter means that Left can guarantee outcome at least a, going first. By "continuously threaten until", we mean that G threatens a, and for every right move there is a left response that reestablishes the threat, until b is achieved. The following recursive definition makes this precise. Recall from Section 3 that $G^{R(L)}$ means a left option of G^R when G^R is composite, or G^R itself when G^R is atomic.

Definition 6.1 (Pivoting fork)**.** A fork G is called *pivoting* if

- $b \leqslant G$, or
- $a \lhd G$ and for all G^R, there exists some $G^{R(L)}$ that is pivoting.

Lemma 6.2. *If G is pivoting and $G \leqslant H$, then H is also pivoting.*

Proof. For the proof, it is convenient to say that G is *pre-pivoting* if there exists $G^{(L)}$ such that $G^{(L)}$ is pivoting. We then prove the following three statements by simultaneous induction: (a) G pivoting and $G \leqslant H$ implies H pivoting; (b) G pivoting and $G \lhd H$ implies H pre-pivoting; (c) G pre-pivoting and $G \leqslant H$ implies H pre-pivoting. With this setup, the rest of the proof is then routine. $\square$

The following proposition characterizes pivoting forks in two different ways. Recall that $G +_j H$ denotes the juxtaposition operation, which was defined in Section 4.3 for 3-terminal positions, but also makes sense for forks. It has the following addition table:

$+_j$	$\bot$	a	b	$\top$
$\bot$	$\bot$	$\bot$	$\bot$	$\bot$
a	$\bot$	$\top$	$\bot$	$\top$
b	$\bot$	$\bot$	$\top$	$\top$
$\top$	$\bot$	$\top$	$\top$	$\top$

Proposition 6.3. *For a passable fork G, the following are equivalent:*

(a) G is pivoting;

(b) $\top \leqslant G +_j \{a, b \mid \{a, b \mid \{a \mid \bot\}\}\}$;

(c) $b \leqslant G$ or $\{a \mid a, \{b \mid \bot\}\} \leqslant G$.

Proof. $(a) \Rightarrow (b)$: Assume G is pivoting. Let $C = \{a, b \mid \{a, b \mid \{a \mid \bot\}\}\}$. We must show that Left has a second player winning strategy in $G +_j C$. We show this by induction on G. If $b \leqslant G$, the claim holds because we can verify by direct calculation that $\top \leqslant b +_j C$. Otherwise, suppose Right moves to $G^R +_j C$. By the pivoting assumption, Left has a response $G^{R(L)}$ (possibly passing) that is again pivoting. Then Left wins the game by the induction hypothesis. Now suppose Right moves to $G +_j C^R = G +_j \{a, b \mid \{a \mid \bot\}\}$. By the pivoting assumption, either $b \leqslant G$, in which case Left wins by moving to $G +_j b$, or $a \lhd G$, in which case $a +_j C^R \lhd G +_j C^R$. Since $a +_j C^R = a +_j \{a, b \mid \{a \mid \bot\}\}$ is easily seen to be a second-player win for Left, it follows that $G +_j C^R$ is a first-player win.

$(b) \Rightarrow (c)$: Let $C = \{a, b \mid \{a, b \mid \{a \mid \bot\}\}\}$, and assume $\top \leqslant G +_j C$. By the fundamental theorem of monotone games, we can assume without loss of generality that G is monotone.

We proceed by induction on G. We first claim that $a \lhd G$ or $b \leqslant G$. Let $C^R = \{a, b \mid \{a \mid \bot\}\}$. By definition of $\leqslant$, we have $\top \lhd G +_j C^R$. By definition of $\lhd$, there is some $(G +_j C^R)^L$ with $\top \leqslant (G +_j C^R)^L$. By definition of the sum, there are two cases: $\top \leqslant G^L +_j C^R$ or $\top \leqslant G +_j C^{RL}$. In the first case, we have $\top \lhd G^L +_j C^{RR}$, i.e., $\top \lhd G^L +_j \{a \mid \bot\}$, i.e., Left has a winning move in $G^L +_j \{a \mid \bot\}$. This move must be to a in the right component, since any other move is losing due to Right's response of $\bot$. Therefore $\top \leqslant G^L +_j a$. This implies $a \leqslant G^L$, hence $a \lhd G$. In the second case, we have $\top \leqslant G +_j C^{RL}$. Since C^{RL} is either a or b, it follows that $\top \leqslant G +_j a$ or $\top \leqslant G +_j b$. In the first case, $a \leqslant G$, which implies $a \lhd G$. In the second case, $b \leqslant G$. In all cases, we are done proving the first claim.

To finish proving (c), we must prove $b \leqslant G$ or $\{a \mid a, \{b \mid \bot\}\} \leqslant G$. In case $b \leqslant G$, we are done, so assume $b \nleqslant G$. We must show $\{a \mid a, \{b \mid \bot\}\} \leqslant G$. By definition of $\leqslant$, we must show two things.

- We must show that $a \lhd G$, but this holds by the first claim (and our assumption that $b \nleqslant G$).
- We must also show that every $G^{(R)}$ satisfies $\{a \mid a, \{b \mid \bot\}\} \lhd G^{(R)}$. So consider some $G^{(R)}$. From assumption (b), we get $\top \lhd G^{(R)} +_j C$. By definition of $\lhd$, there exists some $(G^{(R)} +_j C)^L$ such that $\top \leqslant (G^{(R)} +_j C)^L$. By definition of the sum, $(G^{(R)} +_j C)^L$ is either $G^{(R)} +_j C^L$ or $G^{(R)L} +_j C$.
  - Case 1: $\top \leqslant G^{(R)} +_j C^L$. Note that C^L is either a or b, and therefore either $\top \leqslant G^{(R)} +_j a$ or $\top \leqslant G^{(R)} +_j b$. This implies $a \leqslant G^{(R)}$ or $b \leqslant G^{(R)}$.
    * Case 1.1: $a \leqslant G^{(R)}$, therefore $\{a \mid a, \{b \mid \bot\}\} \lhd G^{(R)}$ directly from the definition of $\lhd$.
    * Case 1.2: $b \leqslant G^{(R)}$. But we had assumed that G is monotone, therefore $b \leqslant G^{(R)} \leqslant G$, contradicting our assumption that $b \nleqslant G$.
  - Case 2: $\top \leqslant G^{(R)L} +_j C$. Then by the induction hypothesis, we have $b \leqslant G^{(R)L}$ or $\{a \mid a, \{b \mid \bot\}\} \leqslant G^{(R)L}$.
    * Case 2.1: $b \leqslant G^{(R)L}$. This implies $b \lhd G^{(R)}$, hence $\{b \mid \bot\} \leqslant G^{(R)}$, hence $\{a \mid a, \{b \mid \bot\}\} \lhd G^{(R)}$ as claimed.
    * Case 2.2: $\{a \mid a, \{b \mid \bot\}\} \leqslant G^{(R)L}$. This directly implies $\{a \mid a, \{b \mid \bot\}\} \lhd G^{(R)}$ by the definition of $\lhd$.

$(c) \Rightarrow (a)$: Clearly both b and $\{a \mid a, \{b \mid \bot\}\}$ are pivoting, and therefore, by Lemma 6.2, so is any G that is at least as good as one of those two. $\square$

6.2 Verifying pivoting templates

Proposition 6.3 suggests a method for verifying the validity of pivoting templates. To prove that some Hex position G has the pivoting property, by Proposition 6.3(b), all we need to do is juxtapose it with a Hex position of value $C = \{a, b \mid \{a, b \mid \{a \mid \bot\}\}\}$ and use a Hex solver to check whether Black has a second-player winning strategy in $G +_j C$. But can we find a Hex fork with value C? Here, the database of Section 5 is helpful. As shown in Figure 15, the value C has a Hex realization with 14 empty cells as $G(1, 2, 3) = P_5^{\text{op}}(4, 5, 1), P_{15}^{\text{op}}(4, 6, 2)^{\times}, P_{25}^{\text{op}}(6, 5, 3)^{\times}$. Following the procedure of Section 5.2, we find that this Hex realization is

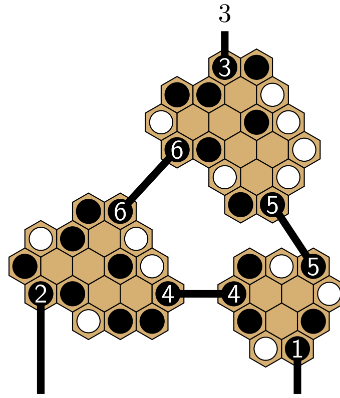

The following is a more compact representation of this position. Note that according to our convention of terminal numbering, outcome a means that terminals 3 and 2 are connected, and outcome b means that terminals 3 and 1 are connected.

