The angel game in the plane

Brian H. Bowditch
School of Mathematics, University of Southampton,
Highfield, Southampton SO17 1BJ, Great Britain.
http://www.maths.soton.ac.uk/staff/Bowditch/

[First draft: August 2006. Revised: September 2006]

0. Introduction.

The angel game, as described in [BeCG], has two players the "angel" and the "devil" who play alternately on the 2-dimensional integer lattice, $\mathbf{Z}^2$. We refer to lattice points as *vertices*, and write $o = (0,0)$ for the origin. The angel has a certain fixed *power* p. (We refer to it as a p*-angel*.) It starts the game at the origin, and at each play moves to another vertex so that the change in each coordinate is at most p in absolute value. On its turn, the devil can "block" any vertex other than that currently occupied by the angel. Once a vertex has been blocked, it remains blocked forever. The angel is not subsequently allowed to visit any blocked vertex. The aim of the devil is to trap the angel so that its only legal move is to remain where it is. (Of course, it would be sufficient to trap it within some bounded set.) Since the game is infinite, we speak of the angel as "escaping" if it never loses, i.e. is never trapped. (There is a finite version played on a square board, where the angel wins if it reaches the edge of the board starting from the centre.)

In [BeCG] it was asked if there is some $p \in \mathbf{N}$ such that the angel of power p can always escape. We shall show that, in fact, the 4-angel has a computable winning strategy. A formal statement for that case is as follows.

Given $n \in \mathbf{N}$, write $I(n) = [-n, n] \cap \mathbf{Z}$ and $W(n) = (I(n))^2$. We write $\sigma(n)$ for the location of the angel at time n, and $\Delta(n)$ for the set of all vertices that have been blocked at time n. Thus, $\sigma(0) = o$, $\sigma(n) \notin \Delta(n)$, and $\sigma(n+1) - \sigma(n) \in W(p)$. We start with $\Delta(0) = \emptyset$, and $\Delta(n+1)$ is obtained from $\Delta(n)$ by adding a single vertex other than $\sigma(n)$. Thus $|\Delta(n)| = n$.

Theorem 0 : *There is a computable function ϕ, which takes as input $n \in \mathbf{N}$ and $A \subseteq W(4n)$, and outputs some $\phi(n, A) \in W(4)$, such that if at time n the angel (of power 4) moves from $\sigma(n)$ to $\sigma(n+1) = \sigma(n) + \phi(n, W(4n) \cap (\Delta(n) - \sigma(n))$, then the angel is never trapped.*

In particular, at any given time, n, the angel need only take into account those vertices that have been blocked up to that point in the square, $\sigma(n) + W(4n)$, centred at its current location (which we have translated back to the origin in the above statement). While the strategy is computable, it is superexponential in n, and thus not a very practical algorithm.

In fact, we will present the argument for a 5-angel which is conceptually a little simpler. In this case, in Theorem 0, $W(4n)$ can be replaced by $W(3n)$. The case of the 4-angel only calls for slight modification as we discuss in Section 1.

It seems quite likely that a variation could also cope with a 3-angel, though this adds some technical complications. We give a brief discussion of how this might work at the end of Section 3.

The game that pits a 1-angel (or "king") against the devil, also known as "quadraphage", has been attributed to David Silverman [S] and Richard Epstein [E], and is discussed in [G]. In [BeCG] it is shown that the devil has a winning strategy against the 1-angel, indeed on any finite square board bigger than 35×35 (see also [KuP]). They observe that it was not known if there exists k such that k-angel can always excape on an infinite board. They state that Körner had shown that for sufficiently large dimension d, in the same game played in $\mathbf{Z}^d$, a sufficiently powerful angel could escape. However, it seems that no account of this has been published. The 2-dimensional game was further publicised by Conway. Some partial results are given in [C]. Winning strategies for the angel in dimension $d = 3$ have recently been described independently by Kutz [Ku] and by Bollobás and Leader [BoL], the former, for an angel of power 13. Though the details are different, both use a strategy involving iterated interpolation. There are reasons to suppose that such a strategy cannot work in dimension $d = 2$, at least without some significant modification (see the discussion in [Ku]). The argument we give here is somewhat different. Some related games, for example, where the devil only blocks a vertex for some specified time, are discussed in [BoL].

To summarise our stategy, we mention an earlier general observation of Conway. If the angel can win, then it can win without ever returning any vertex previously visited. The basic idea behind this is that, if the angel is following some particular winning strategy, and notices that at some point that some hypothetical future sequence of devil moves would force it to return to its current location, then it simply pretends that this sequence has already been played, and proceeds from that new position instead. Turning this idea around, one could modify the rules to explicitly allow the angel to return to any vertex it had previously visited. This makes the game easier. The angel need no longer worry about traps, and one can give a relatively simple strategy. Basically, the angel heads north whenever it can, but skirts around any obstacle it encounters, keeping it on its left hand side. The angel may well have to return to previously visited vertices, but a principle similar to that Conway allows it to shortcut any such loops. In so doing, it will also be playing according to the rules of the original game.

We note that two other entirely independent solutions to the angel problem have been produced, by Kloster [Kl] and by Máthé [M]. Both of these give winning strategies for the angel of power 2.

I am grateful to Imre Leader, Oddvar Kloster and the referee for their helpful comments on earlier drafts of this paper.

1. Some variations.

Rather than tackle the angel game directly as described, we deal first with a couple of variations. An angel strategy for these will give rise to an angel strategy for the original. The first two observations are fairly trivial, and can be thought of as just as expository tools.

