

Transfinite game values in infinite games



Davide Leonessi

Supervised by Prof. Joel David Hamkins

MSc in Mathematics and Foundations of Computer Science

Academic year 2020-21

Abstract

The object of this study are countably infinite games with perfect information that allow players to choose among arbitrarily many moves in a turn; in particular, we focus on the generalisations of the finite board games of Hex and Draughts.

In Chapter 1 we develop the theory of transfinite ordinal game values for open infinite games following [Evans and Hamkins, 2014], and we focus on the properties of the *omega one*, that is the supremum of the possible game values, of classes of open games; we moreover design the class of climbing-through- T games as a tool to study the omega one of given game classes.

The original contributions of this research are presented in the following two chapters.

In Chapter 2 we prove classical results about finite Hex and present Infinite Hex, a well-defined infinite generalisation of Hex.

We then introduce the class of *stone-placing* games, which captures the key features of Infinite Hex and further generalises the class of positional games already studied in the literature within the finite setting of Combinatorial Game Theory.

The main result of this research is the characterization of open stone-placing games in terms of the property of *essential locality*, which leads to the conclusion that the omega one of any class of open stone-placing games is at most ω . In particular, we obtain that the class of open games of Infinite Hex has the smallest infinite omega one, that is $\omega_1^{\text{Hex}} = \omega$.

In Chapter 3 we show a dual result; we define the class of games of Infinite Draughts and explicitly construct open games of arbitrarily high game value with the tools of Chapter 1, concluding that the omega one of the class of open games of Infinite Draughts is as high as possible, that is $\omega_1^{\text{Draughts}} = \omega_1$.

*To my Mother,
who gave me the strength to embark on my journey.*

Contents

1	Infinite Games	1
1.1	Game values	6
1.2	The omega one	11
2	Hex	16
2.1	Prelude: finite Hex	16
2.1.1	Fundamental results	16
2.2	Infinite Hex	20
2.2.1	Winning condition	21
2.2.2	Normal Infinite Hex is a draw	25
2.3	Positional and stone-placing games	30
2.3.1	Positional games	30
2.3.2	Biased Infinite Hex may not be a draw	33
2.3.3	Stone-placing games	34
2.3.4	Impossibility of infinite game values	37
2.4	Other generalisations of Hex	42
3	Draughts	44
3.1	Infinite Draughts	46
3.1.1	Embedding trees in the draughtboard	47
3.1.2	Further results	55
A	Appendix	60
A.1	2-colourability in Maker-Breaker games	60
A.2	Strategy-stealing in stone-placing games	61
	Bibliography	63

Chapter 1

Infinite Games

The object of the present study are games played on infinite structures that generalise finite board games and that can thus allow players to choose among arbitrarily many moves in a turn.

In game-theoretic terms, we will consider two-player zero-sum infinite games with perfect information and no chance moves; we will define such games in extensive form, so to represent them as their induced game tree.¹

We will focus on the class of games in which the order type of each play is at most ω , so that plays consist of at most countably many turns, and, at each turn, players can choose among at most countably many moves. The game trees of such games are naturally seen as sub-trees of the following.

Definition 1.1. The *full countable tree* $\omega^{<\omega}$ is a directed tree whose vertex set is the set of finite sequences of natural numbers, that is

$$\omega^{<\omega} = \bigcup_{n \in \omega} \omega^n = \bigcup_{n \in \omega} \{v : n \rightarrow \omega\},$$

and whose directed edges order such finite sequences by inclusion.

Note that we can label the edges of the full countable tree by ordinals smaller than ω so that each node v is exactly the sequence of edge labels that compose the unique path that ends at v and starts at the root node \emptyset , i.e. the empty sequence.

Observe that the set of *branches* of $\omega^{<\omega}$, which are the maximal paths in $\omega^{<\omega}$ starting at the root node \emptyset , is the collection of infinite sequences $\omega^\omega = \{s : \omega \rightarrow \omega\}$.

For the rest of this work, we will say that a collection of objects $T = (V, E, v_0)$ is a *tree* if T is a directed tree with vertex set V , edge set E , and root node $v_0 \in V$, which can be embedded into the full countable tree $\omega^{<\omega}$; that is, there is a

¹Infinite games in extensive form were first studied by [Gale and Stewart, 1953].

map $\tilde{\iota} : T \rightarrow \omega^{<\omega}$ induced by an injection $\iota : V \hookrightarrow \omega^{<\omega}$ which preserves the tree structure.²

If $(u, v) \in E$ is an edge of T , then we say that v is a *child* node, or *immediate successor*, of u ; conversely, we say that u is the *parent* node, or *immediate predecessor*, of v ; we write $u \rightarrow_T v$.

With this notation we can say that the injection $\iota : V \hookrightarrow \omega^{<\omega}$ induces a structure-preserving embedding in the sense that it preserves the root node and the edges, so that $\iota(v_0) = \emptyset$ and $\iota(u) \rightarrow_{\omega^{<\omega}} \iota(v)$ if and only if $u \rightarrow_T v$.

We can now define the games we consider in extensive form.

Definition 1.2. A *game* \mathcal{G} is a collection of objects $((V, E, v_0), S, S_1, S_2)$, such that;

- i. V is the set of *positions*, and $v_0 \in V$ is the *initial position*;
- ii. (V, E, v_0) is a tree, which we call the *game tree*;
- iii. At each position, a player chooses a child node of that position in the game tree as the next position; the players alternate making moves and the turn is determined by the parity of the number of moves already made;
- iv. S is the collection of branches of (V, E, v_0) , which we call the set of *plays*;
- v. $S_1, S_2 \subset S$ are the *winning conditions* of, respectively, the first and second player, such that $S_1 \cap S_2 = \emptyset$; a play $s \in S$ is a *win* for the first or second player if, respectively, $s \in S_1$ or $s \in S_2$; moreover, a play $s \in S$ is a *draw* if $s \notin S_1 \cup S_2$.

Remark 1.3. As the game tree T of condition ii. is a directed rooted tree, we are assuming that every non-initial position has a unique predecessor; since each branch of such T has order type at most ω , then each position defines a unique history of finitely many moves from the initial position v_0 . In other words, any partial play consists of only finitely many moves.

Hence, we can see that condition iii. has a well-defined meaning in the sense that the first player is the next to move at all positions whose unique path from the initial position comprises an odd number of vertices; this includes the initial position v_0 . Conversely, the second player moves at all positions whose unique path from the initial position consists of an even number of positions.

Remark 1.4. Note that in the majority of popular board games occurs some *board position*, or *board configuration*, which can be achieved by several distinct move sequences; we consider each move sequence as leading to distinct *positions*.

²Note that $\omega^{<\omega}$ refers to both the vertex set of the full countable tree and the tree itself.

Thus, we can see that positions carry more information than board positions; in addition to determining whose turn it is, a position uniquely specifies the history of moves from the initial position, and so all the board positions already achieved over the play.

Remark 1.5. In the games that we consider in this work we will show that not both players can win, satisfying the disjointness of S_1, S_2 required in condition v.

More in general, plays that allow both players to win can be considered draws without loss of generality.

Remark 1.6. As mentioned above, we are interested in games which feature positions with infinitely many successors, that is, with some positions $v \in V$ which have infinitely many children nodes.

It is important to note that the games defined here allow each position to have at most *countably* many successors; this follows from our definition of trees as subtrees of the full countable tree.

Fixing a game \mathcal{G} as in definition 1.2, we define the following common game-theoretic concepts; the definitions for the first player have a natural analogue for the second player.

Definition 1.7. Say that $V_1 \subset V$ is the collection of positions at which the first player is the next to move. A *strategy* for the first player of \mathcal{G} is a function $\sigma : V_1 \rightarrow V$ such that $\sigma(v)$ is a child node of v in the game tree of \mathcal{G} for all $v \in V_1$.

We say that a play $s \in S$ *conforms to* a strategy σ for the first player if, for any position v in s at which the first player is the next to move, i.e. $v \in V_1$, we have that the successor of v in s is $\sigma(v)$; in other words, the second player has a way to play the game that leads to the play s , when the first player follows σ .

Definition 1.8. A *winning strategy* for the first player of \mathcal{G} is a strategy σ such that, for any play $s \in S$ conforming to the strategy σ for the first player, we have that $s \in S_1$; i.e. any play that conforms to such strategy is a win for the first player.

Recall that the set of branches of the full countable tree is ω^ω ; we can turn this set of sequences into a space by imposing on it the product topology induced by the discrete topology on ω .

A base for such topology is given by the collection of position neighbourhoods $\mathcal{U}(v)$ for $v \in \omega^{<\omega}$, where $\mathcal{U}(v) = \{s \in \omega^\omega : v \subset s\}$, i.e. $\mathcal{U}(v)$ is the collection of all branches which have the path from v_0 to v as an initial segment.³

We can similarly impose a topology on the set of branches S of any tree $T = (V, E, v_0)$. Let $\{\mathcal{U}(v)\}_{v \in V}$, the collection of neighbourhoods of finite paths in T , be the base for the topology on S .

³Clearly, $\mathcal{U}(\emptyset) = \omega^\omega$.

Observe that this topology is equivalent to the subspace topology induced on S from ω^ω via any structure-preserving embedding $\tilde{\iota} : T \rightarrow \omega^{<\omega}$.

In particular, if T is the game tree of some game \mathcal{G} , then we can impose this topology on the set of plays S ; in this case we can interpret the neighbourhood $\mathcal{U}(v)$ of a position v to be the collection of plays that complete the partial play up to v .

Now, if the winning condition $S_1 \subset S$ is open, then S_1 is equal to a union of position neighbourhoods $\bigcup_{i \in \mathcal{I}} \mathcal{U}(v_i)$ for some nodes $v_i \in V$; hence, for a given play $s \in S_1$, there is some node v_j such that $s \in \mathcal{U}(v_j) \subset S_1$.

Moreover, any play s' that contains v_j is necessarily a win for the first player, as $s' \in \mathcal{U}(v_j) \subset S_1$; that is, the outcome of any play that at some point reaches v_j is determined at that position and cannot be changed by any move made after v_j by either player. We can then say that any such play s' , including s , is ended in v_j , and is thus finite or *essentially finite*.

We can formally define an essentially finite play as a an infinite play whose outcome is the same of all the plays in the neighbourhood of one of its positions; that is, a play $s \in S$ is essentially finite if $s \in S_1$ (or, respectively, S_2 , $(S_1 \cup S_2)^c$) and there is some position v in s such that $\mathcal{U}(v) \subset S_1$ (or, S_2 , $(S_1 \cup S_2)^c$). We can observe that there is no need to distinguish between such plays and finite plays, so that we will not.

In fact, consider a finite play s of, say, usual Chess; the leaf node in s is the ending position, which occurs after finitely many turns, according to the standard rules. Then, we can extend s to an infinite sequence of moves by allowing the players to make trivial moves in which they acknowledge the end of the game by alternately saying “OK” to each other ad infinitum.

However, we can naturally define the “actual” length of a (essentially) finite play s as the minimal⁴ n for which there is a sequence of n moves from the initial position to a node $v \in s$ and $s \subset \mathcal{U}(v)$.

We will focus on games which, even though infinite, enjoy the following finiteness property.

Definition 1.9. A game is *open* for one player if any win by that player takes place in finitely many moves; that is,⁵ if the winning condition for that player is open in the space of plays.

We similarly define the following.

⁴It is immediate to see that if $v \in s$ and $s \subset \mathcal{U}(v)$, then also $s \subset \mathcal{U}(w)$ for all the descendants w of v in s .

⁵As remarked by [Evans et al., 2017, §1].

Definition 1.10. A game is *closed* for some player if his winning condition is closed in the space of plays; that is, all the plays which are not wins for that player are determined to be so in finitely many moves.⁶

A game is *clopen* for some player if it is open and closed for that player.

Remark 1.11. If a game has no draws, then it is clopen for some player if and only if it is clopen for that player's opponent.

More in general, this implication holds when the set of draws is clopen in the space of plays.

In the majority of board games, players have essentially analogous goals, so that the direct infinite generalisations of such games are often open for both players,⁷ as we can see for instance in Chess, which was studied in detail by [Evans and Hamkins, 2014].

Example 1.12. Infinite Chess.

This game is played on the infinite chessboard, which is a square tiling of the plane in bijection with $\mathbb{Z} \times \mathbb{Z}$, and features the usual Chess pieces, generalising the long-distance ones, i.e. rook, bishop, and queen, to move for arbitrarily large, but finite, distances on the board.

The winner of a game of Infinite Chess is naturally defined as the first player that places the opponent's king in a checkmate position, which can only happen after finitely many moves. It follows that all the plays that involve infinitely many turns are draws.

As a consequence of the definition of winning condition, Infinite Chess is an open game for both players, White and Black, given any starting position.

We can modify the winning conditions to make Infinite Chess not open for one player. Suppose we have an initial position in which Black has one king, while White has none; allow White to win if he satisfies the previous checkmate condition, and let Black win if the game continues for infinitely many turns,⁸ so that his king is never in checkmate. This game is clearly open for White, but not for Black.

This example highlights the importance of both the winning conditions and the initial position in the definition of a game.

We will see later another game open for only one player; the Angel and Devil game in Example 1.26.

⁶However, it may not be determined in only finitely many moves whether such plays are draws or wins for the player's opponent.

⁷A game open for both players is not necessarily clopen, as it may allow draws to occur with infinite plays, so that the game may not be closed for either player.

⁸Such an infinite play would be a tie according to the previous winning conditions.

1.1 Game values

The class of open games allows for the definition of ordinal game values within the framework of infinitary game theory.

We present the theory of transfinite game values in infinite games as developed by [Evans and Hamkins, 2014].

We approach the concept of game value by borrowing another idea from Chess.

Example 1.13. The mate-in- n problem.

Often presented as a puzzle for small values of n , a Chess position is defined to be *mate-in- n* for White if there is a winning strategy for White such that all plays conforming with it are wins for White after at most n moves⁹ and, moreover, there is a strategy for Black such that all plays conforming with it comprise at least n moves;¹⁰ we call this strategy for White a *win-in- n* strategy.

In such a situation we can think of White as being at most n moves away from winning; if n is minimal, then the position is said to have game value n for White.

For instance, observe that White has already won in a mate-in-0 position, and White can force a checkmate from a mate-in-1 position on his next move; in the latter case, White's move may depend on Black's move, if White plays second.

We can generalise this concept to the infinite case for the smallest infinite ordinal, ω .

An Infinite Chess position is defined to have game value ω for White if White has a strategy to win in finitely many moves, but Black is next to move and his initial choice determines a finite lower bound for the number of turns that need to be played before Black's inevitable defeat.

