

A Solution to the Angel Problem

Oddvar Kloster

SINTEF ICT, Postboks 124 Blindern,
0314 Oslo, Norway, okl@sintef.no

Abstract. We solve the Angel Problem, by describing a strategy that guarantees an Angel of power 2 or greater to win. Basically, the Angel should move north as quickly as possible. However, he should detour around eaten squares, as long as the extra distance does not exceed twice the number of eaten squares evaded. We show that an Angel following this strategy will always spot a trap early enough to avoid it.

Keywords: Angel problem, combinatorial games

1. Introduction

The Angel Problem is described by John Conway in [4]: The Angel plays a game against the Devil on an infinite chessboard, whose squares correspond to pairs of integers (x, y) . The Angel starts out on a square of the board. In each turn, the Devil may eat any single square on the board. The Angel can then fly to any square that the Devil has not yet eaten, and whose distance from the Angel's current coordinates is at most k in the infinity metric (i.e. he can fly from square (x, y) to (x', y') if $|x - x'| \leq k$ and $|y - y'| \leq k$). The integer k is the Angel's *power*, a parameter of the game. The Devil wins if he can render the Angel unable to move, all squares within distance k of the Angel's current position having been eaten. The Angel wins by being able to move forever.

The Angel Problem is to determine whether, for a sufficiently large k , there exists a winning strategy for the Angel.

The game was apparently first presented in [1], where it is shown that the Devil can defeat an Angel of power 1, i.e. a chess king. However, it has long been an open problem as to who has a winning strategy when the Angel has power 2 or greater. Some progress has been made by Kutz and Pór [6, 8], who define the α -King and prove that for $\alpha < 2$, the Devil can catch an α -King (who is less powerful than a 2-Angel, but more powerful than a 1-Angel if $\alpha > 1$). Also, the fact that a sufficiently strong Angel can win in three dimensions has been long known and recently published [2, 6, 7].

The purpose of this paper is to demonstrate a strategy that guarantees the Angel of power 2 to win, so for the rest of the paper, we will fix $k = 2$. The fact that an Angel of any higher power can win follows immediately. We will also show a slight refinement that yields a winning strategy for the 2-King, who can make two consecutive chess king's moves in each turn, each of which must bring him to an uneaten square.

Other, independent proofs that the Angel wins have recently appeared. Bowditch [3] shows that the 4-Angel wins, while Máthé's proof [9] works for the 2-Angel. The paper by Gács [5] gives no specific power for the Angel.

2. Border curves

We start by visiting the space of *border curves*, on which the Angel's strategy is based. As the proofs of most lemmas in this section are tedious technicalities, we defer them to the Appendix.

Define a *segment* as the border between two adjacent squares on the board. We will consider continuous curves that are built from an infinite (in both directions) sequence of segments. The curves are directed, from their *past* to their *future* part. As part of a curve, a segment is similarly directed; regarded as part of the board, a segment is undirected.

Definition 1. Let s be a segment in a curve. The *right square* of s is the adjacent square that is on the right of s when looking in the future direction. The other adjacent square is the *left square* of s .

Definition 2. Let κ be a continuous, directed curve that consists of an infinite sequence of segments. We say that κ is a *border curve* if there exists a set V_κ of squares on the board such that:

- i) No segment on the board occurs more than twice in κ .
- ii) If a segment occurs exactly once in κ , then its left square is in V_κ , while its right square is not.
- iii) If a segment occurs twice in κ , then the occurrences have opposite directions, and both adjacent squares are in V_κ .
- iv) If a segment does not occur in κ , then its adjacent squares are either both in V_κ or both not in V_κ .
- v) Both V_κ and its complement (in the set of all squares on the board) are infinite.
- vi) V_κ is connected under separation by κ . That is, the squares in V_κ form a single connected component when we consider two squares neighbours if and only if they have a common border that that does not occur as a segment of κ .

We call V_κ the *left set* of κ . Its complement is the *right set* of κ .

Lemma 1. *The left and right sets of a border curve κ are unique.*

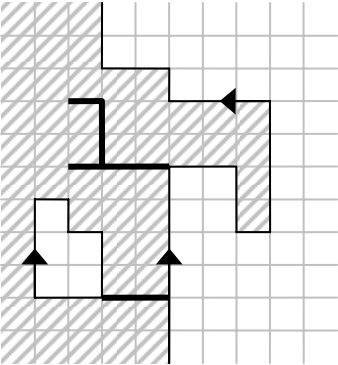


Figure 1. A border curve. Squares in the curve’s left set are shaded. Segments that are traversed twice are shown fatter.

