

CHAPTER 1

The Angel Problem

1. Angels, Kings, and Fools

Two players, the *angel* and the *devil*, play a game on an infinite chess board whose squares be indexed by pairs of integers. The angel is an actual “person” moving across the board like some chess piece, while his opponent does not live on the board but only manipulates it. In each move, the devil blocks an arbitrary square of the board such that this location may no longer be stepped upon by the angel. The angel in turn, flies in each move from his current position $(x, y) \in \mathbb{Z}^2$ to some unblocked square at distance at most k for some fixed integer k , i.e., to some position $(x', y') \neq (x, y)$ with $|x' - x|, |y' - y| \leq k$. Note that devil moves are not restricted to the angel’s proximity or limited by any other distance bounds; he can pick squares at completely arbitrary locations.

The devil wins if he can stop the angel, that is, if he manages to get him in a position with all squares in the $(2k + 1) \times (2k + 1)$ area around him blocked. The angel wins simply if he succeeds to fly on forever. The open question is, whether for some sufficiently large integer k the angel with distance bound k , called the *k-angel*, can win this game.

First variants of this game were discussed by Martin Gardner [17], who names D. Silverman and R. Epstein as original inventors. Though his article deals mainly with finite configurations, i.e., the question whether a chess king (which is simply a 1-angel) can reach the boundary of a given rectangular board, he also asks for a strategy against a chess knight on an infinite board, possibly with a devil who gets to play more than just one block per move. In its present form the angel game first appeared in Berlekamp, Conway, and Guy’s classic [8] (Chapter 19). Amongst detailed analyzes of games with kings and other chess pieces on finite boards against devils with certain additional restrictions, the authors coin the names “angel” and “devil” for the two competitors and give a thorough proof that the chess king can be caught on a 33 by 33 board. Then Conway [11] focused entirely on the infinite angel game, trying to explain possible pitfalls with certain natural escape attempts and pointing out the hardness of the problem. Besides all variants, the central open question remains whether some angel of sufficient power can escape forever. In his overview article [14], Demaine cites it as a difficult unsolved problem of combinatorial game theory.

In this present work, we present modest advances on the current best known devil strategy. Therefore we introduce a slight reformulation of the original game, which allows us to focus on speed as the important parameter. In a further part, we treat a higher-dimensional analog of the angel game, showing that an angel of sufficiently large power can escape in 3D.

Catching the chess king. The only case for which the k -angel problem is solved is $k = 1$, the ordinary chess king. We like to sketch a winning strategy for the devil, which is motivated by the analysis in [8]. This shall get us some feel for the game and make us familiar with some basic principles that will turn up every now and then. The basic ideas are quite simple. Maybe the reader likes to stop reading here for a while and enjoy figuring out such a strategy on his own.

Assume the devil wants to prevent the king from crossing a certain horizontal line. With three squares above the king already blocked on that line, like in Figure 1, this is easily achieved. The devil simply answers a king's move a to the right by an extension of that triple block by a play at u . A further move to b is countered by v and likewise, any left movement to a' is blocked at u' . Pushing along in this simple fashion ensures that wherever the king goes, the three squares above him will always be blocked, making a crossing impossible.




FIGURE 1. Pushing the chess king along a line.

It is not difficult to get the three initial blocks placed on a blank line when a king is just approaching. In the left drawing of Figure 2, the king is only five steps away from the desired line along the upper rim, where the devil has just played his first block. We claim that however the king now approaches that line, the devil will always manage to get his triple block in place.

If the king makes one step forward to a , the devil replies at u . After that the moves b' and b'' both lead directly to a triple block by the devil answers v' and v'' , respectively. So we only have to consider a king's move to b . The devil plays at v , after which the king's moves c' and c'' are both blocked by v' . This leaves only a step to c , which is countered by w . Now the king's right-most option d can be blocked at x and the moves to d' or d'' again lead to a triple by v' ,

The second option for the king's first move is a in the right drawing of Figure 2. (A move to a' being symmetric to this case.) The devil plays at u . Against a step to b' , the devil immediately forms a triple block by playing v' . The two moves to b and b'' lead to symmetric configurations, so we need only consider the remaining option b . The devil replies v , after which c is countered by w , and c' and c'' can both be blocked at v' .

So five preparation steps suffice for the devil to get his triple block in place against an approaching king. Figure 3 shows in a slightly non-proportional drawing how to turn the above wall-pushing argument into a successful devil strategy for catching the chess king. With his first 44 moves, the devil blocks some squares in the four corners of an imaginary box around




FIGURE 2. Getting the triple block in place.

the king. The box must be chosen large enough to ensure that during this preparation phase the king does not get too close to the boundary of that box. After that, the devil plays the above wall-pushing strategy along the dotted lines whenever the king approaches such a line. The four corners are there to ensure that the devil can never be forced to play on two fronts at the same time.




FIGURE 3. Catching the chess king.

We leave the argument at this informal state, hoping that the reader has grasped the idea. We are headed for a stronger result, which we shall then prove in full detail. As we already said, a deeper analysis of the chess king, very similar to the above discussion, can be found in Chapter 19 of [8].

The fool argument. The first general idea for an escape with a k -angel might be to run away in one direction. If the power k is large enough, shouldn't the angel somehow be able to go faster than the devil putting any serious obstacles in his way? Maybe the angel can simply run away in one direction. The answer is *no*! Conway defines a k -fool to be a k -angel who commits himself to strictly increasing his y -coordinate in every move. He shows that a fool of any power k can be caught [11]. The argument is simple, so we take the time to recapitulate it here.

The restriction on the y -movement implies that from a fixed position the k -fool can reach only squares within a cone of slope n . The devil sets out to build a long barrier at a far distance across this cone so that the fool, once he gets there, will stand in front of an impenetrable wall. Note that in order to block a k -fool effectively, we need a thick wall of k consecutive lines. We clearly cannot build such a solid wall across the whole width of the cone because already a 2-fool would arrive at the construction site much

earlier than the devil could finish his work. Conway's trick is the following dynamic refinement strategy.

Say, our desired barrier shall be h units to the north. There the cone of possible future fool positions has width $2hk + 1$, so that a complete wall of thickness k at that distance would consist of about $2hk^2$ squares. The devil begins filling this wall partly. With his first $h/(2k)$ moves, while the angel gets half way to the distant line, he blocks about 1 out of $4k^3$ squares there, distributing his moves evenly over the full width. Once the fool reaches the center line, the devil determines the new cone of potential fool positions, which by simple geometry, covers only half of the original wall. The devil then spreads his next $h/(4k)$ moves evenly on that segment of the construction site that can still be reached by the fool. See Figure 4.




FIGURE 4. Catching a fool.

He obviously gets the same proportion of about 1 out of $4k^3$ squares blocked there until the fool has reduced his distance to the wall to $h/4$. If the initial distance h was chosen large enough, we can iterate this process often enough (about $4k^3$ times) to finish the relevant part of our barrier before the fool arrives.

This argument generalizes to non-strict fools, i.e., angels who are not allowed to make a step in negative y -direction, and it is not limited to one direction. There is also a radial variant where the angel never decreases his distance to the origin. The detailed arguments are given in [11].

Conway's fool counter already indicates that devising an escape strategy for some angel might be a very difficult task. By a simple dove-tailing argument this result can even be turned into the following surprising fact [11].

1. THEOREM ("Blass-Conway diverting strategy"). *There is a strategy for the devil with the following property. For each point p of the plane and each distance d , no matter how the angel moves there will be two times $t_1 < t_2$ such that at time t_2 the angel will be d units nearer to p than at time t_1 .*

This diverting strategy does not imply, however, that the angel must run in a wild zig-zag across the board. Concrete bounds on t_1 and t_2 are astronomical, so that the angel has plenty of time to comply with those requirements. But Theorem 1 can be used to immediately disqualify a variety of ad-hoc angel strategies, like refinements of the fool approach, that do not allow for sufficient freedom of movement in all directions. After all, Conway himself believes that some angel can escape. He awards \$100 for an escape strategy for an angel of some sufficiently high power k .