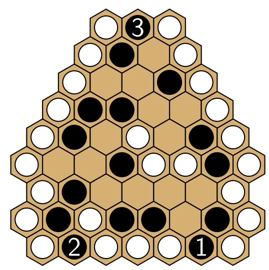

(9)

To check the validity of a pivoting template, it then suffices to juxtapose the template with the position (9) and let a Hex solver check that Black has a second-player win. For example, Figure 16 shows a position that can be used to verify the validity of the pivoting template (8).

As usual, minimality can be checked by placing a white stone on one of the empty cells and re-running the solver, much like we did for other types of templates.

6.3 Sente pivoting forks

Sente is a Japanese Go term that roughly means “keeping the initiative”. Its opposite is *gote*, which means “losing the initiative”. Concretely, suppose there is a goal that a player wants to achieve. To achieve the goal in sente means to achieve it in such a way that it’s the player’s turn immediately afterwards. If instead it is the opponent’s turn after the player achieves the goal, then the player has achieved it in gote.

Our definition of a pivoting fork in Section 6.1 is such that when Left achieves the outcome b, Left achieves it in gote. We can also define a stronger version of a pivoting fork where Left achieves b in sente. The following definition is analogous to Definition 6.1, except that we have moved the condition $b \leqslant G$ to when it is Left’s turn.

Definition 6.4 (Sente pivoting fork)**.** A fork G is called *sente pivoting* if

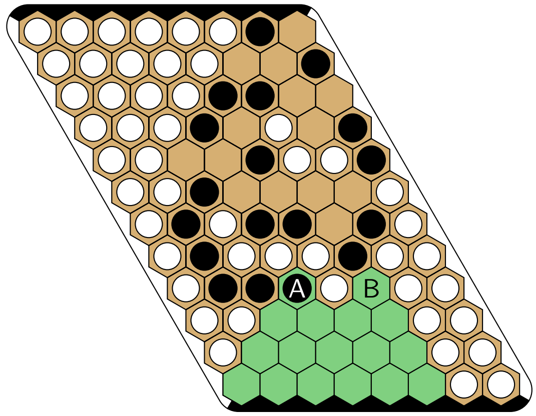

Figure 16: Verifying a pivoting template

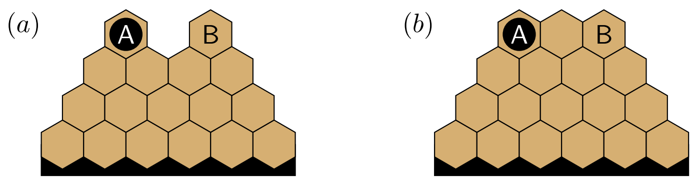

Figure 17: A gote pivoting template and a sente pivoting template

- $a \lhd G$ and for all G^R, either $b \leqslant G$ or there exists some $G^{R(L)}$ that is sente pivoting.

The following characterization of sente pivoting forks is analogous to Proposition 6.3.

Proposition 6.5. *For a passable fork G, the following are equivalent:*

(a) G is sente pivoting;

(b) $\top \leqslant G +_j \{a, b \mid a\}$;

(c) $\{a \mid a, b\} \leqslant G$.

Proof. The proof follows along similar lines as that of Proposition 6.3, and we omit the details. □

An example of a sente pivoting template is shown in Figure 17(b). It is obtained from the (gote) pivoting template in Figure 17(a) by extending the template's carrier with an additional cell between A and B. The idea is that even after Black occupies B, Black still threatens to connect to A, forcing White to spend one more move defending against that threat. So when Black finally stops threatening A, it is Black's turn.

The following proposition shows that this works in general, i.e., extending any (gote) pivoting template with a semi-connection from A to B results in a sente pivoting template.

Proposition 6.6. *As before, let $P_2 = \{\bot, a, b, \top\}$ be the outcome poset for a fork. Also consider the 2-element poset $\mathbb{B} = \{\top, \bot\}$ and any addition operation $+ : P_2 \times \mathbb{B} \to P_2$ satisfying $b + \top \geqslant a$ and $x + \bot \geqslant x$ for all $x \in P_2$. If G is a pivoting fork, then $G + \{\top \mid \bot\}$ is a sente pivoting fork.*

Proof. By Proposition 6.3, we know that $b \leqslant G$ or $\{a \mid a, \{b \mid \bot\}\} \leqslant G$. By Proposition 6.5, we must show that $\{a \mid a, b\} \leqslant G + \{\top \mid \bot\}$. Therefore, it suffices to show that $\{a \mid a, b\} \leqslant b + \{\top \mid \bot\}$ and $\{a \mid a, b\} \leqslant \{a \mid a, \{b \mid \bot\}\} + \{\top \mid \bot\}$. Both are easy to show from the definition of $\leqslant$ and the assumptions about $+$. □

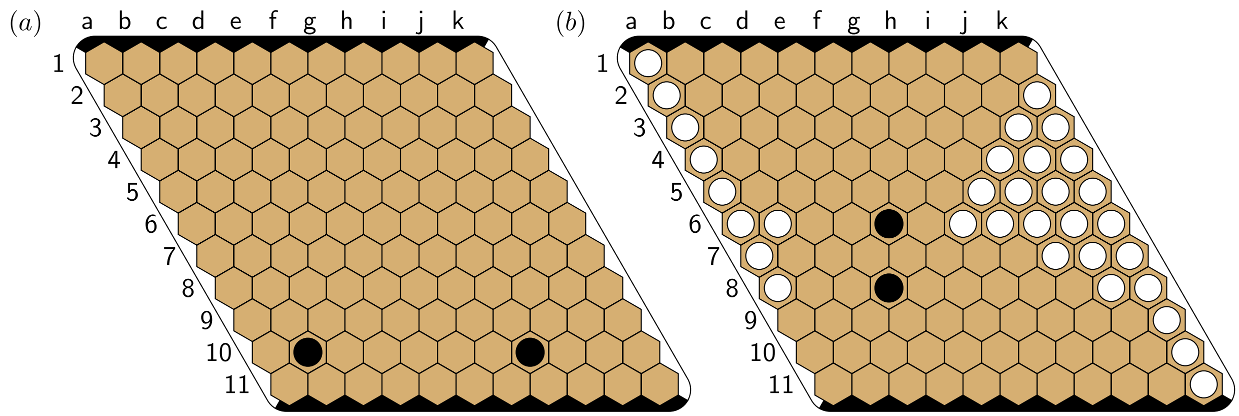

Figure 18: Two handicap winning positions for Black. (a) From Henderson and Hayward [10]. (b) Using a pivoting template.

6.4 Application: A new handicap strategy for 11×11 Hex

In [10], Henderson and Hayward described an explicit winning strategy for Black in 11×11 Hex if Black is allowed to start by playing two stones. (More generally, they described a winning strategy on $n \times n$ Hex, provided that Black is allowed to start with k stones where $6k - 1 \geqslant n$). Figure 18(a) shows Henderson and Hayward's winning opening.

Interestingly, our characterization of pivoting templates allows us to give another such handicap strategy for 11×11. Specifically, we claim that the two black stones shown in Figure 18(b) are a winning opening for Black. The white stones are of course not required in an actual game, but we have included them in the figure to show the area that Black needs to carry out the win (or more precisely, the area that we need to carry out our *proof* of Black's win).

