

Chain-making games in grid-like posets

DANIEL W. CRANSTON, WILLIAM B. KINNERSLEY*,
KEVIN G. MILANS*, GREGORY J. PULEO* AND DOUGLAS B. WEST†

We study the Maker-Breaker game on chains in a poset. In a chain-product poset, the maximum size of a chain that Maker can guarantee building is $k - \lfloor r/2 \rfloor$, where k is the maximum size of a chain in the poset and r is the maximum size of a factor chain. We also study a variant where Maker must build a chain in increasing order, called the *ordered chain game*. Within the bottom k levels of a product of d chains of size at least k , Walker can guarantee a chain that hits all levels if $d \geq 14$; this result uses a solution to Conway’s Angel-Devil game. When $d = 2$, the maximum that Walker can guarantee is only $2/3$ of the levels; when $d = 3$, Walker cannot guarantee all levels, as shown by Clarke, Finbow, Fitzpatrick, Messenger, and Nowakowski by studying a related game. It is unknown whether Walker can guarantee all levels when $4 \leq d \leq 13$. In the product of two chains of equal size, Walker can guarantee $2/3$ of the levels asymptotically.

1. Introduction

The *Maker-Breaker* game on a hypergraph \mathcal{H} is played by Maker and Breaker, who alternate turns (beginning with Maker). A move acquires a previously unchosen vertex of \mathcal{H} . Maker wins by acquiring all vertices of some edge of \mathcal{H} ; Breaker wins by preventing this.

Maker-Breaker games have been studied for many hypergraphs, particularly when the vertices are the edges of an n -vertex complete graph. In that setting, when n is large enough Maker can build a spanning cycle, a complete subgraph with q vertices, a spanning k -connected subgraph, or various other structures. For an introduction to Maker-Breaker games (and more general positional games), see the surveys [1] and [2] or the book [3]. Recent work has also considered how quickly Maker can win [8] and what bias Breaker

*Research of Kinnersley, Milans, and Puleo supported by NSF grant DMS 08-38434, “EMSW21-MCTP: Research Experience for Graduate Students”.

†Research partially supported by the National Security Agency under Award No. H98230-10-1-0363.

needs to win [4, 5, 10], where bias b means that Breaker makes b moves for each move by Maker.

In this paper, we study the *chain game* on posets, where the winning sets are the chains with a given size. For every poset there is a maximum size of chain that Maker can build against optimal play by Breaker; we seek this value. For special posets whose elements are integer d -tuples, we give efficient strategies for Maker that do not waste any move (every element that Maker selects is in the chain constructed).

Chains in posets are ordered from bottom to top, so a natural variant of the chain game is the *ordered chain game*, in which the chain must be built from bottom to top. The rules for moving are the same as in the chain game; the difference is that the “score” that counts for Walker is the maximum size of a subsequence of played elements forming a chain played in increasing order. To distinguish this game from the (unordered) chain game, we call its players *Walker* and *Blocker*.

We use $\mathbf{r}, \mathbf{k}, \mathbf{2}$ to denote the posets that are chains of sizes $r, k, 2$, etc. The *product* of d chains with sizes r_1, \dots, r_d , written $\prod_{i=1}^d \mathbf{r}_i$, is the set of d -tuples x such that $0 \leq x_i < r_i$ for $1 \leq i \leq d$, ordered by $x \preceq y$ if and only if $x_i \leq y_i$ for $1 \leq i \leq d$. (Similarly, $x \prec y$ if $x \preceq y$ and $x \neq y$.) The *d -dimensional k -wedge*, written \mathcal{W}_k^d , is the subposet of \mathbf{k}^d consisting of the nonnegative-integer d -tuples with sum less than k , under the coordinate-wise order.

For the chain game and ordered chain game on these posets, we prove two main results.

Theorem 1.1. *Let P be a product of d chains, with r being the maximum size among these chains and k being the maximum size of a chain in P . In the chain game on P , Maker can build a chain of size $k - \lfloor r/2 \rfloor$, and Breaker can prevent Maker from building a larger chain.*

Theorem 1.2. *If $d \geq 14$, and $k \in \mathbb{N}$, then in the ordered chain game on \mathcal{W}_k^d , Walker can build a chain of size k (hitting all levels).*

In \mathcal{W}_k^d with $d = 2$, Walker can guarantee hitting $\lceil 2k/3 \rceil$ levels but no more. Clarke, Finbow, Fitzpatrick, Messenger, and Nowakowski [6] showed that when $d = 3$, Walker cannot hit all levels. Somewhere between dimension 4 and dimension 14 is a least dimension in which Walker can hit all levels. We conjecture that this happens already in dimension 4.

We begin in Section 2 with the Maker-Breaker chain game on products of chains. The remainder of the paper addresses the ordered chain game. In Section 3 we study the 2-dimensional case for both wedges (proving the $\lceil 2k/3 \rceil$ result mentioned above) and grids (products of equal chains). In the

product of two chains of equal size, Walker can guarantee asymptotically $2/3$ of the levels; we prove this by a “potential function” argument. In Section 4 we briefly discuss the Seepage game of Clarke et al. [6]; this game (formulated independently of this paper) generalizes the ordered chain game to acyclic digraphs. Section 4 also relates the ordered chain game to Conway’s famous Angel-Devil game. We apply this relationship in Section 5 to prove Theorem 1.2; there we also consider the biased game in which Blocker makes b moves after each move by Walker.

