

The Angel Problem

JOHN H. CONWAY

ABSTRACT. Can the Devil, who removes one square per move from an infinite chessboard, strand the Angel, who can jump up to 1000 squares per move? It seems unlikely, but the answer is unknown. Andreas Blass and I have proved that the Devil *can* strand an Angel who's handicapped in one of several ways. I end with a challenge for the solution the general problem.

1. Introduction

The Angel and the Devil play their game on an infinite chessboard, with one square for each ordered pair of integers (x, y) . On his turn, the Devil may eat any square of the board whatsoever; this square is then no longer available to the Angel. The Angel is a “chess piece” that can move to any uneaten square (X, Y) that is at most 1000 king’s moves away from its present position (x, y) —in other words, for which $|X - x|$ and $|Y - y|$ are at most 1000. Angels have wings, so that it does not matter if any intervening squares have already been eaten.

The Devil wins if he can strand the Angel, that is, surround him by a moat of eaten squares of width at least 1000. The Angel wins just if he can continue to move forever.

What we have described is more precisely called an *Angel of power 1000*. The *Angel Problem* is this:

Determine whether an Angel of some power can defeat the Devil.

Berlekamp showed that the Devil can beat an Angel of power one (a chess King) on any board of size at least 32×33 . However, it seems that it is impossible for the Devil to beat a Knight, and that would imply that an Angel of power two (which is considerably stronger than a Knight) will win. But can you prove it?

Well, nobody’s managed to prove it yet, even when we make it much easier by making the Angel have power 1000 or some larger number. The main difficulty seems to be that the Devil cannot ever make a mistake: once he’s made some moves, no matter how foolish, he is in a strictly better position than he was at the start of the game, since the squares he’s eaten can only hinder the Angel.

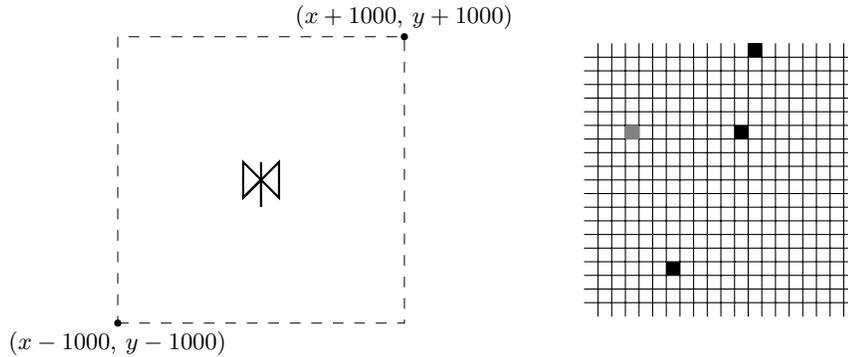


Figure 1. Left: An Angel about to fly; he can fly to any point in the square. Right: A Devil eating a square; three have already been eaten.

2. What's Wrong with Some Potential Strategies

Most people who think about this problem produce trial strategies for the Angel of the following general type. They advise the Angel to compute a certain “potential function” that depends on the eaten squares, and then to make a move that minimizes (or maximizes) this function.

The Devil usually finds it easy to beat such strategies, by finding some feature to which this function is too sensitive. Suppose, for instance, that the function is very sensitive to eaten squares close to the Angel. Then the Devil should build a horseshoe-shaped trap a light-year across, far to the north of the starting position. This will, of course, cause the Angel to fly southwards as fast as possible, but no matter. When the trap is set up, the Devil just goads the Angel into it by repeatedly eating the square just to the south of him. The Angel feels the pain of these “jabs at his behind” so strongly that he blindly stumbles right into the trap!

If the Angel switches to a new function that places great emphasis on eaten squares that are very far away, the Devil might build a mile-wide horseshoe and then scare the Angel into it by eating just a single square a megaparsec to the south of him.