2. From Finite to Infinite Games

Before we go on to devise strategies for angel and devil, let us pause a while to discuss some fundamental aspects of infinite games in general. Such games may behave a little weird: It may be that neither player can force a win, i.e., there exist no winning strategies, even though the game does not allow for draws.

Formally, an infinite *game* is simply a subset \mathcal{A} of $\mathbb{N}^{\mathbb{N}}$. A *play* is an infinite sequence $\tau = (x_0, y_0, x_1, y_1, \dots)$ of natural numbers where Players 1 and 2 choose the x_i and y_i , respectively, in turns, and Player 1 wins iff $\tau \in \mathcal{A}$. A *strategy* is a mapping from all possible finite initial segments of a play to the next move, i.e., a mapping from the set \mathbb{N}^* of finite words to \mathbb{N} and it is a *winning strategy* if it wins against all possible opponent plays.

It is well-known that the axiom of choice allows the construction of games in which neither player has a winning strategy [23, Sec. 43]; but Martin [32] proved that for games that are Borel sets this cannot happen, such games are *determined*: one of the two players must have a winning strategy. This result covers essentially all games that can be defined in simple ways. Any “reasonable” game will be determined. And so is the angel-devil game. However, we do not need the full power of Martin’s deep theorem. The following Lemma is an easy adaptation of an earlier, simpler result of Gale and Stewart [16]. I want to thank Stefan Geschke for introducing me to these set-theoretic foundations of infinite games and for helpful discussions about the arguments in this section.

2. LEMMA. *The angel-devil game is determined. That is, either the angel or the devil has a winning strategy.*

PROOF. Assume the devil has no winning strategy. The angel can play as follows. In each turn he makes a move after which the devil does not have a winning strategy. By induction, such a move must always exist since otherwise the devil would have a winning strategy. The resulting angel strategy is obviously a winning strategy, simply because it allows the angel to play forever. \square

Of course, one could define the above strategy for any given infinite game. The decisive point is that usually such a strategy does not automatically yield a win as is the case with the angel-devil game.

A further observation, which is useful when thinking about our game, is that in a sense it is infinite only from the point of the angel. If the devil wins, the game ends, by definition, after finitely many moves. So it seems that if the devil can win at all against the k -angel, there should exist some constant N_k such that the devil can catch the k -angel in at most that many moves. Equivalently, if some angel should be able to survive M moves, for any arbitrarily large number M that is fixed at the beginning of the game, then he should also be able to escape forever.

These seemingly obvious implications bear a subtlety. It could in principle be possible that the angel would have to choose his strategy dependent on the given M , so that he can in deed escape for M moves as required but will be caught a little later. If he had wanted to survive longer he might have had to choose a different strategy. Ultimately, there might not exist

a strategy that works for all M at the same time. Seen from the devil's perspective this would mean that while he is sure to catch the angel after a finite number of moves, there might not be a universal bound on the time that is required to catch him. Fortunately, our concerns are needless.

3. LEMMA. *If the devil has a winning strategy against some angel then there exists a bound N such that the devil can stop that angel in at most N moves. Conversely, if the angel can survive for any arbitrarily large, previously given number of steps then he can escape forever.*

PROOF. Assume that the devil has a winning strategy. Consider the game tree of all possible plays under such a devil strategy σ . It has a bounded number of options at each angel node (no more than $(2k + 1)^2$) and just one option at each devil node, namely the one prescribed by σ . The leaves are exactly those positions in which the angel cannot move anymore and thus has lost. This tree contains no infinite paths because such a path would directly give the angel an infinite sequence of moves, in contradiction to our assumption that σ is a winning strategy.

Since the degree of the tree is bounded and it contains no infinite paths, it is finite by König's lemma and therefore has finite depth, N , say. This means that the strategy σ allows no more than N moves before the angel is stuck, independent of how the angel plays.

The second statement is equivalent to the first. If the angel can escape as long as required by the beginning of the game, the devil cannot have a strategy that catches him after a fixed number of moves. Hence, the devil has no winning strategy at all, which means by Lemma 2 that the angel can escape forever. \square

3. The Need for Speed

There is pretty little known about even very weak angels. Already the destiny of the 2-angel is not settled and even more, it is unknown whether a chess knight, i.e., a piece that jumps in each move to one of the eight squares at Euclidean distance exactly $\sqrt{5}$, can be caught.

We do not have a solution for the 2-angel, either, but we make a first step in this direction by devising devil strategies against opponents whose power lies somewhere between that of a 1-angel and the strength of 2-angel. The improvement is rather modest but the new concepts we need to introduce in order to obtain them or even state them, reveal details of the game that seem to lie hidden with Conway's original angels.

Let us take a closer look at what happens when we upgrade the original chess king to a 2-angel. This is already a large step. The improvement is actually two-fold. Not only does the 2-angel move at twice the speed, any barriers must also be twice as thick to hold him back. In a sense, the 2-angel can be said to be 4 times stronger than the 1-angel. We focus on the first aspect: *speed!* We would like to suppress the ability to jump over obstacles as an undesired side effect. Define a k -king as a player who in each turn makes exactly k ordinary king's moves, while the devil still gets to place one block per turn. The point is that now every single king's move must be valid, the k -king cannot fly across obstacles.

If we want to use kings for the study of the angel problem, they should, in some qualitative sense at least, be equivalent to angels. Obviously, a k -angel is stronger than a k -king. An escape strategy for a king can be used for an angel of the same power as well. The converse is, of course, not true—not for trivial reasons at least—but we can show that if you can catch kings of arbitrary power k then you can also catch any angel. Before we come to this reduction, let us first remark on a subtlety in the above argument.

A k -king could in principle use a sequence of k steps to run a circle and return to his starting position, thereby simulating a pass between two consecutive devil moves. An angel is formally not allowed to pass. So our trivial transformation from above had a little flaw. The following basic no-return lemma by Conway [11] repairs this defect. It works for k -kings as well as for k -angels and will be needed once more later on.

4. LEMMA. *If the k -angel or k -king can escape then he can also escape without ever visiting any square twice; where in the case of the king we only consider the last step of a sequence of k steps between two devil moves.*

The restriction to the last step in a sequence of k king steps is natural because that final location is always the one that the devil sees when it's his turn. For the intermediate positions the following argument would not work.

PROOF OF LEMMA 4. We assume that we have a winning devil strategy σ against a non-returning k -angel or k -king and derive from that a winning strategy against the non-restricted versions. The idea is simple. When the angel/king revisits a location, the situation is always worse than at his first visit. The set of blocks has only grown. We turn this observation into a formal proof.

The devil plays according to σ until the angel/king lands on a square p he has already visited before. In this case the devil blocks an arbitrary square from the $(2k + 1) \times (2k + 1)$ area around p . Now he simply forgets all moves since his opponent first visited p and resumes the strategy σ from that position. The point is that his reply to the angel's/king's move to p has already been played when he had answered to p the first time, so that his intermediate move was really for free and he did not fall behind with σ .

We must be precise about what we mean by “forget.” The intermediate moves, since the first play of the revisited square p , are really erased from the devil's memory. So when the angel steps on a square he had been before but the devil has forgotten about that move, he plays on without backtracking. Otherwise we would have to show how forward jumps in σ , that is, jumps to a location that has been visited in the forgotten future, should be treated consistently.

The result of the described devil play is, of course, that the angel/king cannot return more than $(2k + 1)^2$ times to the same location because then he would not be able to leave it again due to its barred environment. Consequently, the derived strategy wins just as σ does. \square

Lemma 4 also shows that a little inaccuracy in our definition of the game is inconsequential. We have not said explicitly whether the devil should be allowed to block the square on which the angel currently sits. Since we may

assume that the angel may never return to that position anyway, the devil will never need to make such a move.

We now establish the announced equivalence of angels and kings. Of course, the reduction from angels to kings requires an increase in speed.

5. PROPOSITION. *If the k -angel can escape then so can the $99k^2$ -king.*

PROOF. We derive an escape strategy for the $99k^2$ -king from an escape strategy for the k -angel. While the king plays against the “real” devil, we set up an additional, imaginary board with an imaginary k -angel, where we simulate the action on the king’s board through appropriate transformations.