2. Maker-Breaker on chain-products

In a product of chains, we use *level* ℓ to denote the set of elements whose entries sum to ℓ . A *successor* of a d -tuple x is a d -tuple y such that $x \prec y$. To evoke familiar terminology from games on physical boards, we refer to an element chosen at a particular time as a *move* and say that the player *plays* that move at that time.

In order to solve the Maker-Breaker game on products of finite chains, we first solve the ordered chain game on products of 2-element chains. We then apply this lemma to build an optimal strategy for Maker in the unordered chain game on arbitrary finite chain-products. Let $[d] = \{1, \dots, d\}$. There is a natural isomorphism from 2^d to the lattice of subsets of $[d]$ in which each binary d -tuple x is mapped to $\{i: x_i = 1\}$.

Lemma 2.1. *For $d \geq 2$, let P' be the poset obtained from 2^d by deleting the top element and the bottom element. The maximum size of a chain in P' is $d - 1$, and Walker can build a chain of size $d - 1$ in P' , even if Blocker moves first.*

Proof. View the elements of P' as subsets of $[d]$. On his k th turn, Walker plays a move S at level k such that (i) S is above all his previous moves and (ii) Blocker has played no successor of S . Successfully executing this strategy for $d - 1$ turns builds a chain of size $d - 1$.

Let S be the previous move by Walker. If Blocker responded with a move not a successor of S , then Walker can add any element to S . Otherwise, since the highest level of P' is $d - 1$, the move by Blocker omits some $e \in [d]$. Walker now plays $S \cup \{e\}$ and restores the property that no successor of the current move has been played. \square

Since Walker can build a chain hitting all levels in P' , we conclude also that Maker can build a chain hitting all levels in the unordered game. The latter statement, along with the freedom to let Breaker move first, is what we need to analyze arbitrary chain-products. Since Maker need not take




Figure 1: Maker strategy for chain-products.

the elements of a chain in order, Maker can build chains independently in different copies of the poset P' in Lemma 2.1; they will combine to form a large chain.

To show optimality of the resulting strategy for Maker, we present a strategy for Breaker. Since every chain in a chain-product is contained in a longest chain, it suffices to give a pairing strategy for Breaker that guarantees blocking enough of every longest chain.

Theorem 2.2. *Let $P = \prod_{i=1}^d r_i$. In the unordered chain game on P with $r = \max_i r_i$, Maker can guarantee building a chain of size $k - \lfloor r/2 \rfloor$, where k is the maximum size of a chain in P , and Breaker can keep Maker from building any larger chain.*

Proof. The elements of P are the d -tuples x such that $0 \leq x_i < r_i$ for $1 \leq i \leq d$. By symmetry, we may assume that $r = r_d = \max_i r_i$. For $0 \leq j < r$, let z_j be the element of P whose i th coordinate is $\min\{j, r_i - 1\}$. For $1 \leq j < r$, let A_j be the subposet of P consisting of all x such that $z_{j-1} \prec x \prec z_j$. See Fig. 1, where $\{z_0, \dots, z_{r-1}\}$ are in bold. Note that $z_0 \prec \dots \prec z_{r-1}$, that A_1, \dots, A_{r-1} are pairwise disjoint, and that each A_j is isomorphic to the poset P' of Lemma 2.1 for some dimension (as j increases beyond some r_i , the dimension decreases). Fig. 1 illustrates A_1, \dots, A_{r-1} .

The key to Maker's strategy is that chains in A_1, \dots, A_{r-1} combine with $\{z_0, \dots, z_{r-1}\}$ to form a chain in P . Let $Z = \{z_0, \dots, z_{r-1}\}$. Maker begins by playing an element in Z . When Breaker plays an element in Z , Maker responds by playing another element in Z if one is available. Maker treats

each subposet A_j as an instance of the poset P' of Lemma 2.1; when Breaker plays in A_j , Maker responds using the strategy of Lemma 2.1. When Breaker plays any other move, Maker plays to increase the chain in some A_j or Z .

By Lemma 2.1, Maker obtains in each A_j a chain hitting all levels. These combine with Z to form a long chain in P . The only levels that Maker misses are those containing an element of Z played by Breaker. Maker's strategy ensures that Breaker plays at most $\lfloor |Z|/2 \rfloor$ such moves. Since $|Z| = r$, the bound follows.

To prove optimality, Breaker uses a pairing strategy. For all j with $0 \leq j < \lfloor r/2 \rfloor$ and all (x_1, \dots, x_{d-1}) , Breaker pairs the elements $(x_1, \dots, x_{d-1}, 2j)$ and $(x_1, \dots, x_{d-1}, 2j+1)$. When Maker plays a paired element, Breaker plays its mate; when Maker plays an unpaired element, Breaker responds arbitrarily. We show that Maker misses at least $\lfloor r/2 \rfloor$ elements from every maximal chain.