In other words, Black makes a move from a position with game value ω to a mate-in- n position, where he can choose $n \in \mathbb{N}$ to be arbitrarily large.

Observe the example¹¹ of a position with game value ω for White in figure 1.1 which features infinitely many pawns, as indicated by the pattern.

In that position, Black is the next to move and the central Black rook is attacked by a pawn; the initial move by Black is moving the rook up the empty corridor for an arbitrarily long, but finite, distance.

White can then capture the Black rook with a pawn and systematically advance the pawns below the resulting empty square, until a White rook is let free to mate the Black king. The Black pawn on the side prevents a stalemate.

We have now given enough motivation to define the following.

⁹Clearly White could play sub-optimally, without winning in n moves; this does not change the fact that the position itself is mate-in- n .

¹⁰In both cases, only the moves by White are counted.

¹¹[Evans and Hamkins, 2014, Figure 5]




Figure 1.1: Infinite Chess position with game value ω for White, Black to move.

Definition 1.14. The *game value* of a position p in a game open for a player O , playing against C , is defined by recursion as follows;

- if O has already won in p , then p has game value 0;
- if O is the next to move in p , and α is the smallest ordinal among the game values of the positions¹² that can be reached from p via a legal move by O , then p has game value $\alpha + 1$;
- if instead C is the next to move in p , and all the positions that can be reached from p via a legal move by C have a defined game value, then the value of p is the supremum of all these values.

Note that this definition may not apply to all positions of a game; the positions with defined game value for some player are exactly the positions from which that player can force a win, which is reached in finitely many moves by openness of the game.

Observe that Definition 1.14 reflects the rational behaviour of the players in a game, from the perspective of player O . The *open player* O aims to reduce the game value as quickly as possible, when it is defined, and to reach a position with defined value otherwise; on the other hand, the *closed player* C has the opposite aim and tries to keep the game value undefined, if possible, or as high as possible, when the value is defined.

Suppose that a position p has ordinal value $\alpha > 0$ for O , so that, by definition of game value, if O is the next to move from p , then O can reach a position with game value strictly smaller than α ; if instead C is the next to move from p , then C can only reach positions with game value $\beta \leq \alpha$, in particular $\beta < \alpha$ if α is a limit ordinal.

Thus, O can play from p so to strictly reduce the game value at each turn, while C cannot reach a position with higher, or undefined, game value; we call

¹²Only the positions for which a game value can be defined.

this the *value-reducing* strategy for the winning player. By well-foundedness of the ordinals, it is clear that a winning player following such a value-reducing strategy does eventually reach a winning position of value 0 in finitely many turns, as there are no infinite descending sequences of ordinals.

Conversely, if p has no defined game value for O , then O cannot make a move to reach another position with defined game value, as otherwise p would have had a game value in the first place. Moreover, starting from such p , there is at least one position with undefined game value that C can reach with one move, again by Definition 1.14; this gives rise to the *value-maintaining* strategy, with which C can keep the game value undefined at all positions reached over the play.

We conclude that the positions without a game value for O are exactly the positions from which C can force either a draw or a win, meaning that C has a drawing or winning strategy for the game starting at p .

Furthermore, if a game has no draws and is open for both players,¹³ then the positions with defined game value for one player have undefined value for the other, and vice versa.¹⁴

We will refer to the game value of a game to mean the game value of its initial position.

We can provide a visual intuition of what it means for a position p to have game value for O by fixing a winning value-reducing strategy for the open player O and thus observing that the way the play unfolds depends entirely on the losing player C ; by considering the “partial game tree” that represents only the choices available to C , we obtain the structure of the following game.

Example 1.15. The climbing-through- T game.

Let T be a tree; recall that T is a rooted directed tree whose branching nodes have at most countable degree, and whose branches have at most order type ω .

Let two players, Climber and Observer, play the following game.

Climber aims to climb up the tree, starting from the root; at each turn, if Climber is on a branching node, then he can move to one of its successor nodes, if instead Climber is standing on a leaf node, then he loses.

On the other hand, Observer does not directly influence the game; Observer simply acknowledges Climber’s moves by saying “OK”, until Climber reaches a leaf node, making him win.

This is clearly a game open only for Observer, as Climber wins exactly when he manages to climb up an infinite branch of T , thus playing for infinitely many turns.

¹³Such a game is also clopen.

¹⁴This is true since we only consider games in which at most one player wins, as per the definition of winning condition in 1.2.v.

In particular, Climber can win if and only if T has an infinite branch, and Observer has a (trivial) winning strategy if and only if T is well-founded, that is, if all the branches of T have finite length.

The simplicity of climbing-through- T games enables us to assign values easily to such games; these will be a key tool in the construction of Infinite Draughts positions in Chapter 3.

Observe that \rightarrow_T is a strict partial order¹⁵ on the vertex set of T ; we can thus apply the theory of strict partial orders on well-founded trees, as described in [Kunen, 2013, §III.5], to establish a link between tree rank and game value.

Define the *rank*, or *height*, of a node of a well-founded tree as $\text{rank}(u) = \sup\{\text{rank}(v) + 1 : u \rightarrow_T v\}$; the rank of a well-founded tree is defined to be the rank of its root.

We can see that the leaves of the tree have rank 0, as they have no children. Nodes whose children are all leaves have rank 1; in general, the rank of a branching node is the supremum of the ordinal successors of the ranks of the nodes it is the parent of.

We can now easily prove by transfinite induction the following.

Proposition 1.16. The rank of a node v of a tree T is exactly equal to the game value for Observer of the position of the climbing-through- T game in which Climber is standing on v .

Proof. The nodes of T with rank 0 are leaves, which are precisely the nodes on which Climber stands when Observer has won.

Say v is a node of T with rank $\beta + 1$; then, the children of v have rank bounded by β , and one of them, say w , has exactly rank β . When standing on v , Climber can only move to positions with defined game value bounded by β , by inductive hypothesis; by climbing one step up to w , Climber reaches a position of game value exactly β , again by hypothesis.

Hence, Climber was at a position with value $\beta + 1$, when standing on v .

Say v is a node with a limit ordinal γ as rank. From v , Climber can only reach nodes of arbitrarily high rank $\alpha < \gamma$, which correspond to positions of value $\alpha < \gamma$ for Observer, by inductive hypothesis.

This completes the inductive proof by definition of game value.

Moreover, note that v has undefined rank exactly when it is contained in an infinite branch of T ; in that case, the corresponding position of climbing-through- T has undefined game value for Observer since Climber has a winning strategy, namely climbing from v up such infinite branch of T . \square

¹⁵A *strict order* is an irreflexive and transitive relation, which is thus antisymmetric.

We present the concrete example of a tree of rank $\omega + 3$ in figure 1.2,¹⁶ where the nodes are labelled by their rank.




Figure 1.2: A well-founded tree with nodes labelled by their ranks.

Note in particular the node with rank ω ; it is a branching node of countable degree whose children have arbitrarily high, but finite, rank $n \in \mathbb{N}$.

We conclude this example justifying the presentation of T as a “partial game tree”; T really is the tree obtained by the game tree of the climbing-through- T game by joining each decision node of Observer to its preceding node, and deleting the edge between them.

Setting the scene for the next section, we now show that the ordinals that can be realised as game values of positions in a game define the initial segment of some ordinal.

Proposition 1.17. Let p be a position in a game with game value α for some player; for any ordinal β such that $\beta \leq \alpha$, there is a finite sequence of legal moves that from p reaches some position q with game value β for the same player.

Proof. Let \mathcal{G} be a game played by O and C , and let it be open for O . The position p has value α for the open player O .

We prove the result by transfinite induction on α ; the base case $\alpha = 0$ is trivial.

Let $\alpha = \delta + 1$ be a successor ordinal and take $\beta \leq \delta < \alpha$.

¹⁶Even though the height of nodes in a tree increases towards the root, we represent trees with the root down and the leaves up, because that is what “they do in nature”, as remarked by [Kunen, 2013, p. 202].

If O is the next to move in p then, by definition of game value, he can make a move to some position with game value δ ; then there is a sequence of moves that leads to a position of value $\beta \leq \delta$ by inductive hypothesis.

If C is the next to move, then $\delta + 1$ is the supremum of the game values of all positions reachable by C ; in particular, she can reach some position also of value $\delta + 1$ in which O is next to move. The result follows from the previous case.

Let $\alpha = \gamma$ be a limit ordinal and take $\beta < \gamma$.

Recall that a limit ordinal can be achieved as game value of some position only if the losing player is the next to move. Hence, C is the next to move in p .

Again by definition, γ is the supremum of the game values of all positions in the reach of C ; in particular, she can reach positions of value arbitrarily high, but strictly bounded by γ . Hence, C can make a move to reach a position q with game value $\delta < \gamma$ such that $\beta \leq \delta$. We can now apply the inductive hypothesis to q , which concludes the proof. \square

1.2 The omega one

We can informally think of high-value positions as evidence of the “strategic complexity” of a specific open game. We formalise this idea in the following definition that we make for a *class of games*, that is a collection of games constructed according to the same rules, but with different initial positions.

Definition 1.18. The *omega one* of a class of games \mathbb{G} , all open for some player, is the supremum of the game values of all the games in \mathbb{G} , for that player. We denote it by $\omega_1^{\mathbb{G}}$.

Recall that we only consider games in which plays are sequences of at most countably many moves and, at each turn, players can choose one of up to countably many moves. In fact, we will only consider games played by moving or placing pieces on boards of countably infinite size.

Observe that, in such a game \mathcal{G} , the game value of all positions will be the supremum of at most countably many ordinals, and if all such ordinals are countable, then ω_1 cannot arise as game value of any position.

The result below follows by transfinite induction.

Proposition 1.19. For any open game \mathcal{G} , we have that the game value of \mathcal{G} is strictly smaller than ω_1 .

Corollary 1.20. For a class of games \mathbb{G} , all open for some designated player, we have that $\omega_1^{\mathbb{G}} \leq \omega_1$.

Remark 1.21. Proposition 1.19 and Corollary 1.20 crucially rely on game positions having at most countably many successors, which is an assumption that we make throughout this work, as highlighted in Remark 1.6.

These results do not hold for classes of more general games, such as the one mentioned later in footnote 17.

We will see now that this bound is sharp for the class of climbing-through- T games, which are open for Observer.

Lemma 1.22. For any ordinal $\alpha < \omega_1$, there is a well-founded tree T of rank α .

Proof. We prove this by transfinite induction.

If $\alpha = 0$, then T is the trivial tree with only one node.

If $\alpha = \beta + 1$, then take a tree T' of rank β by inductive hypothesis and simply let T be a tree obtained by adding an extra node v to T' and let v have the root of T' as only child; v is the root of T , which thus has rank $\beta + 1 = \alpha$.

If α is a limit ordinal, then let T_β be a tree of rank β for all $\beta < \alpha$. Define $T = \bigoplus_{\beta < \alpha} T_\beta$ similarly to the previous construction; T is obtained by letting an extra node v have as children the roots of each T_β , for $\beta < \alpha$. If T has v as root, then it is clearly as needed.

Note that this construction of T satisfies our earlier definition of tree because α is a countable ordinal, and so the root node v of T has countable degree. \square

Remark 1.23. A climbing-through- T game has a defined game value for Observer if and only if T is well-founded, otherwise the game would be a win for Climber, as we saw in the proof of Proposition 1.16.

Proposition 1.24. The omega one of the class of open climbing-through- T games is ω_1 .

Proof. By Lemma 1.22 and Proposition 1.16, there are climbing-through- T games of arbitrarily high value $\alpha < \omega_1$; therefore, the omega one of the class of climbing-through- T games is at least ω_1 .

If there were some climbing-through- T' game of value ω_1 , then there would be some well-founded tree T' of rank ω_1 by Proposition 1.16, so that the root of T' would have uncountable degree; this is a contradiction because we assume trees to be sub-trees of the full countable tree.

Invoking the contrapositive of Proposition 1.17, we confirm the bound given by Corollary 1.20 and complete the proof. \square

Remark 1.25. Observe that no climbing-through- T game can have value ω_1 since the ranks of well-founded trees, which correspond to the possible initial positions, are countable.

Therefore, a necessary condition to have uncountable omega one for the class of climbing-through- T games, as for any other game class,¹⁷ is to have uncountably many possible initial positions.

In order to stress the importance of infinite branching nodes in achieving infinite game values, we propose the following game defined by [Conway, 1996], which is an interesting example of game open for only one player.

Example 1.26. Angel and Devil game.

Let two players, Angel and Devil, play on an infinite chessboard, which is in bijection with $\mathbb{Z} \times \mathbb{Z}$, as usual.

On his turn, the Devil can mark, i.e. *burn*, any square of the chessboard; such square becomes unavailable to the Angel.

At the initial position, the Angel stands on the origin (0,0) and, at each turn, moves by at most p chess king's moves away from the previous square; that is, the Angel can move from (x, y) to (x', y') if $|x - x'|, |y - y'| \leq p$ and (x', y') is not burnt by the Devil; we say that the Angel has *power* p .

The Devil wins if the Angel remains trapped, that is, at some turn the Angel has no legal moves; conversely the Angel wins by remaining free and by having a legal move to make at all stages of the play.

Observe how Conway's game, which is played on an infinite board, does not allow for infinite game values for the Devil because of the finite choice of moves of the Angel.

The game is clearly open only for the Devil, while the Angel wins exactly when both players make infinitely many moves. Thus, assume that some position has infinite game value, where the Devil is winning; by Proposition 1.17, it is enough to take it to be ω .

By definition of game value, the Angel, who is the losing player, has to make the first move. But the Angel can only choose one of finitely many moves; namely, the Angel can only choose among the at most $(2p + 1)^2$ cells within the square allowed by its power p which have not yet been burnt by the Devil.

But then, ω would be the supremum of at most $(2p + 1)^2$ finite values, which is a contradiction.

The only way to obviate to this issue would be giving 'infinite power' to the Angel, allowing it to move to an arbitrary cell of the infinite board, but that would lead the Angel to trivially win the game against a Devil that only burns finitely many cells at each turn; such an Angel of infinite power would win even against a

¹⁷We could consider a climbing-through- T *hypergame* in which Climber makes the first move and chooses a well-founded tree T on which to play; it follows from our discussion that such a hypergame would have uncountable game value.

Devil who burns coinfinately many cells at each turn, that is, if the Devil always leaves infinitely many cells available for the Angel to move on.

We conclude that, generalising the choice of initial position, the class of games induced by the Angel and Devil game has omega one at most ω .