The king’s board is partitioned into a regular grid of sidelength- $18k^2$ boxes. Likewise, the angel’s board is segmented into blocks of sidelength k . The boxes of the two worlds are in one-to-one correspondence with each other, in the obvious fashion: the box containing the king’s starting point corresponds to the angel’s initial box and further all adjacencies are preserved. These partitions and the correspondences are fixed once and for all at the beginning of the game.

We play as follows. When the devil blocks some square in the king’s world, we cross out an arbitrary empty square from the corresponding box in the angel’s world or from one of the eight adjacent boxes there.

When it’s the king’s turn, we use our escape strategy for the angel to get a move in the imaginary world. This move is then translated into the king’s plane by a movement of the king into the corresponding box there. For example, in Figure 5 the angel jumps from his current box into the next box to the north; then the king runs into the northern box in his world, too. The precise position within that box is completely independent of the angel’s position in his box, however. It will depend on the following technical details.




FIGURE 5. Simulating a king by an angel.

We have to describe precisely how the king should run and also prove that the required movement will always be possible. Observe that when the angel can jump into some box, the devil cannot have blocked all k^2 of its squares. From our simulation rule for devil moves, we conclude that the corresponding box in the king’s world and also the eight surrounding boxes there contain less than $9k^2$ devil blocks each.

Knowing that a target box in the king’s world contains less than $9k^2$ blocks, we can now find a route for the king. We introduce an invariant for

king positions: the king only stops at locations from where the four lines into the four axis parallel directions within the current box are completely free. We maintain this invariant to always ensure a free passage for the king into his target box. Assume the king needs to go one box to the east. The density bound of $9k^2$ guarantees that in both orientations, vertical and horizontal, strictly more than half of the $18k^2$ lines are completely empty, in each box. This implies that there are at least two free long horizontal lines through both boxes. It is easy to see that from his good position with free roads in at least three directions (the fourth direction possibly blocked by the last devil move) the king can reach such a line in less than $9k^2 + 18k^2 = 27k^2$ steps, as shown in Figure 6. (The additional $9k^2$ moves result from a possible detour, which won't take longer, because there are at most $9k^2$ blocks in the whole square.) Then it takes no more than another $36k^2$ steps to reach a good position in the target box. If the king is headed for one of the four diagonally connected boxes, he first makes a stop-over in the horizontal or vertical direction and proceeds from there with at most another $36k^2$ steps. This gives a total of less than $99k^2$ steps.




FIGURE 6. The king runs into a neighboring box.

Again a remark on passing. Since the king is forced to use up all his 99 moves in each turn, he might in principle get in troubles when he arrives at his destination too early. However, zugzwang is not really an issue here because the king finds enough empty squares along the side of his road to waste arbitrary numbers of moves by running in little circles. \square

We emphasize again that the quantitative proportion of the above reduction is not our main concern. The purpose of Proposition 5 is only to establish the qualitative equivalence between angels and kings, as a legitimation to use kings as a tool to attack the angel problem.

Preparing fences. Let us have a closer look at the devil strategy against the 1-king from the beginning. It seems we wasted some potential there. After the preparation of the corners, the devil simply sits and waits for the king to arrive at one of the four sides. Couldn't he perhaps use this time for some further preparations so that he can catch a faster king, a 2-king, maybe.

The basic idea for the king counter was our dynamic-wall argument, where we had the king pushing along a line without ever letting him break through. Can we extend this method to the 2-king? Since the 2-king makes two steps for each devil move, it would suffice to have every second square along the desired frontier already in place. Starting from the initial position

in Figure 7 with only two additional squares blocked, the devil can push along with the 2-king by answering the double move a_1, a_2 at u , then b_1, b_2 at v , and so on.




FIGURE 7. A wall against the 2-king.

How long would it take the devil to prepare such a density-1/2 wall against the 2-king? Since he needs to block 1 square out of 2, he can set up such a wall at an absolute speed of 2, which is exactly the speed of the 2-king. In other words, the devil can build such fences against the 2-king at the same speed the 2-king runs. For example, a 2-king who sits at the bottom of a square box of sidelength R with solid walls to the east, west, and south but completely open to the north is lost. We just learnt that the devil can build a fence of density 1/2 across that open gate in the north, in just the time it takes the 2-king to get there. Hence, the 2-king can never leave that box.

Can we extend these ideas to encircle the 2-king completely? The answer is *yes*—almost. We shall present successful devil strategies against any king of speed $2 - \varepsilon$ for any fixed real $\varepsilon > 0$. First, of course, we have to say what such a statement shall mean. We need a definition of fractional, or even irrational speed.

Real kings. What is a 3/2-king? On average he should get to make three king's steps for 2 devil steps, which we could realize by a move sequence like $KKKDDKKKDD\dots$, which shall mean that the king makes 3 steps, then the devil blocks 2 squares, and so on. However, such a concept would depend on the actual representation of a rational number. The 6/4-king would get a different sequence. We could get around this by demanding reduced fractions but then a 1001/8-king would behave completely different from a 1000/8-king, who should simply be the 125-king. What's worse, the grouping of devil moves can be lethal for the king. For example, the eight consecutive devil moves in the sequence $K^{1001}D^8K^{1001}D^8\dots$ could be used to encircle the king completely, even though his average speed would be greater than 125.

What we want are move sequences that approximate a given speed $\alpha \in \mathbb{Q}^+$ as fair as possible, avoiding unnecessarily large chunks of moves for either side. The sequence $(u_n)_{n \in \mathbb{N}}$ defined by [2]

$$(1) \quad u_n = \lfloor (n+1)\gamma + \phi \rfloor - \lfloor n\gamma + \phi \rfloor \in \{0, 1\} \quad \text{with } \gamma = \frac{\alpha}{\alpha+1} \in (0, 1)$$

and some constant offset $\phi \in \mathbb{Q}$ shows this behavior—if we interpret 1's in the sequence as the king's and 0's as the devil's moves.

This sequence (u_n) is easy to understand; it simply compares consecutive elements of the arithmetic progression $(n\gamma + \phi)$. Whenever there lies an

integer between the n th and the $(n + 1)$ st element of $(n\gamma + \phi)$, we have $u_n = 1$, otherwise, when the two elements fall in a common integer gap, (1) evaluates to $u_n = 0$. We conclude that the frequency of 1's in (u_n) is γ , hence the frequency of 0's is $1 - \gamma$ and we get (cf. [2])

$$(2) \quad \lim_{n \rightarrow \infty} \frac{|\{i \leq n : u_i = 1\}|}{|\{i \leq n : u_i = 0\}|} = \frac{\gamma}{1 - \gamma} = \alpha.$$

The sequences (u_n) are called *Sturmian sequences* if α is irrational, and they are well-studied. See [2] for a broad treatment and for historic references.

6. DEFINITION. For $\alpha \in \mathbb{Q}^+$ we define the α -king to be a king whose move sequence is given by (1) with $\phi = 0$. This means that in the n -th time step the king moves by one square if $u_n = 1$ and the devil gets to block a new square if $u_n = 0$.

The choice of the offset ϕ looks arbitrary. For a natural definition it would be desirable that the chances of the α -king in the game do not depend on this parameter. And in fact, they don't.

7. LEMMA. *Any two kings with move sequences generated by (1) with the same speed parameter α but different ϕ 's either can both escape or can both be caught.*

PROOF. Let (u_n) be a sequence defined by (1) with offset ϕ and let (u'_n) be another sequence defined with some other offset ϕ' , both with the same α , though. We distinguish rational and irrational α .