Given a maximal chain X , let X_j be the subchain of X consisting of all elements whose last coordinate has value j , for $0 \leq j \leq r-1$. Let x be the last element of X in X_{2m} , and let y be the first element of X in X_{2m+1} (note that X cannot skip any X_j). Since Breaker has paired x with y , Maker misses at least one of these. Thus in every maximal chain Maker misses elements from at least $\lfloor r/2 \rfloor$ levels. \square

3. Walker vs. Blocker in two dimensions

For the ordered chain game, wedges are easier to analyze than chain-products, because the poset formed by the successors of any element is isomorphic to a truncation of the same wedge by discarding the highest levels. This allows Walker to define a strategy locally. In a chain-product, when the top of a growing chain reaches the maximum in a given coordinate, no further moves in that direction are possible. Blocker may be able to take advantage of Walker being "trapped in a corner". To overcome this, Walker may need to look farther ahead to plan a strategy.

In this section, we first give an exact solution for the ordered chain game on \mathcal{W}_k^2 . We then express Walker's strategy in a more general way using a "potential function" to show that asymptotically as big a chain can be built in the game on the product of two $(k+1)$ -element chains as on the wedge \mathcal{W}_{2k+1}^2 that contains it.

Theorem 3.1. *In the ordered chain game on \mathcal{W}_k^2 , Walker can build a chain of size $\lfloor 2k/3 \rfloor$ in $\lfloor 2k/3 \rfloor$ moves. Also, Blocker can prevent Walker from building a larger chain.*

Proof. We present a strategy in which Walker follows a single chain, with no wasted moves. Walker first plays $(0, 0)$. For each subsequent move, Walker plays a successor of his previous move; among all unchosen successors, he plays one at the lowest level. A level containing a move by Walker is “won by Walker”. After a move (a, b) , Walker next plays at level $a + b + 1$ unless Blocker has already played both $(a + 1, b)$ and $(a, b + 1)$. We then say that Blocker wins level $a + b + 1$.

Since Blocker spends at least two moves in a level to win it, while Walker spends only one move per level won, the number of levels that have been won by Walker is always at least twice the number won by Blocker. At the end of the game all k levels have been won by one player or the other; hence Walker has won at least $\lceil 2k/3 \rceil$ levels.

For the upper bound, we present a strategy for Blocker to keep Walker from building a larger chain. If Walker’s previous move was at (a, b) and exactly one of $(a + 1, b)$ and $(a, b + 1)$ have been played, then Blocker plays the other. If neither of them has been played, then Blocker plays $(a + 1, b + 1)$, if available. Otherwise, Blocker plays an arbitrary move.

Once the game has finished, let (a, b) be an element on a largest chain C that was occupied by Walker in order. If when Walker played (a, b) one of $(a + 1, b)$ and $(a, b + 1)$ had already been played, then C has no element from level $a + b + 1$. If C has an element x from level $a + b + 1$, then before x was played the element $(a + 1, b + 1)$ was played by one of the players. Blocker next ensures that the other successor of x at level $a + b + 2$ is occupied, thus preventing C from having an element from level $a + b + 2$.

We have shown that Blocker’s strategy prevents Walker from building a chain in order that hits three consecutive levels. Hence Walker wins at most $\lceil 2k/3 \rceil$ levels. \square

In efficient strategies, where all moves by Walker form a chain, we refer to the most recent move by Walker as the *head* of the chain. A more global view of the strategy for Walker uses a potential function to measure the difficulty that Walker faces in the levels above the head.

We define a potential function to measure the future levels that Walker may need to skip. Thus Walker wants to keep the potential small. When Walker skips a level, the potential will decrease by 1. Other moves by Walker will not increase the potential. A move by Blocker will increase the potential by at most $1/2$. We design such a potential and strategy for \mathcal{W}_{2k+1}^2 and use it to show that even when the game is restricted to the subposet $(\mathbf{k} + \mathbf{1})^2$, Walker can still win asymptotically $2/3$ of the levels.

Blocker’s move at a position (c, d) makes it harder for Walker to win level $c + d$. To measure this difficulty when the head is at (a, b) , we define

the *influence* of (c, d) on (a, b) , where $a \leq c$ and $b \leq d$, to be $\binom{c'+d'}{d'} 2^{-(c'+d')}$, where $(c', d') = (c, d) - (a, b)$ (the influence is 0 if $a > c$ or $b > d$). We write $f_{a,b}(c, d)$ for the influence of (c, d) on (a, b) . Define the *potential* $\Phi_{a,b}$ to be the total influence on (a, b) of the moves Blocker has played. Large potential is good for Blocker.

To motivate these definitions, note that the influence of (c, d) on (a, b) equals the probability that a random walk from (a, b) to level $c + d$ will end at position (c, d) , where the walk iteratively increases a randomly chosen coordinate by 1. Let (a, b) be the current head of the chain, and let (c, d) be another position. The average of the influences of (c, d) on $(a + 1, b)$ and $(a, b + 1)$ equals the influence of (c, d) on (a, b) . Walker will want to choose the option that produces the smaller potential.