Figure 1 shows an example of a border curve and its left set. Despite a slightly laborious definition, border curves are easily grasped intuitively. Basically, they trace the infinite border between a left and a right set of squares. While the left set must be connected, the right set may contain isolated components that are enclosed in the left set. If this is the case, the border curve must at some point leave the main border and carve a channel through the left set, circle the enclosed component clockwise and then return to the main border along the same channel. The curve may also trace additional dead end channels, as long as the left set is not split into multiple connected components.

We define two operations for transforming a border curve, κ , into another:

- *Extension.* Select a segment of κ whose right square is $q \notin V_\kappa$. Replace this segment with the other three borders of q , thus oriented that q is their left square.
- *Contraction.* Select two consecutive segments of κ that traverse the same segment on the board in opposite directions, thus forming a dead end. Erase these two segments from the curve.

Figure 2 shows some examples of these operations.

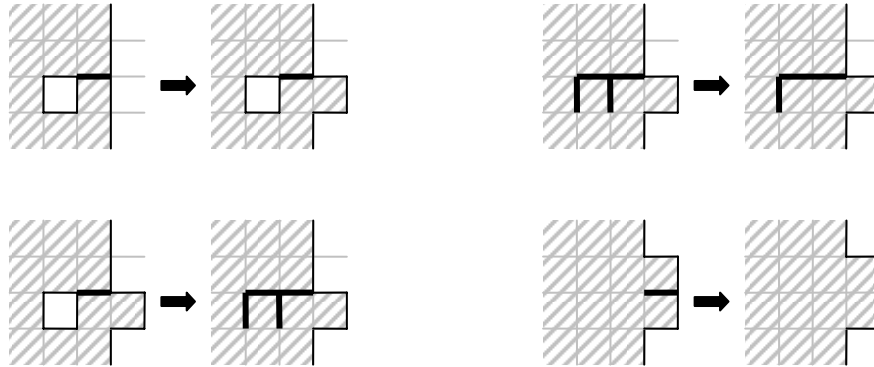


Figure 2. Left: Examples of extensions. Right: Examples of contractions.

Lemma 2. Let κ be a border curve, μ an extension of κ involving square q , and ν a contraction of κ . Then μ and ν are border curves, with $V_\nu = V_\kappa$ and $V_\mu = V_\kappa \cup \{q\}$.

Definition 3. Let κ and ν be border curves. If κ can be turned into ν by some (possibly empty) finite sequence of extensions and/or contractions, we call ν a *descendant* of κ .

Lemma 3. Let κ be a border curve and ν a descendant of κ . Then $V_\kappa \subseteq V_\nu$.

Proof. No extension or contraction can remove a square from the curve's left set. Thus any square in V_κ must also be present in V_ν . \square

Lemma 4. Let κ be a border curve and s a segment that occurs twice in κ . Let ν be the curve produced by erasing the section between both occurrences of s , inclusive, from κ . Then ν is a descendant of κ , and consequently a border curve.

3. The Angel's Strategy

We now turn to describing the Angel's strategy. While executing the winning strategy, the Angel maintains a *path* that represents his past movements and his current plans for the future. At any time, the path is a border curve. One of the path's segments is called the *perch*, and the Angel sits on the right square of the perch. On his turn, the Angel moves the perch two segments along the path toward its future part and alights on the right square of the new perch. Figure 3 shows a few of the Angel's turns. It is easy to see that the new square is within the Angel's power to reach: each time the perch is advanced one segment, the Angel moves at most one square in each dimension, or may even stay put (if the path turns clockwise). In fact, the Angel does not need to exert all his powers. Because he moves diagonally only when the path turns counter-clockwise, the largest move required of the Angel is that of a chess knight.

At the start of the game, the path is the infinite straight line from south to north that passes just west of the Angel's starting square. This line is easily seen to be a border curve, whose left set is the half-board west of the starting square.