First note that we allow the devil to pass (i.e. not play anywhere on its turn). This cannot be to the devil's advantage, but the option is introduced to simplify later discussions.

Conversely, the angel can employ the services of a "demon" who, on the angel's turn, may block any set of vertices at the angel's behest (not including the vertex currently occupied by the angel). This seems to play into the hands of the devil, though the policy may serve to remind the angel of any devil traps the angel may have noticed, thereby simplifying the description of its subsequent strategy.

Slightly less trivially, we can also consider a devil of *weakness* w (or w-*devil*) for any $w \in \mathbf{N}$. Such a devil must visit a vertex w times (on separate plays) before it becomes blocked. If $w > 1$, such a devil may visit the vertex currently occupied by the angel, provided it has been there at most $w-2$ times before. Thus, at any given moment, each vertex has a weight associated to it which counts the number of devil visits there.

We note:

Lemma 1.1 : *If a 1-angel can always escape from a k^2-devil, then a $(2k-1)$-angel can always escape from in against a 1-devil.*

Proof : Divide the lattice $\mathbf{Z}^2$ into square arrays of $k \times k$ vertices, which become the vertices of a new lattice. We regard all the vertices in a given square as equivalent. Such a square is only blocked if all the vertices there are blocked. Thus a 1-devil in the old game acts as a k^2-devil in the new. Any 1-angel move in the new game gives a $(2k-1)$-angel move in the old. $\diamondsuit$

To simplify our account, we will further restrict the angel moves so that one coordinate is always fixed while the other is either incremented or decremented by 1, i.e. it takes exactly one step in one of the four "cardinal" directions. Such a move is sometimes called a *duke* move. We therefore arrive at our first variant game:

Game 1 : The angel moves as a duke against a devil of weakness w, for some fixed $w \in \mathbf{N}$. The goal of the angel is to avoid being trapped.

As before, the angel can employ a demon, and we can allow the devil to pass.

We will show that the angel wins Game 1. To make our analysis work, we will need $w \geq 5$. Since $2^2 < 5 < 3^2$, this means taking $k = 3$, and so corresponds to an angel of power $p = 2.3 - 1 = 5$ in the original game. Given this, we may as well take $w = 9$, which gives slightly better inequalities.

In fact, $w = 5$ can also cope with a 4-angel. In other words, if a duke can always escape from a 5-devil, then a 4-angel can escape from a 1-devil. To see this, instead of cutting the integer lattice into square arrays, we cut it into crosses. In other words, take the set $\{(0,0),(0,1),(0,-1),(1,0),(-1,0)\}$, and all its translates under the the group generated by $[(x,y) \mapsto (x+2,y+1)]$ and $[(x,y) \mapsto (x-1,y+2)]$. Now duke moves turn into 4-angel moves. However, since the square array is conceptually simpler, we will stick with that for the purposes of exposition.

We now move on to a further modification, whose rules allow the angel to backtrack out of any trap.

Definition : At any given time, a vertex is *safe* if it was visited by the angel before it was visited w times by either the devil or the demon.

Once a vertex is made safe, it remains safe forever. This does not prevent the devil or demon returning there, but such a visit will have no effect on legality of any subsequent angel move in the new game.

We now have:

Game 2 : The rules are as with Game 1, except that now, the angel is allowed to return to any vertex already made safe. The goal of the angel is now more challenging: it must trace out a "circuitous trail" satisfying (L) below.

We postpone the definition of a "circuitous trail" until Section 2. For the moment, we just make the following remarks. In general, a *trial* will mean formally a sequence of vertices with consecutive vertices adjacent. (Though we will frequently view it as a continuous path with the edges filled in.) An *arc* is an injective trail, i.e. no vertex is visited more than once. If a trail starts and finishes at the same point, we refer to it as a *loop*. It will be the case that any subtrail of a circuitous trail is circuitous (so the devil could definitely win at this game if it forced the angel into a non-circuitous trail). Among other properties, any semi-infinite circuitous trail, π, has the form $\pi = \sigma \cup \bigcup_{i=1}^{\infty} \gamma_i$, where σ is a semi-infinite arc, and where γ_i is a loop based at the ith vertex, $\sigma(i)$, of σ. Moreover, $\gamma_i \cap \gamma_j = \emptyset$ if $i \neq j$. (We allow for $\gamma_i = \{\sigma(i)\}$.) Here, the notation $\sigma \cup \bigcup_i \gamma_i$ means that each γ_i is traversed before proceeding to the next vertex of σ. We see that such a decomposition is unique, and refer to σ as the *spine* of π.

Condition (L) gives a metric constraint on the loops.

(L) For each i, the length of γ_i is at most i.

The actual bound, i, is an artifact of the proof. Any computable bound would serve to give some angel strategy in the original game.

We note that the initial segment of π from $\sigma(0)$ to $\sigma(i)$, including the loop, γ_i, lies inside the square $\sigma(i) + W(i)$. (This observation is used in Section 4.)

We aim to show:

Proposition 1.2 : *In Game 2, the angel has a winning strategy (against a 5-devil).* $\diamondsuit$

To get from here back to Game 1, the idea is that the angel employs a "phantom" who plays to the rules of Game 2. The phantom warns the angel of traps by performing loops in some circuitous trail. The angel then shortcuts these loops, following the spine of the phantom's trail. Since the spine is an arc, none of these vertices were made safe by the phantom, and so remain unblocked by the devil and demon. The angel moves are thus in accordance with the rules of Game 1. (This is analogous to the observation of Conway that, mentioned in the introduction, that an angel need never repeat a vertex.) The details of this are described in Section 4.