Proposition 6.7. *The position shown in Figure 18(b) is winning for Black, with White to move.*

Proof. In principle, we could show that this position is winning for Black by directly inputting it into a Hex solver. However, the region is too large to be efficiently solvable. Instead, we prove it using a divide-and-conquer method. We divide the board into two regions as shown in Figure 19. We also insert some thin white lines, and claim that Black can win even without the winning path crossing those lines.

First, we claim that region 1, minus the cell marked "$*$", is a pivoting template. As explained in Section 6.2, this can be checked by juxtaposing it with an appropriate context, as in Figure 20(a). The position in Figure 20(a) is simple enough to be solved by Mohex, and is a second-player win for Black as claimed.

Second, we claim that region 1, including the cell marked "$*$", is a sente pivoting template. This follows by Proposition 6.6. Therefore, by Proposition 6.5, the value of region 1 (including the cell marked "$*$") is at least $\{a \mid a, b\}$.

Third, we claim that region 2, when juxtaposed with a fork of value $\{a \mid a, b\}$, is a second-player win for Black. Since the dual superswitch of Section 4.3 is known to have exactly value $\{a \mid a, b\}$, the claim can be checked by juxtaposing region 2 with such a dual superswitch, as in Figure 20(b). This position is simple enough to be solved by Mohex, and is a second-player win for Black as claimed.

Putting all claims together, since region 2 is winning in the context of $\{a \mid a, b\}$, and region 1 is at least as good for Black as $\{a \mid a, b\}$, region 2 is also winning in the context of region 1, as claimed. $\square$

The idea of using divide-and-conquer methods for finding winning strategies in Hex is not new. Yang et al. [17] decomposed the board into smaller regions to identify winning moves on boards up to size 9×9.

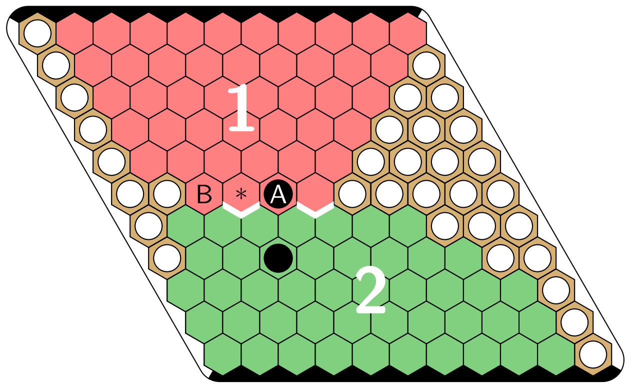

Figure 19: Divide. . .

However, they only considered decompositions into simple (2-terminal) templates. By contrast, our regions have much more intricate combinatorial properties.

6.5 Corollary: A winning strategy for 1-move handicap

A system for measuring the strength of handicap in Hex was proposed by one of the authors, and has found some acceptance in the Hex community [3]. The idea is to measure handicap as a nominal "number of extra moves" given to Black at the start of the game. Through selective use of the swap rule, this "number of moves" can be adjusted in increments of 0.5. Specifically, when playing Hex with the swap rule, the game is approximately fair (we ignore the second player's theoretical advantage because it is very small in practice). This is considered a handicap of 0 moves. On the other hand, when playing without the swap rule, the difference between going first and going second is exactly one extra move at the start of the game. So compared to a theoretically fair game, it makes sense to say that Black, who goes first, has an advantage worth 0.5 moves when playing without swap. An advantage of exactly 1 move can be achieved by playing with swap, but giving Black one extra move at the earliest opportunity (i.e., right after White decides to swap or not). This means: Black plays, White swaps, and Black makes two consecutive moves, or: Black plays, White declines to swap, and Black makes an additional move. A handicap of 1.5 moves can be achieved by playing without swap and letting Black start with two moves, and so on.

In this terminology, both Henderson and Hayward's strategy and our pivoting strategy (Figure 18(a) and (b)) are for 1.5-move handicap: the swap rule is not used and Black gets to start with one extra move.

Interestingly, the two strategies can be combined to yield a guaranteed winning strategy for Black with 1-move handicap. The strategy is as follows: Black opens at `h10`. If White doesn't swap, Black plays a free move at `b10` and then follows the Henderson-Hayward strategy (see Figure 18(a)). If White does swap, the black stone at `h10` effectively becomes a white stone at `j8`, which is outside of the carrier of our pivoting strategy. Black plays two stones at `e6` and `d8` and then follows the pivoting strategy (see Figure 18(b)).

7 Application: Non-inferiority of probes in Hex templates

7.1 Inferiority of probes

Consider the following edge template, which is called the *ziggurat* or the *4-3-2 template* [16, 8]:

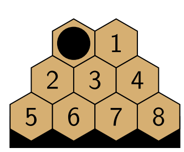

(10)

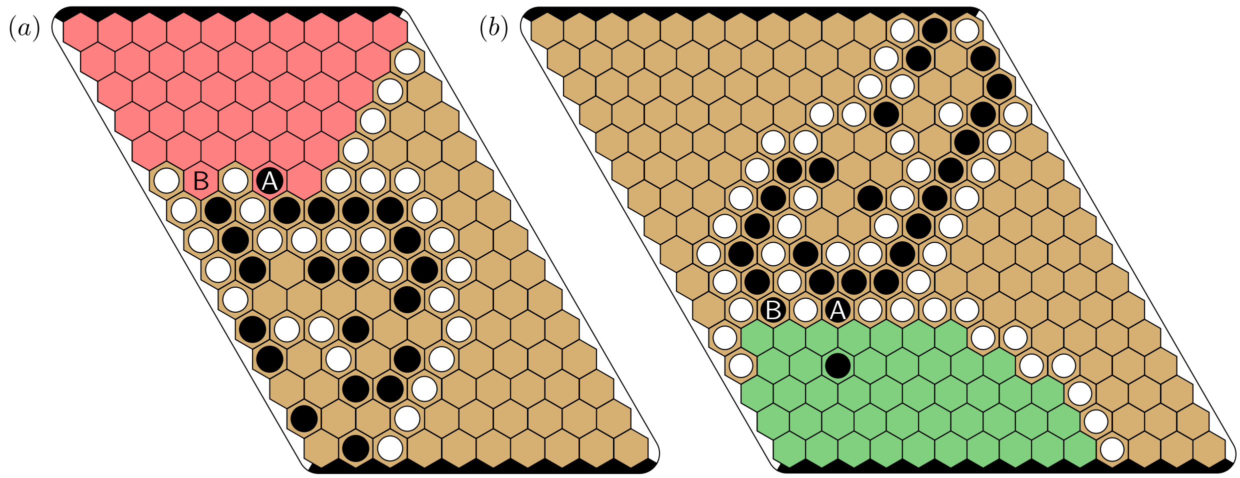

Figure 20: . . . and conquer. Both positions are second-player wins for Black and can be efficiently solved by Mohex.

It is easy to verify that Black can indeed connect the stone to the edge. One way for Black to accomplish this is to play the pairing strategy $\{\{1,3\},\{2,4\},\{5,6\},\{7,8\}\}$: if White plays in any numbered cell in the template, Black plays in the other numbered cell of the same pair.

Nevertheless, the fact that Black can defend the connection does not mean that it is never useful for White to play in the template. By judiciously playing in the template, White can force Black to respond in a way that gives White an advantage. For example, consider the following situation, with White to move. It is easy to see that White's move at 4 is winning, and every other move is losing:

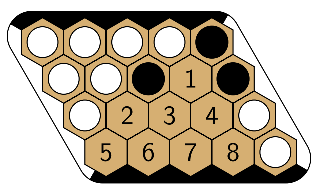

(11)

A white move in Black's template is called an *intrusion* or a *probe* [16, 8]. A probe x is *inferior* if it can never be the unique winning move. More precisely, for every way of embedding the template in a Hex board, if playing at x is a winning first move for White, then there also exists some other winning first move (inside or outside the template). Conversely, a probe is *non-inferior* if there exists some board position where x is the unique winning move. Such a board position is called a *witness* of the non-inferiority of the probe. Thus, (11) is a witness of the non-inferiority of probe 4 in the ziggurat.