Every time the Devil has eaten a square, the Angel will survey the board around the future part of the path to see if any traps are lurking. If he finds a sufficient number of eaten squares sufficiently close to the path, he chooses a new path whose left set includes those squares, thus evading them. The new path must be a descendant of the current path and may only differ from it future of the current perch. We will show that after each path update, the right square of essentially every future segment of the path is uneaten. This guarantees that the Angel can move forever by updating and following the path.

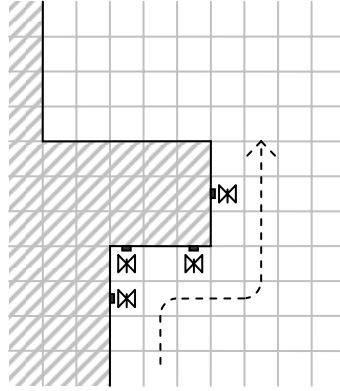


Figure 3. The Angel hugs the right side of the path, moving the perch two segments in each turn.

Definition 4. At any time, the squares of the board can be partitioned into three sets: *Free*, *Eaten* and *Evaded*. The Evaded set is equal to the current path's left set. An Eaten square is one that is in the current path's right set and has been eaten by the Devil. The uneaten squares in the path's right set are Free.

Note that 'eaten' and 'Eaten' have different meanings: an eaten square is Eaten only if it is in the path's right set; otherwise it is Evaded.

Initially, the entire western half-board is Evaded, while the eastern half-board, including the Angel's starting square, is Free. When the Angel updates the path, some Free and Eaten squares may become Evaded, and the Evaded set branches out into the eastern half-board. By eating a Free square, the Devil converts it to an Eaten square. Neither the Angel nor the Devil can change an Evaded square back to Free or Eaten. If the Devil should choose to eat an Evaded square, it stays Evaded, and the Devil has wasted a move.

Definitions 5. We enumerate the Angel's turns, each consisting of updating the path and then moving, by 1, 2, Define λ_i to be the path after being updated in turn i ; λ_0 is the initial path. Let κ be a descendant of λ_0 . Since λ_0 can be turned into κ by a finite sequence of extensions and contractions, each of which affects only a finite number of segments, λ_0 and κ must be equal sufficiently far in the past and future directions. Thus we can define L_κ , the *length* of κ , as the number of segments in κ minus the number of segments in λ_0 after equal infinite past and future parts have been removed from both curves.

Let j be a turn and κ a border curve. Define $n_\kappa(j)$ to be the number of squares in V_κ that the Devil has converted from Free to Eaten at some point before turn j . In other words, $n_\kappa(j)$ is the number of eaten squares in V_κ at turn j , but not counting those that were already Evaded at the time they became eaten.

For any turns i and j , we write $L_i = L_{\lambda_i}$ and $n_i(j) = n_{\lambda_i}(j)$ as shorthand notation.

Define p_i to be the Angel's perch after the Angel moves in turn i ; p_0 is the initial perch.

Informally, the rule by which the Angel updates the path can be stated as follows: The Angel must make the path's future part as short as possible, but is allowed to increase the length of the future path if this evades an additional Eaten square for every two segments added. Subject to this constraint, he must evade as many Eaten squares as possible. Figure 4 illustrates how the path may be updated in a few situations (but in each case, there are other valid ways to perform the update).

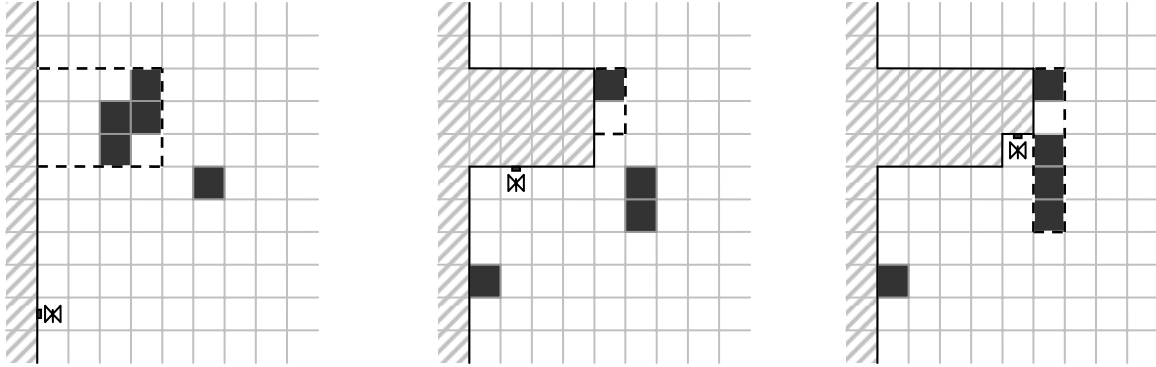


Figure 4. Updating the path. Black squares are Eaten, shaded squares are Evaded before the update.