Remark 1.27. Conway originally asked for a proof of the intuitive fact the Angel and Devil game starting from an empty board is a win for the Angel.

It was proved by Berlekamp¹⁸ that an Angel of power 1, that is a chess king, can be trapped on any sufficiently large finite board; the question remained open for a decade, until [Máthé, 2007] proved that the Angel of power 2 wins against the Devil.

Proposition 1.28. Let \mathbb{G} be a class of games open for player O , and say that his opponent C has only finitely many moves to choose from at each turn.¹⁹ Then, $\omega_1^{\mathbb{G}} \leq \omega$.

Moreover, if O has a winning strategy for some game \mathcal{G} in \mathbb{G} , then O has a win-in- $n_{\mathcal{G}}$ strategy²⁰ for some $n_{\mathcal{G}} \in \mathbb{N}$.

Proof. The discussion in Example 1.26 shows that no infinite ordinal can be realised as the value of a position of the Angel and Devil game, and so the omega one of its induced class of games is at most ω .

That argument relies only on Proposition 1.17 and on the fact that the Angel has finitely many moves to choose from when at a position of value ω , so we can directly generalise it and obtain that $\omega_1^{\mathbb{G}} \leq \omega$.

Suppose now that O plays according to some winning strategy σ in a game \mathcal{G} of \mathbb{G} . Let p be the first position in which C is the next to move.²¹

Note that C can reach only finitely many positions from p , and all such positions have finite game value for O ; let m be the supremum of all such values.

It is clear that m is finite and that σ is a win-in- n strategy for $n = m$ or $(m + 1)$, depending on who moved first, so that the longest play of \mathcal{G} ends after n moves by O . \square

We conclude by focussing our attention more narrowly than as stated at the beginning of this chapter; the games which are interesting for our analysis are open for some player and allow that player's opponent to have a strategy ensuring that she can choose among infinitely many possible moves at least once in any play.

¹⁸According to [Conway, 1996, §1].

¹⁹Observe that a finite bound for the number of possible moves for C across all turns is not required.

²⁰See the mate-in- n problem in Example 1.13.

²¹ p is either the initial position, or the one reached by the first move of O .

We will prove in Chapter 2 that in Infinite Hex the freedom of infinite choice granted to the losing player is not enough to construct a proper obstruction to the winning player's main line of play. The class of open games of Infinite Hex has the smallest possible omega one, that is $\omega_1^{\text{Hex}} = \omega$; we can interpret this as saying that the complexity of the infinite generalisation of Hex is lost when we restrict our analysis to games open for some player.

In Chapter 3 we will show a dual result; the class of open games of Infinite Draughts features positions of arbitrarily high game value, so that the losing player's obstructing power is put to good use, and hence $\omega_1^{\text{Draughts}} = \omega_1$. This means that we can directly appreciate the strategic complexity of the infinite version of Draughts in games open for some player.

Chapter 2

Hex

2.1 Prelude: finite Hex

The game of Hex was first designed by the inventor and polymath Piet Hein, who published it as “Polygon” and made it become a popular board game in Denmark, starting in 1942.¹

Hex was then rediscovered by Nobel laureate John Nash when still a student at Princeton in 1948; he described it as a “matter of connecting topology and game theory”.²

From the common room of the Department of Mathematics at Princeton, the then-called game of “Nash” became popular in the mathematical community, and eventually published as the mainstream board game “Hex” in 1952.

Hex is played on a rhomboidal board of $n \times n$ hexagonal tiles by two players, Red and Blue, who alternate placing stones on the empty tiles of the board; to each player is assigned a pair of opposite board sides and the player who joins his assigned sides with a connected chain of stones of his respective colour is the winner.

See in figure 2.1a the initial position and in 2.1b the ending position of a play won by Red.

2.1.1 Fundamental results

Remark 2.1. Observe that the following Theorems 2.2 and 2.5 rely on the topological properties of the boundary of the finite Hex board; namely, that the boundary is made of 4 contiguous segments alternatively assigned to the two players.

¹[Hein, 1942]

²According to a private communication by [Björnsson et al., 2006, §3].




Figure 2.1: The 6×6 Hex board.

Theorem 2.2 (Hex Theorem). At a finite Hex position in which all the tiles are marked by either Red or Blue, at least one of the players has won.

The Hex Theorem is the result that carries the highest mathematical significance of finite Hex; [Gale, 1979] showed it to be equivalent to the Brouwer Fixed Point Theorem, which in turn was shown to be equivalent to the Jordan Curve Theorem by [Maehara, 1984] and [Adler et al., 2016].

In order to present the argument by [Gale, 1979], we need to mention the following simple graph-theoretic result.

Lemma 2.3. A finite graph whose vertices have degree at most 2 is the union of disjoint subgraphs, each of which is either (i) an isolated vertex, (ii) a simple cycle, or (iii) a simple path.

Proof of Theorem 2.2. Let Γ be the edge graph of the Hex board with the addition of 4 edges ending in the vertices n, e, s, w , as in figure 2.2a.

Now, given a completely marked board, we present an algorithm to find a winning chain of stones.

Starting from s , we construct a tour along Γ following the rule of always proceeding along an edge which is the common boundary of a cell with a Red stone and a cell with a Blue stone; observe that the edge starting from s does have this property, by considering the quadrants outside of the board as assigned to the players, see figure 2.2b for an application of this touring rule starting from all the vertices of Γ .

Observe that this touring rule determines a unique path; suppose that the tour has proceeded along some edge d and arrives at some vertex v . Two of the three tiles incident to v are those of which d is the common boundary, hence one has a




Figure 2.2: The argument for Theorem 2.2.

Red stone and the other a Blue stone; the third tile incident to v may have a Red or a Blue stone, but in either case there is exactly one edge d' which satisfies the rule described.

As a consequence of starting from s , we have that such a tour will never revisit any vertex of Γ ; however, Γ is finite and so the tour must terminate.

It is clear from the above that such a tour must describe a subgraph of Γ whose vertices have degree at most 2; hence, by Lemma 2.3, the tour starting at s can only be a simple path, thus ending at either w , n , or e . Observe that each of these vertices is incident to either the NW Blue quadrant or the NE Red quadrant.

If, say, the terminal vertex of the described tour is incident to the NE quadrant, then Red has a winning chain which is incident to the edge tour starting in s and joins the NE side with the SW side of the board; the other case for Blue is analogous. \square

Remark 2.4. Observe that this proof does not require any “right-left” orientation of the tour.

Moreover, we do not even require that the fully marked board is the ending position of some legal play.

Theorem 2.5. There is no finite Hex position in which both players have won.

This Theorem is usually proved via an application of the Jordan Curve Theorem; we propose another visual proof by [Schachner, 2019].

Proof of Theorem 2.5. Similarly to the previous proof, we add one distinguished vertex for each side of the board and an extra one.

Supposing for a contradiction that both Red and Blue win, we can use the two winning paths in the ending board position to construct the full graph with 5 vertices, as in figure 2.3.

However, such a graph cannot be planar; contradiction.




Figure 2.3: Obtaining the graph K_5 from a Hex position.

□

Historical Remark 2.6. According to [Hayward and Toft, 2019, pp. 7–8], Piet Hein was aware of this alternative proof; he realised that the (then unproved) Four Colour Theorem implied that there is no planar full graph with 5 vertices. He mentioned this as the foundation for his intuition behind the design of Hex.

From Theorems 2.2 and 2.5 we deduce the following;

Corollary 2.7 (No-tie property). All plays of finite Hex end with exactly one winner.³

Proposition 2.8 (Strategy-stealing). In finite Hex, the second player does not have a winning strategy.

Remark 2.9. We observe the following basic properties of finite Hex.

- i. It is never disadvantageous for any player to have an extra stone on the board,⁴ as first observed by the discoverer in [Nash, 1952].
- ii. There exists a rigid motion of the board B into its *opposite* – that is, B can be rotated, translated, and possibly reflected, into the board consisting of B with each boundary segment assigned to the opposite player.

³We can unambiguously say that plays end because they can last only for as many turns as there are tiles on the finite Hex board.

⁴Note how in a *dynamic game* in which pieces can change position like Chess, for instance, having an extra piece that obstructs the escape of one's king can facilitate a checkmate. Or observe that it is possible to construct dynamic game positions in which making more than one move in a turn can put a player at a disadvantage.

Proof of Proposition 2.8. Suppose for a contradiction that Blue, the second player, has a winning strategy σ .

Let σ' be the *opposite* strategy of σ , that is the corresponding strategy on the opposite board, in the sense of Remark 2.9.ii.

Now, let Red, the first player, place a stone on an arbitrary tile t and then play according to σ' ignoring the extra tile t , so to pretend to play as second. If, at some turn, σ' prescribes Red to play on t , then Red can place a stone on another arbitrary tile t' and continue to play ignoring it.

Note that the strategy for Red is winning, exactly because σ is a second-player winning strategy and because ignoring his initial move can not obstruct his win, as observed in Remark 2.9.i.

But now both players can have a winning strategy; this is a contradiction by Corollary 2.7. \square

Observe how the symmetry property of Remark 2.9.ii is crucial for the argument above; it would not be possible to let Red use Blue's strategy otherwise, as the board, and so the relevant winning conditions, would have “looked different” to each of them.

A natural example of boards that satisfy condition 2.9.i. but fail to satisfy condition 2.9.ii. is that of asymmetric $m \times n$ Hex boards.

Aiming to balance the game of Hex in favour of the otherwise losing second player, it seems reasonable to use asymmetric boards; however, it turns out that the player playing along the short edge of the board has such a strong advantage to be able to win even when playing second.

In fact, it is possible to pair the tiles of any $(n + 1) \times n$ Hex board so that the player joining the long edges wins by choosing, at each turn, the cell paired to the last cell taken by the opponent, as observed by [Gardner, 1988, §8]. When making the initial move, the player following this strategy can just make an arbitrary move as in the proof of Proposition 2.8.

This is clearer in the example of figure 2.4.

We will see in section 2.2.2 that a direct generalisation of this strategy will provide a drawing strategy for the second player in the game of Infinite Hex starting from an empty board.

2.2 Infinite Hex

It is immediate to generalise the board of Hex to the infinite plane; Infinite Hex is a class of games in which the two players, still Red and Blue, alternate placing their respective stones on the hexagons that tile the infinite plane.




Figure 2.4: The pairing strategy on the 5×4 Hex board.

Note that the hexagonal tiles of the Infinite Hex board are in bijection with $\mathbb{Z} \times \mathbb{Z}$, so that the players, at each turn, can choose to mark one of at most countably many tiles; hence, the game tree of any game of Infinite Hex can always be embedded in the full countable tree $\omega^{<\omega}$, as we need.

In other words, any position of Infinite Hex, including the positions with infinitely many stones already placed, is the initial position of a well-defined game of Infinite Hex.

We next discuss the definition of winning condition, which is less straightforward.

2.2.1 Winning condition

Since the Infinite Hex board has no edges to join, it is not immediately clear how to define a winning condition; we would like to “join the opposite infinite ends of the plane”, so we will now formalise this intuition.

Given a position p , we aim to determine whether some given player has won. Moreover, we need to ensure that not both players win in a single play, so to comply with condition 1.2.v. in the definition of game.

We first need some definitions, which can be easily repeated for Blue; let \tilde{H} be the Infinite Hex board and fix some position p , that is an assignment of each cell of \tilde{H} to either or no player, together with a turn indicator.

An infinite Red path, or Red \mathbb{Z} -chain, is a function $r : \mathbb{Z} \rightarrow \tilde{H}$ such that all the cells in the image of r are marked by Red at p and $r(n)$, $r(n+1)$ are adjacent in \tilde{H} for all $n \in \mathbb{Z}$.

We can now approach the condition to determine whether an infinite Red path is *winning*.⁵

Pick some cell $h_0 \in \tilde{H}$ and impose an orthogonal system of reference with origin in h_0 and axes directed from South to North and from West to East, as in figure 2.5a.

Observe that, with some flexibility in the identification of the South-North columns, we can use these axes to obtain unique coordinates for each tile in \tilde{H} ; this is best described by figure 2.5b.

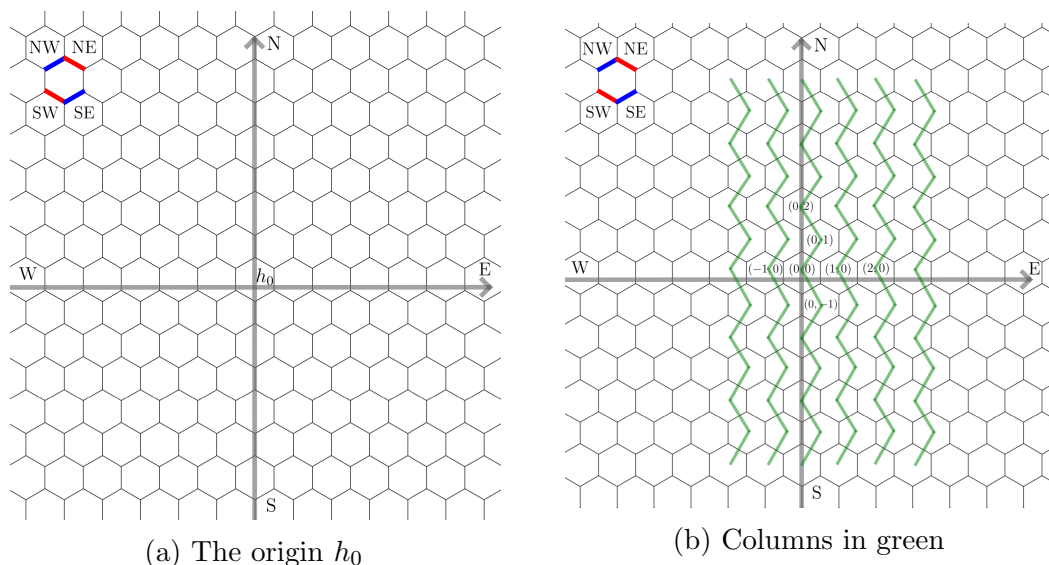


Figure 2.5: The reference system for the Infinite Hex board.

It can be useful to keep in mind the sketch of a general Red winning path in figure 2.6, while reading the following definitions.

Let r be an infinite Red path.

Definition 2.10. Given a system of reference centered at h_0 as above, we say that r is *winning with respect to h_0* if $r[\mathbb{Z}^+]$ is eventually contained in the NE-quadrant and $r[\mathbb{Z}^-]$ is eventually contained in the SW-quadrant of such reference system.