The exact details of the Devil's strategy will of course depend on exact structure of the Angel's potential function. It seems to be impossible to get the function just right, and if it isn't just right, the Devil builds a trap that the potential function is insensitive to, and then steers the Angel into it using something to which the function is over-sensitive! A friend of mine, after hearing my talk of horseshoe-shaped traps, dreamed up a wonderful potential function that cleverly searched for horseshoe-shaped traps of all sizes and avoided them. I defeated it trivially by having the Devil use a modest little horseshoe that easily frightened the Angel into a trap of another shape!

3. Fools Rush On Where Angels Fear to Tread

About twenty years ago in Angell Hall, Ann Arbor, Andreas Blass and I spent some time with the Angel. It is reasonable to think that the Angel should be able to win by moving northward as fast as he can, combined with occasional zigzags to the east or west to avoid any obvious traps. So Blass and I made the following definition:

DEFINITION. A *Fool* is an Angel who is required always strictly to increase his y coordinate. So a Fool can make precisely those Angel moves from (x, y) to (X, Y) for which $Y > y$.

THEOREM 3.1. *The Devil can catch a Fool.*

PROOF. If the Fool is ever at some point P , he will be at all subsequent times in the “upward cone” from P , whose boundary is defined by the two upward rays of slope $\pm \frac{1}{1000}$ through P . Then we counsel the Devil to act as follows (Figure 2): he should truncate this cone by a horizontal line AB at a very large height H above the Fool’s starting position, and use his first few moves to eat one out of every M squares along AB , where M is chosen so that this task will be comfortably finished when the Angel reaches a point Q on the halfway line that’s distant $\frac{1}{2}H$ below AB (we’ll arrange H to be exactly divisible by a large power of two).

At subsequent times, the Devil knows that the Fool will be safely ensconced in the smaller cone QCD , where CD is a subinterval of AB of exactly half its length, and for the next few of his moves, he should eat the second one of every M squares along the segment CD . He will have finished this by the time the Fool reaches a point R on the horizontal line $\frac{1}{4}H$ below AB . At later times, the

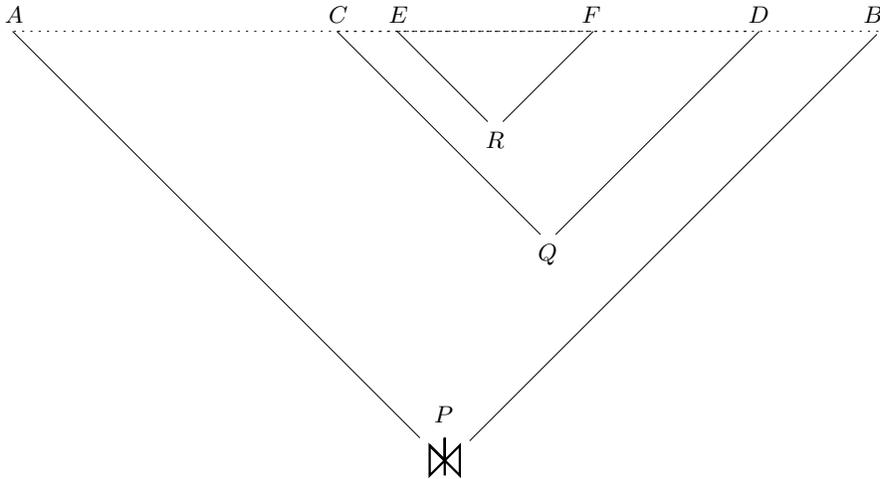


Figure 2. A Fool travelling north with the Devil eating along AB .

Fool will be trapped inside a still smaller cone REF , with $EF = \frac{1}{2}CD = \frac{1}{2}AB$, and the Devil should proceed to eat the third one of every M squares along the segment EF of AB .