2. Circuitous trails.

We want to describe the basic properties expected of the trail traced out by the angel in Game 2. We begin with some elementary graph theory.

Let Π be a connected locally finite graph. For us, an *arc* is an injective trail, and a *circuit* is an injective cycle (closed trail). We generally think of a cycle as a cyclically ordered set of vertices, where the basepoint is unimportant. If we want to stress that is has a preferred basepoint, we will refer to it as a *loop*.

Recall that a *cut vertex* is a vertex that separates Π. A subgraph is 2-*vertex connected* if it has no cut vertex. A *block* is a maximal 2-vertex connected subgraph. A *bridge* is an edge that separates Π. Note that a bridge is a block connecting two cut vertices. Two blocks meet, if at all, in a single cut vertex. The blocks are arranged in a treelike manner (more precisely the blocks and cut vertices together form a bipartite tree, where the membership relation is interpreted as adjacency), see, for example, [Bo] for details.

We note:

Lemma 2.1 : *The following are equivalent:*
(1) Each block of Π *is either a bridge or a circuit.*
(2) Each edge of Π *is contained in at most one circuit.*
(3) Any pair of distinct vertices of Π *are separated by some set of at most two edges.* $\diamondsuit$

Definition : We refer to a graph satisfying any of the above as a *cactus*.

In practice it is mainly (1) that we will be using.

We note that any connected subgraph of a cactus is a cactus, and that a graph is a cactus if and only if every finite connected subgraph is.

Definition : We say that two vertices of a graph are *uniquely connected* if there is precisely one arc that connects them.

Now let Γ be any connected locally finite graph, with vertex set $V(\Gamma)$. (In practice, it will be the infinite planar grid.) Suppose that π is a (possibily infinite) trail in Γ. We write Π for its image in Γ. Consider the condition:

(P1) No edge of Π is traversed twice in the same direction by π.

For a trail satisfying (P1), we refer to those edges of Π that are traversed twice (in opposite directions) as *double edges*. The remainder are directed.

Definition : We say that a trail, π, satisfying (P1) is *circuitous* if its image Π is a cactus and every double edge of Π is a bridge.

It is easily verified that:

Lemma 2.2 : *Any subtrail of a circuitous trail is circuitous.* $\diamondsuit$

It is also worth noting that a trail is circuitous if and only if every finite subtrail is.

Definition : The *spine* of a circuitous trail, π, is the union of all directed bridges of Π.

Now since, by assumption, the edges of each circuit of Π are directed, it is easily seen that π must eventually leave any circuit for the last time by the same vertex at which it first entered. It follows that the spine, σ, must be either empty, or an arc. In the latter case, it must have the same endpoints as π.

We also note that the edges of each block must be consistently oriented (to give a directed circuit). It follows easily that:

Lemma 2.3 : *If π is finite, then its endpoints are uniquely connected in* Π. $\diamondsuit$

Indeed, its spine, σ, is the unique arc connecting them. Note that (with the convention of Section 1), we can write $\pi = \sigma \cup \bigcup_i \gamma_i$, where γ_i is a (possibly trivial) loop based at $\sigma(i)$. Indeed, the same is true of a semi-infinite circuitous trail. Note that a loop, γ_i, might pass more than once through its basepoint (in which case the basepoint will be a cut vertex of its image).

The following is needed in Section 4. Here $\vee$ is used to denote concatenation of two trails.

Lemma 2.4 : *Suppose that $\pi = \pi' \vee \rho$ is a finite or semi-infinite circuitous trail where π' and ρ are subtrails concatenated at some vertex x. Let σ and σ' be the spines of π and π' respectively, and suppose that ρ never returns to σ' after leaving x. Then ρ never returns to π', and $\sigma' \subseteq \sigma$.*

Proof : It is enough to show that $x \in \sigma$. We know that x lies in some loop, γ, of π. Let $y \in \sigma$ be its basepoint. Now any trail from x back to the initial vertex must pass via y, and so, in particular, $y \in \sigma'$. Also, the endpoint of σ cannot lie strictly inside a loop of π, and so ρ must also pass by y. Since ρ never returns to σ', it follows that $x = y$. Thus, $x \in \sigma$, and the rest is easy. $\diamondsuit$

Note, in particular, that every loop of π' is also a loop of π.

We now consider specifically the case of the integer lattice, $\mathbf{Z}^2$, in the plane, $\mathbf{R}^2$. This is the vertex set, $V(\Gamma) = \mathbf{Z}^2$, of the the 1-skeleton, Γ, of a square tessellation of the plane. We think of Γ as a regular 4-valent graph.

We shall refer the first and second coordinates as "Easterly" and "Northerly" respectively.

Definition : The *height* of a vertex of Γ is its Northerly coordinate.

Note that any $x \in V(\Gamma)$ has four neighbours to the East, North, West and South, going around the "cardinal directions" in an "anticlockwise" sense. We shall denote these neighbours respectively by $x + \mathrm{E}$, $x + \mathrm{N}$, $x + \mathrm{W}$ and $x + \mathrm{S}$, or simply write E, N, W and S, where the reference point is clear from context. We write $D(x) = \{\mathrm{E}, \mathrm{N}, \mathrm{W}, \mathrm{S}\}$.

If $A \subseteq D(x)$, then the cyclic orientation allows us to speak of *clockwise* and *anticlockwise successors* in A. For example, if $A = \{\mathrm{N}, \mathrm{W}, \mathrm{S}\}$, then N has clockwise successor S and anticlockwise successor W, etc. If A is a singleton, then this is both the clockwise and anticlockwise successor of itself.