Theorem 3.2. *In the ordered chain game on $(\mathbf{k} + \mathbf{1})^2$, Walker can build a chain hitting more than $\frac{2}{3}(2k + 1) - 4\sqrt{k \ln k}$ of the $2k + 1$ levels, and this is asymptotically sharp.*

Proof. Since Blocker can limit Walker to winning $\lceil (2/3)(2k + 1) \rceil$ levels in \mathcal{W}_{2k+1}^2 , Walker can do no better on the subposet $(\mathbf{k} + \mathbf{1})^2$. Hence it suffices to prove the lower bound.

Consider the game on \mathcal{W}_{2k+1}^2 . At a given time, let $S(a, b)$ denote the set of elements at or above (a, b) that Blocker has played. Recall that the potential $\Phi_{a,b}$ at a point (a, b) is $\sum_{(c,d) \in S(a,b)} f_{a,b}(c, d)$.

We have noted that always $f_{a,b}(c, d) = \frac{1}{2} [f_{a+1,b}(c, d) + f_{a,b+1}(c, d)]$. When the head of the chain is at (a, b) , we have $(a, b) \notin S(a, b)$, and hence summing over $S(a, b)$ yields $\Phi_{(a,b)} = \frac{1}{2} (\Phi_{(a+1,b)} + \Phi_{(a,b+1)})$. To keep the potential small, Walker wants to move to whichever of $(a + 1, b)$ and $(a, b + 1)$ has smaller potential.

If this strategy directs Walker to play a move (a', b') that Blocker already played (that is, $(a', b') \in S(a, b)$), then Walker simply computes the choice the strategy would make from (a', b') instead. The influence of (a', b') on $\Phi_{a',b'}$ is 1, and by skipping (a', b') this influence is lost. Thus when Walker chooses a successor of (a', b') , the potential decreases by (at least) 1. Walker may skip several levels before the preferred option is available, decreasing the potential by 1 for each level skipped.

Since Blocker cannot play where Walker just played, the increase in potential from Blocker's move is at most $1/2$. Since the potential is 0 at the start and the end of the game, Blocker must make at least two moves for every level skipped by Walker. Walker wins a level for each move played, so Walker wins at least twice as many levels as are skipped.

In order to restrict play to $(\mathbf{k} + \mathbf{1})^2$, which is a subsubset of \mathcal{W}_{2k+1}^2 , we grant Blocker initially all moves that are outside $(\mathbf{k} + \mathbf{1})^2$. The key observation, which we formalize below, is that all of these free moves for Blocker increase the potential only by $o(k)$. Since again the potential decreases to 0 at the end, Walker loses at most $o(k)$ more levels than in the original game on \mathcal{W}_{2k+1}^2 , and hence Walker gets at least the fraction $2/3 - o(1)$ of the $2k + 1$ levels. (Here $o(g(k))$ denotes any function of k whose ratio to $g(k)$ tends to 0 as $k \rightarrow \infty$.)

To bound the initial influence of the forbidden moves, recall that $f(c, d)$ is the probability that a random walk from $(0, 0)$ to level $c + d$ ends at (c, d) . The distribution of the endpoint is the standard binomial distribution for $c + d$ trials. Let X_n be the binomial random variable counting the heads in n flips of a fair coin. The initial value of the potential function is $\sum_{n=k+1}^{2k} \Pr(|X_n - n/2| > k - n/2)$. Let $k' = \lceil \sqrt{2k \ln k} \rceil$ and $n_0 = 2k - k'$. For each n greater than n_0 , the probability is at most 1. For $n \leq n_0$, we use the well-known Chernoff bound.

The Chernoff bound states that $\Pr(|X_n - n/2| > nt) \leq 2e^{-2nt^2}$; we apply it with $t = k/n - 1/2$. Since $2nt^2$ increases as n decreases, we may assume $n = n_0$ and use $2e^{-2n_0t^2}$ as a bound on the contribution from these terms. We have $n_0 = 2k(1 - x)$, where $x = k'/2k > \sqrt{\ln k / (2k)}$. Also, $2t = \frac{1}{1-x} - 1 = \frac{x}{1-x}$. We compute

$$2n_0t^2 = \frac{n_0}{2}(2t)^2 = k(1 - x)\frac{x^2}{(1 - x)^2} > kx^2 > \frac{1}{2} \ln k.$$

Thus $2e^{-2n_0t^2} < 2k^{-1/2}$, bounding the total contribution from these terms by $2\sqrt{k}$. From the k' terms with largest n , the bound on the total is $\lceil \sqrt{2k \ln k} \rceil$. Hence the initial potential is less than $4\sqrt{k \ln k}$. As a result, Walker misses fewer than $4\sqrt{k \ln k}$ levels in addition to the $(1/3)(2k + 1)$ levels missed by the earlier argument. □

4. Seepage and the Angel-Devil game

In this section, we study two games: the Seepage game, defined by Clarke, Finbow, Fitzpatrick, Messenger, and Nowakowski [6]; and Conway’s Angel-Devil game. We relate these two games to the ordered chain game.