For $\alpha \in \mathbb{Q}$ we write $\gamma = p/q$ in reduced form. Obviously both sequences are periodic with period q . All we have to do is to align them in the right way. Partition the unit interval into congruent half open intervals of length $1/q$ and let $r \in \{0, \dots, q - 1\}$ be the index of the residue class with

$$\phi \in \left[m + \frac{r}{q}, m + \frac{r+1}{q} \right) \quad \text{for some } m \in \mathbb{Z}.$$

We look for some index j , where the sequence $(jp/q + \phi')$ hits the same residue class; i.e.,

$$j \frac{p}{q} + \phi' \in \left[m' + \frac{r}{q}, m' + \frac{r+1}{q} \right) \quad \text{for some } m' \in \mathbb{Z},$$

which clearly exists because p and q are coprime. It is easy to see that

$$u_n = u'_{n+j} \quad \text{for all } n \in \mathbb{N}.$$

If the devil has a winning strategy on the move sequence (u_n) , he can therefore also win on (u'_n) by simply waiting j time steps and then starting to play according to the strategy for (u_n) . By exchanging (u_n) and (u'_n) we get the converse simulation.

For irrational α the sequences (u_n) and (u'_n) are non-periodic. We use a deeper result from the theory of such Sturmian sequences: The set of contiguous subwords of the sequence (u_n) depends only on α and not on the offset ϕ [2]. (Even more, there are exactly $n + 1$ different subwords of

length n and each of them occurs infinitely often in u_n .) Therefore, any initial segment of (u_n) can also be found somewhere in (u'_n) .

Since Lemma 3 tells us that if the devil can win on the move sequence (u_n) , he can do so in a bounded number of N steps, say, he can use a strategy on (u_n) to win on (u'_n) by simply waiting until a copy of the N -prefix of (u_n) starts in (u'_n) and then pursuing this strategy. \square

For $\alpha \in \mathbb{N}$, the above definition of an α -king obviously coincides with the previous one that was restricted to integral speed. For $\alpha = k \in \mathbb{N}^+$ the defining sequence (1) produces exactly k 1's between any two consecutive 0's, just as expected. It is also clear that our notion of an α -king fulfills our wish for fairness, large chunks of devil moves cannot occur. One easily checks that for $\alpha \geq 1$, the devil never gets to block two squares at a time. On the other hand, we can guarantee that not only in the long run but also locally, the devil always gets his share of moves.

8. DEFINITION. A 0/1-sequence is (s, t) -bounded, $s, t \in \mathbb{N}^+$, if every contiguous subword that contains strictly more than s occurrences of 1's contains at least t occurrences of 0. We call a king with a given move sequence (s, t) -bounded if the sequence is (s, t) -bounded. (Where we interpret 1's as king's moves and 0's as devil moves.)

9. LEMMA. An α -king, $\alpha \in \mathbb{Q}^+$, is (s, t) -bounded for every pair $s, t \in \mathbb{N}^+$ with $\alpha \leq s/t$.

The ‘‘strictly’’ in the definition appears for a technical reason; it does not mean that we get only t devil moves per $s + 1$ king moves on average. Namely, starting from any 1 in the sequence, we count 0's until we reach the $(s + 1)$ st 1. By then we have passed at least t 0's. When we read on until the $(2s + 1)$ st 1 shows up, we are sure to have counted at least $2t$ 0's. And so on. Before the $(rs + 1)$ st 1 appears, we are guaranteed to read at least rt many 0's.

PROOF OF LEMMA 9. Assume we have $s + 1$ many 1's between two positions a and b (inclusively) in the sequence (u_n) . Telescoping (1) yields

$$s + 1 \leq \sum_{a \leq i \leq b} u_i = \lfloor (b + 1)\gamma + \phi \rfloor - \lfloor a\gamma + \phi \rfloor < (b - a + 1)\gamma + 1,$$

where the terminal 1 accounts for the error that might result from the deletion of the floors. For the number of 0's in this interval we thus get

$$b - a + 1 - (s + 1) > \frac{s}{\gamma} - (s + 1) = \frac{s - \alpha}{\alpha} \geq t - 1. \quad \square$$

4. Catching a $(2 - \varepsilon)$ -King

In this section we develop a devil strategy to catch all kings of speed less than 2. The following main theorem emerged from joint work with Attila Pór.

10. THEOREM. The devil can catch any α -king with $\alpha < 2$.

Have a look at Figure 7 again, where the devil pushed a 2-king along a line of density $1/2$. With every second square already in place, the 2-king could never break through. We generalize this idea to kings of arbitrary speed.

11. DEFINITION. An *infinite (s, t) -fence* is an infinite horizontal or vertical strip in the plane with some squares blocked such that when an (s, t) -bounded king enters the strip from one side, the devil can play in a way that prevents the king from leaving it on the other side. Formally, such a fence is just a map $F: \mathbb{Z} \times [1..w] \rightarrow \{0, 1\}$, where $F^{-1}(1)$ is the set of blocked squares. The integer w is called the *width* of F .

We call such a fence *periodic* if there exists some integer λ such that $F(x, y) = F(x + \lambda, y)$ for all $x \in \mathbb{Z}$. Call the minimal such λ the *period* of F . In this case we also define the *density* of the fence, as the ratio

$$\frac{1}{\lambda} \left| \{(x, y) \mid 1 \leq x \leq \lambda, 1 \leq y \leq w, F(x, y) = 1\} \right|.$$

Note that density is measured with respect to length, not area. Width is not the crucial quantity, it appears for merely technical reasons.

12. LEMMA. *Against an (s, t) -bounded king, $s/t \leq 2$, there exists a periodic infinite fence of density $1 - t/s$ and width $10s + 1$.*

PROOF. We provide a periodic map $F: \mathbb{Z} \times [1..10s + 1] \rightarrow \{0, 1\}$ with the desired properties. Let F be everywhere zero except for those points (x, y) with

$$0 \leq x \bmod s < s - t \quad \text{and} \quad y = 5s + 1.$$

In other words, we group the central horizontal line $y = 12 + 1$ into segments of s squares and place $s - t$ blocks in each segment. See Figure 8. The density of this pattern is obviously the claimed $(s - t)/s$.




FIGURE 8. An infinite (s, t) -fence.

We now show how the devil keeps the king from crossing F by making sure that he can never step on the central line. By symmetry we may assume that the king enters the strip from the bottom.

Like in the case of the 2-king, we make sure that in the proximity of the king the central line is always filled completely. Precisely, if the segment S_0 above the king's current position has already been filled completely and the one to the left and right, S_{-1} and S_1 , too, then the devil acts as follows. As soon as the king steps into the area below the segment S_{-1} to the left, the devil uses his next t moves to fill up the segment S_{-2} , further to the left. By (s, t) -boundedness, this is finished before the king gets to play his $(s + 1)$ st move (counting the move that entered S_{-1} as the first). Hence,

by that time the king must be somewhere below the segments S_{-1} and S_0 , and because S_{-2} is now filled we are in the situation as before: the three segments directly above the king are blocked. If the king started running to the right, the devil would have filled the segment S_2 , of course. The devil can iterate this recipe forever, never letting the king step on the central line.

To obtain the above configuration, we reuse the procedure for the 1-king from section 1, where we managed to get a block of three consecutive squares in the king's way. Interpreting a whole segment S_i as a single square in which we must play t moves, we immediately see that 5 such meta moves suffice to get three segments prepared. Since any sequence of t devil moves yields no more than s king's moves, this gives a total of $5s$ approach moves, which is just the width of the strip below the central line. \square

The devil cannot build infinite structures in finite time. Infinite fences serve as a mere theoretical concept, which is easier to handle than finite fences, whose existence can be easily derived from the infinite ones.

13. DEFINITION. A *finite (s, t) -fence* is a rectangular box of size $\ell \times w$ in the plane with some squares blocked, such that when an (s, t) -bounded king enters through one of the length- ℓ sides he can only leave through that side again and such that all squares along the two length- w sides are blocked. Formally, such a fence is a map $F: [1.. \ell] \times [1.. w] \rightarrow \{0, 1\}$, where $F^{-1}(1)$ is the set of blocked squares. The integers ℓ and w are called the *length* respectively *width* of F . The *density* of the fence is the ratio

$$\frac{1}{\ell} |\{(x, y) \mid 1 \leq x \leq \ell, 1 \leq y \leq w, F(x, y) = 1\}|.$$

The following lemma provides the trivial transformation of an infinite fence into a finite fence.

14. LEMMA. *If there exists a periodic infinite (s, t) -fence of density σ then there exist finite (s, t) -fences of the same width and of density no more than*

$$\sigma + \frac{2w}{\ell}$$

for any length $\ell \geq 1$.