Henderson and Hayward conjectured that in the ziggurat, probes 3, 5, 6, 7 and 8 are inferior [8, Conj. 1]. Here, we show that the conjecture is false: in fact, we prove that none of the ziggurat's probes are inferior. (Impatient readers can skip directly to Figure 22 to see the witnesses).

Proving the non-inferiority of some probe in a template, or more generally, of some move in a Hex region, is not an easy task, because it requires finding a witnessing context in which that move, and only that move, is winning. It is not usually possible to find such witnesses by trial and error or by a brute force search, because the witnesses can be exceedingly rare and subtle. Instead, we need a finely-tuned tool. It turns out that combinatorial game theory, along with our database of Hex-realizable 3-terminal values from Section 5, is the right tool for the job.

7.2 Abstract probes

Our method for constructing witnesses of non-inferiority is best illustrated in an example. Our running example will be probe 7 in the ziggurat. We introduce the required machinery in this section, and then give a step-by-step demonstration of the method in Section 7.3.

First, some terminology. Consider the ziggurat in (10). Let A be the outcome poset for the ziggurat, as defined in Section 3. Note that the ziggurat has a complicated boundary that is not an n-terminal region, and so its outcome poset is likely to be complicated. To carry out the computations below, we must be able to do computations in this outcome poset, but fortunately, we do not need to explicitly understand what its elements are.

Let G be the game form of the ziggurat (10) over the poset A. We have

$$G \simeq \{G_1^L, \ldots, G_8^L \mid G_1^R, \ldots, G_8^R\}, \tag{12}$$

where G_i^L is the value of the position obtained by placing a black stone on cell i of the ziggurat, and G_i^R is the value of the analogous position with a white stone on cell i. Since we are interested in situations where G_7^R is the only winning move for White, we can ask what happens if we modify G so that White cannot play 7 as the first move. Let H be the game form that is just like G, except that the right option G_7^R has been removed:

$$H = \{G_1^L, \ldots, G_8^L \mid G_1^R, \ldots, G_6^R, G_8^R\}. \tag{13}$$

(Note that we only remove the white move at 7 from G itself, not from any proper followers of G; in other words, White will be allowed to play in cell 7 as long as it is not the first move in the region). We call the pair (G, H) an *abstract probe*. More generally:

Definition 7.1 (Abstract probe)**.** An *abstract probe* over an outcome poset A is a pair (G, H) of games over A such that $G \leqslant H$. We say that the abstract probe is *viable* if $G < H$.

When an abstract probe is viable, it means that there *potentially* exists some context in which White is winning in G and losing in H. In particular, in the intended situation where H is obtained from G by removing a single white option, it means that there potentially exists a context where the removed option was the only winning move for White. On the other hand, when an abstract probe is not viable, there is no such context.

Our general strategy is to start with a viable abstract probe, and then gradually refine and extend the context while keeping the probe viable at every step. To carry out this refinement, we must be able to add a context to a probe. We do this in a componentwise way, by defining $(G, H) +_f X = (G +_f X, H +_f X)$.

Our refinement process ends when we reach a 2-terminal position, i.e., a probe $(G, H) +_f X$ over the boolean poset $\mathbb{B}$. At this point, the following lemma will give us a witness of non-inferiority.

Lemma 7.2. *Let (G, H) be a probe over a poset A, such that G and H have the same left options, and such that G has exactly one right option K that is not a right option of H. Let X be a game over a poset B, and let $f : A \times B \to \mathbb{B}$ be a monotone function. If the probe $(G, H) +_f X$ is viable, then there exists some follower Y of X such that $K +_f Y$ is the unique winning move for White in $G +_f Y$.*

Proof. By induction on X. Since $G +_f X < H +_f X$ and both games are over $\mathbb{B}$, we know that either as first player or as second player, White has a winning strategy in $G +_f X$ but not in $H +_f X$. Case 1: White has a first-player winning strategy in $G +_f X$ but not in $H +_f X$, and none of White's winning moves in $G +_f X$ are in X. In this case, White must have a winning move of the form $G^R +_f X$. Since White does not have a winning move in $H +_f X$, G^R cannot be a right option of H, and therefore, G^R must be K, proving the claim of the lemma. Case 2: White has a first-player winning strategy in $G +_f X$ but not in $H +_f X$, and at least one of White's winning moves in $G +_f X$ is of the form $G +_f X^R$. By assumption, $H +_f X^R$ is not winning, so $G +_f X^R < H +_f X^R$, and the claim follows by the induction hypothesis. Case 3: White has a second-player winning strategy in $G +_f X$ but not in $H +_f X$. Then Left has some winning move in $H +_f X$. Since G and H have the same left options and Left has no winning move in $G +_f X$, Left's winning move in $H +_f X$ must be some $H +_f X^L$, and the corresponding $G +_f X^L$ is not winning for Left. Then $G +_f X^L < H +_f X^L$, and the claim follows by the induction hypothesis. $\square$

7.3 Step-by-step computation of witnesses

Let us now apply this method to probe 7 in the ziggurat. Note that each of the below steps is explorative; if any step fails, one must redo the previous step.

Step 1. Let $P = (G, H)$ be the abstract probe defined by (12) and (13) above, i.e., the one corresponding to probe 7 in the ziggurat. Since the ziggurat only has 8 empty cells, the values G and H can be computed reasonably efficiently by traversing the game tree. The first thing we must check is that the probe P is viable, i.e., that $G < H$. This can be verified by a computation.

Step 2. Next, we try to convert P into an n-terminal probe by surrounding it with black and white stones. There are potentially many ways to do so. By trial and error, we find that the following boundary keeps the probe viable.

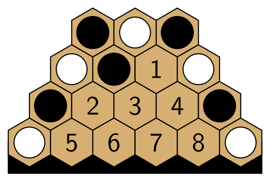

Let P' be the value of this new probe over the 5-terminal outcome poset. We can schematically embed the probe P' inside a Hex board as follows:

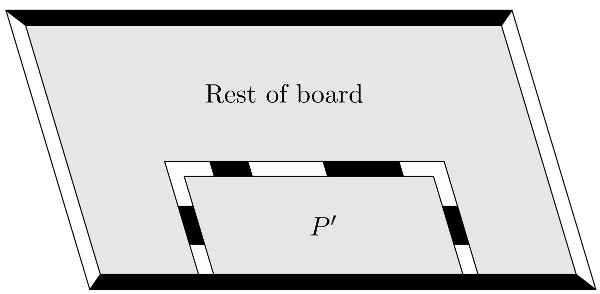

Our goal is to find a Hex-realizable value for "Rest of board" that keeps the probe viable.

Step 3. We divide the rest of the board into simpler regions. By trial and error, we find that a good first step is to slice off two corners A and B, as in the following diagram:

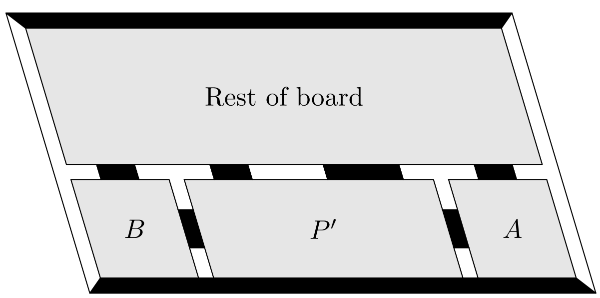

Using the database from Remark 5.2, we make a list of Hex-realizable corner values A such that the probe $P' + A$ remains viable. Anecdotally, we find that the probe is kept viable by about 1.5 percent of the possible values of A. Similarly, we make a list of Hex-realizable corners B such that the probe $P' + B$ remains viable. Here, we find that about 64 percent of the candidate values of B work. Next, we combine the two separate lists of viable values for A and B into a single list of pairs (A, B) such that the probe $P' + A + B$ is viable. The vast majority of such pairs (A, B) work.