We proceed to define the update rule formally. As the Angel starts turn i , λ_{i-1} is the current path and p_{i-1} is the current perch. Let P_i^1 be the set of border curves μ that satisfy conditions 1 and 2:

1. μ is a descendant of λ_{i-1} .
2. μ is equal to λ_{i-1} in its past part, up to and including p_{i-1} .

Then, let P_i^2 be the set of those $\mu \in P_i^1$ that satisfy condition 3:

3. For any $\kappa \in P_i^1$, $L_\mu - 2n_\mu(i) \leq L_\kappa - 2n_\kappa(i)$.

Finally, let P_i^3 be the set of those $\mu \in P_i^2$ that satisfy condition 4:

4. For any $\kappa \in P_i^2$, $n_\mu(i) \geq n_\kappa(i)$.

The new path λ_i may be selected arbitrarily among the members of P_i^3 .

We note that the update rule will tend to produce loop-less paths (in the sense that no segment occurs twice), since condition 3 encourages short-circuiting loops to reduce the length. Indeed, the future part of the path will never contain an entire loop. However, a loop that the Angel is already inside cannot be short-circuited, due to condition 2. Thus a major concern in the proof of the strategy will be to narrow down the circumstances in which such a loop may arise.

While P_i^1 , P_i^2 and P_i^3 may easily be infinite, the number of border curves that the Angel needs to examine just to find a member of P_i^3 is limited by two considerations. Firstly, $\lambda_{i-1} \in P_i^1$, and $n_\mu(i)$ cannot exceed i for any μ . So a border curve that satisfies condition 3 cannot be longer than $L_{\lambda_{i-1}} - 2n_{\lambda_{i-1}}(i) + 2i$. Secondly, there is no point in letting λ_i differ from λ_0 north of the northernmost square so far eaten by the Devil, as this cannot yield a border curve that is shorter nor has more eaten squares in its left set than a similar curve that equals λ_0 north of that square. Since the number of border curves that can be created by replacing a given finite section of λ_0 by a section of bounded length, is finite, updating the path can be achieved with a finite computation, which is good news for the Angel.

The fact that $\lambda_{i-1} \in P_i^1$ also guarantees that P_i^2 and P_i^3 cannot be empty. Thus the update rule can always be followed successfully.

Lemma 5. *If i and j are turns with $j > i$, then $L_j - 2n_j(j) \leq L_i - 2n_i(i)$.*

Proof. Since $\lambda_{i-1} \in P_i^1$ and $\lambda_i \in P_i^2$, condition 3 yields

$$L_i - 2n_i(i) \leq L_{i-1} - 2n_{i-1}(i). \quad (1)$$

After turn $i - 1$, the Devil cannot produce any new Eaten squares in the left set of λ_{i-1} , so

$$n_{i-1}(i) = n_{i-1}(i - 1). \quad (2)$$

Combining (1) and (2), we get $L_i - 2n_i(i) \leq L_{i-1} - 2n_{i-1}(i - 1)$, and a simple induction yields the lemma. \square

4. Proof that the Angel Wins

We shall now fulfil our promise to demonstrate that after each update, the right square of essentially every future segment of the path is uneaten, and in fact Free. The only exception is that the right square of the very next path segment future of the perch can be Evaded. Before presenting the formal argument, we summarize it informally.

Consider a future path segment s and its right square q . We first find that q cannot be Eaten, since the update rule would have preferred to extend the path around q to make it Evaded. This extension is possible, as every Eaten square is in the path's right set.

We then suppose that q is Evaded. This implies that s occurs twice in the path, forming a loop. This loop can either lie entirely future of the perch, or it can include the perch, thus enclosing the Angel. The former case is not possible because the update rule would simply short-circuit the loop. In the latter case, we go back in time to when the Angel was about to enter the loop-to-be. We find that, but for the exception mentioned above, the update rule does not allow the Angel to enter the loop at all, but prefers a path that evades the entire region that would contain the loop. This arises because in order to be able to close the loop later, the Devil must have so many Eaten squares already in place at this earlier time that the Angel is scared off.