A *rooted digraph* is a digraph with a single vertex designated as the *root*. The *Seepage game* is played by Sludge and Green on an acyclic rooted digraph G . Sludge and Green take turns claiming vertices of G . A vertex claimed by Sludge is said to be *polluted*, while a vertex claimed by Green

is said to be *protected*. In each round, first Sludge pollutes an unprotected out-neighbor of a currently polluted vertex (on the first turn, Sludge instead pollutes the root). Next, Green protects some unpolluted vertex. On a finite digraph (the only kind we consider), the game ends when a sink becomes polluted, in which case Sludge wins, or when the active player has no legal move, in which case Green wins provided that no sink is polluted.

Acyclic rooted digraphs are closely related to *rooted posets*, which are posets having a unique minimal element called the *root*. A *k-prefix* in a rooted poset is a chain of size k consisting of elements x_0, \dots, x_{k-1} such that x_0 is the root and x_i covers x_{i-1} for $1 \leq i \leq k-1$. The top element of a prefix is its *head*; *climbing* a prefix means following it in order.

The *k-prefix game* on a rooted poset P is a variant of the ordered chain game in which Walker must climb a k -prefix to win. Winning does not require the k -prefix to be followed consecutively; there may be dead ends along the way. If Walker wins the k -prefix game on P , then Walker can build a chain of size k in the ordered chain game, since a k -prefix is a k -chain with the additional requirements of skipping no levels and starting at the bottom.

Following [6], when G is a rooted acyclic digraph we write (G, d) to refer to the subgraph of G induced by the vertices having distance at most d from the root. The following proposition establishes a straightforward correspondence between the Seepage game and the k -prefix game:

Proposition 4.1. *Let P be a rooted poset, and let G be the cover digraph of P , rooted at the minimal element. Walker wins the k -prefix game on P if and only if Sludge wins the Seepage game on (G, k) .*

Proof. It is clear that a Sludge win in G translates into a Walker win in P : translating Blocker moves directly into Green moves, Walker plays the corresponding winning strategy of Sludge, and when Sludge wins the constraints on Sludge's legal moves imply that some k -prefix has been climbed in the correct order.

Conversely, if Walker wins the k -prefix game in P , then Walker has a winning strategy that never claims a vertex whose in-neighbors are all unclaimed: such a vertex can clearly never be part of a winning k -prefix. Thus any winning strategy for Walker can be modified to eliminate such moves and then translated into a winning strategy for Sludge. \square

In light of this result, we can translate Theorem 4.6 of [6] into the following theorem about k -prefix games:

Theorem 4.2. *If $k \geq 10$, then Blocker wins the k -prefix game on \mathcal{W}_k^3 .*

Turning from the Seepage game, we now define a slightly more general version of Conway's Angel-Devil game [7]. The *Angel-Devil game* is played on an infinite rooted digraph G by two players, Angel and Devil. Angel and Devil alternate turns. In each round, Angel moves to an available out-neighbor of the current position (except that Angel's first move is to occupy the root vertex), and then Devil *burns* at most one vertex, preventing Angel from using it in future moves.

Devil wins if Angel is ever unable to move, meaning that all out-neighbors of the current position are burned. Angel avoids losing by having a strategy to continue moving forever. When every vertex of G has finite outdegree, an equivalent condition for Angel to win is that for every natural number n , Angel has a strategy to guarantee moving for n turns.

The classical Angel-Devil game is played on the points of an infinite grid, with Angel starting at the origin and also moving first. We model this variant by translating the root to the vertex Angel reaches before Devil's first move.

Remark 4.3. As noted by Conway, one may grant that Devil burns Angel's current position and its out-neighbors not chosen for the next move. If there is a way for some winning strategy to return to the current vertex or its out-neighborhood, then Angel does at least as well by making the subsequent move to an out-neighbor now. Similarly, Devil never *needs* to burn the vertex Angel is leaving.

The k -prefix game on a poset can be translated directly into the Angel-Devil game on an appropriate digraph, with Angel required to survive for k turns. The difficulty is showing that Angel does not actually need Walker's ability to backtrack. Theorem 2.4 of [6] allows us to avoid explicitly presenting this argument.

Proposition 4.4. *Let P be a rooted poset with minimal element x_0 , and assume that every element of P is covered by finitely many other elements. Let G be the rooted digraph with root x_0 that is the cover digraph of P . Devil wins the Angel-Devil game on G if and only if Blocker wins the k -prefix game on P for all k .*

Proof. Clearly a winning Angel strategy on G can be translated directly into a winning Walker strategy on P .

For the other direction, let any k be given, and observe that (G, k) is an acyclic digraph with a single source vertex. Theorem 2.4 of [6] states that if Sludge wins the Seepage game on (G, k) , then Sludge has a winning strategy that only pollutes the vertices of a path. This winning strategy can be directly translated into a strategy for Angel to survive k turns. \square

In the next section, we will apply the Angel-Devil game to the k -prefix game on \mathcal{W}_k^d .