We can also speak about relative directions. Suppose we are at x and facing in some cardinal direction, F, which we call Forward. The cardinal directions going anticlockwise are then Forward, Left, Back and Right, or abbreviated to F, L, B and R.

If e is an edge of Γ directed from x to y, we write $e = xy$, and refer to x and y and the *tail* and *head* of e respectively. If we are facing in this direction, then $y = x + \mathrm{F}$.

Suppose that π is a trail in Γ satisfying (P1) above. To save worrying about endpoints, we will assume, for the moment, that π is bi-infinite. Let $\Pi \subseteq \Gamma$ be the image of π. Given $x \in V(\Gamma)$, let $D(x, \pi) = D(x) \cap V(\Pi)$. In other words, those neighbours of x visited at some time by π (not necessarily directly to or from x).

Consider the following condition:

(P2) Each time π arrives at some vertex $x \in V(\Gamma)$ from some $y \in D(x, \pi)$, then it immediately leaves by the clockwise successor of y in $D(x, \pi)$.

Here is a more intuitive description of (P2). Let $\hat{\pi}$ be the trail in $\mathbf{R}^2$ obtained by pushing π slightly to the Left. In particular, $\hat{\pi}$, does not pass through any vertex of $\mathbf{Z}^2$, but rather skirts around it on the left. (For example, if π immediately doubles back at some vertex, x, then $\hat{\pi}$ will make a semicircular clockwise turn around x before returning on the other side.) The condition (P2) is now equivalent to saying that $\hat{\pi}$ is embedded. Indeed it is a boundary component of a small regular neighbourhood of Π in $\mathbf{R}^2$.

Lemma 2.5 : *Any bi-infinite trail satisfying (P1) and (P2) is circuitous.*

Proof : Let π be such a trail, and Π its image. Let C be a complementary region, i.e. a component of $\mathbf{R}^2 \setminus \Pi$. Its boundary determines a trail, γ, in Γ. (Consider the corresponding boundary component of a small regular neighbourhood of Π in $\mathbf{R}^2$.) Of course, a-priori, γ might not be embedded in Π, and π may enter and leave the image of γ several times.) Note that the orientation of $\mathbf{R}^2$ determines "clockwise" and "anticlockwise" directions for γ. If C is bounded, then γ is a closed cycle.

Suppose that e is an edge of γ, and is directed so that C is on its left, i.e. consistent with the anticlockwise orientation of γ. (At this stage, we can allow e to be a double edge.) Now proceeding around γ, applying (P2), we see that all the edges of γ have an anticlockwise direction (or maybe some are also double edges). But it now follows that consecutive edges of γ must also be consecutive in π, since there is nowhere else for π to go, without violating (P2). Since π is not a closed trail, this is a contradiction. From this, we conclude that the boundary of any bounded complementary component is a clockwise cycle. Again, by (P2), we see that its image has no cut vertices, and so it must be a circuit. (Of course, π, might enter and leave this circuit several times.)

In particular, this means that the region to the left of any directed (or double) edge is unbounded. The regions on either side of any double edge are unbounded, and so any double edge is a bridge. (Retrospectively, we see that it was the same component on both sides.)

Suppose that γ is a clockwise circuit bounding a region C. Now each complementary region meeting C in at least one edge is unbounded, and so there can be no arc outside C connecting two distinct vertices of γ. It follows that any connected subgraph strictly containing γ must have a cut vertex. In other words, γ is a block of Π.

This shows that any edge of Π that is not a bridge lies in such a block. In other words, every block of Π is either a bridge or a circuit, so Π is a cactus. $\diamondsuit$

Of course, the hypotheses (P1) and (P2) together are somewhat stronger than being circuitous. What we eventually require of π will be that is it circuitous and satisfies the metric constraint (L) described in Section 1.

In order to deduce (L) we will need some more lemmas about a trail satisfying (P1) and (P2).

First, we give some definitions. Let α be a finite subtrail of π. (The initial discussion makes sense for any trail α in Γ.)

Definition : The *Northward displacement*, $v(\alpha)$, of α is the height of the terminal point minus the height of the initial point.

We call an edge *horizontal* if it is directed either W or E, and vertical otherwise.

The *Southward migration* of α is defined as $m(\alpha) = r(\alpha) + 2s(\alpha)$, where $r(\alpha)$ is the number of horizontal edges in α, and $s(\alpha)$ is the number of edges directed South.

The *length*, $l(\alpha)$, of α is the total number of edges in α.

We record the following simple observation for future reference:

Lemma 2.6 : $v(\alpha) = l(\alpha) - m(\alpha)$. $\diamondsuit$

We borrow the follow terminology, rather liberally, from go. A *liberty* of a directed edge $e = xy$ of π is the set $\{x + \mathrm{L}, y + \mathrm{L}\} \setminus V(\Pi)$ (where $y = x + \mathrm{F}$). If π immediately doubles back along e after reaching y, then we refer to $y + \mathrm{F}$ as a *terminal liberty*. Note that if π satisfies (P2), then this cannot lie in $V(\Pi)$.

Definition : If α is a subtrail of π, then the set of *liberties* of α is the union of the liberties of all edges of α that are directed S, W, or E, together with the terminal liberties of any edge directed S, W or E along which α immediately doubles back.

We write $\Lambda(\alpha)$ for the set of liberties of α. Thus, $\Lambda(\alpha) \cap V(\Pi) = \emptyset$. (Note that the definition of $\Lambda(\alpha)$ depends on the ambient trail, π.) Let $\lambda = \lambda(\alpha) = |\Lambda(\alpha)|$.