More precisely, we say that r is winning with respect to h_0 if there is some $M \in \mathbb{N}$ such that if $m \geq M$, then the coordinates of $r(m)$ are both positive, while the coordinates of $r(-m)$ are both negative.

Definition 2.11. We say that r is a *winning* Red path if it is winning with respect to every choice of centre h .

⁵We could require a winning path r to be *geodesic*, so that the sub-path between any two tiles in r is the shortest Red path between such tiles; however, it will be clear that if Red has a winning path, then Red also has a geodesic winning path.




Figure 2.6: Sketch of a general Red winning path.

The winning condition for a Blue path is analogous, substituting NW for NE and SE for SW in Definition 2.10, as hinted by the key in figure 2.5.

Observe that the arbitrary choice of the origin for the reference system is necessary to avoid considering as wins paths which are bounded in one of the axes' directions; otherwise we would not be able to assign consistently an outcome to the position in figure 2.7, in which the Red and Blue paths are unbounded towards North and South, but bounded in the WE-direction.

In fact, each path is winning for the relevant player with respect to several cells on the paths themselves, while neither path is winning with respect to the starred tile, for instance.




Figure 2.7: Infinite paths bounded in the WE-direction.

Remark 2.12. By insisting in Definition 2.10 that the coordinates of $r(m)$ have to be strictly positive, and of $r(-m)$ strictly negative, then we can define a path to be winning, equivalently to Definition 2.11, if it is winning with respect to each of its tiles; such a path cannot then be bounded in either of the axes' directions.

We now show that the condition above is well-defined in the sense of defining disjoint winning conditions for the two players, as per Definition 1.2.v.

Proposition 2.13. Not both players can win in a position of Infinite Hex.

Proof. Suppose for a contradiction that p is a position of Infinite Hex such that both Red and Blue have winning paths, respectively, r and b .

Pick some origin $h_0 \in \tilde{H}$ and let $x, y : \tilde{H} \rightarrow \mathbb{Z}$ be the projections on the WE- and SN-axes centered at h_0 , so that the coordinates of a tile $h \in \tilde{H}$ are $(x(h), y(h))$.

Note that both paths r and b are winning, for the relevant players, with respect to h_0 ; hence there is some $M \in \mathbb{N}$ such that for $m \geq M$ we have that⁶

$$x(r(m)), y(r(m)) > 0; \quad x(r(-m)), y(r(-m)) < 0;$$

$$x(b(m)) < 0, \quad y(b(m)) > 0; \quad x(b(-m)) > 0, \quad y(b(-m)) < 0.$$

Consider now the finite paths $r([-M, M])$ and $b([-M, M])$; we claim that they cannot be both connected.

Let H be the smallest finite Hex board centered at h_0 and with the usual orientation, such that it contains both $r([-M, M])$ and $b([-M, M])$.

Now, let r', b' be the maximally connected finite sub-paths of r, b that extend, respectively, $r([-M, M]), b([-M, M])$ in H ; we claim that they are both winning in H , according to the usual finite Hex winning conditions, leading to a contradiction.

Observe that, by choice of M , the ends of the path $r([-M, M])$, i.e. $r(-M)$ and $r(M)$, are included in, respectively, the NE- and SW-quadrants of the reference system centered at h_0 .

As r' extends $r([-M, M])$ without getting out of the relevant quadrants, then r' is a Red winning path in the game of finite Hex played on H .

An analogous argument can be made to show that b' is a Blue winning path in the game of finite Hex on H , which leads to the contradiction needed, by Theorem 2.5. \square

We conclude this section by presenting two “extreme” positions that motivated the choice of the definition of winning condition given and allowed to discard several other possible definitions. These positions are draws according to our definition, but can easily be considered wins by either or both players with weaker conditions.

In particular, observe how the Red path in figure 2.8 is unbounded in the NE-SW-direction, as desired, but also goes through the East-West pink band infinitely many times.

⁶ M is the maximum of the two bounds given by Definition 2.11 for r and b .




Figure 2.8: Infinite Red path unbounded towards East, West, NE, and SW.

Moreover, the Red and Blue paths in figure 2.9 are unbounded in all directions; any coherent winning condition that identifies one of them as winning has to do the same with the other.




Figure 2.9: Two spiralling Red and Blue paths.

2.2.2 Normal Infinite Hex is a draw

In section 2.2.1 we provided a definition of winning positions, rather than plays; this is because the theory of transfinite game values developed in section 1.1 only applies to open games, and so we are mostly interested in applying such theory to the sub-class of games of Infinite Hex which are open for some designated player.

Nevertheless, we can consistently extend Definition 2.11 to all plays of Infinite Hex by dropping the finiteness condition; a player wins if he has a strategy with which he can eventually construct⁷ an infinite path which is winning according to Definition 2.11, however his opponent plays.

Equipped with this new definition, we can ask the question: who wins from the initially empty position? We are particularly interested in clarifying whether the advantage of the first player in finite Hex carries on to Normal Infinite Hex, the game of Infinite Hex with the empty board as initial position.

We will show that this is not the case and that both players can force at least a draw from the empty board. However, we later conjecture in section 2.3.2 that the first player may still win if given some advantage, possibly of only one tile.

Let θ be a pairing of the tiles of the Infinite Hex board as depicted in figure 2.10; cells in the same pair have the same label. Note that θ pairs each cell with exactly one other cell.

Let μ be the second-player strategy, called *mirroring strategy*, induced by θ as follows; whenever Red, the first player, places a stone on a tile h , then μ prescribes that Blue places a stone on the tile paired with h , according to θ .




Figure 2.10: The pairing θ for μ , with the mirroring axis in green.

Observe that μ is a well-defined strategy for Normal Infinite Hex. If Blue follows it since her first move, then Red will not be able to mark over the course of the play both tiles of any pair given by θ ; hence, at each turn, the tile indicated by μ to Blue will always be empty. Moreover, note that, following μ , Blue always places a stone in the opposite half-plane to where Red played immediately before, with respect to the East-West mirroring axis.

⁷In at most countably many moves, as plays still have order type at most ω .

Remark 2.14. Similarly to the Strategy-stealing argument of Proposition 2.8, we can adapt μ to be a first-player strategy.

Let θ' be the pairing given by reflecting the pairing θ across some North-South axis.⁸ Let μ' be the strategy that prescribes Red to make an arbitrary initial move and then follow the (second-player) mirroring strategy induced by θ' , possibly making arbitrary moves when the mirroring strategy prescribed to place a stone on a tile already marked by Red; this is because the mirroring strategy ignores the initial move and the subsequent arbitrary moves by Red.

Since the reflection of all possible winning paths for Red across a North-South axis gives exactly all possible winning paths for Blue, it is then clear that showing that μ is a drawing strategy for Blue implies that μ' is a drawing strategy for Red.

Remark 2.15. The pairing in figure 2.10 is called mirroring because it is not a symmetry and does not represent a rigid motion of the board. In particular, we obtain the Northern half-plane as the reflection across the mirroring axis of the Southern half-plane, followed by a translation by half a cell towards East; this translation is in effect “tearing” the plane.

Remark 2.16. As already observed, when Blue plays according to μ , Red can mark at most one tile in each pair given by θ .

It follows that if Red constructs some winning path r , then r contains at most one tile in each pair given by θ , so that Blue can mark all the cells of $\theta(r)$, the *mirroring* of r with respect to θ ; we call this the *mirroring condition* on r .

Before proving the main result of this section, we make two observations.

Firstly, any Red path that crosses the mirroring axis of μ can only do so with two adjacent cells in the NW-SE-direction, that is the “wrong” direction for Red. This will be key for the argument, and is clear by inspection of figure 2.11.




Figure 2.11: Single Red intersection of the mirroring axis.

⁸Observe that θ' is essentially unique, given θ .

Secondly, μ allows Blue to win, as long as Red plays appropriately bad. Observe that the position in figure 2.12 is winning for Blue and satisfies the mirroring condition of Remark 2.16 for the Red path.




Figure 2.12: Blue wins following the mirroring strategy μ .

Proposition 2.17. In Normal Infinite Hex, μ is a drawing strategy for the second player.

Proof. Suppose for a contradiction that Red has a winning strategy σ .

Letting Red play according to σ and Blue according to μ , Red will be able to construct a winning path⁹ r which satisfies the mirroring condition of Remark 2.16, so that r does not contain both tiles of any of the pairs of θ .

Now, we claim that r cannot be a winning path. Observe that r has to cross the East-West mirroring axis at least once. So, suppose that r crosses the mirroring axis exactly once. We then expect a single Red intersection as in figure 2.11.

Since r is winning for Red, then the semi-infinite sub-path of r in the Northern half-plane will eventually go towards North-East; by definition of μ , Blue will correspondingly construct the mirrored sub-path in the Southern half-plane that eventually goes towards South-East, as in figure 2.13a.

However, we also have that the sub-path of r in the Southern half-plane must eventually go towards South-West, but if that is the case, then we can draw a finite Red path¹⁰ as in figure 2.13b so that we can identify a simple closed curve determined by adjacent Red stones, such curve is a Jordan curve which is intersected by a *connected* Blue path; this is a contradiction by the Jordan Curve Theorem.

⁹Such path is only constructed after infinitely many turns, as described at the beginning of this section.

¹⁰This extra Red finite path is drawn directly on the board for the sake of the geometric argument; it is not part of the play.

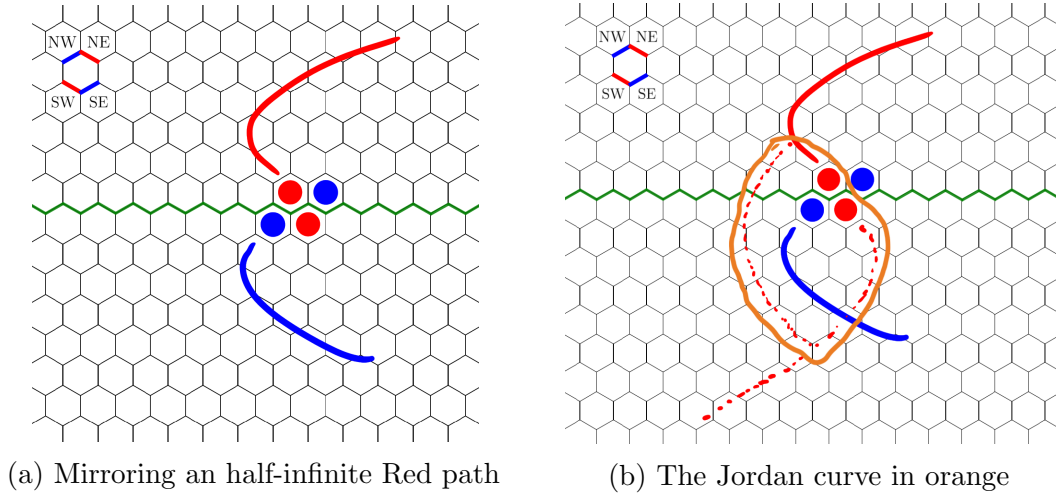


Figure 2.13: Application of the Jordan Curve Theorem.

Suppose now that r crosses the mirroring axis more than once; r cannot cross it exactly twice because it is a winning path, and in particular a connected path, thus assume that r crosses the mirroring axis exactly thrice.

We expect to see three Red intersections of the mirroring axis as in figure 2.14a; these were discussed before in figure 2.11.

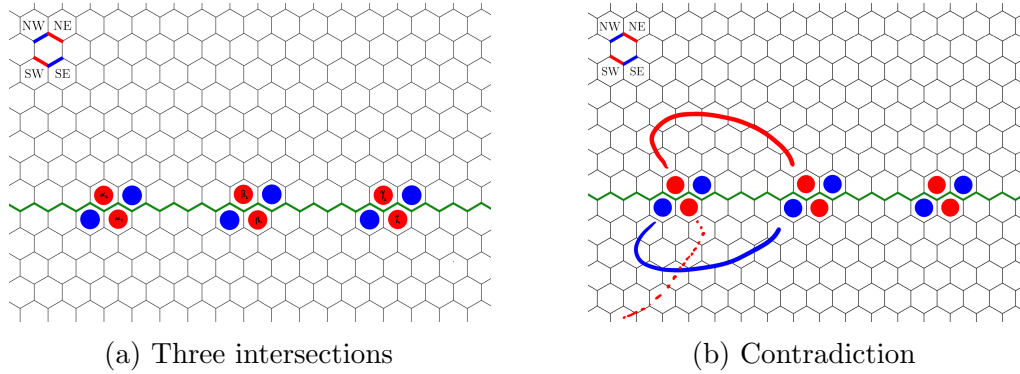


Figure 2.14: A Red path cannot cross the mirroring axis three times.

Observe that the only way to construct such a path r is to join with Red finite paths the ends labelled α_2 with β_2 , and β_1 with γ_1 , and joining appropriate semi-infinite Red paths to the ends labelled α_1 and γ_2 .

Such a construction is not possible due to the mirroring condition of Remark 2.16; observe in figure 2.14b how joining two ends in the same half-plane prevents one end in the opposite half-plane to be extended indefinitely.

These observations clearly hold by applying the Jordan Curve Theorem similarly to the above.

Observe that the argument for r intersecting the mirroring axis three times can be easily adapted to any position in which r intersects the mirroring axis a larger (odd) number of times.

Since r can only intersect the mirroring axis finitely many times, as it is a winning path, we then conclude that r cannot intersect the mirroring axis more than thrice, either. The proof is complete. \square

Remark 2.18. It is possible to prove Proposition 2.17 via a more involved argument by contradiction which reduces a Red winning path to a winning path for the disadvantaged player of a game of finite Hex on an asymmetric board.

In this sense, we can think of Infinite Hex as generalising the asymmetric variant of finite Hex, rather than the standard finite version.

2.3 Positional and stone-placing games

2.3.1 Positional games

Many popular board games belong to the class of positional games, which share the property of extra moves being always advantageous as in Hex,¹¹ and whose simplest representative is the game of Tic-tac-toe.

Even though universally popular, we mention the rules of this game; in Tic-tac-toe two players, Naughts and Crosses, alternatively place their marks on a 3×3 square board. The first player to mark three squares on a straight line wins, i.e. the first player to mark all 3 elements of any of 8 distinguished subsets of the board wins.

Note that the sets of board cells that allow the players to win are the same for both players.

We define such “Tic-tac-toe-like games” as in [Hefetz et al., 2014]; in order to additionally allow for infinite games, we define the following structure.

Definition 2.19. A *generalised hypergraph* is a pair of sets (B, \mathcal{F}) , such that B is the collection of *vertices*, and $\mathcal{F} \subset \mathcal{P}(B) \setminus \emptyset$ is the collection of *hyperedges*.