The argument does not entirely exclude the possibility of a loop being formed, but it shows that if this happens, then the Angel is perched at the very end of the loop, immediately before the repeated segment, and will jump out of the loop in the same turn.

Lemma 6. *Let s be a segment of λ_j future of p_{j-1} , and let q be the right square of s . Then at the end of turn j , q is not Eaten.*

Proof. By way of contradiction, assume that q is Eaten at the end of turn j . Let μ be the extension of λ_j where s is replaced with the other three borders of q . Then $L_\mu = L_j + 2$ and $n_\mu(j) = n_j(j) + 1$. Using $\lambda_j \in P_j^2$ and $L_\mu - 2n_\mu(j) = L_j - 2n_j(j)$, we easily verify that $\mu \in P_j^2$. But since $\lambda_j \in P_j^3$, condition 4 implies that $n_j(j) \geq n_\mu(j)$, contradicting the above. \square

Lemma 7. *Let s be a segment of λ_j future of p_{j-1} , and let q be the right square of s . Assume that q is Evaded at the end of turn j . Then s is the very next segment of λ_j after p_{j-1} .*

Proof. As a segment on the board, s occurs at least once in λ_j , and both its adjacent squares are in the left set of λ_j . Then Definition 2 implies that it must occur exactly twice in λ_j , in opposite directions. Call the first occurrence s_1 and the second s_2 . Define κ to be the curve

produced from λ_j by erasing the section from s_1 to s_2 , inclusive. By Lemma 4, κ is a descendant of λ_j . Its length is

$$L_\kappa = L_j - l, \quad (3)$$

where $l \geq 2$ is the number of segments in λ_j from s_1 to s_2 , inclusive. By Lemma 3, $V_{\lambda_j} \subseteq V_\kappa$, giving us

$$n_j(j) \leq n_\kappa(j). \quad (4)$$

If s_1 were to lie future of p_{j-1} , we have $\kappa \in P_j^1$. Since $\lambda_j \in P_j^2$, condition 3 would then imply that $L_j - 2n_j(j) \leq L_\kappa - 2n_\kappa(j)$, which contradicts (3) and (4). So we conclude that s_1 is in the past of or coincident with p_{j-1} . It follows that $s = s_2$. Figure 5 exemplifies the situation, where the Devil has managed to close the path around the Angel.

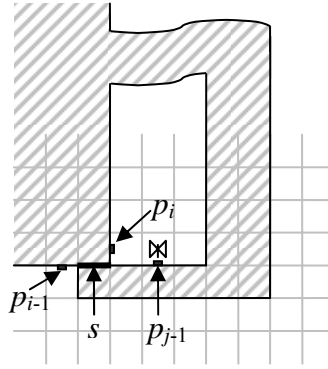


Figure 5. A trapped Angel.

Let i be the turn in which the Angel moves the perch to or beyond s_1 . p_i is at or after s_1 , while p_{j-1} is before s_2 (along λ_j). From p_i to p_{j-1} , the Angel moves the perch at most $l - 2$ segments, at two segments per turn. Thus we have $l - 2 \geq 2(j - i - 1)$, or equivalently

$$2(j - i) \leq l, \quad (5)$$

with equality only if $p_i = s_1$ and p_{j-1} is immediately before s_2 along λ_j . From Lemma 5, we have

$$L_j - 2n_j(j) \leq L_i - 2n_i(i). \quad (6)$$

Between turns i and j , the Devil has eaten only $j - i$ squares, and consequently

$$n_\kappa(j) - n_\kappa(i) \leq j - i. \quad (7)$$

Taking the sum of equations (3) + 2*(4) + (5) + (6) + 2*(7), we get

$$L_\kappa - 2n_\kappa(i) \leq L_i - 2n_i(i). \quad (8)$$

In turn i , we have $\kappa \in P_i^1$, and condition 3 implies

$$L_i - 2n_i(i) \leq L_\kappa - 2n_\kappa(i).$$

So, if the Angel obeyed the update rule in turn i , (8) must hold with equality. But then equations (4) – (7) must also hold with equality. As noted above, equality in (5) yields the conclusion of the lemma. \square

The attentive reader may have noticed that the above argument contains an as yet unjustified assumption, namely the existence of turn i . What if s_1 is so far in the path's past that it actually lies before the Angel's starting perch and has never been passed? We can get around this complication by translating the game in time and space. Suppose that for the first m turns, we require the Devil to effectively pass, by only eating squares in the western half board, while the Angel plods two squares north in each turn. After turn m , the game unfolds as normal. Due to the symmetry of the initial path, this translated game is equivalent to the original game for any m . Thus we can assume, without loss of generality, turn j to be arbitrarily late in the game, and are thereby entitled to posit the existence of turn i .