Lemma 2.7 : *Suppose that π is a bi-infinite trail satisfying (P1) and (P2) and that α is a finite subtrail beginning with a Northerly edge, from which we either turn right or double back. Then $m(\alpha) \leq 4\lambda(\alpha)$.*

(If we remove the clause about starting on a Northerly edge, we would have to include an additive constant.) The actual bound $4|\Lambda(\alpha)|$ is a bit delicate. However, it is very easy to give some linear bound. Any linear bound would suffice to make our analysis work against a sufficiently weak devil.

Lemma 2.7 is easier to see if we consider a closed cycle, rather than a trail. Since there are an equal number of N and S edges, we can count each just once. Moreover, suppose for the moment we are also counting liberties of Northerly edges. In this case, the inequality becomes $m \leq 4\lambda + 4$. To see this, note that m is just the lotal length of the trail. Moreover each liberty is adjacent, on the left, to at most 4 vertices along our trail. Now counting liberties to the left of a vertex, we get 0, 1, 2 or 3, depending on whether we turned L, F, R or B, at that point. Writing l, f, r and b respectively for the numbers of such vertices, we get $m = l + f + r + b$ and $f + 2r + 3b \leq 4\lambda$. Now $t = l - r - 2b$ is the total left turning of the trail, which for a closed cycle equals 4. Thus $m = (f + 2r + 3b) + 4 \leq 4\lambda + t = 4\lambda + 4$, as claimed. Now, suppose we rule out any liberty that arises only from a Northerly edge (as in the defintion given). We get the same inequality: since going due West from such a liberty, we will come to another liberty before we come a vertex of our trail, which allows us to do a different trade off. We trade one quarter of this new liberty instead of one quarter of the old against the same vertex.

The case of a trail can be dealt with by a similar argument, though the details get a bit more involved. (We need to worry a bit more about Northerly liberties, and total turning.) In fact, when we come to apply the result, we could allow ourselves to include Northerly liberties other than those blocked by the demon. Moreover, it is easy to see that at most half of the vertices in the trail can be Left turns. The latter observation gives rise to a very simply argument to give a multiplicative factor of 8, say, which would allow our proof to work (with slightly worse inequalites) against a devil of weakness 9. For these reasons, we omit the details here.

3. A strategy for Game 2.

We give a strategy for the angel in Game 2, so that it traces out a trail satisfying (P1) and (P2) of Section 2, and so, in particular, is circuitous. Recall that, in Game 2, a vertex is blocked if it is visited by the devil w or blocked by the demon, before being visited by the angel. For our inequalities, we will assume $w \geq 9$, though only $w \geq 5$ is needed for the general argument to work.

Let σ_- be the ray heading directly South of the origin, o, i.e. $\sigma_- = \{0\} \times (-\infty, 0]$. We deem these vertices to be safe, and all of the vertices immediately to the West (i.e.

$\{-1\} \times (-\infty, 0])$ to be blocked by the demon. In other words, we imagine the angel to have traversed σ_- from time $-\infty$, arriving at the origin at time 0. (This is compatible with the algorithm below.)

Of course, this is an artificial construction which serves only to streamline the exposition. In retrospect, we will see that the angel remains in the positive quadrant, $[0, \infty)^2$, so that the safety of σ_- is of no practical use.

Once the devil has made its first move (at time 1), the angel proceeds according to the following algorithm. Let us suppose that the angel has arrived at its current position x, and is facing in the direction of its last move (deemed to be N at the start). It now proceeds as follows:

(A1) If L unblocked, it goes L.

If L is blocked, but F is not, it goes F.

If L and F are both blocked, but R is not, it goes R.

If L, F and R are all blocked, it goes to B.

(A2) If at the end of its move it is at some point y, and has never been strictly further North than y, then it instructs the demon to block $y + \mathrm{W}$ if this point is not already safe or already blocked by the devil.

In (A2) by "strictly further North", we mean any point with strictly greater Northerly coordinate, regardless of its Easterly coordinate. Note that (A1) is telling us to go to the clockwise successor of its previous position, among those adjacent vertices that are not blocked. The final option of doubling back is always available as a last resort in this game. Note that if (A2) is applied, this must have been our first arrival at y, and the move must in fact, have been to the North. It cannot have been to the West, since had we previously visited $y + \mathrm{E}$, then y would have been blocked on our first arrival there.

Let π_+ be the trail traced out by an angel following the above strategy, for some set of devil moves, and let $\pi = \sigma_- \cup \pi_+$, be the bi-infinite trail.

As we have formulated it, we are allowing the angel to play first. To get us started, we note that the first w moves will take us directly North.

The general idea is that our trail, π, will always have blocked vertices to its left. These prevent π from looping around in an anticlockwise sense and rejoining itself from the left. It can still loop around in a clockwise sense, and rejoin itself from the right. When it does, it must immediately turn left, and possibly retrace its steps some distance, before branching out on a new adventure to the left (relative to the direction we are currently facing). While we are heading North, the vertices on the left may have been blocked by the demon, but while we are going in other directions, they must all have been blocked by the devil. Since it takes the devil several visits to block any given vertex, we can quickly overtake the devil, thereby controlling the length of such meanders.