PROOF. The basic idea is obv: we cut a length- ℓ segment out of the infinite fence S . We only face a little inconvenience. Unless the desired length ℓ is a multiple of the period λ of S , our chosen segment might contain more than the average density due to local inhomogeneities. This problem is easily overcome by looking at a sequence of λ aligned length- ℓ segments of S . Since their total length is an exact multiple of λ , the total mass in all of them is exactly $\sigma\lambda\ell$. Now at least one of those segments contains no more than the average $\sigma\ell$ blocks. To turn this segment into a finite fence, we have to fill the length- w sides up completely, which costs the additional $2w$ squares. \square

Lemma 12 provides us with an infinite fence of density $1 - t/s$, which is strictly smaller than $1/2$ for an α -king with $\alpha < 2$. This does not seem to suffice to catch any such king, yet, but for $\alpha < 9/8$ we already get a devil win as follows. By Lemma 9 this speed bound grants us (s, t) -boundedness

with $s/t < 9/8$. So there exist infinite fences of density $\sigma \leq 1 - t/s < 1/9$. Choosing sufficiently long finite subfences of such an infinite strip, we can make the additional cost of $2w/\ell$ in Lemma 14 arbitrarily small, so that it gets absorbed by the small gap between σ and $1/9$. Altogether there exist finite (s, t) -fences of density at most $1/9$. This is all we need against our α -king. We simply build a square box of four such fences around him; in such a way that these fences touch but don't overlap. For a sidelength of ℓ this takes $4\ell/9$ devil moves, which in turn yield less than $9/8 \cdot 4\ell/9 = \ell/2$ king's moves. That means, all four fences will be finished by the time the king reaches the boundary of the box. Hence, he will be caught.

In prospect of the proof of Theorem 10, we forgo more formal details of this argument because the $9/8$ -king is covered by that result. The strategy behind Theorem 10 starts of just like the $9/8$ case, by obtaining some fence of density below $1/2$. The trick then is to assemble many such fences into a huge new fence of slightly smaller density. Iterating this process, we will eventually produce fences of arbitrarily small density. The key tool is the following lemma, whose proof describes this construction. (Observe that the bound $(s/t)\sigma^2$ in this lemma is strictly less than σ , which means that the density is really decreased.)

15. LEMMA. *If there exist finite (s, t) -fences, $s/t \leq 2$, of any length above some value ℓ_0 , all of the same width w and with density bounded by a common $\sigma < 1/2$, then there also exists a periodic infinite (s, t) -fence with density below*

$$\frac{s}{t} \sigma^2.$$

PROOF. The basic idea is to assemble infinitely many identical vertical finite density- σ fences to a wide horizontal fence of the desired density. As the length ℓ of those finite fences we pick any multiple of s larger than ℓ_0 and w . (Actually ℓ_0 should be much bigger than w anyway, but let us demand $\ell \geq w$ here for the sake of rigor.) As the distance between those fences we choose

$$m := \left\lfloor \frac{t\ell}{s\sigma} \right\rfloor \geq \ell.$$

Let the width of the infinite fence L we want to construct be 7ℓ , i.e.,

$$L: \mathbb{Z} \times [1 .. 7\ell] \rightarrow \{0, 1\}.$$

Figure 9 shows how the vertical fences of length ℓ and width w are placed in the central ℓ -strip of L . Precisely, the region

$$[\nu(w + m) + 1 .. \nu(w + m) + w] \times [3\ell + 1 .. 4\ell]$$

forms a fence for each $\nu \in \mathbb{Z}$.

Before we start to play on L , let us compute its density. The period is $w + m$ and each segment of this length receives no more than $\sigma\ell$ blocks, so we can bound L 's density by

$$\frac{\ell\sigma}{m + w} \leq \frac{\ell\sigma}{\left\lfloor \frac{t\ell}{s\sigma} \right\rfloor + 1} \leq \frac{s}{t} \sigma^2,$$

which is what we claimed.




FIGURE 9. Assembling many finite vertical fences into one big infinite horizontal fence.

We turn to the more difficult part: showing that L is indeed an (s, t) -fence. Assume the king enters L from the south, so we have to keep him from reaching the upper border. The basic idea is to build a horizontal fence between the upper ends of two vertical fences whenever the king runs north between them. Such a horizontal fence will be of length m to make it fit nicely in the gap. It will be placed in the rectangle

$$[\nu(w + m) + w + 1 .. (\nu + 1)(w + m) - 1] \times [4\ell .. 4\ell + w - 1]$$

for the respective $\nu \in \mathbb{Z}$. This arrangement is displayed in Figure 10 (the shaded area between the fences will soon be addressed). Note the vertical one-point overlap with the vertical fences on line 4ℓ . To avoid confusion: Those horizontal fences will be created *dynamically* by the devil when necessary, they are not part of the original strip L when the king enters.




FIGURE 10. A horizontal fence between two vertical fences and the shaded slot between them.

We now describe the essential aspect of the devil strategy starting from a standard situation, postponing the matter how to reach that situation for later. Therefore we give the shaded area between two vertical fences and below the (potential) horizontal fence a name: call such a rectangle of the form

$$[\nu(w + m) + w + 1 .. (\nu + 1)(w + m) - 1] \times [1 .. 4\ell - 1], \quad \nu \in \mathbb{Z},$$

a *slot*. We say that the king is in *standard position* if he is located within a slot whose upper border is already closed with a horizontal fence or he sits between two such blocked slots, perhaps within the vertical fence between them.

Let us assume the king is in standard position. We claim that if he leaves the slot then the devil can force him into standard position again by playing as follows. When the king enters one of the three surrounding fences, he follows the strategy of that respective fence to make sure that the king does

not break through to the other side of that fence. Note that we use the fact here that those fences do not overlap so that the devil is not forced to play in two fences simultaneously. Since there are no gaps where the three fences touch, this play guarantees that the king cannot leave the current slot above line 3ℓ without rebounding from the fences.

If the king leaves the slot that way below, to the left, say, the devil starts constructing the horizontal fence across the slot to the left. This takes no more than

$$(3) \quad m\sigma = \left\lfloor \frac{t\ell}{s\sigma} \right\rfloor \sigma \leq \frac{t\ell}{s}$$

devil moves. During this time the devil completely ignores the king play. In particular, he does *not* respond to the possible king's crossing of any fences, thus rendering them ineffective. Where can the king get while the devil is off at work? By (s, t) -boundedness the king gets no more than s steps per t devil moves. Counting the step out of the slot as the king's first move, we reckon that until the $(\ell + 1)$ st king move, the devil has made at least $t\ell/s$ moves, i.e., the $(\ell + 1)$ st king move comes after that many devil moves. Since this figure is just what we have computed in (3), the king gets no more than ℓ moves before the new fence is finished. Note that we used that ℓ is a multiple of s when we applied (s, t) -boundedness.

A look back at Figure 9 reveals how far the king can have run in ℓ moves. Since the first move lead strictly below the $(3\ell + 1)$ st line, he cannot have reached the $(4\ell + 1)$ st line, where the horizontal fences start. Neither can he have crossed completely the slot to the left, nor the original slot because those areas are each $m \geq \ell$ points wide. Consequently, the king ended up somewhere inside the old slot or the new slot, which is fine because both now have a fence above them, or he sits somewhere in between them. That means he is in standard position again.

It remains to show how to reach a standard position from the initial situation when the king enters the unmodified strip on line 1. The argument is again very similar to the respective part of the strategy against the 1-king. Here it is actually even simpler because we need only one horizontal fence instead of a triple block. However, the notion of standard position requires a little extra attention.