If he proceeds in this way, then by the time the Fool reaches the horizontal line at distance $H' = 2^{-M}H$ below AB , the Devil will have eaten every square of the subsegment of AB that might still be reached by the Fool. The Devil should then continue, by eating the first out of every M squares on the segment $A'B'$ just below this one, a task which will be finished before the Fool reaches the horizontal line distant $\frac{1}{2}H'$ below $A'B'$, when he should start eating the second of every M squares on the portion $C'D'$ of $A'B'$ that is still accessible, and so on. We see that if we take H of the form 1000×2^N , where $N > 1000M$, then before the Fool crosses the horizontal line that is 1000 units below AB , the Devil will have eaten all squares between this line and AB that the Fool might reach, and so the Fool will be unable to move. \square

4. Lax Fools and Relaxed Fools

The kind of Fool we've been considering up to now is more precisely called a Plain Fool, since he may make only those Angel moves for which Y is strictly greater than y . Blass and I also showed how the Devil can beat some not-quite-so foolish Fools.

DEFINITION. A *Lax Fool* is a piece that can make precisely those Angel moves for which $Y \geq y$.

THEOREM 4.1. *The Devil can catch a Lax Fool.*

PROOF. The Devil uses his odd-numbered moves to convert the Lax Fool into a Plain Fool (of a much higher power). He chooses two squares L and R a suitable distance D to the left and right of the Fool's starting position, and eats inward alternately from L and R , so long as the Fool stays on the starting line (Figure 3).

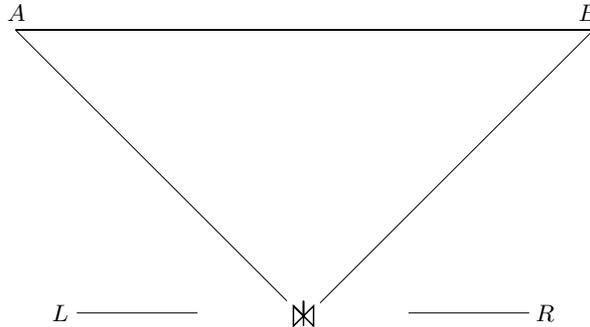


Figure 3. A Lax Fool travelling north.

Whenever the Fool makes an upwards move, the Devil takes the horizontal line through his new position as a new starting line.

Suppose that the Fool stays on the starting line for quite some time. Then he can use the four moves he has between two consecutive bites of the Devil at the left end of this line to move at most 4000 places. So, if we take $D = 4,000,000$, the Devil will have eaten 1000 consecutive squares at the left end of this line before the Fool can reach them. We now see that the Fool can stay on the starting line for at most 8,000,000 moves, since in 8,000,001 moves the Devil can eat every square between L and R .

If the Devil adopts this strategy, then the Fool must strictly increase his y coordinate at least once in every 8,000,000 moves. If we look at him only at these instants, we see him behaving like a Plain Fool of power 8,000,000,000, since each of his moves can have carried him at most a 1000 places. So if the Devil uses the strategy that catches an 8,000,000,000-power Plain Fool on his even-numbered moves, he'll safely trap the 1000-power Lax Fool. \square

Suppose now that the Angel does occasionally allow himself to decrease his y -coordinate, but promises never to make any sequence of moves that decreases it by more than a million.

DEFINITION. A *Relaxed Fool* (of laxity 1,000,000) is an Angel who promises that if he's ever at a position (X, Y) , he will never be later at a position (x, y) with $y > Y - 1,000,000$.