At least, that's the basic idea. Some of the details of the local analysis are a bit messy. We begin with:

Lemma 3.1 : *If the angel turns to the left at some point x, then the diagonal point $x + \mathrm{B} + \mathrm{L}$ was blocked at the time it arrived at x (from $x + \mathrm{B}$).*

Proof : Let $y = x + \mathrm{B}$. Suppose, for contradiction, that $x + \mathrm{B} + \mathrm{L} = y + \mathrm{L}$ was unblocked at the time of our arrival at x. The fact that we did not go to $y + \mathrm{L}$ (which is unblocked) immediately after y means that we have must have just come from there. In other words, we made a left turn at y. But now, by the same argument applied to the previous vertex, we see that we must have turned left there as well. Iterating back in time, we see that we must have been in an infinite anticlockwise circuit around the four points, $y, x, x + \mathrm{L}, y + \mathrm{L}$, again giving a contradiction. $\diamondsuit$

Lemma 3.2 : *Suppose that the angel has arrived for the first time at some point $x \in V(\Gamma)$, and that it has never previously been strictly further North of x. Then the previous move was either E or N. If it was N, then the point $x + \mathrm{W}$ cannot have been safe at the time the angel arrived at x (and so was blocked by the demon on the arrival at x if it was not already blocked by the devil).*

We are allowing the angel to have previously visited other points at the same height as x.

Proof : We cannot have come directly from $x + \mathrm{E}$, otherwise x would have been blocked by the demon on our arrival at $x + \mathrm{E}$. Thus, our previous move was E or N.

Suppose it was N. In other words, we have just come from $x + \mathrm{S}$. Suppose, for contradiction, that $x + \mathrm{W}$ was safe at the time of our arrival at x. Consider the first time at which the angel arrived at $y = x + \mathrm{W}$. We cannot have arrived at y from $y + \mathrm{W}$, or we would have gone immediately on to $y + \mathrm{N}$ or to $x = y + \mathrm{E}$, either way contradicting our assumptions. Thus, we got there from $y + \mathrm{S}$. Now $y + \mathrm{W}$ cannot have been blocked at this time, or again we would have gone on to $y + \mathrm{N}$ or to x. In particular, it was not blocked by the demon on our arrival at y. Thus, we must have been to $z = y + \mathrm{W}$ at some even earlier time. Moreover, we must have made a left turn at y, so Lemma 3.1 tells us that $z + \mathrm{S} = y + \mathrm{S} + \mathrm{W} = y + \mathrm{B} + \mathrm{L}$ must have been blocked at the time of our arrival at y, and hence also at our first arrival time at z. In particular, we cannot have got to z from there, and so we must have got to z first from $z + \mathrm{W}$. But now, since we did not go N from z, $z + \mathrm{N}$ must have been blocked. Thus, we should have proceeded on to y. But we have already observed that we got to y first from $y + \mathrm{S}$, thereby giving a contradiction. $\diamondsuit$

Definition : A move is *progressive* if it takes the angel strictly further North than it has previously been.

An immediate consequence of Lemma 3.2, is that the angel will make a progressive move whenever it is possible, i.e. if N is not blocked and strictly further North than the angel has ever been.

Lemma 3.3 : *π satisfies (P1).*

Proof : If not, let f' be the first edge of π to be superimposed on some previous edge, f, of π oriented in the same direction. We write f_0 for the underlying edge of Π. Let e' and e be the edges of π immediately prior to f' and f respectively. By assumption, these are distinct in Π. Let x be the common vertex. Now π can never have doubled back at x (since there are at least two unblocked exits), and so e, e' and f_0 are all distinct. Now if e' is the clockwise successor of e among $\{e, e', f_0\}$, then we should have followed e immediately by e', or perhaps by the fourth incident edge, if this were also unblocked. Similarly, if e is the clockwise successor of e' among $\{e, e', f_0\}$, we get the same contradiction on swapping e and e'. $\diamondsuit$

Lemma 3.4 : *π satisfies (P2).*

Proof : Let x be a vertex of π. Since all the adjacent vertices of x in Π are unblocked (perhaps made safe by the angel), then, by rule (A1) each exiting edge is a clockwise successor of the corresponding entering edge (among the incident edges of Π). This is precisely (P2). $\diamondsuit$

Lemma 3.5 : *π never enters any vertex directly to the West of the head of any progressive edge (i.e. with the same Northerly coordinate, but strictly smaller Easterly coordinate).*

Proof : If it does, let z be the first vertex of π that lies to the West of the head of some progressive edge of π. Let e be this progressive edge. Let v be a most Northerly point visited by the angel before arriving at z. Among such most Northerly points, we take v to be most Westerly. Thus π leaves v to the East or South. Let α be the subtrail of π from o to v and let β be the subtrail from v to z. Now, v is at least as far North as z, which in turn is as far North as e. Thus, e lies in α. Since π makes a right turn or doubles back at v, in order to get to z, β must cross $\sigma_- \cup \alpha$. In other words, the trail $\hat{\pi}$ must cross itself, in contradiction to (P2). $\diamondsuit$

(In the above, we a formally using a principle of planar separation that we will use again later. Let ρ be the embedded bi-infinite trail obtained from $\hat{\pi}$ by heading strictly North from v rather than following the remainder of $\hat{\pi}$. Now there is a closed trail in the plane that meets ρ only once: the horizonal trail from z to the head of e meets it once, and then we can follow π around from e to v and then on to z without meeting it at all. This gives the contradiction that ρ is non-trivial in locally finite $\mathbf{Z}_2$ homology.)

Corollary 3.6 : *The angel never goes strictly North West of the tail of any progressive edge.*

By "strictly North West", we mean any point with strictly greater Northerly coordinate, and strictly smaller Easterly coordinate.

Proof : Let x be the tail of a progressive edge. By continuing inductively with Lemma 3.5, we see that all subsequent progressive edges must be to North East of x. But any point to strictly North West of x is strictly West of the head of such and edge. $\diamondsuit$

In particular, this tells us that the angel always remains in the Easterly half-plane.