Call the slot in which the king's first position lies S_0 , in case he enters just between two slots, just pick any of them; and label the four neighboring slots correspondingly $S_{-2}, S_{-1}, S_1,$ and S_2 , from left to right. The devil first constructs the horizontal fence above slot S_0 . We already know from the previous computation that such an endeavor grants the king at most ℓ steps. So afterwards, the king sits in one of the slots S_{-1}, S_0, S_1 or in one of the two gaps between them. Inside S_0 he is already in standard position. If he sits in a gap, the one between S_0 and S_1 , say, then the devil builds a fence above S_1 , after which the king can only be in S_0, S_1 or the gap where he already was before. So we have reached standard position. It remains to consider a king in slot S_{-1} or S_1 with S_0 blocked; in S_1 , say, by symmetry. Then the devil builds the fence above S_2 , squeezing the king between S_0 and S_1 . Blocking S_1 with the next m moves, then leads into standard position.

Altogether, the total number of devil moves is in no case larger than $3m$, so that by the time we attain standard position, the king will not have reached the $(3\ell + 1)$ st line, yet. \square

PROOF OF THEOREM 10. Pick positive integers s and t with $\alpha \leq s/t < 2$, so that the α -king is (s, t) -bounded by Lemma 9. Then Lemma 12 provides us with an infinite periodic (s, t) -fence of density $\sigma < 1/2$.

For an application of Lemma 15 we have to fix a suitable lower bound ℓ_0 on the length of the finite fences we allow for the construction of the new fence. Therefore we write $s/t = 2/(1 + \delta)$ with some, possibly very small $\delta > 0$ and choose ℓ_0 large enough to ensure that the density of the finite fences longer than ℓ_0 , as obtained by Lemma 14, is bounded through

$$\sigma + \frac{2w}{\ell_0} \leq \sqrt{1 + \delta} \sigma.$$

Now Lemma 15 gives us an infinite (s, t) -fence of density

$$\sigma' \leq \frac{s}{t} \left(\sigma + \frac{2w}{\ell_0} \right)^2 \leq \frac{2}{1 + \delta} (\sqrt{1 + \delta} \sigma)^2 = 2\sigma^2.$$

Repeated application of this procedure yields a sequence $\sigma_0, \sigma_1, \sigma_2, \dots$ with $\sigma_n \leq 2\sigma_{n-1}^2$ and $\sigma_0 < 1/2$. The resulting bound

$$\sigma_n \leq \frac{1}{2} (2\sigma_0)^{2^n}$$

is easily verified, so that we see that the sequence (σ_n) converges to 0.

In a game against the α -king, the devil can now arrange four finite (s, t) -fences of density smaller than $1/16$, say, along the four sides of a huge square around the king. With α bounded by 2, the devil builds such fences more than 8 times faster than the α -king runs and thus finishes them before the king reaches any of them. Hence, the king will never leave that big box. Note that the fences have to be arranged in a non-overlapping way to ensure that the devil can play in each of them independently. And maybe we should also remark that the king cannot run around in his cage forever. After some time, when the fences are filled to the rim with devil moves, the devil simply starts flooding the central region with blocks until the king eventually gets stuck. \square

As the proof has shown, the 2 in Theorem 10 maxes out the potential of our fences. We have already indicated in our discussion of the 2-king on page 10 that a speed of 2 can be considered “fair” with respect to fence building. If one used fences in the described way against faster kings, their construction would be more expensive than the gain through the resulting king’s detour. This can perhaps be seen as some very weak indication that a $(2 + \varepsilon)$ -king cannot be caught anymore, but fences could, of course, just be one technical tool, without any deeper meaning for the game.

Anyway, since Conway’s article of 1996, there has apparently not been any progress on the angel problem. Maybe Theorem 10 stimulates interest in this game again, since the concept of α -kings allows for arbitrarily small improvements in devil strategies. Perhaps we can learn something new about our two antagonists from the voyages of very slow kings. Join the game!

5. An Escape into Space

So we do not have any escape strategy for any k -angel in the plane. Maybe we can obtain some positive result in higher dimensions, where an escape should be potentially easier. And in fact we can. 3D-angels live in a 3-dimensional world of cubes, indexed by coordinates in \mathbb{Z}^3 . Just like in the plane, in every move the k -angel jumps from his current position (x, y, z) to some other cube (x', y', z') with $|x' - x|, |y' - y|, |z' - z| \leq k$, and in turn, the devil blocks some cube of his choice. We prove the following.

16. THEOREM. *On the three-dimensional board, the 13-angel can escape forever.*

The problem at hand has only been mentioned once in the literature, also in [8], where the authors actually report to know escape strategies for angels in higher dimensions. However, a proof has apparently never been published.

Theorem 16 should be seen in the proper light. It is not a breakthrough on the way towards a solution of the two-dimensional case but rather confirms that the original question by Conway, Berlekamp, and Guy addresses the right problem. Moreover, it will become clear that our solution, a density-sensitive path-search method based on a hierarchical space partition, will *not* carry over to the two-dimensional game—not without major modifications, at least. So we not only provide a first constructive escape strategy for a variant of the angel problem but also want to point out the intrinsic obstacles for similar strategies in dimension two, to further emphasize the hardness of the original angel problem.

The box hierarchy. Our escape strategy divides the world into an infinite hierarchy of larger and larger boxes. The angel will have to make sure that on each level, his current box contains not too many devil blocks. This shall then guarantee his free travel.

A remark on terminology. Our usage of the word “cube” might get a little confusing when we speak about our hierarchy, since higher-level boxes will themselves be cubes—of cubes of cubes of cubes, etc. We shall use the expression *elementary cube* to emphasize that we mean the basic locations of the board, while the term *box* be reserved for collections of such objects. With other expressions the intended meaning should in general be clear from the context.

On the first level, the world is regularly partitioned into cubes of side-length 13, such that the origin $0 \in \mathbb{Z}^3$, where the angel starts, lies at the very center of one of these boxes. Formally, the first level H_1 is the collection of all boxes

$$H_1^{(u,v,w)} := \left\{ (x, y, z) \in \mathbb{Z}^3 \mid \begin{array}{l} 13u - 6 \leq x \leq 13u + 6, \\ 13v - 6 \leq y \leq 13v + 6, \\ 13w - 6 \leq z \leq 13w + 6 \end{array} \right\},$$

with $u, v, w \in \mathbb{Z}$, where we reference elementary cubes of the world via their coordinates $(x, y, z) \in \mathbb{Z}^3$.

The sidelength 13 corresponds to the power of the 13-angel. From level 2 on, sidelengths grow by a factor of 29 per step, where there is no deeper

reason for the choice of this particular value except that it makes the forthcoming computations work. On each level we again demand that the origin lie at the very center of the one box that contains it. Technically, for $j \geq 2$ the j th level H_j of our hierarchy is the collection of all boxes

$$H_j^{(u,v,w)} := \left\{ H_{j-1}^{(a,b,c)} \mid \begin{array}{l} 29u - 14 \leq a \leq 29u + 14, \\ 29v - 14 \leq b \leq 29v + 14, \\ 29w - 14 \leq c \leq 29w + 14 \end{array} \right\},$$

with $u, v, w \in \mathbb{Z}$.

So any box on level $j \geq 2$ contains 29^3 boxes on level $j - 1$ and the whole hierarchy is symmetric to the origin. Note that formally the elements of a higher-level box are again boxes, which is what we want. But with a certain laxness we shall also consider a level- j box simply as the set of the $(13 \cdot 29^{j-1})^3$ elementary cubes that lie inside it. In this vein we define the *level- j box of a cube $a \in \mathbb{Z}^3$* to be the unique box in H_j that “contains” the elementary cube a and denote it by

$$Q_j(a).$$

Further we define a mass function μ for all boxes A on all levels of our hierarchy, letting

$$\mu(A)$$

count the number of elementary cubes inside A that have already been blocked.

Clear roads ahead. Globally, the angel’s route through our hierarchy of boxes will be guided by simple mass constraints, in a quite elegant way. The basic step, the transition between two adjacent boxes, however, requires some dirty work. We need to introduce a few technical notions to ensure that locally the angel does not get stuck in unfortunate arrangements of blocks. The ideas are similar to the invariant from the proof of Proposition 5 from page 8.

17. DEFINITION. Let E be a quadratic grid of 29×29 cubes with some cubes marked *forbidden*. We say that a cube a of E *lies clear in E* if

- no more than 12 of the $29^2 = 841$ cubes in E are forbidden,
- a lies in the central 13 by 13 square of E ,¹ and
- the two axis-parallel lines through a in E contain no forbidden points.