Theorem 8. *The presented strategy permits the Angel to play indefinitely without ever landing on an eaten square.*

Proof. Let j be any turn and q the right square of p_j . Recall that p_j is two segments future of p_{j-1} along λ_j . At the end of turn j , q must be Free, since by Lemma 6, it cannot be Eaten, and by Lemma 7, it cannot be Evaded. Thus q is uneaten as the Angel lands there in turn j . \square

5. Discussion

Figure 6 shows an example of the situation addressed in Lemma 7, where the Devil almost manages to trap an Angel that uses the presented strategy, by getting the Angel inside a loop in the path. The Devil sets up a clockwise turn in the Angel's path, and when the Angel is perching right before the turn, the Devil eats away the Angel's current square (assuming this is allowed by the rules). The Angel must then update the curve to evade this square, and finds himself in a dead end, but one that is short enough for him to immediately fly out of.

This trivial loop is actually the longest that the Devil can produce. However, we do not devote effort to proving it in this paper.

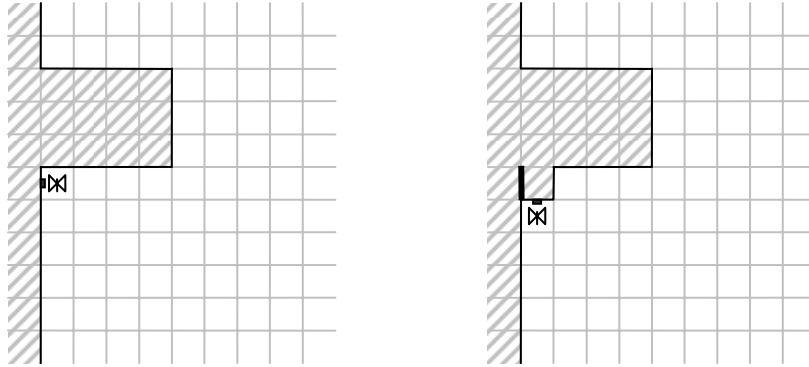


Figure 6. Almost trapping the Angel. Left: The Devil is about to eat the Angel's current square. Right: The Angel has escaped the trap.

The rules as given by Conway do not explicitly state whether or not the Angel is allowed to remain on his current square, i.e. not move at all, if the Devil did not eat away that square on his last turn. Our strategy directs the Angel to stay put if the path makes two consecutive clockwise turns just after the perch. But if this is not permitted, we can amend the strategy to say that in each turn, the Angel moves the perch two segments, *not counting clockwise turns*. This does no damage to our arguments, and ensures that the Angel always flies to a new square on each turn.

Under this amended strategy, we prove that the Devil can never create a loop longer than the trivial one shown above. For if the loop is longer, $p_i \neq p_{j-1}$, and the Angel must traverse at least one clockwise turn of the path on his way from the right square of p_i to the right square of p_{j-1} (defining i and j as in Lemma 7). At this turn, he travels faster than in the original strategy, so equation (5) cannot hold with equality. Consequently, the Angel can never wind up inside a non-trivial loop, and will never have to jump over an eaten square. Thus, this amended strategy is actually a winning strategy for the 2-King. A quick inspection of Figure 6 shows that the 2-King has no problem dealing with the trivial loop either.

As explained in [4], the Devil has some power over the behaviour of an Angel that uses a winning strategy. In particular, the Devil can force the Angel to detour arbitrarily far in any

chosen direction, and to follow a path that winds arbitrarily many times around some point on the board. We conclude by showing that this is indeed the case for the presented strategy.

Assume that the Angel is in a position where the path leads straight north for the next n squares. The Devil can then use his moves before the Angel reaches the n 'th square to build a wall that protrudes east from this square. This wall is 1 square thick and $\lfloor n/2 \rfloor$ squares long.