5. Walker vs. Blocker on high-dimensional wedges

It is difficult to devise winning strategies in Angel-Devil games. To benefit from the few explicit strategies that are known, we seek ways to transfer these strategies between games. We define a type of map from one rooted digraph to another that facilitates such a transfer.

Definition 5.1. Let G and H be digraphs with roots g_0 and h_0 . A *robust map* from G to H is a map $\phi: V(G) \rightarrow V(H)$ with $\phi(g_0) = h_0$ such that whenever there is an edge from $\phi(v)$ to w in H , there is also some vertex z in G such that $\phi(z) = w$ and $vz \in E(G)$.

Informally, a map is robust if, whenever the image $\phi(P)$ in H of a path P in G can be extended, P can also be extended to P' in G so that $\phi(P')$ is the extended path in H .

Theorem 5.2. *Let G and H be two rooted digraphs, and let $\phi: G \rightarrow H$ be a robust map from G to H . If Angel wins the Angel-Devil game in H , then Angel also wins in G .*

Proof. Given a winning strategy for Angel in H , we define a winning strategy in G . We play an imaginary game in H to track and simulate the actual game in G . The G -Angel will maintain a position in G that maps under ϕ to the current position of the H -Angel in H . This holds initially, since they both start at the root.

At some time later, let v be the location of the G -Angel, so the imagined H -Angel is at $\phi(v)$. The G -Devil moves by burning some vertex y in G . The imagined H -Devil burns the corresponding vertex $\phi(y)$. The imagined H -Angel has a winning response w for this move.

Since w must be an out-neighbor of the current vertex $\phi(v)$ in H , the robustness of ϕ guarantees a vertex z in G such that $vz \in E(G)$ and $\phi(z) = w$. The vertex z cannot previously have been burned by the G -Devil, because the imagined H -Devil would have immediately burned w in H to copy that move. Since z is available and $vz \in E(G)$, the G -Angel can move to z . This preserves the property that the H -Angel is on the image of the position of the G -Angel, and the game continues. Since the H -Angel can move forever, the G -Angel also can move forever. \square

Abusing notation, we will write \mathcal{W}^d to denote both the infinite wedge poset in d dimensions and its cover digraph, rooted at the minimal element.

Thus the vertices of \mathcal{W}^d are nonnegative integer d -tuples, with $xy \in E(\mathcal{W}^d)$ if y is obtained from x by increasing one coordinate by 1.

We compare the wedge with the “power-2” Angel-Devil game on \mathbb{Z}^2 . An Angel of power k can move to any unburned square that is at most k units away in each horizontal or vertical direction. Thus in the digraph for the power-1 game each vertex has outdegree 8, while in the power-2 game the vertices have outdegree 24. It is known that Devil wins the power-1 game (see Conway [7]), while Angel wins the power-2 game (proved independently by Kloster [9] and Máthé [11]).

To prove our result, we give a robust map to the digraph for the power-2 Angel on \mathbb{Z}^2 . Since Angel wins that game, Theorem 5.2 implies that Angel wins on \mathcal{W}^d when $d \geq 24$. Wästlund’s refinement of the argument for the power-2 Angel gives a winning Angel with only 14 moves [12], lowering our bound to $d \geq 14$. By Proposition 4.4, Walker then wins the k -prefix game on \mathcal{W}_k^d for every finite k .

The construction of our robust map uses the following intuitive idea: if Angel has d different “types of move” in some digraph, and these moves commute, then we can introduce a (highly redundant) coordinate system on the graph by counting how many times Angel has made each type of move. This coordinate system induces a robust map from \mathcal{W}^d into the digraph. The following lemma formalizes the idea.

Lemma 5.3. *If H is a rooted digraph with $V(H) \subset \mathbb{Z}^n$, and $M \subset \mathbb{Z}^n$ is a finite set such that $xy \in E(H)$ implies $y - x \in M$, then there is a robust map from $\mathcal{W}^{|M|}$ to H .*

Proof. Let $d = |M|$, and let m_1, \dots, m_d be the elements of M . Define $\phi: \mathcal{W}^d \rightarrow V(H)$ by $\phi(x_1, \dots, x_d) = h_0 + \sum_{i=1}^d x_i m_i$, where h_0 is the root of H . Since $\phi(0, \dots, 0) = h_0$, the start condition is satisfied. Now consider $(\phi(v), w) \in E(H)$. We must find $z \in V(\mathcal{W}^d)$ such that $\phi(z) = w$ and $vz \in E(G)$. Since $(\phi(x), v) \in E(H)$, the hypothesis guarantees existence of $m_i \in M$ such that $\phi(x) + m_i = v$. With e_i denoting the unit vector with 1 in coordinate i , we have $\phi(x + e_i) = \phi(x) + m_i = v$, and $(x, x + e_i) \in E(\mathcal{W}^d)$. Hence ϕ is robust. \square