Step 4. We divide the rest of the board into smaller regions again. Consider the regions C and D shown in the following diagram:

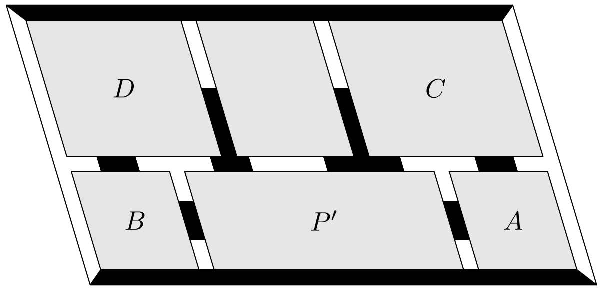

Similarly to what we did in the previous step, we make a list of pairs (A, C) keeping the probe $P' + A + C$ viable (about 24 percent of the pairs work). We also make a list of pairs (B, D) keeping the probe $P'+B+D$ viable (only about 0.8 of such pairs work). Next, we look for triples (A, B, D) keeping the probe $P'+A+B+D$ viable.

Here, we hit a snag. We do not find any viable triples (A, B, D).

Step 5. Since we were not able to complete the previous step, we try a different subdivision of the board. Using more trial and error, we arrive at the following subdivision:

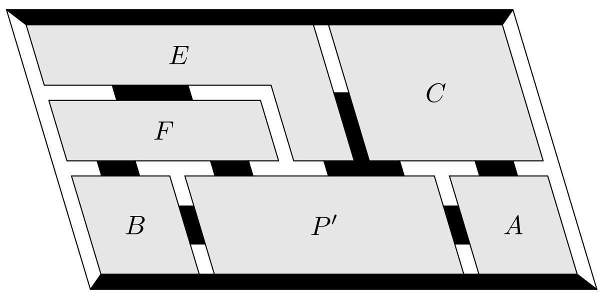

Note that regions A, B, C, and E are corners, but F is a general 3-terminal position. Also note that region C is the same as in step 4, so we can re-use the list of viable pairs (C, A) we have already computed. Continuing as above, we eventually find a tuple (A, B, C, E, F) of Hex-realizable values for these regions, such that the probe $P' + A + B + C + E + F$ is viable. The simplest such tuple we found, as measured by the total number of empty cells in its Hex realization, is

$$\begin{array}{rcl} A & = & \{a, c \mid \{a \mid \bot\}\}, \\ B & = & \{a, c \mid \{a \mid \bot\}\}, \\ C & = & \{\{\top \mid a\}, \{\top \mid c, \{b \mid \bot\}\} \mid \bot\}, \\ E & = & \{\{\top \mid b\} \mid \bot\}, \\ F & = & \{\{\top \mid b\} \mid b, \{c \mid \bot\}\}. \end{array}$$

Step 6. The trial-and-error phase of the method is now completed. What remains to be done is to assemble the witness position. Looking up the Hex realizations of A, B, C, E, and F in the database of Section 5, and connecting the appropriate terminals, we obtain the following position:

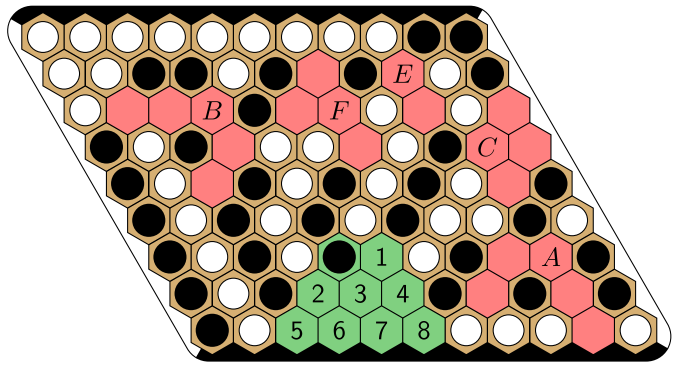

(14)

Here, we have labelled the regions A–F to make them more easily recognizable. By construction, the probe $P' + A + B + C + E + F$ is viable; therefore, Lemma 7.2 guarantees that some follower of $A + B + C + E + F$ is the desired witness of non-inferiority, in which probe 7 is the unique winning move for White. In practice, since we typically enumerate the possible contexts in order of increasing depth, it is almost always the case that $A + B + C + E + F$ itself is the witness.

But can one trust this to work? The required computations took several days, and we could have easily made a mistake somewhere. Luckily, one does not have to take our word for it; it is easy to check the final answer using a Hex solver such as Mohex. Indeed, Mohex easily confirms that probe 7 is the only winning move in (14).

Discussion. The method described in this section requires both heavy computation and significant human guidance. Computations are used for such tasks as searching large databases of Hex-realizable values, computing outcomes and sums, and determining the viability of probes. Human judgement is needed for tasks such as proposing subdivisions of the board, deciding in which order to fill them, when to backtrack, etc. On the whole, this is a labor intensive process. It may be possible to better automate this process in the future, but we have not attempted to do so.

We also note that the witness position (14) is very complex (certainly much more complex than any position that is likely to arise naturally during game play). While one can computationally verify that probe 7 is the only winning move for White, we do not have a human-level explanation for why this is so. Indeed, finding the winning move in this position is a difficult puzzle that would likely stump most Hex masters.

7.4 Results

7.4.1 Probes in some common Hex templates

We investigated the inferiority or non-inferiority of probes into many common Hex templates. Figure 21 shows a number of named templates, and we have numbered the probes in each template.

To prove that a probe is non-inferior, it suffices to provide a witness position as discussed above. For proving that a probe x is inferior, there are several methods. The most common such methods, and the only ones that we need here, are to show that x is *dominated* by some other move in the template, or to show that x is *strongly reversible*. Here, a white move x is strongly reversible if there exists a black response y so that if White has a winning move in the resulting position, then White already had a winning move other than x before x and y were played. For a more detailed discussion of inferiority proofs and additional methods and examples, see Henderson and Hayward [8, 9].

We found that in the templates of Figure 21, the following probes are inferior:

- In the trapezoid, probe 2 is inferior. It is strongly reversed by 1.
- In the crescent, probes 2 and 4 are inferior. Both are strongly reversed by 3. Probe 2 is also strongly reversed by 1.

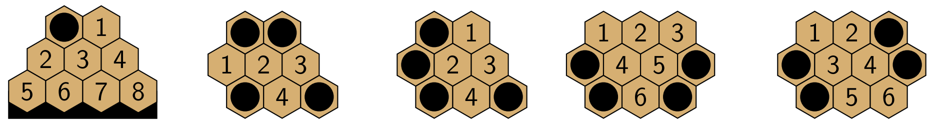

The ziggurat. The trapezoid. The crescent. The span. The parallelogram.

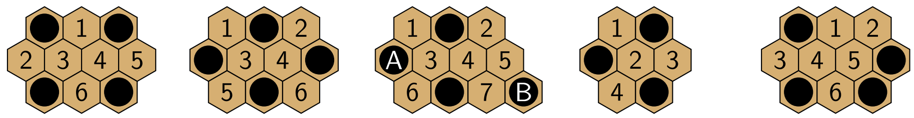

The box. The diamond. The hammock. The wheel. The shopping cart.

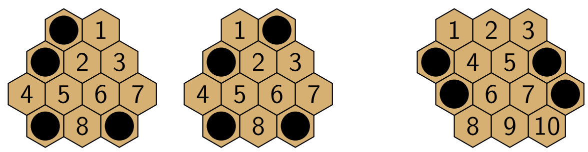

The scooter. The bicycle. The wide parallelogram.

Figure 21: Some Hex templates. The ziggurat is an edge template, and the others are interior templates. In the hammock, Black is guaranteed to connect the stones labelled A and B to each other. In the remaining interior templates, Black can connect all of the black stones.