Note that from now on we will not distinguish between generalised hypergraphs and *hypergraphs*, which are usually defined stressing the finiteness of both the vertex and hyperedge sets.

¹¹See Remark 2.9.i.

Definition 2.20. A *positional game* played on a hypergraph (B, \mathcal{F}) is a game such that the players alternatively mark previously unchosen elements of the vertex set B , which represents the set of cells of the game board that are unmarked at the initial position.

The first player to mark all the vertices of some hyperedge in \mathcal{F} wins; here $\mathcal{F} \subset \mathcal{P}(B) \setminus \emptyset$ is called the collection of *winning sets*.

Recall from Remark 1.3 that a *position* in a positional game is an assignment of some of the board cells to the players, together with the history of moves and a turn indicator.

Remark 2.21. Observe that, in Definition 2.20, specifying that only the first player to satisfy the winning condition is the winner ensures openness of the game for both players and is necessary to avoid ambiguity.

See for instance the Tic-tac-toe board position in figure 2.15; both players seem to be winning, but assuming that Crosses is the first player of the game, then we can say that, for any history of legal moves leading to that board position, the resulting position is a win for Crosses.




Figure 2.15: An ambiguous Tic-tac-toe board position.

Remark 2.22 (Strategy-stealing). Observe that Nash's Strategy-stealing argument¹² can be applied to show that any positional game played on (B, \mathcal{F}) is either a draw or a first-player win.

In fact, if the second player has a winning strategy, then the first player can use it after having made an arbitrary initial move, leading to a contradiction; this holds because both players have exactly the same winning conditions, which are represented by \mathcal{F} , and because in all positional games it is always advantageous to make an extra move, i.e. to choose one extra vertex.

Thus, the best outcome for the second player in a positional game is achieving a draw. For this reason, we do not lose generality when considering the class of games defined below.

¹²See proof of Proposition 2.8.

Definition 2.23. A *Maker-Breaker game* played on a hypergraph (B, \mathcal{F}) is a game in which Maker wins by choosing all the vertices of some winning set in \mathcal{F} , while Breaker wins by choosing at least one vertex of each winning set in \mathcal{F} .¹³

Remark 2.24. Clearly, Breaker's goal is precisely to prevent Maker from winning; it follows that a play of any Maker-Breaker game cannot allow both players to win.

Moreover, we can see that there is a duality between Maker and Breaker that preserves strategies, but not game openness.

Namely, take a Maker-Breaker game \mathcal{G} open for Maker and played on some hypergraph (B, \mathcal{F}) ; observe that such a game is equivalent¹⁴ to the Maker-Breaker game \mathcal{G}' not necessarily open for Breaker and played on the hypergraph (B, \mathcal{F}') , but with the players' roles reversed, where

$$\mathcal{F}' = \{F' \in \mathcal{P}(B) : |F' \cap F| \geq 1 \text{ for all } F \in \mathcal{F}\}.$$

Observe that Breaker aims to mark at least one vertex in each of the winning sets in \mathcal{F} , so that \mathcal{F}' does represent the correct winning condition in the game with the roles reversed.

Remark 2.25. Recalling the no-tie property of finite Hex,¹⁵ we observe that a finite Hex player who prevents the opponent from winning is equivalent to the player himself winning. In other words, the best offensive moves in finite Hex are exactly the best defensive moves, and vice versa.

Hence, we see that finite Hex is a Maker-Breaker game played on the hypergraph (H, \mathcal{C}) such that the vertex set H is in bijection with the empty finite Hex board and \mathcal{C} is the collection of all the winning paths for Red, who corresponds to Maker.

What we mentioned in section 2.1.1 for asymmetric Hex boards is that no winning component for the player who is joining the short edges can avoid containing two cells paired with respect to the opponent's strategy.

A weaker graph-theoretic property that follows from this pairing is that of *proper 2-colourability* of $(H_{asym}, \mathcal{C}_{disadv})$, where H_{asym} represents the relevant empty asymmetric board, and \mathcal{C}_{disadv} is the collection of winning paths for the disadvantaged player, who is losing; that is, H_{asym} can be coloured with two colours, say

¹³A different definition could let Breaker win whenever Maker does not win; such a no-tie variant is not equivalent to our definition in the setting of infinite games. Not all vertices in B may be marked over some plays and neither player may satisfy their original winning condition over such plays, so that we would consider such plays ties with our definition and wins for Breaker with the no-tie variant.

¹⁴Equivalent from the point of view of the closed player.

¹⁵See Corollary 2.7.

black and white, such that no collection of tiles that represents a winning path in \mathcal{C}_{disadv} is monochromatic.

Indeed, the cell pairing given by the mirroring strategy described in figure 2.4 is a particular case of such a *proper 2-colouring*, which is obtained by colouring with distinct colours the cells in each pair.

We prove the following result in appendix A.1.

Proposition 2.26. Take some Maker-Breaker game \mathcal{G} played on a hypergraph (B, \mathcal{F}) such that \mathcal{G} is open for Breaker, and Maker is the first player to move. If Breaker has a winning strategy, so that he can play to take at least one vertex in each of Maker’s winning sets in \mathcal{F} , then (B, \mathcal{F}) is properly 2-colourable.

2.3.2 Biased Infinite Hex may not be a draw

Informally, our intuition of Infinite Hex is that, since the infinite board is “too large”, the advantage of placing the first stone is “too diluted”, and so an initial position with finitely many cells already marked by Red¹⁶ should not give a consistent advantage to Red, compared to the initially empty board, so that the game should result in a draw, as we saw in section 2.2.2.

However, adaptations of the mirroring strategy μ do not seem to be sufficient even in the case of just one extra stone for Red.¹⁷

It seems reasonable to claim that no pairing strategy can be formulated. A possible proof could model Biased Infinite Hex as a Maker-Breaker game played on $(\tilde{H}_{adv}, \mathcal{C}_1)$, where \tilde{H}_{adv} is the infinite board without the cells given to Red as advantage, and \mathcal{C}_1 represents the winning condition for Red, the first player, and then assume for a contradiction that Blue has a drawing strategy, which would imply the proper 2-colourability of $(\tilde{H}, \mathcal{C}_1)$ by Proposition 2.26; it is not clear how to get a contradiction from here.

The non-openness of the games of Biased Infinite Hex makes such problem less tractable.

We then formulate the following.

Conjecture 2.27. A game of Infinite Hex with exactly one tile of the board assigned to Red at the initial position, and none to Blue, is a win for Red.

As an apparently easier problem, we also propose the following.

Conjecture 2.28. A game of Infinite Hex with initial position as in figure 2.16 is a win for Red.




Figure 2.16: Infinite East-West Red path.

It is not difficult to see that from such a position Red can construct arbitrarily long, but finite, “prongs” that shoot out of the East-West infinite path; however, this is short of a definitive answer.

2.3.3 Stone-placing games

We will now prove that no infinite game values can be achieved in the following wider class of vertex-colouring games, that retain the property of extra moves being always advantageous for the player making them, which we first observed for finite Hex in Remark 2.9.i.

Definition 2.29. A *stone-placing game* played on $(B, \mathcal{F}, \mathcal{S})$, where both (B, \mathcal{F}) and (B, \mathcal{S}) are hypergraphs, is a game in which the two players alternatively choose previously unchosen elements of the common vertex set B ; at the initial position, all the vertices of the board B are unmarked.

The hyperedge sets $\mathcal{F}, \mathcal{S} \subset \mathcal{P}(B) \setminus \emptyset$ represent the winning conditions of, respectively, the first and the second player; that is, the winner is the player who chooses all the vertices of some winning set specified by his respective winning condition first.

Thus, we impose that both \mathcal{F} and \mathcal{S} are closed under taking supersets, without loss of generality.

Remark 2.30. Note that Definition 2.29 does indeed generalise the other games defined previously in the current section 2.3.

Say that \mathcal{G} is a stone-placing game played on $(B, \mathcal{F}, \mathcal{S})$; if $\mathcal{F} = \mathcal{S}$, then \mathcal{G} is a positional game.¹⁸

We can see that if $\mathcal{S} = \mathcal{F}'$, where

$$\mathcal{F}' = \{F' \in \mathcal{P}(B) : |F' \cap F| \geq 1 \text{ for all } F \in \mathcal{F}\},$$

¹⁶Or even infinitely many, if placed “badly” enough.

¹⁷Of course, games of Infinite Hex biased in favour of the second player present essentially the same problem.

¹⁸This clarifies why positional games are also known as *Maker-Maker games*.

then \mathcal{G} is a Maker-Breaker game in which Maker plays first; an analogous condition allows \mathcal{G} to be a Maker-Breaker game with Breaker playing first.

Remark 2.31. If \mathcal{G} is a stone-placing game played on $(B, \mathcal{F}, \mathcal{S})$ open for the first player, then every win by the first player is essentially finite, and so for any $f \in \mathcal{F}$ there is some finite $f' \in \mathcal{F}$ such that $f' \subset f$. Thus, we can say that \mathcal{F} has a basis of finite sets; that is, there is $\{f_i\}_{i \in \mathcal{I}} \subset \mathcal{F}$ such that f_i is finite for each $i \in \mathcal{I}$, which is not necessarily finite, and, for any $f \in \mathcal{F}$, there is $f_i \subset f$ for some $i \in \mathcal{I}$.

To express this more explicitly, say that a play s is a win for the first player because he manages to mark all the vertices in some infinite winning set $f \in \mathcal{F}$, then there is a finite $f_i \subset f$, for some $i \in \mathcal{I}$, which is completely marked by the first player over s ; since each vertex is marked at some finite stage of the play,¹⁹ then the whole of f_i is marked by the first player after finitely many turns and s is an essentially finite play, as required by the openness of \mathcal{G} .

Hence, if \mathcal{G} is open for the first player, then we can regard the elements of the basis for \mathcal{F} as the “meaningful” winning sets, and we can thus assume that there are no “redundant” infinite elements in \mathcal{F} , without essentially changing \mathcal{G} . In particular, if all the elements of \mathcal{F} are finite, then \mathcal{G} is open for the first player.

If \mathcal{G} is open for the second player, then \mathcal{S} satisfies analogous conditions.

Analogously, we see that if a Maker-Breaker game \mathcal{G} played on (B, \mathcal{F}) is open for Maker, then \mathcal{F} has a basis of finite sets.

Dually, if \mathcal{G} is open for Breaker, then \mathcal{F}' has a basis of finite sets, where \mathcal{F}' is defined as above.²⁰

In particular, if both \mathcal{F} and its elements are finite, then \mathcal{G} is open for both players.

Remark 2.32. We can generalise the mirroring strategy described in section 2.2.2 for Normal Infinite Hex, and at page 20 for finite asymmetric Hex boards, to give a condition that enables the second player²¹ to force a draw.

We propose another way to see that argument; the pairing θ for the the mirroring strategy μ determines an involution²² of the infinite board $\theta : \tilde{H} \rightarrow \tilde{H}$; note that θ has no fixed point, as each cell is paired with a different cell.

Take a winning path for Red and take its image under θ ; this will not necessarily be a winning path for Blue, exactly because the original tile pairing θ does not

¹⁹This is a consequence of plays having order type at most ω .

²⁰Recall the duality between Maker and Breaker discussed at page 32.

²¹Note that an analogous result holds for the first player, who additionally makes an arbitrary initial move, as usual. We state the result for the second player because she is a priori at a disadvantage.

²²Given a pairing of the tiles in the Infinite Hex board, we need the Axiom of Choice to define such an involution.

determine a rigid movement of the board; however, we have shown in section 2.2.2 that it will intersect Red's path.

Observe that the intersection consists only of paired cells; hence, if the intersection is finite, then it has even size.

Proposition 2.33 (Mirroring). Let \mathcal{G} be a stone-placing game played on $(B, \mathcal{F}, \mathcal{S})$. Suppose that there is a fixed-point free involution θ of the board, that is a bijective function $\theta : B \rightarrow B$ such that $\theta^2(b) = b$ and $\theta(b) \neq b$ for all $b \in B$, and say that, for each $f \in \mathcal{F}$, we have $f \cap \theta[f] \neq \emptyset$.

Then, the second player in \mathcal{G} can force a draw.

Proof. Let the second player play as follows; at each turn, when the first player marks some vertex $v \in B$, then the second player chooses $\theta(v) \neq v$ immediately after. Note that if $\theta(v)$ was already chosen by the first player, then $\theta(\theta(v)) = v$ would have been previously taken by the second player, leading to a contradiction; thus, this strategy is well-defined.

Say that $v \in f \cap \theta[f]$. Since $v \in f$, then $\theta(v) \in \theta[f]$; moreover, since $v \in \theta[f]$, then $\theta(v) \in \theta^2[f] = f$. Thus, v and $\theta(v)$ are distinct elements of $f \cap \theta[f]$; we deduce that $f \cap \theta[f]$, if finite, has even size.

Following this strategy, the second player prevents the first player from choosing all the vertices in some hyperedge $f \in \mathcal{F}$; since $|f \cap \theta[f]| \geq 2$, then the second player will always be able to counteract the moves of the first player, so that the first player will not be able to mark all the vertices in $f \cap \theta[f]$.

Therefore, the second player can force a draw. \square

Remark 2.34. Observe that a board involution θ as in Proposition 2.33 specifies a vertex pairing that can be used to describe a proper 2-colouring of (B, \mathcal{F}) , as we had seen in Proposition 2.26.

Remark 2.35. Note that the Strategy-Stealing argument²³ does not generally apply to this class of games, in fact, there is not always an automorphism of the vertices in B such that each winning set in \mathcal{S} is sent to some winning set in \mathcal{F} .

However, we can generalise the symmetry condition of the finite Hex board highlighted in Remark 2.9.ii to construct the Strategy-stealing argument for stone-placing games which are not open; we provide an argument in appendix A.2.

²³See Proposition 2.8.

2.3.4 Impossibility of infinite game values

Finite game values in Infinite Hex

Firstly, we show that all finite ordinals can be realised as the game values of some Infinite Hex positions through the use of *virtual links*.²⁴

The position in figure 2.17a is a *bridge*, as defined by [Browne, 2000, §12.3], which ensures that Red can join the two marked cells by choosing either of the two cells separating them. Blue can threaten this construction by occupying one of these two cells, and Red can simply respond by picking the remaining one, thus “realising” the virtual connection.




Figure 2.17: The realisation of a virtual link by Red.

Joining two half-infinite lines by finitely many interlinked bridges allows to construct positions of any finite game value.