The underlying digraphs of the classical Angel-Devil game fit the hypothesis for H in the lemma, yielding the following corollary:

Corollary 5.4. *For $d \geq 14$, Angel wins the Angel-Devil game on \mathcal{W}^d (and hence also Walker wins the k -prefix game on \mathcal{W}^d for all k).*

Proof. Since Angel wins the power-2 Angel-Devil game [9, 11], in which Angel always has 24 possible moves expressed as coordinate vectors, Lemma 5.3

and Theorem 5.2 together imply that Angel wins in \mathcal{W}^{24} . Furthermore, Angel can win that power-2 Angel-Devil game using only moves changing the horizontal coordinate by at most 2 and the vertical coordinate by at most 1 (proved by Wästlund [12]); hence Angel wins in \mathcal{W}^{14} (and thus in all higher-dimensional wedges). \square

In Section 3, we showed that Blocker wins the ordered chain game on the wedge \mathcal{W}^2 , and the result of [6] quoted in Theorem 4.2 implies that Blocker also wins on \mathcal{W}^3 . The question remains: Who wins when $4 \leq d \leq 13$? We conjecture the following.

Conjecture 5.5. *For $d \geq 4$ and for all k , Walker wins the k -prefix game on \mathcal{W}^d .*

Since Walker wins the ordered chain game on wedges except in very small dimensions, it seems fair to give Blocker a chance by introducing bias; in the b -biased game, Blocker occupies b elements after each move by Walker. Introducing bias is a typical approach to studying Maker-Breaker games won by Maker; how much bias is needed to enable Breaker to win? Here we have another parameter, the dimension of the wedge, so we ask the question a bit differently. The effect is to generalize the question asked in Conjecture 5.5.

Definition 5.6. For $b \in \mathbb{N}$, let $f(b)$ be the least d such that for all k , Walker wins the k -prefix game with bias b on \mathcal{W}^d .

Clearly $f(b) > b$, since in b dimensions Blocker can immediately occupy all elements covering the first move by Walker. In fact, the function is strictly increasing.

Proposition 5.7. *If Blocker wins the b -biased k -prefix game on \mathcal{W}^d , then Blocker wins the $(b + 1)$ -biased k -prefix game on \mathcal{W}^{d+1} .*

Proof. On each turn Blocker plays the out-neighbor of Walker's previous move with coordinate $d + 1$ increased and uses the remaining b moves to play the strategy for d dimensions. Any move by Walker that increases the last coordinate immediately loses a level, so Walker does best by playing the game in d dimensions. \square

To obtain upper bounds, we discuss the effect of bias in the Angel-Devil game. Bias b allows Devil to burn b vertices on each turn.

Proposition 5.8. *For $b \in \mathbb{N}$, Angel wins the power- $2b$ Angel-Devil game with bias b on \mathbb{Z}^2 . Moreover, Angel can do so while changing the horizontal coordinate by at most $2b$ and the vertical coordinate by at most b on each turn.*

Proof. Angel uses a winning strategy for the power-2 game (with bias 1) on \mathbb{Z}^2 . On each turn in the true game, Devil makes b moves. Angel imagines these moves in some order and responds, in the imagined power-2 game, to each move in turn. This leaves Angel at some vertex v ; Angel moves directly to v in the true game. After each turn in the true game, Angel sits on the same vertex in both games, and the same vertices have been burned. A technical requirement is that after b rounds of the imagined game, Angel has not returned to the original vertex. This follows from Remark 4.3, where a winning Angel can burn visited vertices and their unused out-neighbors without cost. Now, since Angel wins the imagined game, Angel also wins the true game.

As proved by Wästlund [12] and used in Corollary 5.4, Angel can win the power-2 game by changing the horizontal coordinate by at most 2 and the vertical coordinate by at most 1 on each turn. Thus, in the true game there is a winning strategy that changes the horizontal coordinate by at most $2b$ and the vertical coordinate by at most b on each turn. \square

Theorem 4.2 shows that $f(1) > 3$, and thus Conjecture 5.5 is the statement that $f(1) = 4$. Propositions 5.7 and 5.8 yield general lower and upper bounds on $f(b)$.

Proposition 5.9. $f(b) > b + 2$ for $b \geq 1$, and $f(b) \leq (4b + 1)(2b + 1) - 15$ for $b \geq 2$.

Proof. From Proposition 5.8, the same arguments used to prove Corollary 5.4 yield $f(b) \leq (4b + 1)(2b + 1) - 1$. However, Remark 4.3 implies that within the $(4b + 1)$ -by- $(2b + 1)$ rectangle around the current position, Angel can guarantee a winning strategy without using the central 5-by-3 rectangle when $b > 1$. When $b > 2$ it is plausible that a larger internal region could be thrown away, yielding a tighter bound, but the calculations become more complex (we cannot simply throw away all spaces reachable in 2 moves, for example).

Now consider lower bounds. Theorem 4.2 states that $f(1) > 3$, and by Proposition 5.7 it follows that $f(b) > b + 2$ for all $b \geq 1$. However, the resulting Blocker strategy is complex, using the Green strategy of [6] as its “base case”. We give a simple Blocker strategy when $b = 3$, resulting in simpler strategies for higher b .