- In the span, probes 4 and 5 are inferior. They are dominated by 2, using star-decomposition domination [9].
- In the parallelogram, probes 3 and 4 are inferior. They are strongly reversed by 5 and 6, respectively.

All other probes in all templates of Figure 21 are non-inferior, with witnesses shown in Figures 22 and 23. In each witness position, the cell marked with a white dot is the unique winning move for White. In choosing these witness positions, we generally tried to minimize the number of empty cells (rather than minimizing the number of stones), because the number of empty cells is a more meaningful measure of a position's complexity.

7.4.2 Probes in long crescents

Some templates form infinite families. An example of this is the *long crescent*, which is a generalization of the crescent that can be of any length:

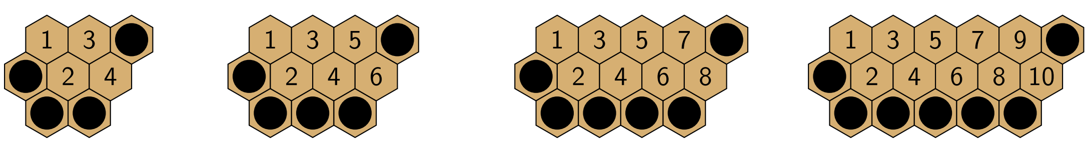

The crescent. The 6-cell long crescent. The 8-cell long crescent. The 10-cell long crescent.

In the long crescent of any length, probes 2 and 4 are inferior for the same reason as in the crescent. The remaining probes (i.e., all odd probes, and all even probes $\geqslant 6$) are non-inferior; witnesses for the 10-cell long crescent are shown in Figure 24. Moreover, these witnesses can be adjusted to all possible lengths, due to the following equivalences, which can be verified by computing their combinatorial values.

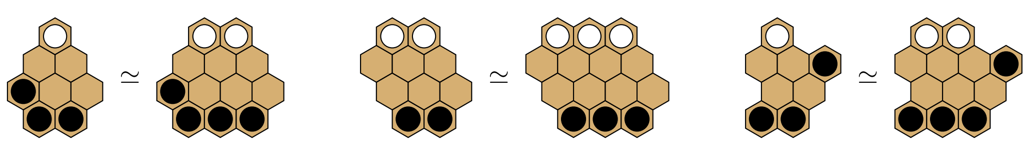

Ziggurat:

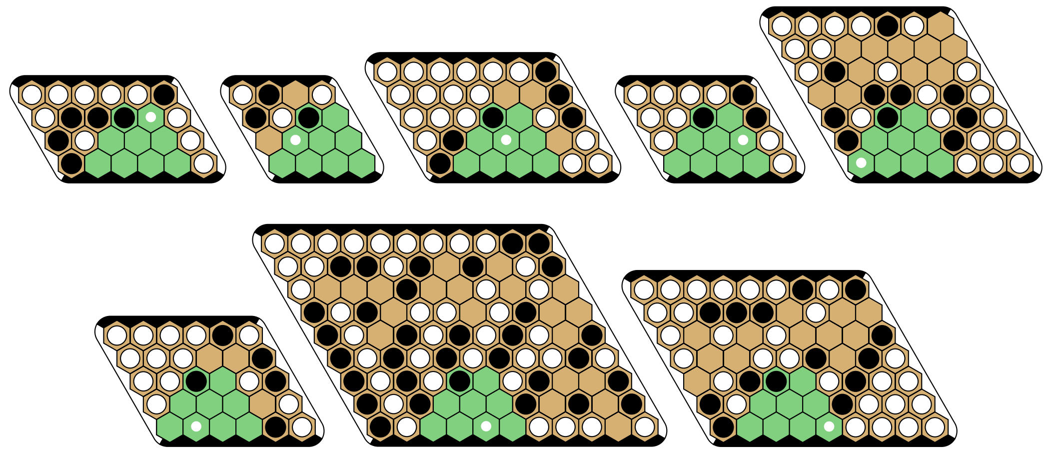

Trapezoid:

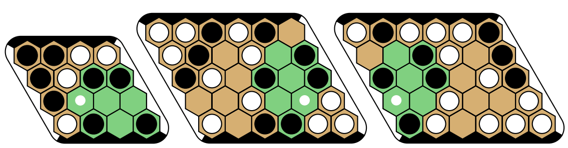

Crescent:

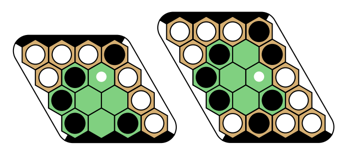

Span:

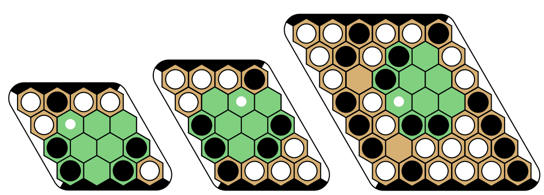

Parallelogram:

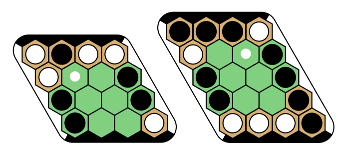

Box:

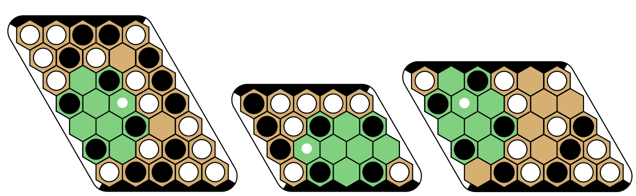

Diamond (and wheel):

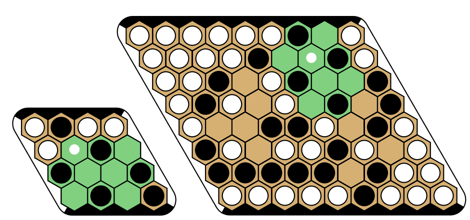

Hammock (and wheel):

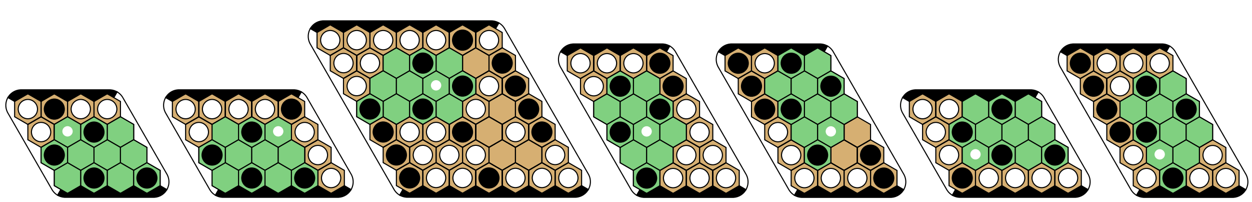

Figure 22: Witnesses for the non-inferiority of probes in various templates, part 1. In each position, the white dot denotes the unique winning move for White.