If the Devil eats the westernmost square in the wall last (and no other eaten squares on the board interfere), the Angel will update the path just as he reaches the wall, and then faces travelling straight east for at least $\lfloor n/2 \rfloor - 2$ squares. The Devil can then start building a shorter wall that protrudes south from the end of the previous wall. When the Angel reaches this wall, the Devil starts building a still shorter wall that protrudes west, and so forth, until some wall becomes so short that the Angel reaches its end before the Devil can build the next one. It is evident that by building this structure for a large enough n , the Devil can cause the Angel both to travel as far as he desires in any direction, and to follow a spiral path that winds as many times as he likes before unwinding again. And as the Angel's path at the start of the game is an infinite straight line, the Devil can indeed do this for an arbitrarily large n .

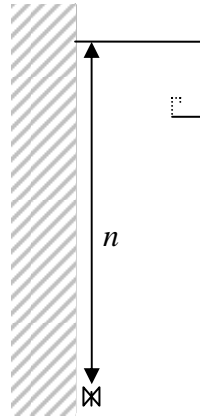


Figure 7. A spiral of walls that the Angel must follow.

6. Acknowledgements

I am grateful to Brian H. Bowditch and the two anonymous referees for helpful comments.

Appendix

This appendix contains the proofs of Lemmas 1, 2 and 4 from section 2.

Proof of Lemma 1. By way of contradiction, assume that U_1 and U_2 are left sets of κ and q is a square that is in U_1 but not in U_2 . Choose t to be a square adjacent to a segment of κ . We can construct a connected sequence of squares starting at q and ending at t . (Here, connected means that any two consecutive squares in the sequence share one border). Let o be the first square in the sequence that is adjacent to a segment of κ . Since no two consecutive squares in the sequence from q to o are separated by a segment of κ , part iv) of Definition 2 implies that all the squares from q to o are in U_1 , and none in U_2 . But by parts ii) and iii), o must be either in both U_1 and U_2 , or in neither, giving a contradiction.

The uniqueness of right sets follows since the complement of a set is unique. \square

Proof of Lemma 2. We need to verify parts i) – vi) of Definition 2 for μ and ν , given that they hold for κ .

Extension. Let s be the segment of κ that is replaced, and r_1 , r_2 and r_3 the segments that replace it. Note that to verify parts i) – iv), we need only to consider the segments s , r_1 , r_2 and r_3 , since all other segments are unchanged with respect to occurrences in κ versus μ and membership of their adjacent squares in the respective left sets. Recall that q is the right square of s and the left square of r_1 , r_2 and r_3 , $q \notin V_\kappa$ and $q \in V_\mu$.

- i) No border of q can occur more than once in κ , since otherwise iii) would imply $q \in V_\kappa$. So these segments occur at most twice in μ .
- ii) Assume that r_1 occurs exactly once in μ , and let o be the right square of r_1 . We need to show that $o \notin V_\mu$ and $q \in V_\mu$. If r_1 occurs just once in μ , it cannot occur in κ . By iv) and $q \notin V_\kappa$, we have $o \notin V_\kappa$ and thus $o \notin V_\mu$. $q \in V_\mu$ is known. The same argument holds for r_2 and r_3 .
By iii) and $q \notin V_\kappa$, s cannot occur twice in κ . Then s cannot occur in μ .
- iii) Assume that r_1 occurs twice in μ . Then it occurs exactly once in κ . By ii), the occurrence in κ must have q as its right square and thus have opposite direction to r_1 . Also by ii), one adjacent square is in V_κ , thus also in V_μ , and the other is $q \in V_\mu$. The same argument holds for r_2 and r_3 . s cannot occur twice in μ .
- iv) s does not occur in μ , but once in κ . By ii), its left square is in V_κ , and thus in V_μ . Its other adjacent square is $q \in V_\mu$. r_1 , r_2 and r_3 occur at least once in μ .
- v) Since L and its complement both are infinite, and V_κ and V_μ differ by a single square, V_μ and its complement must also both be infinite.
- vi) V_κ is connected under separation by κ . Neither r_1 , r_2 nor r_3 introduces additional disconnections in V_κ , since each is adjacent to $q \notin V_\kappa$, so V_κ is connected under separation by μ as well. $V_\mu = V_\kappa \cup \{q\}$ is connected under separation by μ if q is adjacent to some square $o \in V_\kappa$ and the segment between them does not occur in μ . Taking o to be the left square of s , we see that this is the case.