Walker starts at the root, the all-0 vector. Blocker wants to limit Walker’s options by playing moves that remain above the head of the chain no matter where Walker moves. We write elements as strings instead of lists.

Proposition 5.7 and Theorem 3.1 together yield $f(2) > 3$. Following those arguments, Blocker starts with $\{001, 110\}$ in response to 000. By symmetry, Walker plays 100, and Blocker wins with $\{101, 200\}$.

Now consider bias 3 in \mathcal{W}^5 . Blocker forces Walker to increase a coordinate that is 0 on each subsequent move to avoid skipping a level, while seizing all elements on one level that avoid coordinate values larger than 1. Below are the moves by Walker and Blocker, up to symmetry. Walker increases a coordinate at the second move that bypasses one of Blocker's first moves, because increasing one of the last two coordinates creates a situation that is isomorphic to this but with the loss of an extra position ahead. After Blocker's second and third moves, the coordinates where Walker can increase are isomorphic.

Walker	Blocker		
00000	01111	10111	11011
10000	20000	11101	11110
11000	21000	12000	–
11100	21100	12100	11200

With the three moves of his first turn and two moves of his second, Blocker occupies all points at the fifth level with largest coordinate 1. Blocker's other moves prevent Walker from moving to an element having a coordinate larger than 1 while not skipping a level. No elements with weight 4 and largest coordinate 1 are available, so Walker loses then. The proof of Proposition 5.7 shows how to extend this strategy to higher b . \square

In fact, in each dimension where Blocker wins the b -biased game, one can ask what fraction of the levels will be obtained by Walker under optimal play. By Theorem 3.1, the fraction is $2/3$ when $d = 2$ and $b = 1$. The arguments in Proposition 5.9 show that the fraction also is at most $2/3$ when $d = 3$ and $b = 2$, and it is at most $4/5$ when $d = 5$ and $b = 3$.

References

- [1] J. Beck (1996). Foundations of positional games. *Random Struct. Alg.* **9** (1–2) 15–47. [MR1611724](#)
- [2] J. Beck (2005). Positional games. *Combin. Probab. Comput.* **14** 649–696. [MR2174650](#)
- [3] J. Beck (2008). *Combinatorial Games: Tic-Tac-Toe Theory*, Cambridge University Press. [MR2402857](#)

- [4] M. Bednarska, T. Łuczak (2000). Biased positional games for which random strategies are nearly optimal. *Combinatorica* **20** 477–488. [MR1804821](#)
- [5] V. Chvátal, P. Erdős (1978). Biased positional games. *Ann. Discrete Math.* **2** 221–229. [MR0500701](#)
- [6] N. E. Clarke, S. Finbow, S. L. Fitzpatrick, M. E. Messenger, R. L. Nowakowski (2009). Seepage in directed acyclic graphs. *Australas. J. Combin* **43** 91–102. [MR2489411](#)
- [7] J. H. Conway (1996). The angel problem. In: *Games of No Chance*, R. Nowakowski ed., MSRI Publications, Vol. 29, pp. 3–12. [MR1427954](#)
- [8] D. Hefetz, M. Krivelevich, M. Stojaković, T. Szabó (2009). Fast winning strategies in Maker-Breaker games. *J. Combin. Theory B* **99** 39–47. [MR2467817](#)
- [9] O. Kloster (2007). A solution to the angel problem. *Theoret. Comput. Sci.* **389**(1–2) 152–161. [MR2363369](#)
- [10] M. Krivelevich (2011). The critical bias for the Hamiltonicity game is $(1 + o(1))n / \ln n$. *J. Amer. Math. Soc.* **24**(1) 125–131. [MR2726601](#)
- [11] A. Máthé (2007). The angel of power 2 wins. *Combin. Probab. Comput.* **16**(3) 363–374. [MR2312432](#)
- [12] J. Wästlund. *A Weaker Winning Angel*. <http://www.math.chalmers.se/~wastlund/weakerAngel.pdf>

DANIEL W. CRANSTON
MATHEMATICS DEPT.
VIRGINIA COMMONWEALTH UNIVERSITY
USA
E-mail address: dcranston@vcu.edu

WILLIAM B. KINNERSLEY
MATHEMATICS DEPT.
RYERSON UNIVERSITY
CANADA
E-mail address: wkidders@gmail.com

KEVIN G. MILANS
MATHEMATICS DEPT.
WEST VIRGINIA UNIVERSITY
USA
E-mail address: milans@math.wvu.edu

GREGORY J. PULEO
MATHEMATICS DEPT.
UNIVERSITY OF ILLINOIS
USA

E-mail address: puleo@illinois.edu

DOUGLAS B. WEST
MATHEMATICS DEPT.
ZHEJIANG NORMAL UNIVERSITY
CHINA

MATHEMATICS DEPT.
UNIVERSITY OF ILLINOIS
USA

E-mail address: west@math.uiuc.edu

RECEIVED DECEMBER 9, 2011