Shopping cart:

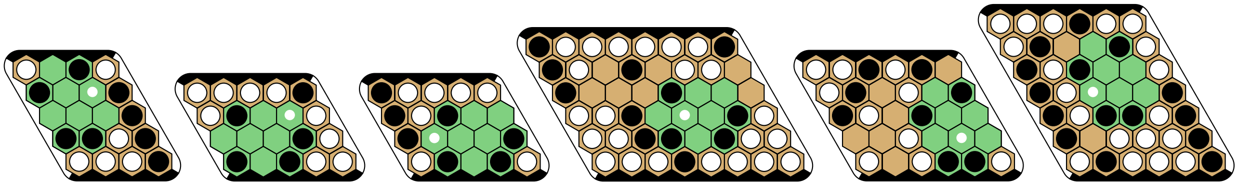

Scooter:

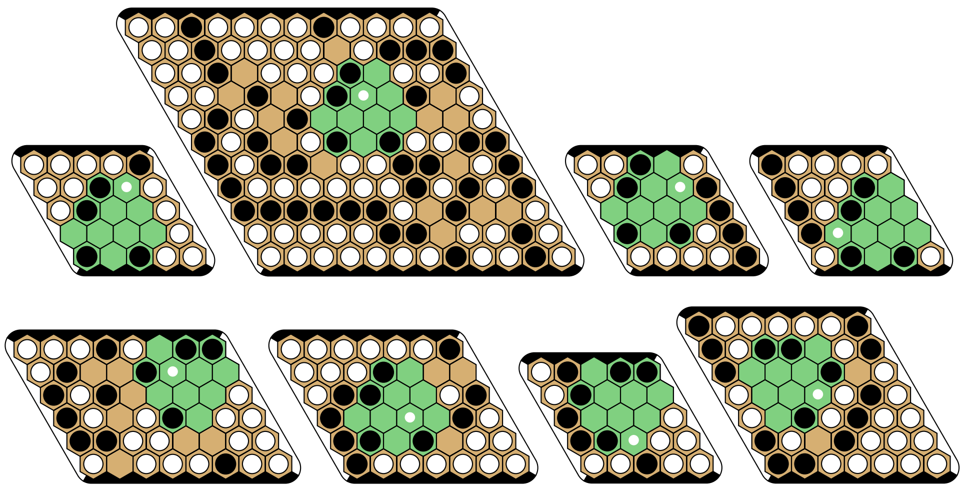

Bicycle:

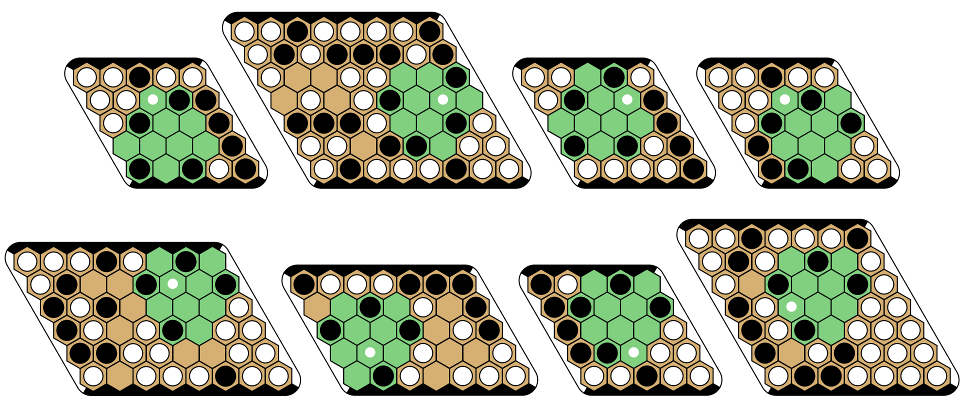

Wide parallelogram:

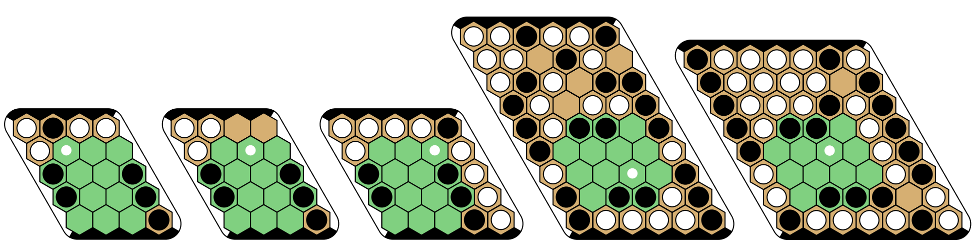

Figure 23: Witnesses for the non-inferiority of probes in various templates, part 2. In each position, the white dot denotes the unique winning move for White.

7.4.3 Probes in bolstered ziggurats

The inferiority of probes in some Hex region can be affected by nearby stones outside of the region. For example, consider the bridge in (a) below. Both probes in the bridge are non-inferior. However, if we *bolster* the bridge by adding a white stone on the left, as in (b), then probe 1 becomes inferior, because it is then strongly reversed by 2. Only the white intrusion at 2 remains useful. If the bridge is bolstered on both sides as in (c), then there are no useful intrusions at all.

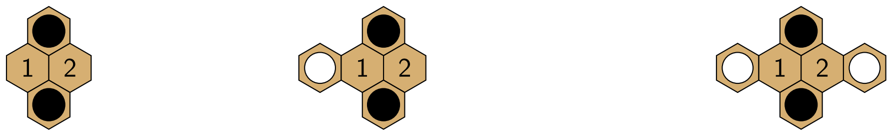

(a) The bridge. (b) A bridge bolstered on the left. (c) A bridge bolstered on both sides.

Let us investigate what happens to a ziggurat in the presence of a single white stone on the perimeter. We label the cells adjacent to the ziggurat with letters $a, \dots, i$, as follows:

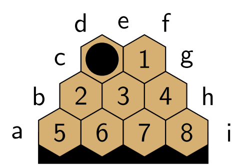

For example, a ziggurat with a white stone at b is the following position, which we call the *half-star*.

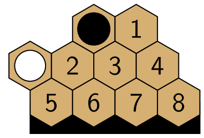

Henderson and Hayward showed in [9] that in the half-star, probe 4 dominates probes 2, 3, 5, and 7. In other words, probes 2, 3, 5, and 7 are inferior in this situation. One may ask whether any of the remaining probes, 1, 4, 6, or 8 are inferior in the half-star. This is not the case; in fact all are non-inferior, and our witnesses for these probes in Figure 22 are half-stars.

The following table summarizes the status of each probe 1–8 in the presence of a white stone at a–i.

	a	b	c	d	e	f	g	h	i
1	O	W	O	W	W	A	W	W	W
2	W	■	W	O	W	W	W	W	W
3	O	■	W	W	O	■	W	■	W
4	W	W	W	W	W	B	B	W	W
5	W	■	W	O	O	W	W	C	W
6	W	W	W	W	O	■	W	D	D
7	W	■	W	E	W	■	W	■	W
8	O	W	F	W	W	F	W	■	W

In this table, "■" means that the probe is provably inferior, and all other entries mean that the probe is non-inferior. "W" indicates that the corresponding witness in Figure 22 works, and "O" indicates that an obvious modification of that witness works. By "obvious modification", we mean one that is isomorphic to the original witness as a set coloring game (see Section 2.1). For example, the following is an obvious modification of the probe 3 witness from Figure 22, which works in the presence of a white stone at e:

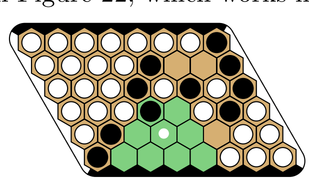

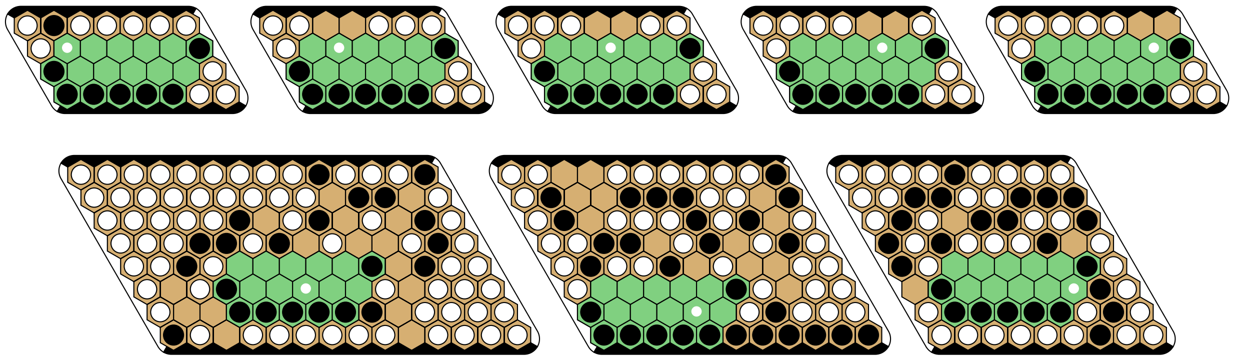

Figure 24: Witnesses for the non-inferiority of all probes except probes 2 and 4 of the long crescent.