We present in figure 2.18 a position with value 5 for Red, two half-infinite linear components joined by 5 bridges, in which the virtual connections are highlighted.

Note that this position has infinitely many pieces already placed. More in general, since the winning player places only finitely many stones when starting from a position with defined game value, then it follows from the definition of winning condition that no position with defined game value can have only finitely many stones already placed.

Infinite game values

It can be reasonable to expect that some well-constructed open games of Infinite Hex can achieve an infinite game value; however, we will show that this is not possible by proving that Hex, as all stone-placing games, is an *essentially local* game.

²⁴These were first called *angle-positions* by [Hein, 1942], and then received this name by Claude Berge in 1977, according to [Björnsson et al., 2006, §7]; they are now part of the standard terminology used by Hex players [HexWiki, 2021].




Figure 2.18: Infinite Hex position with value 5 for Red.

Definition 2.36. A *dead region* for some player in a stone-placing game, which is played on $(B, \mathcal{F}, \mathcal{S})$ and has defined game value α for the player's opponent,²⁵ is some subset of the board $R \subset B$ such that, if all the vertices in R are assigned to the (losing) player, then the game with the resulting initial position, which has the same turn indicator, still has value α for the (winning) player's opponent.²⁶

In other words, modifying the initial position of the game by assigning the cells of a dead region as advantage to the losing player does not substantially alter the game; even though the strategy of the winning player may change, the game value for such player does not vary.

Definition 2.37. A stone-placing game, which has defined game value for some player, is *essentially local* for that player if there is a cofinite dead region of the board for that player's opponent.

Remark 2.38. An analogous definition can be made for essentially local positional games. Also, it is obvious that all finite games are essentially local.

²⁵So that the game is clearly open for the player's opponent.

²⁶This definition is inspired by the concept of dead cells, which are usually defined in a more restrictive sense as cells that cannot influence the outcome of the game from a given position; see for instance [Björnsson et al., 2006, 3.1].

For instance, see that all the empty tiles in figure 2.18, except for the 10 empty tiles which allow Red to realise the virtual links, belong to a dead region for Blue, hence that is the initial position of an essentially local game.

We aim to prove the following.

Theorem 2.39. Any game of Infinite Hex with defined value for some player is essentially local for that player; thus, it has finite game value.

Observe that a game \mathcal{H} of Infinite Hex can be represented as a stone-placing game.

Suppose that Red is the first player to move in \mathcal{H} and say that $(\tilde{H}, \tilde{\mathcal{R}})$ and $(\tilde{H}, \tilde{\mathcal{B}})$ are hypergraphs in which the vertex set \tilde{H} is in bijection with the board cells unchosen at the initial position of \mathcal{H} , and the hyperedges in $\tilde{\mathcal{R}}, \tilde{\mathcal{B}} \subset \mathcal{P}(\tilde{H})$ represent all the possible winning paths for, respectively, Red and Blue. Note that the vertices of all hyperedges in $\tilde{\mathcal{R}}$ and $\tilde{\mathcal{B}}$ are unmarked at the initial position of \mathcal{H} .

We can see that \mathcal{H} can be interpreted as a stone-placing game played on $(\tilde{H}, \tilde{\mathcal{R}}, \tilde{\mathcal{B}})$, where the first player is Red and the second player is Blue.

It is now clear that Theorem 2.39 is a consequence of the following more general result.

Theorem 2.40. Any stone-placing game with defined value for some player is essentially local for that player; thus, it has finite game value.

In order to highlight the symmetry between the first and second player, in the rest of this section stone-placing games will be played on $(B, \mathcal{O}, \mathcal{C})$ by the players Open and Closed, who have winning conditions \mathcal{O} and \mathcal{C} , respectively.

Lemma 2.41. Given any stone-placing game \mathcal{G} played on $(B, \mathcal{O}, \mathcal{C})$ which has finite game value n , we have that \mathcal{G} is essentially local for Open.

That is, there is a finite subset of the board $D \subset B$ such that Open has a strategy that allows him to win in at most n turns by choosing vertices only in D , and $B \setminus D$ is a dead region for Closed.

Proof. We prove this by induction on n .

If $n = 0$, then there is nothing to prove, as Open has already won at the initial position and does not need to make any move. Trivially, $D = \emptyset$.

Suppose that $n = 1$; we consider this case for clarity of argument. Let σ be a winning strategy for Open realising this game value.

If Open is the next to move, then he can simply win by choosing a single vertex $v \in B$, as prescribed by σ ; in this case, $D = \{v\}$ is sufficient.

If instead Closed is the next to move, then suppose that he chooses some vertex w , without winning. Then Open can again make a winning move by choosing some $v' \neq w$. If instead Closed decides to initially choose v' , then Open can still win by choosing the vertex $v'' \neq v'$ prescribed by σ in that position. In either case, $D = \{v', v'', w\}$ is as needed.

Suppose now that the Lemma is true for $n = N$. Let \mathcal{G} have game value $N + 1$ for Open, and let τ be a value-reducing strategy that realises this game value.

If Open is the next to move from the initial position, then, by definition of game value,²⁷ he can choose a vertex $v \in B$ so to reach a position with game value N . Then, by inductive hypothesis, there is a finite subset $D' \subset B$ such that Open can win in at most N moves by choosing vertices only from D' . In this case, $D = D' \cup \{v\}$ is as needed.

If instead Closed is the next to move from the initial position, then suppose that he chooses some $w_0 \in B$. Following τ , Open can mark some other $v_0 \in B$ to reach a position with game value at most N ; by inductive hypothesis, there is a finite sub-board $D'_0 \subset B$ such that Open wins by choosing vertices only from it. Let $D_0 = D'_0 \cup \{v_0, w_0\}$.

Consider again the initial position of \mathcal{G} ; suppose that there are k unchosen vertices²⁸ in $D'_0 \cup \{v_0\}$ at the start, say $\{w_i\}_{i=1}^k$. Now, take some $1 \leq i \leq k$ and suppose that Closed decides to initially choose w_i , to which Open replies by choosing some v_i determined by τ , reaching a position of game value at most N ; by inductive hypothesis, there is a finite sub-board $D'_i \subset B$ such that Open wins by playing only on its vertices. Note that, in this case, the line of play of Open is only influenced by the moves made on the sub-board $D_i = D'_i \cup \{v_i, w_i\}$.

We claim that $D = \bigcup_{i=0}^k D_i$ is a finite sub-board that satisfies the condition of the Lemma at the initial position of \mathcal{G} .

Suppose that Closed initially chooses an available vertex in D_0 , that is, he chooses w_i for some $0 \leq i \leq k$. Open can reply by choosing the corresponding vertex v_i , and subsequently win in at most N more moves by playing on D_i ; Open thus ends the game having chosen vertices only in $D_i \subset D$, by construction.

Suppose now that Closed makes an initial move outside of D_0 , choosing some vertex $w' \notin D_0$ (but possibly in D). Open can pretend to see w' as still unchosen and he can play as if Closed had made the *phantom move* of choosing w_0 . Hence Open can win by playing only on $D_0 \subset D$, as before. Observe that in this case the initial move by Closed does not obstruct Open's play since Open wins by marking all the vertices of some $o \in \mathcal{O}$ such that $o \subset D_0$, and that certainly does not depend on $w' \notin D_0$.

²⁷Definition 1.14.

²⁸Note that both v_0 and w_0 are still unchosen at the initial position.

This concludes the proof of the Lemma. \square

Lemma 2.42. No stone-placing game open for some player has infinite game value for that player.

Proof. Let \mathcal{G} be a stone-placing game played on $(B, \mathcal{O}, \mathcal{C})$, which is open for Open. Suppose for a contradiction that \mathcal{G} has infinite game value for Open; by Proposition 1.17, it is enough to assume that \mathcal{G} has game value ω for Open.

By definition of game value, Closed is the first to move in \mathcal{G} and can only reach positions with finite game value from the initial position.

So, suppose that Closed initially chooses some $w \in B$, reaching a position with finite value N for Open. Now, by Lemma 2.41, there is a finite sub-board $D' \subset B$ such that Open can win in at most N turns by marking vertices only in D' . Let σ be a value-reducing strategy for Open that realises this in \mathcal{G} ; note that the moves prescribed by σ for Open only depend on the moves that take place on $D = D' \cup \{w\}$.

Suppose now that, from the initial position, Closed makes a different move $w' \notin D$.

Then, Open can ignore such move by Closed, and imagine instead that Closed had made the phantom move of choosing w , and then play according to σ from such position imagined by Open.

Similarly to the situation in the proof of Lemma 2.41, Open can safely ignore Closed's move in w' because σ allows him to win by playing only on D , and so by marking all the vertices of some $o \in \mathcal{O}$ such that $o \subset D$; it is clear that Closed's moves outside of D cannot obstruct this goal.

It is important to note that, by playing outside of D , Closed cannot aim to win, as he can only reach positions with defined game value for Open.

We have shown that, for an arbitrary initial move by Closed, Open can force a win in at most N turns; thus, the position achieved by the initial move of Closed has game value bounded by N . But the initial position had game value ω ; this is a contradiction. \square

Proof of Theorem 2.40. A stone-placing game with value defined for some player can only have finite game value, by Lemma 2.42; such game is thus essentially local for that player by Lemma 2.41. \square

Remark 2.43. Observe that the arguments above can be adapted to generalise the results to n -player stone-placing games, once the definition of game value is extended to n -player games.

By definition of omega one, we have the following immediate consequence of Theorem 2.40.

Theorem 2.44. The omega one of any class of stone-placing games open for some player is at most ω .

We see that also Theorem 2.39 is now proved; following our previous discussion of Infinite Hex positions with finite game value, we can conclude by stating the result which answers the main question that motivated the present research.

Theorem 2.45. The omega one of the class of open games of Infinite Hex is the smallest infinite ordinal, that is,

$$\omega_1^{\text{Hex}} = \omega.$$

2.4 Other generalisations of Hex

It is not easy to intuitively capture the complexity of finite board games, so we now present explicitly some measurement of complexity for the finite games we mention in this work.

We mention the following estimates²⁹ of the state-space and game-tree complexities of the different board games we are considering.

Game	State-space complexity	Game-tree complexity
Checkers (8×8)	10^{21}	10^{31}
Draughts (10×10)	10^{30}	10^{54}
Chess	10^{46}	10^{123}
Hex (11×11)	10^{57}	10^{98}
TwixT (24×24)	$>> 10^{140}$	10^{159}

The Shannon game

The Shannon game is played on a graph Γ with two distinguished vertices.

Rephrasing the definition of [Björnsson et al., 2006, §2], we can define this game as a Maker-Breaker game played on (B, \mathcal{F}) , such that B is the set of non-distinguished vertices of Γ and \mathcal{F} is the collection of paths that join the distinguished vertices of Γ .

Usually, Maker is called Short, and Breaker is called Cut.

It should come as no surprise that Hex can be represented as a Shannon game.

²⁹From [van den Herik et al., 2002], and [Uiterwijk and Moesker, 2009] for TwixT.

TwixT

TwixT is a connection game played on a 24×24 square lattice. The two players place stones on the vertices of such lattice and join two of their stones with an edge when they are opposite vertices of a 2×1 rectangle.

The peculiarity of this game is that edges cannot intersect, so that the order in which moves are made have a great importance in determining the outcome of a game.

Moreover, even though TwixT looks rather different from Hex, it is actually a generalization of it. Thanks to the constructions provided by [Bonnet et al., 2016], we can simulate any Hex game on an appropriately filled TwixT board.

Observe that TwixT can be described as a particular edge-colouring game on the Knight's graph, i.e. the graph obtained by joining the squares of a chessboard separated by a Chess knight's move.

However, TwixT cannot be directly studied as a stone-placing game, and so we cannot apply our results to an infinite version of TwixT.

Chapter 3

Draughts

Also known as Checkers, Draughts is a board game that has a long tradition and many variants, all with the same core rules.

Some pieces, called *pawns*,¹ are initially placed on the white squares² of a chessboard, or *draughtboard*, usually of size 8×8 , which we orient as usual. At the initial position, the pieces of Black are placed on the Southern side of the board and the pieces of White on the Northern side, as in figure 3.1.




Figure 3.1: Initial position of finite Draughts.

At each turn, a player can move one of his pieces to one of the empty squares diagonally adjacent to it.

If a player's piece is adjacent to an opponent's piece, which in turn is adjacent to an empty square aligned with such two pieces, then the player can make a *simple jump* move letting his piece “jump” to the empty square, removing the “jumped” opponent's piece from the board; see figure 3.2.

¹Pawns are usually called *pieces*, which is a word that we will use to refer to both pawns and kings.

²We follow the common custom of displaying Draughts positions as played on the white squares of the board for clarity of presentation; we state the rules accordingly.




Figure 3.2: A simple jump by Black.

Moreover, if the player's piece is, immediately after this jump, in a similar position, then the player can extend his move by making another jump; the resulting move is called an *iterated jump*. It is important to note that an iteration of jumps constitutes a single move, rather than multiple moves.

We refer to these two last types of move as *jump* moves.

In particular, Black pawns are only able to make such moves going strictly towards North; when a Black pawn reaches the Northernmost row of the chessboard, it is promoted to *king*, so that it can then make the moves described above in all directions. An analogous rule applies to White pieces.

The first player that is not able to make a move, either because he has no pieces left on the board or because he has no available legal moves, loses.

It is required in Draughts competitions that a player who can make a jump move at some turn must do so at that turn. Moreover, after a jump move, the piece just moved cannot be in a position from which another jump move could be possible.

In other words, if a player makes a piece jump, then he has to make it keep jumping, choosing which way to jump, until it reaches a position from which no more jumps are possible.

We can formalise this as follows.

Rule 3.1 (Forced Jump). Each player must make jump moves, when they are legal.

Rule 3.2 (Forced Iteration). No player can make a jump move which can be extended as an iterated jump move.

Observe that rule 3.2 allows players to make only maximal jump moves; in particular, it allows players to choose among distinct maximal jump moves. See for instance the position in figure 3.3, where Black can make a simple jump to the square labelled with a dagger, or make either of the iterated jumps to the starred squares; however, Black cannot make a simple jump to the square labelled with “ \otimes ” by rule 3.2.




Figure 3.3: Distinct maximal jumps for Black.

Remark 3.3. Rule 3.2 is justified by the fact that if a player makes one of his pieces jump without completing an available jump iteration, then that piece will be threatened, and likely captured at the next turn,³ by an opponent's piece, as in the case of Black stopping on the square labelled “*” in figure 3.3.⁴

3.1 Infinite Draughts

Analogously to Infinite Chess,⁵ Infinite Draughts is a game played on the infinite chessboard; in particular on its white squares, which are also in bijection with $\mathbb{Z} \times \mathbb{Z}$.