THEOREM 4.2. *The Devil can catch a Relaxed Fool (of any laxity).*

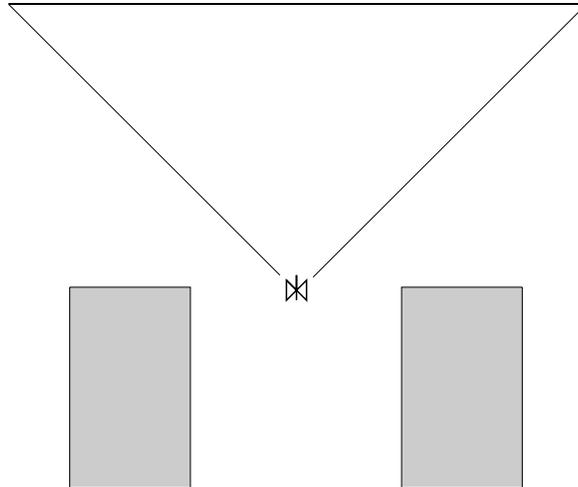


Figure 4. Funnelling the Relaxed Fool.

PROOF. This time, the Devil eats away sunken caissons of depth 1,000,000 at a suitable distance to the Left and Right of the Fool (see Figure 4). \square

5. Variations on the Problem

Because our problem was inspired by Berlekamp's proof that the Devil can catch a chess King on all large enough boards, we've been using the chess King metric, in which the unit balls are squares aligned with the axes. However, the arguments that one meets early on in topology courses show that (provided one is allowed to change the number 1000) this really doesn't matter. It makes no difference to our problem if we suppose that the Angel is positioned at any time on an arbitrary point of the Euclidean plane, and can fly in one move to any other point at most 1000 units away, while the Devil may eat out all points of an arbitrary closed unit disc in any one of his moves.

In a similar way, we can prove that our problem is invariant under certain distortions. This happens for instance when the distorted image of every ball of radius R contains an undistorted ball of radius $\frac{1}{2}R$, and is contained in one of radius $2R$.

Suppose, for example, that the Angel promises always to stay inside the circular sector of Figure 5, and also that he will always increase his r -coordinate (his distance from the vertex O of this sector). Then the Devil can catch him, by distorting this sector into the triangle of the same figure, in such a way that the r -coordinate turns into the y -coordinate. But the transformation turns the Angel into a Plain Fool.

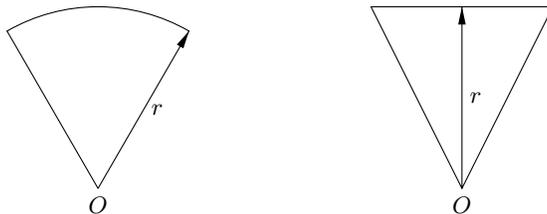


Figure 5. The transformation of regions.

6. Kaleidoscopes Catch Out-and-Out Fools.

DEFINITION. An *Out-and-Out Fool* is an Angel who promises always to increase his distance from the origin.

THEOREM 6.1. *The Devil can catch an Out-and-Out Fool.*

PROOF. The Devil should divide the plane into $2N$ sectors radiating from the origin, and pretend that the walls of these sectors are the mirrors of a kaleidoscope, as in Figure 6. There will then usually be $2N$ images of the Fool, one

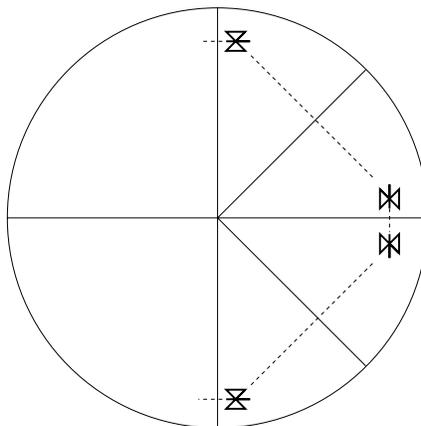


Figure 6. The kaleidoscope strategy.

in each sector. The Devil correspondingly distributes his moves between $2N$ demons, bidding each one of these to take care of the Fool's image in his own sector, which behaves like a Plain Fool of power $2000N$ confined as in Figure 5, and so can be caught by a slightly distorted Plain Fool strategy. \square

Of course, the Devil can also catch the Relaxed Out-and-Out Fool who merely promises never to make any sequence of moves that reduces his distance from the origin by more than 1,000,000 units.