We may strengthen Corollary 3.6 as follows:

Lemma 3.7 : *Suppose that, when the angel first arrives at some point $x \in V(\Gamma)$ it has never been strictly further North. Then it never subsequently goes strictly further North West of x.*

Proof : If it does, let e be the first progressive edge taking us strictly North of x. By Corollary 3.6, it is enough to show that the tail, z, of e does not lie strictly to the West of x.

Suppose that it did (so in particular, $x + \mathrm{N}$ must have been blocked). Let α be the subtrail of π from o to x, and β be the subtrail from x to z. By Lemma 3.2, we must have arrived at x from $x + \mathrm{W}$ or $x + \mathrm{S}$. In the former case, we next go $x + \mathrm{W}$, $x + \mathrm{S}$ or $x + \mathrm{E}$. In the latter case, again by Lemma 3.2, $x + \mathrm{W}$ was not safe at the time of our first arrival at x, and so was blocked by the devil or demon. Since $x + \mathrm{N}$ must also have been blocked, our subsequent move must have been either E or S. We now see, as in the proof of Lemma 3.5, that in order to get from x to z, β must cross α. In other words, $\hat{\pi}$ must cross itself, giving a contradiction. $\diamondsuit$

Continuing with the same line of argument we obtain:

Lemma 3.8 : *Suppose that when the angel visits some $x \in V(\Gamma)$, the vertex $y = x + \mathrm{W}$ is blocked by the demon according to rule (A2). Then the angel never visits $y + \mathrm{W}$ or $y + \mathrm{S}$ (before or after).*

Proof : Note that this must be the first time the angel has visited x, and it has never been strictly further North than x. We first show that it never visited $y + \mathrm{W}$.

Suppose that the angel visits the vertex $z = y + \mathrm{W}$ for the first time before it gets to x. Since it had never been further North, $z + \mathrm{N}$ must have been blocked (see the remark after Lemma 3.2). However, $z + \mathrm{E} = y$ must have been unblocked, since, by assumption it is subsequently blocked by the demon. Moreover, it cannot progress to y, or y would have been safe on our arrival at x. Now by Lemma 3.2, the angel must have arrived at z from $z + \mathrm{W}$ or $z + \mathrm{S}$. But the former is not possible, or it would have progressed immediately to y. Thus it must have arrived from $z + \mathrm{S}$, and so by Lemma 3.2, $z + \mathrm{W}$ would have been blocked by the demon at that time. Again the angel should have progressed to y, giving the same contradiction.

Thus, if z is visited, it must be after our first arrival at x. In this case, let γ be the subtrail of π from o to z. If this never goes strictly North of x, then set $v = x$. If it does, let v be the most Westerly point among the most Northerly points visited by γ (cf. the proof of Lemma 3.5). Either way, we see that π must turn right or double back at v. Now v cuts γ into two subtrails, $\gamma = \alpha \cup \beta$, and $x \in \alpha$. In order to get to z, we see that β would have to cross α, giving a contradiction, as usual.

We now show that it never visits $w = y + \mathrm{S}$. If it did, this point is never blocked. It follows, that immediately before going from $x + \mathrm{S}$ to x, we must have come from w. But however we arrived at w (from $w + \mathrm{E}$, $w + \mathrm{S}$ or $w + \mathrm{W}$), we should have gone to y in preference to $x + \mathrm{S} = w + \mathrm{E}$ (cf. Lemma 3.1), giving a contradiction. $\diamondsuit$

Putting Lemmas 3.7 and 3.8 together, we see that if a point x is ever blocked by the demon, then the angel never passes by $x + \mathrm{N}$, $x + \mathrm{W}$ or $x + \mathrm{E}$ (before or after).

Now suppose that $\alpha \subseteq \pi$ is a subtrail. Let $\Lambda(\alpha)$ be the set of liberties of α. By definition, each liberty is to the N, E or S of some vertex of α. Thus, by the above observation, no such liberty can ever have been blocked by the demon.

In fact:

Lemma 3.9 : *If $\alpha \subseteq \pi$ is a subtrail, then each element of $\Lambda(\alpha)$ was blocked by the devil by the time the angel arrives at the terminal vertex.*

Proof : We have already observed that no liberty was blocked by the demon, so it suffices to show that they must all be blocked. This is a simple consequence of rule (A1). For example, if $e = xy$ is such an edge, if $y + \mathrm{L}$ is not blocked, then, we should go there on the next move, so it cannot be a liberty. (We need some special consideration for the initial and final edges.) Similary each terminal liberty must be blocked. $\diamondsuit$

Since the devil must visit any vertex w times before it is blocked, the following is an immediate consequence:

Corollary 3.10 : *If $\alpha \subseteq \pi$ is the subtrail of length n starting at o, then $w|\Lambda(\alpha)| \leq n$.* $\diamondsuit$

Lemma 3.11 : *At time n, the height of the angel is at least $\left(1 - \frac{4}{w}\right) n$.*

Proof : The height, h, is the Northward displacement of the trail α of length n from o. By Lemma 2.6, $h = l(\alpha) - m(\alpha)$, where $l(\alpha)$ and $m(\alpha)$ are, respectively, the length and Southward migration of α. But $l(\alpha) = n$, and by Lemma 2.7, $m(\alpha) \leq 4|\Lambda(\alpha)|$. Thus, $m(\alpha) \leq \frac{4}{w} n$, and so $h \geq n - \frac{4}{w} n = \left(1 - \frac{4}{w}\right) n$, as claimed. $\diamondsuit$

Thus, if $w \geq 5$, we get $h \geq \frac{n}{5}$. This is good enough to make subsequent arguments work. We get better inequalities if we take $w \geq 9$, so that $h \geq \frac{5n}{9} > \frac{n}{2}$.