We can give a North-South orientation to the board so to maintain the difference between pawns and kings; however, all the pieces in the positions that follow are kings, unless otherwise specified.

Even though the rules of the game are simple and immediately generalise to the infinite case, we encounter a difficulty with the forced-iteration rule 3.2; observe the position in figure 3.4.

In fact, assuming only rule 3.1, if Black makes a jump move up the White infinite *ladder* in figure 3.4 and captures only finitely many White pieces, then White wins on the next move, as Black would stop his only piece in a square already threatened by White. Thus, if we applied strictly the forced-iteration rule 3.2, we would say that Black must make an infinitely iterated jump, capturing infinitely many White pieces in a single move.

Note that, on the next turn, White can move his only surviving piece at the bottom right, but, following that, Black has no pieces left on the board and thus loses; in fact, having made an infinitely iterated jump, the Black king is not on any of the squares of the infinite chessboard, thus it is not on the board anymore.⁶

³Or certainly captured, in case that is the only jump move available to the opponent.

⁴Moreover, it is usually felt as an unfair advantage the possibility of a king to end a jump at the “back” of an opponent's pawn, by which the king cannot be threatened.

⁵See Example 1.12.

⁶In other words, the Black king effectively went to infinity.




Figure 3.4: Black piece at the bottom of an infinite White ladder, Black to move.

It is clear that, for Black, making an infinitely iterated jump in the position of figure 3.4 is equivalent to admitting defeat; as we will consider positions in which Black also loses if he decides to make an infinitely iterated jump and we will want Black to delay his defeat as much as possible, we generalise the forced-iteration rule of Draughts as follows, in order to simplify the exposition in the next section 3.1.1.

Rule 3.4 (Finite Iteration). If a player can make an iterated jump move, then he can only make a finite iteration of jumps.

Remark 3.5. Observe that rule 3.4 is still coherent with rules 3.1 and 3.2 also when considering the position of figure 3.4. In that position, making a jump move up the White ladder would not be a legal move for Black, who can thus move his only king in one of the other directions.

3.1.1 Embedding trees in the draughtboard

In this section, we will consider games of Infinite Draughts assuming rules 3.1 and 3.4, so that players must make jump moves when they can, and they can choose to iterate such jumps for at most finitely many times.

This will allow for simpler constructions in the rest of this section; we will consider different sets of rules in section 3.1.2.

Observe that, similarly to Infinite Hex, games of Infinite Draughts are not in general open, as it is possible to construct positions from which a player can win over an infinite play, but not at any finite stage of such play.

We therefore focus on the class of games of Infinite Draughts which are open for some designated player; we will prove the following.

Theorem 3.6. Every countable ordinal arises as the value of an open game of Infinite Draughts. Hence,

$$\omega_1^{\text{Draughts}} = \omega_1.$$

We will prove this Theorem by constructing an argument closely related to the proof that the omega one of 3-dimensional Infinite Chess is ω_1 , as shown by [Evans and Hamkins, 2014, §4].

We have already met a class of games with uncountable omega one; recall the class of climbing-through- T games and Proposition 1.24.

We will now describe a way to construct, given any tree T_0 , a game of Infinite Draughts \mathcal{D}_{T_0} equivalent to the game of climbing-through- T_0 , so that a closed-player strategy for one game determines a closed-player strategy for the other game.

First, we ensure that the infinite draughtboard does provide enough space for our construction.

Keeping in mind the bijection between the white squares of the infinite chessboard and $\mathbb{Z} \times \mathbb{Z}$ as specified by the system of reference presented in figure 3.5, we prove the following.




Figure 3.5: Orthogonal axes on the Infinite draughtboard.

Proposition 3.7. The full binary tree $2^{<\omega}$ can be realised in $\mathbb{Z}^+ \times \mathbb{Z}^+$.

Proof. Recall that the full binary tree is the tree whose nodes all have exactly two children.

We now describe an algorithm to construct recursively the full binary tree.

For $n = 0$, we start the recursion with a single vertex.

For $n = 1$, we stack two copies of the same vertex, in red in figure 3.6, at a sufficient distance in order to join them to the original vertex.




Figure 3.6: First recursive step.

We then proceed similarly, given a subtree T_n of $2^{<\omega}$, we construct T_{n+1} by attaching a copy of T_n to each of the leaves of the original T_n we stack two copies of it to the right of the previous subtree at a sufficient distance from the original subtree and join each root of the new copies of T_n to some leaf of the original T_n ; this is always possible by stacking the new copies of T_n far enough in the positive x -direction; this is clear in the second recursive step pictured in figure 3.7.

We can see that the trees T_n so constructed are binary and have 2^{2^n} leaves.




Figure 3.7: Second recursive step.

□

Remark 3.8. We have just shown that $2^{<\omega}$ can be realised in a quadrant of the integer lattice $\mathbb{Z} \times \mathbb{Z}$; however, at each step of the previous recursion, the new finite tree copies can be placed arbitrarily far in the positive direction of the x -axis, so that the whole tree $2^{<\omega}$ can be contained in a cone centered at the origin with an arbitrarily small non-zero angle at the centre.

We now define an arrangement of pieces on the draughtboard which will be key in the next argument.

Definition 3.9. A White *king tree* T_W is an arrangement of White kings on the infinite draughtboard, such that;

- i. if a Black king B is placed on a distinguished square, the *root node* of T_W , diagonally adjacent to exactly one White king of the king tree, then each king of T_W could be captured via a single jump move made by B ; this ensures that the kings of T_W are suitably alternated to empty squares and T_W is *connected*;

- ii. the *nodes* of T_W are the empty squares of the board on which B could end a, possibly not maximal, jump move; such jump move is unique for each node of T_W ; this ensures that T_W has no *loops*, and so really looks like a tree;
- iii. there is no empty square of the board diagonally adjacent to 4 kings of T_W ; so that the *nodes* of T_W have degree at most 3, and so we can say that T_W is a binary tree.

Following the usual graph-theoretic terminology, we say that a node of a White king tree T_W is a *leaf* if it is the square reached by a Black king that makes a maximal jump move starting from the root of T_W .

Theorem 3.10. For any well-founded tree T such that the game of climbing-through- T has value α for Observer, there is an open game \mathcal{D}_T of Infinite Draughts which has value α for White; moreover, any strategy for Climber in the climbing-through- T game determines a strategy for Black in \mathcal{D}_T , and vice versa.⁷

The initial position of \mathcal{D}_T comprises of a White king tree T_W , a Black king on the root of T_W , and White kings on all the leaves of T_W ; the Black king is placed on the Southernmost non-empty square of the board and Black is the first to move.

Proof. We will prove by transfinite induction on the rank α of the tree T that there exists a game \mathcal{D}_T of Infinite Draughts as described.

If $\alpha = 0$, then White has already won and the empty board can be the trivial initial position of \mathcal{D}_T ; Black loses by being the first player.

If $\alpha = 1$, then the children of the root node v_0 of T are all leaves. If v_0 has only one child, then we can choose the position in figure 3.8 as the initial position of \mathcal{D}_T ; Black must make an initial jump move that ends on the square labelled with a dagger, and his only piece is then captured by White in the following turn.

Observe that figure 3.8 also shows what the leaves of the White king tree embedded in the draughtboard look like.




Figure 3.8: Position with value 1 for White, Black to move.

⁷Such strategy correspondence could be made unique by imposing arbitrary conditions on them; it is not an issue in our discussion.

Similarly, if the root node of T has any other finite number of children, we can iterate this construction; see figure 3.9 for the initial position of \mathcal{D}_T when v_0 has exactly 3 children.




Figure 3.9: Black to move and choose one of 3 jump moves; value 1 for White.

Furthermore, if v_0 has (countably) infinitely many children, then we can let the initial position of \mathcal{D}_T be as in figure 3.10.

Observe that, in the last 3 figures presented, the Black king can only make legal maximal jump moves ending on the squares labelled with daggers.

In particular, in figure 3.10, Black must choose one of such squares; in fact, rule 3.4 prevents the Black king from making an infinitely iterated jump and thus remaining on the White ladder highlighted in green.

We can see that, given a strategy for Climber, that is a choice of a child node of v_0 , we can determine a strategy for Black in \mathcal{D}_T , that is a choice of jump move to the corresponding square labelled with a dagger.

Here, the fundamental observation is that we can see the initial position of \mathcal{D}_T , in which the Black king stands at the root node of the king tree, as corresponding to the initial position of the climbing-through- T game, in which Climber stands on the root of T ; moreover, we can see the Black king reaching a square labelled with a dagger as corresponding to the Climber moving up to the relevant node.

Hence, this construction allows us to determine the strategy correspondence of the Theorem.




Figure 3.10: Position with value 1 for White; Black to move.

Now, suppose that for all trees of rank $\beta < \alpha$ there are games of Infinite Draughts with value β as described.⁸

First, suppose for simplicity that the root v of T has a single child v' ; if the maximal sub-tree T' of T with root v' has rank δ ,⁹ then T has rank $\delta + 1 = \alpha$.

By induction, there exists some Infinite Draughts game $\mathcal{D}_{T'}$ of value δ as described by the Theorem; let \tilde{T}'_W be the arrangement of White kings at the initial position of $\mathcal{D}_{T'}$, so that \tilde{T}'_W consists of a White king tree T'_W with White kings on the leaf nodes.

Let the initial position of \mathcal{D}_T be as in figure 3.11, in which we embed two copies of \tilde{T}'_W , so that each distinct copy of the king tree T'_W has one of the two starred squares as a root node.

Observe that this initial position of \mathcal{D}_T is constructible thanks to Proposition 3.7. The full binary tree can be embedded in a quadrant of the integer lattice and the king tree T'_W is binary, so the position of figure 3.11 can be constructed so that the Black king stands on the Southernmost non-empty square of the draughtboard. Observe that the embedded copies of T'_W may have more White kings than the original arrangement because Proposition 3.7 ensures constructibility, provided flexibility in the length of edges; this is not an issue because it is never advantageous for the Black king to end jump moves on nodes of the king tree which are not

⁸Recall that the rank of T equals the game value of climbing-through- T by Proposition 1.16.

⁹ T' is maximal in the sense that its vertex set is the same as the one of T , except for v .

binary,¹⁰ such “stretched” copies of T'_W will be equivalent to the original T'_W with respect to the possible strategies available to Black.

In other words, Proposition 3.7 allows us to embed in the draughtboard king trees with the same graph-theoretic features of the original T'_W .




Figure 3.11: Inductive step.

We can now describe the main line of play of Black, who is the first to move in \mathcal{D}_T .

The initial move by black is a jump move to the tile labelled with a dagger; we call it a *resting* square because the Black king can stop there without being captured by White at the next turn.¹¹

Then, we can see that the Black king is standing on the root of two White king trees, highlighted in orange, which each contain a copy of T'_W ; at the next turn, White can only move one piece from either of such king trees.¹²

At the next turn, Black can let his king make a jump move up the king tree whose pieces were not moved by White, and then play as in $\mathcal{D}_{T'}$.

¹⁰According to our definition of king tree, all the empty squares between white kings are *nodes*; however, the not properly branching nodes cannot be resting squares.

¹¹Observe that the Black king reaches the resting square by jumping over a White piece, hence it ends the move being adjacent to exactly 2 White pieces.

¹²Or one of the kings on the leaves of such king trees.

Since Black can then play in $\mathcal{D}_{T'}$ so to realise the game value δ , then the strategy just described realises the value $\delta + 1 = \alpha$. In fact, if the Black king does not stop on the described resting square, then it could stop on a non-resting square, leading to Black's defeat on the next turn, or move past the first resting square and stop on some resting square higher up in one of the embedded copies of $\mathcal{D}_{T'}$; however, recalling the correspondence with the climbing-through- T game, such a move would correspond to the Climber climbing up more than one edge per turn, causing a decrease in game value larger than necessary. Moreover, observe that White's moves cannot really obstruct Black in the same way that the Observer does not obstruct the Climber by saying "OK".

Hence, the main line of play for Black describes a well-defined strategy.

We have now all that we need to complete the proof; if the root of the tree T has finitely many children, then we can iterate the previous construction so that, for each child v of the root of T , letting T_v be the maximal sub-tree of T rooted at v , we can have, for each v , a distinct resting square which is the root of two White king trees, both realising the White king tree of \mathcal{D}_{T_v} .

It should be clear that we can inductively construct the initial position of \mathcal{D}_T when the root of T is infinitely branching, as in figure 3.12; a position for \mathcal{D}_T when the root of T has finitely many children can be obtained by truncation of the green branch.




Figure 3.12: \mathcal{D}_T when the root of T is infinitely branching.

Observe that the discussion following figure 3.11 regarding the correspondence

between the resting squares of the initial position of \mathcal{D}_T and the nodes of T in climbing-through- T easily extends to the finite and infinite branching positions of figures 3.11-3.12, with the only difference that White may respond to Black's moves by moving pieces which Black does not intend to capture.¹³

The proof is now complete. □

Observe that Theorem 3.6 is a direct consequence of Theorem 3.10 and Proposition 1.24; we can appreciate that Theorem 3.10 is a stronger result because it does not just prove the existence of Infinite Draughts games of arbitrarily high (countable) value, but also constructs them explicitly.

3.1.2 Further results

Computable positions

Note that we can define a position of Infinite Draughts to be a function that assigns to each white square of the chessboard a value to mark it as free or occupied by a piece, specifying if it is a pawn or a king and whether it is Black or White, in addition to specifying a turn indicator.

It is then immediate to consider computable positions as computable functions.

The following result is an adaptation of [Evans and Hamkins, 2014, Theorem 6] and is a consequence of Theorem 3.10.

Corollary 3.11. There is a computable position of Infinite Draughts which is a win for a designated player if both players are constrained to play according to computable strategies, and is a draw otherwise.

We first recall a classic graph-theoretic result.

Lemma 3.12 (König's Lemma). Any infinite finitely-branching tree has an infinite branch.

Proof of Corollary 3.11. Let $T \subset 2^{<\omega}$ be a computable infinite tree with no computable infinite branch; such trees exist by elementary results of computability theory.

Since $T \subset 2^{<\omega}$ is clearly finitely-branching, then T is an infinite finitely-branching tree, and so T has an infinite branch by König's Lemma; such branch is not computable by definition of T .