7. An Extremely Diverting Strategy

THEOREM 7.1. *There is a strategy for the Devil that, when used, has 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 at the latter of which the Angel will be at least D units nearer to P than he was at the former.*

PROOF. It suffices to restrict P to the integral points of the plane, and the distances D to the integers. So we need only deal with countably many pairs (P, D) , say (P_0, D_0) , (P_1, D_1) , (P_2, D_2) , \dots . The Devil uses his odd-numbered moves to defeat an Angel that behaves like a Relaxed Out-and-Out Fool who has promised never to move more than D_0 units closer to P_0 , his moves at twice odd numbers to catch one who promises never to move more than D_1 units closer to P_1 , his moves at four times odd numbers to catch one who promises never to move more than D_2 units closer to P_2 , and so on. Any sequence of moves by the real Angel that does not satisfy the conclusion of the theorem will lead to him being caught. \square

It can be shown that, if the Devil employs our diverting strategy, there will be some very wiggly points on the Angel's trajectory, in the sense that he must cross any line through such a point at least 1000 times—no matter how he plays.

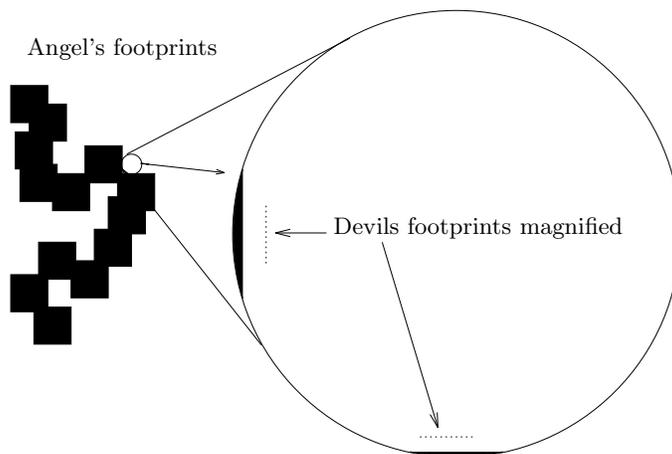


Figure 7. The Angel's large and the Devil's small footprints.

8. The Angel Is His Own Worst Enemy

The Devil only eats one square for each move of the Angel: we shall show that there is a sense in which the Angel himself eats millions!

THEOREM 8.1. *It does no damage to the Angel's prospects if we require him after each move to eat away every square that he could have landed on at that move, but didn't.*

PROOF. Consider the 2001-by-2001 array of squares that is accessible to the Angel from his starting position. We first show that the Devil, without harming his own prospects, can ensure that the Angel only returns to this array a bounded number of times.

If not, there will be indefinitely many times at which the Angel lands on some square of this array that he has already visited previously. At each such time, the Devil essentially has a free move, and he can use these moves to eat away a moat of width 1000 all around the array.

Now suppose that the Angel has any winning strategy whatsoever, say W . Then we shall convert W into another strategy that does not require him to land on any square on which he could have landed at a previous move. When he is at any point P , the Angel considers all the circumstances under which his winning strategy W causes him to repeatedly reenter the 2001-by-2001 array of squares accessible from P . He imagines himself in one of those circumstances that makes him re-enter this array the maximal number of times, and then "short-circuits" W by making right now the moves he would make on and after the last of these times. (He is in a better position now than he would be then, because the Devil has eaten fewer squares.) \square