See the left-hand side of Figure 11.

Let C be a cubic grid of $29 \times 29 \times 29$ cubes with some cubes marked *forbidden*. We say that a cube a of C *lies clear in C* if

- no more than 333 of the the $29^3 = 24,389$ cubes in C are forbidden and
- a lies clear in one of the three axis-parallel 29×29 planes through a in C .

See the cube in Figure 11.

¹The occurrence of the number 13 here is coincidental. This is a “different” 13 than the one from Theorem 16.




FIGURE 11. Clear positions.

The idea behind the above definitions is, as we said before, to guarantee free navigation from a clear cube within a sidelength-29 box to somewhere outside this box. A cube that lies clear will have enough free space around it to guarantee an easy route out. The forbidden cubes may, of course, not be used for travel. We do not speak of blocked cubes in Definition 17 because the little cubes will usually themselves be boxes of smaller cubes. But forbidden cubes will almost be blocked, meaning that their mass exceeds a certain threshold.

For *paths* through such boxes we allow axis parallel steps of unit distance only. That is, a single step of a path is a change of ± 1 in just one coordinate—in contrast to basic angel moves. This restriction is due to the hierarchical structure of our argument. We will be able to travel between two little cubes inside the big cube in Definition 17 only if these cubes share a face which may be used for a transition on the next lower level.

From a purist's point of view, the grids E and A of Definition 17 could, of course, just be called grid graphs, with “cubes” replaced by “vertices.” Then a path would just be a paths in the graph theoretic sense and the following lemmas are in fact just statements about such grid graphs. However, we like to keep with our view of cubes and boxes, hoping that this does not cause any confusion.

The following lemma about planes only serves as a tool for the three-dimensional case. Our actual interest will be in paths through boxes.

18. LEMMA. *Let q be a cube lying clear in a 29×29 grid E . Then at least $763 = 29^2 - 78$ cubes of E are reachable from q in at most 40 steps each.*

PROOF. Any cube on the two lines through a is by assumption reachable directly through that respective line. For every other point $p \in E$ we consider the two potential paths that run parallel to the axes with exactly one turn. A cube p may not be reachable on either of these two paths for two reasons: both paths are blocked or p is a forbidden cube itself. Since by the special choice of our paths, a single pair of forbidden cubes covers at most one cube of E , the first situation can happen for at most $\binom{12}{2} = 66$ cubes, the second, by definition, for at most 12; which makes 78 inaccessible places altogether. One easily computes that any of the remaining $29^2 - 78 = 763$ cubes is reachable in at most 40 steps since the distance from any location in the central region to any side of E is at most 20. \square

19. LEMMA. *Let q be a cube lying clear in a $29 \times 29 \times 29$ box grid C and let D be another $29 \times 29 \times 29$ box aligned with C along one face of C , also with no more than 333 points marked forbidden. Then there exists a cube r lying clear in D such that there is a path of length at most 165 from q to r , which after the first 96 steps uses no more cubes in C .*


PROOF. Let E denote the plane within C in which p lies clear as required by Definition 17. The basic idea for the path construction is to pick a suitable plane F in D , which will contain the target point r , and then to find many disjoint paths from E to F not all of which can be blocked by forbidden cubes.

Observe that by the pigeon-hole principle, among the 29 axis-parallel planes in D that lie parallel to that face of D which borders on C , at least one contains no more than 12 forbidden cubes ($29 \cdot 13 = 377 > 333$). Choose F to be such a plane. For both dimensions of F , at most 12 of the 13 axis-parallel lines passing through the central 13×13 region of F are blocked by forbidden cubes, which leaves at least one clear line in each direction. We choose b as the intersection of two such lines, which makes it lie clear in D . We now distinguish two different cases: when the planes E and F are parallel and when they are not.

First case: E parallel to F . Partition the union of C and D into the $29^2 = 841$ disjoint lines that intersect E and F orthogonally. By Lemma 18, all but 78 of these lines intersect F in cubes that are reachable from q in at most 40 steps and likewise all but 78 lines intersect F in cubes that are reachable in 40 steps from r . This leaves $841 - 2 \cdot 78 = 685$ lines whose intersections with E and F are reachable in 40 steps from a respectively b . By assumption, there are no more than 666 forbidden cubes in C and D altogether, so several of those lines are completely free. Since the distance between the planes E and F is bounded by twice the sidelength of the boxes C and D , we get a path from q to r of no more than $2 \cdot (40 + 29) - 1 = 137$ steps.

The second case, where E and F are not parallel, can be treated similarly. Only the connecting lines must be chosen in a more complicated way. Partition the union of C and D into 29 parallel planes of size 29×58 such that each plane intersects E and F in exactly one line. Within each of these planes we match the 29 cubes of C with the 29 cubes of D by 29 disjoint paths as displayed in Figure 12. As in the first case, we thus get a positive amount of paths connecting locations in E reachable from q to locations in F reachable from r and free of forbidden cubes. The length bound is a little worse, however. Paths in Figure 12 can require up to $28 + 29 + 28 = 85$ steps, which together with the paths within the planes E and F yields an upper bound of 165 steps from q to r . It is also easily checked that in either configuration we spend no more than 96 steps inside C . \square

We want to apply the box-travel lemma to boxes of our hierarchy (H_j) . Therefore we have to define which level- $(j - 1)$ subboxes inside a level- j box should be considered forbidden. This shall, for now, depend on a simple mass constraint. (Later we will also need a slightly modified definition.)

FIGURE 12. Traveling between non-parallel planes E and F .

20. DEFINITION. Call a box $A' \in H_{j-1}$, $j \geq 2$, *light* if

$$(4) \quad \mu(A') \leq \frac{17}{3} \cdot 165^{j-1}$$

and *heavy* otherwise.² We then say that the angel's position a is *nice on level j* if the subbox $Q_{j-1}(a)$ lies clear in $Q_j(a)$, with exactly the heavy level- $(j-1)$ boxes forbidden. The position is *nice on level 1* simply if

$$(5) \quad \mu(Q_1(a)) \leq 1157.$$

We say that a position is *nice up to level j* if it is nice on all levels from 1 through j .

The notion of niceness will be suitable to guarantee an escape route out of the current level- j box $Q_j(a)$. Recall that the constant 165 is exactly the step bound provided by Lemma 19. Level 1 receives a special treatment because it will be used in the induction basis, founding our hierarchy argument on actual angel moves.

The main induction—escaping from larger and larger boxes. With the notion of niceness at hand, it is actually rather straightforward to formulate an appropriate induction hypothesis for angel strategies that allow to travel between arbitrarily large boxes. Only a few constants remain to be chosen thoroughly. And of course, we have to make some assumption on the target box we want to run into. Actually, a simple mass constraint will do.

21. PROPOSITION. *Let B be one of the six level- j boxes neighboring the angel's current box $A \in H_j$, $j \geq 1$. If his current position is nice up to level j and the mass of B is bounded by*

$$(6) \quad \mu(B) \leq 7 \cdot 165^j$$

then the 13-angel can get in no more than

$$(7) \quad 2 \cdot 165^{j-1}$$

elementary moves from his actual position in A to some location in B such that after he has arrived there, his position will be nice up to level j again.

²We prefer to write $j-1$ instead of simply j to emphasize that although lightness is a property of a single box, it shall always be used in reference to the containing box on level j .

Note that the coefficient 7 in (6) is slightly larger than the $17/3$ in (4). So for the box B in Proposition 21, we impose a weaker mass constraint than would be required for being considered light as a subbox in the containing box on level $j + 1$. We also remark that 165^j lies somewhere in between the sidelength of a level- j box and the number of points in a face of such a box. One could say that with increasing level, the mass bound (6) grows strictly faster than one-dimensional objects but strictly slower than two-dimensional objects. Likewise the path length (7); compared to the diameter of a level- j box, it gets arbitrarily large, hence, seen from a far distance, the angel slows down to almost zero speed. Compared to surface growth, however, and this is the crucial measure because potential devil obstacles must be two-dimensional, the speed can actually be seen to *increase* by $29^2/165 > 5$ per level.