Even though T is not well-founded, it still has branches of order type at most ω , so that we can apply the construction in the proof of Theorem 3.10 to construct a game \mathcal{D}_T of Infinite Draughts equivalent to the climbing-through- T game.

¹³In other words, White could move pieces in irrelevant far branches of the king tree.

It is clear that Climber has a winning strategy for the climbing-through- T game, i.e. climbing up an infinite branch of T ; however, T has only non-computable infinite branches, so such strategy for Climber must be not computable.

By the strategic equivalence of climbing-through- T with \mathcal{D}_T , we have that also Black has a non-computable drawing strategy in \mathcal{D}_T .

However, if Climber plays according to a computable strategy in climbing-through- T , then he can only climb up a computable branch of T , which must be finite by definition of T . Correspondingly, if Black plays according to a computable strategy, then he will eventually reach in finitely many turns a leaf of the king tree in \mathcal{D}_T ; in such case, Black reaches an arrangement of White kings as in figure 3.8 and thus loses.

It is clear that any strategy is winning for White, as it is enough to require that White captures the Black king at the first mistake of Black, and make any other uninfluential move, otherwise; hence, we can say that White has a computable strategy which allows him to win in \mathcal{D}_T when Black plays according to any computable strategy.

This completes the proof. □

Tree nodes for different rules

We now focus on extending the previous results to games of Infinite Draughts that assume a less strict set of rules.

In particular, we will design nodes for the White king tree constructed in Theorem 3.10 with more structure than the one in figure 3.11.

All the positions in this chapter were kindly reviewed by Sergio Scarpetta, current World Champion of English Draughts (3-move variant), who confirmed the interpretations presented here, in particular the one regarding figure 3.15.

We observe that by allowing nodes to have fixed finite size in Proposition 3.7, we conclude that the full n -ary tree $n^{<\omega}$ can be embedded in the integer lattice $\mathbb{Z} \times \mathbb{Z}$.

Now that we have ensured to have enough space on the infinite draughtboard, we can present the node needed when dropping the forced jump rule 3.1, but maintaining the finite iteration 3.4.

Firstly, we need to ensure that the Black king actually starts jumping on the White king tree, so that the root of the king tree needs two extra *guardian* White kings, as in figure 3.13, where Black immediately loses by making an initial non-jump move; this was not needed in the previous root configuration.

Even though the Black king will still find it disadvantageous to stop on non-resting squares, we still need to ensure that the Black king does not get *off-track*, that is out of the king tree, when making a move from a resting square; that is




Figure 3.13: Root with guardians.

exactly why we add two guardian White kings to the king tree branching node, which are highlighted in yellow in figure 3.14.




Figure 3.14: Branching node with guardians.

Remark 3.13. Observe that when the Black king is on some resting square, White could move one of the guardian kings to a square adjacent to the Black king, offering it the possibility to go off-track.

The interesting thing to note is that White would make such a move only when, following the Black king going off-track, White could reach a position of value lower than the ones reachable by allowing Black to play as usual; such a situation means exactly that it is not advantageous for Black to go off-track.

This is enough to say that, with optimal play, Black will not capture a guardian king moved by White. In other words, White moves guardian kings only when

that is uninfluential, exactly as the Observer says “OK” in the climbing-through- T game.

Recalling that we stated rule 3.4 in order to avoid discussing infinitely iterated jumps, we are now ready to drop it and only assume the forced iteration rule 3.2 with its strict interpretation in Infinite Draughts.

In other words, we now assume that players are not obliged to make jump moves, but, if they do make a jump move, then they can only make maximal jumps, which include infinitely iterated jumps.

This complicates the structure of the nodes of, what we should now call, *extended king tree* constructed as in the proof of Theorem 3.10; in fact, we will require the Black king to reach a resting square and possibly make a non-jump move, before continuing to climb the extended king tree.

Observe in figure 3.15 the first node that the Black king can decide to pick; it is a 7-branching node, so that each branch leads to an equivalent extended tree, as for the binary case in the original argument.




Figure 3.15: 7-branching node.

Remark 3.14. It may be possible to further simplify the node in figure 3.15; however, the one presented is likely to be the minimal node when reintroducing the forced jump rule.

Observe that the branches are coloured red or orange depending on their row parity.

Suppose that the Black king decides to pick that first node, rather than continuing on the green ladder, and so ends the initial jump move on the resting square labelled with a dagger.

The main line of play for Black is as follows. After White's non-jump move, move the Black king to the starred square adjacent to an external triplet of branches which was not modified by White; White can then move a piece from only one of the red branches next to the Black king, which can then make a jump move to climb up the other, going to another node.

There are several ways in which White could interfere with this strategy, and we will see that they are uninfluential in terms of reducing the game value.

Even though players are not forced to make jumps when legal, White might offer the possibility of a jump move to Black.

Recalling that jumps must be maximal, we have that, if White moves a king so that Black can make a jump towards North-West, then Black can decide to deviate from the main line of play to go up the orange branch, in which case Black leaves the node reducing the value by 1 rather than 2, which is still a positive outcome for Black.

White could also offer possible jump moves towards NE or SW by moving a White king to a starred tile; if Black takes up the offer, then he simply jumps up the relevant external orange branch.

Moreover, if White moves a king to the South-East of the Black king, then it is likely to be advantageous for Black to make such a jump move and go to choose another node up the green branch; this would correspond to the Observer allowing the Climber to go back down to the tree root and make a different choice.

Following White's response to Black's initial move, Black may decide to move his king to a non-starred square.

If the Black king is moved to the North-West of the resting square, then it would be captured by White on the next move.

If instead the Black king is moved towards South-East, that is "backwards", then White can place a king to the North-West of the (original) resting square, and then we can see that the Black king is trapped, also thanks to the guardian kings highlighted in yellow.

Finally, the only other option available to White when Black is standing on a starred square is moving the piece at the bottom of the external orange branch next to the Black king; by an analogous argument to Remark 3.13, we have that White would make that move only if that does not affect Black's play.

With this, we conclude our discussion of Infinite Draughts.

Appendix A

Appendix

A.1 2-colourability in Maker-Breaker games

We prove the following result, which is adapted for any generalised hypergraph (B, \mathcal{F}) from [Hefetz et al., 2014, Proposition 2.3.1].

Proposition A.1. Take some Maker-Breaker game \mathcal{G} played on a hypergraph (B, \mathcal{F}) such that \mathcal{G} is open for Breaker, and Maker is the first player to move.¹ If Breaker has a winning strategy, so that he can play to take at least one vertex in each of Maker's winning subsets in \mathcal{F} , then (B, \mathcal{F}) is properly 2-colourable.

Proof. Suppose that σ is a winning strategy for Breaker, and let Breaker play according to it, so that at least one vertex in each hyperedge in \mathcal{F} is chosen by Breaker, who does that after finitely many turns.

Moreover, let Maker make an arbitrary initial move, which cannot be disadvantageous, and then play according to σ , as if the arbitrary move was not placed on the board. Whenever σ prescribes to choose a vertex already marked by Maker himself, then he can analogously make another arbitrary move. Note that Maker's extra move is the only one that the strategy σ does not consider, thus σ will never prescribe Maker to choose a vertex already marked by Breaker.

Then, Maker will satisfy Breaker's winning condition; that is, given any hyperedge in \mathcal{F} , Maker will choose one vertex from it, and will do that after finitely many turns.

Note that if B is infinite, then not all vertices will necessarily be chosen by either player at any finite stage; if a vertex does not belong to any hyperedge, then it will never be chosen.

¹Note that this argument relies on the fact that Breaker's winning strategy is a second-player strategy.

Moreover, if B is uncountable, then at most countably many vertices will be marked by either player, as all plays are countable sequences of moves, by Definition 1.2.iv.

However, by the definition of σ , all hyperedges comprise one vertex marked by Maker and one by Breaker, both chosen at some finite stages of the game.

Assign the colour white to all the vertices marked by Maker, and the colour black to all the vertices marked by Breaker or left unchosen.²

This assignment is a proper 2-colouring of (B, \mathcal{F}) , as needed. \square

A.2 Strategy-stealing in stone-placing games

We generalise the symmetry condition of the finite Hex board highlighted in Remark 2.9.ii to construct the Strategy-stealing argument for stone-placing games.

Definition A.2. A game is *strictly not open* for one player if that player can only win by playing for infinitely many turns; that is, if the winning condition of that player has empty interior, and so does not contain any basic open set, within the space of plays.

Proposition A.3 (Strategy-stealing). Let \mathcal{G} be a stone-placing game played on $(B, \mathcal{F}, \mathcal{S})$, strictly not open for both players. Say that there is a fixed-point free involution g of the board B . Suppose that g is such that for each $s \in \mathcal{S}$ there is some $f \in \mathcal{F}$ for which $f \subset g[s]$.

Then, the second player in \mathcal{G} does not have a winning strategy.

Hence, we can informally interpret the involution g as a symmetry (or better, a reassemblage) of the board that allows the first player to act as the second player on the new board.

Proof of Proposition A.3. Let the first player make an arbitrary initial move, which is advantageous, and then play in the following way; when the second player chooses a vertex $v \in B$, then the first player, in the next turn, marks $g(v)$.

Note that $s \in \mathcal{S}$ is infinite by strict non-openness of \mathcal{G} , and that the second player will mark all the vertices of s , each after finitely many turns. Thus, the first player will eventually mark all the vertices of $g[s]$.

Therefore, the first player will also win by marking all the vertices of some $f \in \mathcal{F}$ such that $f \subset g[s]$, which is a contradiction. \square

Remark A.4. Observe that we cannot relax the assumption of strict non-openness in Proposition A.3.

²So that we do not need the Axiom of Choice here.

Otherwise, the second player could pick a finite $s \in \mathcal{S}$, which preferably does not contain the initial move of the first player, and then could proceed to mark all the vertices in s , so that the first player would be able to fully satisfy his winning condition $f \subset g[s]$ only with his move that follows the last move of the second player; that is a move too late, so that the second player actually wins.

Note that strengthening the condition by imposing that the involution g is such that for each $s \in \mathcal{S}$ there is some $f \in \mathcal{F}$ for which $f \subsetneq g[s]$ is still not enough; there could be an appropriate order in which the second player could mark the vertices of a finite $s \in \mathcal{S}$ that defeats the opponent.

Bibliography

- [Adler et al., 2016] Adler, A., Daskalakis, K., and Demaine, E. (2016). The complexity of hex and the jordan curve theorem. In Chatzigiannakis, I., Mitzenmacher, M., Rabani, Y., and Sangiorgi, D., editors, *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, number XXX in Leibniz International Proceedings in Informatics. Dagstuhl Publishing.
- [Björnsson et al., 2006] Björnsson, Y., Hayward, R., Johanson, M., and van Rijswijk, J. (2006). Dead cell analysis in hex and the shannon game. In Bondy, A., Fonlupt, J., Fouquet, J., Fournier, J., and Ramírez Alfonsín, J., editors, *Graph Theory in Paris*, Trends in Mathematics, pages 45–59. Birkhäuser Basel.
- [Bonnet et al., 2016] Bonnet, É., Jamain, F., and Saffidine, A. (2016). On the complexity of connection games. *Theoretical computer science*, 644(C):2–28.
- [Browne, 2000] Browne, C. (2000). *Hex strategy: making the right connections*. CRC Press.
- [Conway, 1996] Conway, J. H. (1996). The angel problem. In Nowakowski, R. J., editor, *Games of No Chance*, volume 29 of *Proceedings of the MSRI Workshop on Combinatorial Games*, pages 3–12. Cambridge University Press.
- [Evans and Hamkins, 2014] Evans, C. D. A. and Hamkins, J. D. (2014). Transfinite game values in infinite chess. *Integers*, 14:Paper No. G2, 36.
- [Evans et al., 2017] Evans, C. D. A., Hamkins, J. D., and Perlmutter, N. L. (2017). A position in infinite chess with game value ω^4 . *Integers*, 17:Paper No. G4, 22.
- [Gale, 1979] Gale, D. (1979). The game of hex and the brouwer fixed-point theorem. *The American Mathematical Monthly*, 86(10):818–827.
- [Gale and Stewart, 1953] Gale, D. and Stewart, F. M. (1953). Infinite games with perfect information. In Kuhn, H. W. and Tucker, A. W., editors, *Contributions to the Theory of Games*, volume II, pages 245–266. Princeton University Press.

- [Gardner, 1988] Gardner, M. (1988). *Hexaflexagons and other mathematical diversions – The First Scientific American Book of Puzzles and Games*. University of Chicago Press.
- [Hayward and Toft, 2019] Hayward, R. B. and Toft, B. (2019). *Hex, Inside and Out: The Full Story*. CRC Press.
- [Hefetz et al., 2014] Hefetz, D., Krivelevich, M., Stojaković, M., and Szabó, T. (2014). *Positional Games*. Oberwolfach Seminars. Springer Basel.
- [Hein, 1942] Hein, P. (1942). Vil de laere polygon? *Politiken*, Copenhagen, Dec. 26. Translated into English by [Hayward and Toft, 2019, §3.1].
- [HexWiki, 2021] HexWiki (2021). Strong connection. https://www.hexwiki.net/index.php/Strong_connection (accessed 1st July 2021).
- [Kunen, 2013] Kunen, K. (2013). *Set theory*, volume 34 of *Studies in Logic*. College Publications, London, revised edition.
- [Maehara, 1984] Maehara, R. (1984). The jordan curve theorem via the brouwer fixed point theorem. *The American mathematical monthly*, 91(10):641–643.
- [Máthé, 2007] Máthé, A. (2007). The angel of power 2 wins. *Combinatorics, Probability and Computing*, 16(3):363–374.
- [Nash, 1952] Nash, J. (1952). Some games and machines for playing them. *Technical Report D-1164*. Rand Corporation. <https://www.rand.org/content/dam/rand/pubs/documents/2015/D1164.pdf>.
- [Schachner, 2019] Schachner, M. (2019). The game of hex: a study in graph theory and algebraic topology. *Mathematics REU 2019*. University of Chicago. <http://math.uchicago.edu/~may/REU2019/REUPapers/Schachner.pdf>.
- [Uiterwijk and Moesker, 2009] Uiterwijk, J. and Moesker, K. (2009). Mathematical modelling in twixt. In Löwe, B., editor, *Proceedings Logic and the Simulation of Interaction and Reasoning 2 Workshop (LSIR2)*, pages 29–35. Pasadena: IJCAI.
- [van den Herik et al., 2002] van den Herik, H., Uiterwijk, J. W., and van Rijswijck, J. (2002). Games solved: Now and in the future. *Artificial Intelligence*, 134(1):277–311.