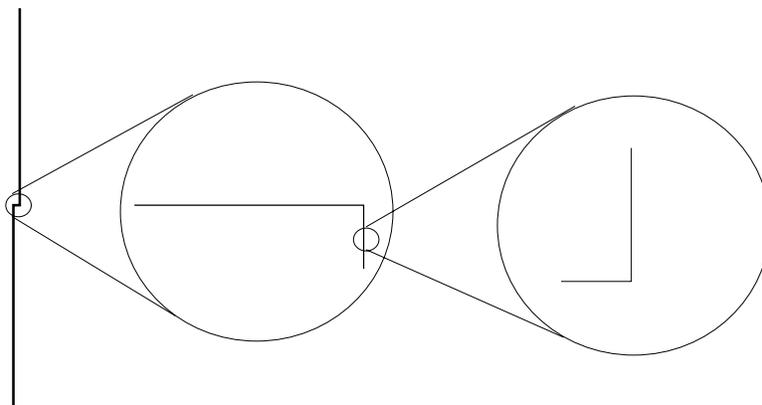


Figure 8. The straight and narrow path.

So the Angel is “really” burning out the path of all points within 1000 of each of his landing points as he moves. The Devil’s consumption seems meager by comparison (see Figure 7).

9. Can the Angel Really Win?

If you’ve got this far, you’ve probably begun to wonder whether it’s really quite safe to trust your fortune to the Angel. But please consider very carefully before selling your soul to the Devil! The Blass–Conway “Diverting Strategy” does indeed force the Angel to travel arbitrarily large distances in directions he might not have wanted to. But it only does so at unpredictable times in the middle of much larger journeys. So for example, at some stage in a megaparsec’s journey north we can divert the Angel (say to the east) for a light-year or so; in the course of this smaller journey we might make him move south for an astronomical unit, with a mile’s westward wandering in that, inside which we might divert him a foot or so to the north, and so on. If he moves like this, then (see Figure 8) we have indeed produced a point about which his path spirals, but the spiral is a very loose one indeed, and an observer at a suitably large scale would only see a very slight wobble on an almost entirely straight Northward path.

So, although we occasionally meet some unbelievers, and even after producing our Diverting Strategy, Blass and I remain firmly on the side of the Angel.

10. Kornering the Devil?

So does Tom Korner, who thinks the Devil might be defeated by the following kind of “military strategy”. The Angel seconds from their places in Heaven a large number of military officers. He then puts a Private in charge of each individual square, and groups these by hundreds into 10-by-10 blocks in charge

of Corporals, which are again gathered by hundreds into 100-by-100 blocks in charge of Majors, 1000-by-1000 blocks in charge of Colonels, and so on.

Each officer is required to make a report to his immediate superior on the status of the square array he is responsible for. He does this on the basis of the similar information he has received from his subordinates. He might for instance classify his array as Safe, Usable, or Destroyed, according to certain rules; perhaps it is Safe only if at least ninety of its immediate sub-arrays are safe, and at most five have been destroyed.

An officer who has just reported that his region is Usable may be immediately telephoned by his superior with a message that a Very Important Person is now entering such-and-such a square on his south boundary and needs to be safely transported to some point on his north or east boundary. In this case he is to plot a suitable route and telephone his subordinates with similar requests. An officer of the n -th rank whose region is Safe may face a stronger demand: perhaps he is only told that the VIP might arrive at any time in the next 10^n time units.

Korner thinks that perhaps the rules can be so designed that these orders can always be safely executed, and that at any time all officers of sufficiently high rank will make Safe reports, so that on a large enough scale the Angel can move completely freely. So do I, but neither of us has been able to turn these ideas into an explicit strategy for the Angel. All we can say is that a “military strategy” of this type could allow for Blass–Conway diversionary tactics, and is not so easy to hoodwink as the potential-function strategies, and so there might well be one that solves our problem.

Join the Fun!

This problem has been alive too long! So I offer \$100 for a proof that a sufficiently high-powered Angel can win, and \$1000 for a proof that the Devil can trap an Angel of any finite power.

JOHN H. CONWAY
JOHN H. CONWAY
DEPARTMENT OF MATHEMATICS
PRINCETON UNIVERSITY
FINE HALL
WASHINGTON ROAD
PRINCETON, NJ 08544