Contraction. Let s be the segment that occurs twice in κ and not in v . To verify parts i) – iv), we need only consider s , since all other segments are unchanged by the contraction.

- i) This holds since s does not occur in v .
- ii) This holds since s does not occur in v .
- iii) This holds since s does not occur in v .
- iv) Since s occurs twice in κ , we know that both adjacent squares are in $V_v = V_\kappa$.
- v) This is satisfied since $V_v = V_\kappa$.
- vi) Every segment of v is also present in κ . Since $V_v = V_\kappa$ is connected under separation by κ , it must also be connected under separation by v . \square

Proof of Lemma 4. We have to show that there exists a finite sequence of extensions and contractions that erases the specified section from κ . We shall construct such a sequence by first using extensions to remove all segments that occur exactly once in the section, and then remove all segments occurring twice by contractions.

Call the first and second occurrence of s in κ s_1 and s_2 , respectively. Let q_1 and q_2 be the squares adjacent to s . From part vi) of Definition 2, we know that we can find a sequence of squares in V_κ that connects q_1 to q_2 without crossing κ , and furthermore, we can choose it such that no square is repeated. Since q_1 and q_2 are adjacent, the squares in this sequence form a closed boundary, broken by no segment in κ except s . Let I be the finite part of the board enclosed by this boundary. κ can have only a finite number of segments in I , and enters or exits I exactly twice, at s ; therefore I must contain exactly the section of κ between s_1 and s_2 .

We now modify the border curve, starting with κ , by choosing a segment in I that occurs exactly once and using an extension to replace this segment. This operation is repeated until there are no more segments in I that occur exactly once. Each extension adds one square in I

to the curve's left set, and since I contains only a finite number of squares, we must stop after a finite number of extensions. Thus we arrive at a curve, μ , where every segment between s_1 and s_2 occurs twice between s_1 and s_2 . This property must continue to hold through the contractions that will follow, as contractions cannot introduce a segment that occurs exactly once.

Let r_1 and r_2 be a pair of segments in μ between s_1 and s_2 that are occurrences of the same segment on the board, and choose them such that the distance from r_1 to r_2 along μ is minimal over all such pairs. We claim that r_1 and r_2 are consecutive segments of μ . For if not, we can repeat the construction of I above and find a part of the board that contains exactly the section of μ between r_1 and r_2 . Any segment in this section occurs twice between r_1 and r_2 , contradicting the distance minimality of r_1 and r_2 .

We now further modify the border curve, starting with μ , by choosing a pair of consecutive segments between s_1 and s_2 that are occurrences of the same segment on the board and removing these by contraction. By the reasoning above, such a pair exists as long as s_1 and s_2 are not consecutive. We repeat this operation until s_1 and s_2 are consecutive. A final contraction to remove s_1 and s_2 completes the construction of v as a descendant of κ . \square

References

- [1] E.R. Berlekamp, J.H. Conway, R.K. Guy, *Winning Ways for your Mathematical Plays*, volume 2: *Games in Particular*, Academic Press, 1982.
- [2] B. Bollobás, I. Leader, *The Angel and the Devil in Three Dimensions*, *Journal of Combinatorial Theory, Series A* 113 (2006) 176-184.
- [3] B.H. Bowditch, *The Angel Game in the Plane*, to appear in *Combinatorics, Probability and Computing*
- [4] J.H. Conway, *The Angel Problem*, in: R. Nowakowski (editor) *Games of No Chance*, volume 29 of *MSRI Publications*, 1996, pp. 3 – 12.
- [5] P. Gács, *The Angel Wins*, <http://www.cs.bu.edu/~gacs/papers/angel.pdf>
- [6] M. Kutz, *The Angel Problem*, Chapter 1 of PhD thesis, Freie Universität Berlin, 2004.
- [7] M. Kutz, *Conway's Angel in Three Dimensions*, *Theoretical Computer Science*, 349 No. 3 (2005) 443-451.
- [8] M. Kutz, A. Pór, *Angel, Devil, and King*, *Proceedings of COCOON 2005*, 925-934.
- [9] A. Máthé, *The Angel of Power 2 Wins*, to appear in *Combinatorics, Probability and Computing*