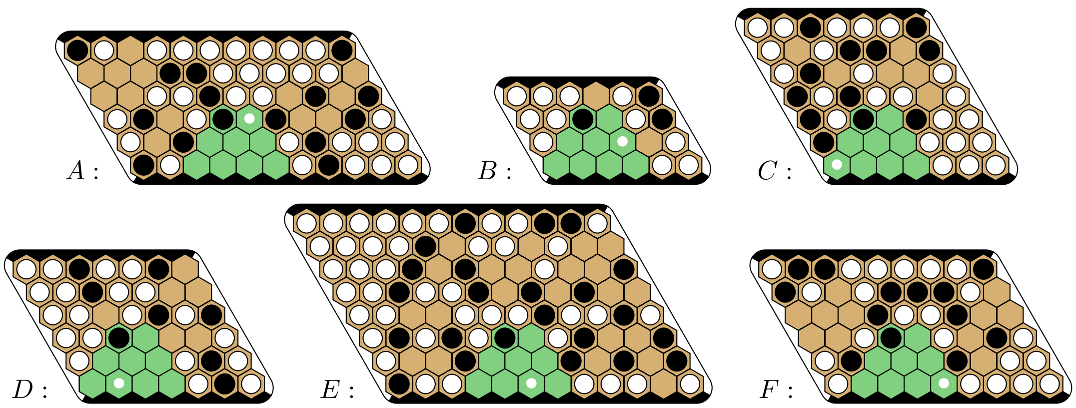

Figure 25: Witnesses for the non-inferiority of some probes in the ziggurat in the presence of certain white boundary stones.

Finally, the letters A–F refer to the six additional witnesses shown in Figure 25.

Conclusions and future work

In this paper, we studied 3-terminal positions in Hex using combinatorial game theory. It turns out that there are many 3-terminal positions with interesting and intricate behaviors. We started by investigating superswitches, an infinite class of positions with remarkable properties. The basic superswitch is a relatively complex Hex position with 12 empty cells, but it has the strikingly simple combinatorial value $\{a, b \mid a\}$. This means that under optimal play, no player can achieve any outcome but a or b, and the left player gets to "set" the switch. This allows the basic superswitch to be used as a building block for making other interesting gadgets, such as higher superswitches and tripleswitches. Among other things, we used this to show that there are infinitely many Hex-realizable and Hex-distinct values for 3-terminal positions. From a theoretical point of view, superswitches are interesting because they form a strictly increasing, cofinal sequence among all games in which Left cannot achieve outcome $\top$ or c. Superswitches (and the related simpleswitches) also have practical applications, such as the verification of connects-both templates.

The most useful tool in our toolbox is a large database of Hex-realizable 3-terminal positions, which we generated by starting from a few simple positions and repeatedly closing them under symmetries, duality, and pinwheel composition. This is how we were able to find Hex implementations of the superswitches and simpleswitches and many other Hex devices. The database is useful because many difficult questions in Hex boil down to constructing Hex positions with carefully crafted combinatorial values. A compelling example of this is a characterization of pivoting templates, a kind of Hex template defined by an intricate balance of threats and answers. We were able to give a combinatorial game theory characterization of two kinds of pivoting templates. Moreover, we found particular 3-terminal contexts that act as perfect tests for pivoting templates: any candidate pivoting template can be mechanically verified by composing it with the corresponding testing context.

In our final application of 3-terminal positions, we settled an open question by Henderson and Hayward on the inferiority of certain probes in the ziggurat template. Inferiority can often be proved using local methods, such as domination and strong reversibility of moves. However, non-inferiority is much harder to show, as it requires the construction of a witnessing position in which the probe in question is the unique winning move. Such witnessing positions are often extremely rare and hard to find. Prior to the present paper, there was no systematic method for constructing them. Our database of 3-terminal positions, along with a process of stepwise refinement (and considerable computational efforts) allowed us to determine the status of all probes in the ziggurat and many other Hex templates.

The paper probably raises as many questions as it answers. Our database provides a practical way of searching for concrete Hex realizations of many combinatorial 3-terminal values. But even the largest such database is necessarily finite. The question remains open whether *every* passable combinatorial 3-terminal value is Hex-realizable (and failing that, how to characterize the ones that are). Progress on this question might allow us to eventually replace our database by a procedure that can compute Hex realizations of arbitrary passable game values.

Another area for future improvement is the stepwise refinement method of Section 7. The theory of abstract probes already proved to be useful in answering open questions, but the process of stepwise refinement is still very labor intensive. The main problem is that, given a probe P and a value A, there is currently no a priori way to predict whether the refined probe $P + A$ will be viable. We also lack a notion of which refinements are "more" viable than others. As a result, there is very little to guide our selection of candidate positions for subdivisions, and we are left to compute a large number of candidates in the hope that enough of them will be viable. It would be nice to have future theorems that can make this process more goal directed.

References

[1] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways for Your Mathematical Plays.* A. K. Peters, 2nd edition, 2001. The first edition was published in 1982.

[2] J. H. Conway. *On Numbers and Games.* A. K. Peters, 2nd edition, 2001. The first edition was published in 1976.

[3] The Demer handicap system. HexWiki article at `https://www.hexwiki.net/index.php/Handicap`, accessed 2025-07-10.

[4] E. Demer and P. Selinger. Database of Hex-realizable 3-terminal values. Available from `http://www.mathstat.dal.ca/~selinger/papers/#hex-3term`. Also available as ancillary material from this paper's arXiv page, 2025.

[5] M. Gardner. Mathematical games: Four mathematical diversions involving concepts of topology. *Scientific American*, 199(4):124–129, Oct. 1958.

[6] R. B. Hayward and B. Toft. *Hex: The Full Story.* CRC Press, 2019.

[7] P. Hein. Vil de laere Polygon? *Politiken*, December 26, 1942. Translated in [6].

[8] P. Henderson and R. B. Hayward. Probing the 4-3-2 edge template in Hex. In H. van den Herik, X. Xu, Z. Ma, and M. Winands, editors, *Computers and Games (CG 2008)*, volume 5131 of *Lecture Notes in Computer Science*. Springer, 2008.

[9] P. Henderson and R. B. Hayward. Captured-reversible moves and star decomposition domination in Hex. *Integers*, 13:G1, 2013.

[10] P. Henderson and R. B. Hayward. A handicap strategy for Hex. In R. J. Nowakowski, editor, *Games of No Chance 4*, volume 63 of *MSRI Publications*, pages 129–136. Cambridge University Press, 2015.

[11] S. Huang, B. Arneson, R. Hayward, M. Mueller, and J. Pawlewicz. Mohex 2.0: a pattern-based MCTS Hex player. In *Proc. Computers and Games CG2013*, volume 8427 of *Lecture Notes in Computer Science*, 2014.

[12] A. Lehman. A solution of the Shannon switching game. *Journal of the Society for Industrial and Applied Mathematics*, 12(4):687–725, 1964.

[13] J. Nash. Some games and machines for playing them. Technical Report D-1164, Rand Corporation, Feb. 1952.

[14] S. Reisch. Hex ist PSPACE-vollständig. *Acta Informatica*, 15:167–191, 1981.

[15] P. Selinger. On the combinatorial value of Hex positions. *Integers*, 22:G3, 2022. Also available from arXiv:2101.06694.

[16] M. Seymour. *Hex: A Strategy Guide*. Online book, available at `http://www.mseymour.ca/hex_book/`, 2019. Accessed: 2022-03-21.

[17] J. Yang, S. Liao, and M. Pawlak. On a decomposition method for finding winning strategy in Hex game. In *Proceedings of the 1st International Conference on Application and Development of Computer Games (ADCOG 2001)*, pages 96–111. City University of Hong Kong, 2001.