PROOF OF PROPOSITION 21. By induction on j . The induction basis is $j = 1$. We have exactly 2 moves to get from the current sidelength-13 box A to an arbitrary elementary cube in B . By niceness, A contains at most 1157 devil blocks and by (6), B contains no more than $7 \cdot 165 = 1155$ blocks. Thus, by the pigeon-hole principle, any 7 planes within the current box A or the target box B contain at least $7 \cdot 13^2 - 1157 = 26$ free locations. Hence, the 13-angel may jump from its current position a to some other place in A at most 7 units away from B . From there he can reach in just one further jump any point within the first 7 layers of B , which still contain some unblocked cubes. He jumps to one of them with his second move. The two devil answers cannot raise the mass of B over 1157, so afterwards the position will be nice on level-1 again, as required.

Induction step from $j - 1$ to $j \geq 2$. Niceness of the current position a guarantees that there are at most 333 heavy subboxes A' in A , all the other boxes satisfying the lightness condition (4). In our target box B we also mark forbidden subboxes, based however, on a slightly stronger mass constraint. Mark a level- $(j - 1)$ subbox B' in B forbidden if does *not* satisfy

$$(8) \quad \mu(B') \leq \frac{11}{3} \cdot 165^{j-1}.$$

So in B , non-forbidden subboxes are “ultra light.” Since 334 such forbidden boxes in B would yield a total mass of

$$334 \cdot \frac{11}{3} \cdot 165^{j-1} > 7 \cdot 165^j,$$

our assumption (6) implies that B contains no more than 333 forbidden boxes.

Now there are two adjacent level- j boxes A and B with at most 666 level- $(j - 1)$ subboxes forbidden altogether, based on two slightly different criteria. By niceness on level j of the current position a , the box $Q_{j-1}(a)$ lies clear within the box $A = Q_j(a)$. Further, the neighboring level- j box B contains fewer than 333 forbidden boxes. Hence Lemma 19 applies to A and B , giving a path (U_0, U_1, \dots, U_t) of level- $(j - 1)$ boxes with $t \leq 165$ from the current box $Q_{j-1}(a) = U_0$ to some U_t that lies clear in B with respect to the ultra-light boxes there. Moreover, the lemma guarantees that from U_{97} on all boxes lie in B .

We use this path of boxes to obtain an actual strategy that gets the angel from a to some point in U_t . Niceness up to level j at his starting position a implies niceness up to level $j - 1$, so we apply our induction hypothesis on level $(j - 1)$ to the pair U_0, U_1 , getting the angel to a position within U_1 that is also nice up to level $j - 1$ and from there to a nice position inside U_2 —and so on, all the way to some b that is nice up to level j in U_t . However, this will only work if the mass constraint (6) is satisfied for the target box U_τ in each single transition between two adjacent boxes $U_{\tau-1}$ and U_τ .

This is easily checked. The whole journey from a to b would grant the devil at most

$$(9) \quad 165 \cdot 2 \cdot 165^{j-2} = 2 \cdot 165^{j-1}$$

moves. Even if he spends all of them on a single box U_τ in B , the mass of this box will remain bounded by

$$(10) \quad \mu(U_\tau) \leq \frac{11}{3} \cdot 165^{j-1} + 2 \cdot 165^{j-1} = \frac{17}{3} \cdot 165^{j-1}.$$

For a box U_τ in A we know that it cannot receive more than $95 \cdot 2 \cdot 165^{j-2}$ devil moves before we want to enter it, so that by the time we invoke Lemma 19 the following mass bound will hold:

$$(11) \quad \mu(U_\tau) \leq \frac{17}{3} \cdot 165^{j-1} + 95 \cdot 2 \cdot 165^{j-2} < 7 \cdot 165^{j-1}.$$

Both bounds, (10) and (11), satisfy the requirement (6) of Proposition 21 with j replaced by the appropriate level $j - 1$ there. Hence, all those transitions between the U_τ will be possible. Also note that the number of moves counted in (9) is exactly what we had to show for (7).

Eventually, the angel reaches a point b in U_t in the required number of elementary moves such that by that time the resulting position is nice up to level $j - 1$. It remains to show niceness on level j . To see this, recall that the relaxed mass bound for the originally ultra-light subboxes in B , which we computed in (10), matches exactly our definition (4) of light boxes. Hence, all subboxes B' of B that are heavy after the angel's trip from a to b , had already been forbidden in the beginning when the box-travel lemma was invoked, and thus the terminal box U_t lies clear in B with respect to those boxes. In other words, b is nice on level j , too. \square

Proposition 21 immediately implies the *existence* of an escape strategy. But since the following argument uses Lemma 3, we do *not* get an explicit strategy, yet.

PROOF OF THEOREM 16 (NON-CONSTRUCTIVE VERSION). At the very beginning of the game all boxes on all levels of our hierarchy are empty and thus light within their respective containing boxes. Because of the symmetry of the hierarchy with respect to the origin, the angel starts at the very center of the box $Q_j(0)$ on every level $j \geq 1$. Therefore, the starting position is nice on every level $j \geq 1$.

By Proposition 21 the angel can thus travel to some adjacent box on any previously given level of the hierarchy, which allows him to escape the devil for any previously chosen amount of time. So by Lemma 3, the angel can escape forever. \square

An explicit infinite escape strategy. If someone really wants to play the angel game for some infinite time, the previous, abstract proof is no big help, telling us just that the angel *can* win—somehow. To obtain an explicit escape strategy, we have to work a little harder and revisit some details of the proof of Proposition 21.

PROOF OF THEOREM 16 (CONSTRUCTIVE VERSION). We start escape strategies on *all* levels of the hierarchy simultaneously—in such a way that on initial segments those strategies are compatible. Therefore we introduce a small technical convention about the paths provided by Lemma 19.

Unrolling the induction in the proof of Proposition 21, we can interpret that result as a concrete strategy for journeys between adjacent boxes of our hierarchy, which on each level invokes Lemma 19 as an algorithm (implicitly given by its proof) for path finding in grid graphs. In this algorithmic view, let us agree that whenever Lemma 19 is invoked to find a path between two boxes that contain no forbidden cubes at all, it returns a path that starts with a step in the direction of the target box.

The angel begins by traveling from the origin 0 to a nice position a_1 in the level-1 box B_1 that lies directly behind (in positive z -direction, say) the initial box $Q_1(0)$. Having arrived at position a_1 , we can now interpret these first steps as the initial sequence of a travel from the box $Q_2(0)$ to a nice position a_2 in the level-1 box B_2 just behind the initial level-2 box $Q_2(0)$. As we already observed in the non-constructive proof above, such a strategy exists by Proposition 21 and by our convention it would have started with a travel to a position in B_1 , just as we did. We now follow the new level-2 strategy until we reach the position a_2 . At that point, we again interpret this journey as the initial sequence of a travel from the origin to a nice position a_3 in the level-3 box behind $Q_3(0)$. Iterating this argument indefinitely, we obtain an infinite escape strategy for the angel. The crucial argument here is that what we have done up to some point, will always fit into strategies on higher levels that we have not considered yet. \square

Why our hierarchy does not work in 2D. One might want to try to transform the hierarchy approach for the three-dimensional case into an escape strategy for the two-dimensional game. Such an attempt would face two major obstacles. First, as we already remarked after the statement of Proposition 21, the step bound (7) grows strictly faster than the sidelengths of the boxes. This effect is due to the detours we are making with each application of Lemma 19. On higher and higher levels, the effective speed of the angel thus gets arbitrarily slow. In the plane, this would allow the devil to completely encircle the angel on a sufficiently large scale since the boundary of a rectangle is proportional to the radius. Hence, we would need an improved path finder that might probably employ some means of charging devil moves against angel moves such that devil plays that force the angel to make detours cannot be counted for wall building far away.

But even if one should succeed in maintaining the “effective speed” of the angel, there would remain a more fundamental problem about hierarchical strategies like the one we presented. While routing out of a level- j rectangle R (or whatever regular shape might be used), the angel must at some point