One immediate consequence is that:

Proposition 3.12 : *The angel remains in the positive (i.e. North East) quadrant.* $\diamondsuit$

In particular, it never returns to σ_-, justifying the assertion that, in retrospect, it plays no role in the angel's strategy.

Now, the trail π_+ from o onward, is circuitous. We write it as $\sigma \cup \bigcup_{i=0}^{\infty} \gamma_i$, where γ_i is a loop based at $\sigma(i)$.

Lemma 3.13 : $l(\gamma_i) \leq i$.

Proof : Let h be the height of the basepoint, $x = \sigma(i)$. Let n and n' be first and last times respectively at which the angel passes by x. Thus $n \geq i$, and $l(\gamma_i) = n' - n$. Clearly, $h \leq i$, and by the above, $n' \leq 2h$, so $n' \leq 2i$. Thus, $l(\gamma_i) = n' - n \leq 2i - n \leq 2i - i = i$. $\diamondsuit$

We have verified property (L) of Section 2, thereby proving Proposition 1.2.

It's also worth noting that, even if we were to abandon the game at some finite stage, the angel will still have followed a circuitous trail satisfying (L). We need not modify our analysis. We can simply continue the angel moves deeming the devil to pass, and note that any subtrail of a circuitous trail is circuitous. The proof of Lemma 3.13 goes through unchanged.

As remarked in the introduction, it is likely that the argument can be modified to deal with $w = 4$, though the details become a bit complicated. The basic idea would be to allow the angel to make diagonal moves as well as duke moves (i.e. like a king). We can think of this as almost equivalent to making two successive duke moves without waiting for the devil to play, if the latter move is L or R. With an appropriately modified definition of "Southward migration" it seems one gets a multiplicative factor of 3, in the analogue of Lemma 2.7, which is then good enough. Of course, the arguments of Section 4 would also need modification. Thus, there is a reasonable hope of dealing with 3-angel. A 2-angel, however, would be a completely different matter.

4. Back to the original game.

The angel wants to use its understanding of Game 2 to play Game 1. To this end, the angel employs a "phantom" and two demons. The phantom plays according to the rules and strategy of Game 2 as already laid out, using the first demon for the purpose of rule (A2). The angel also controls the second demon, using it to block the phantom in a similar manner to the devil. The idea is that the phantom will warn the angel in advance of any traps the devil may have set. It does this by performing a loop of its circuitous trail. The angel can then follow the spine of such a trail, without fear of being trapped. Moreover, the angel moves will be in accordance with the rules of Game 1.

For the purposes of planning its strategy and measuring time, the phantom imagines that it is playing against both the devil and the second demon. It then chooses its moves according to rules (A1) and (A2). (For the purposes of rule (A2), it borrows the first demon, so these don't count as devil moves for the phantom.) Now the analysis of Section 3 applies, with words "angel" replaced by "phantom" and "devil" replaced by "devil and second demon". In particular, we know that the phantom will trace out a circuitous trail satisfying (L). Note that the angel and phantom are working to different times. It is the angel's time that we are really interested in.

To be more precise, we proceed inductively. The angel and phantom both start at o. Let us suppose that at (angel) time i, the angel and phantom have both arrived at the same point x. We assume that the phantom has been following the strategy of Game 2 against the devil and second demon. In particular, it must have followed a circuitous trail.

We also assume that the angel has followed the spine of the phantom's trail up to that point. Moreover, we assume that no subsequent sequence of devil moves could ever return the phantom to any vertex previously visited by the angel (other than x).

The devil now plays, and the angel has to decide where to go. It is possible that some future sequence of devil moves will return the phantom to x. But by (L), any such sequence must have length at most i. It must also therefore remain within a distance of $i/2$ of x. Only (hypothetical) devil moves within this region will influence the phantom's strategy. Thus, the angel can search for all possible such sequences (where the phantom moves first, in response to the move the devil just made). It chooses one such devil sequence of maximal length (taking the empty sequence if no non-trivial sequence returns the phantom to x). To get an explicit algorithm, the angel could choose the first such sequence in some fixed a-priori enumeration of such sequences. The angel now instructs the second demon to make precisely these moves. Meanwhile it sends the phantom off to continue its Game 2 strategy, confident in the knowledge that it will return to x once this process is complete. (It does not really matter whether or not the second demon is given a final move after the phantom has returned.) Once it does, the angel and phantom move together, applying the Game 2 strategy (without waiting for devil to make another move). By the inductive hypothesis, this cannot return us to any vertex previously visited by the angel. In fact, by Lemma 2.4, it cannot have been previously visited by the phantom either. In other words it was never made safe under the rules of Game 2, and so remains unblocked by the devil. It is thus a legal move for the angel, by the rules of Game 1. Moreover, the angel has once again followed the spine of the phantom's trail, and we are back in the same situation as before, with the angel's time incremented by 1.

Note that at time i, the phantom will have traced out a trail of radius at most i about its current location $x = \sigma(i)$. (This is a simple consequence of (L), as remarked in Section 1.) In particular, only devil moves in the square $\sigma(i) + W(i)$ have influenced our moves up to this point. (The fact that the first $w \geq 5$ moves were North, means we don't need to bother add 1 to our radius bound.)

Referring back to the original game, where we divided our lattice into 3×3 squares, this gives a square $\sigma(i) + W(3i)$.

This does not quite prove Theorem 0 as we stated it, since it obliges the angel to remember both where it started, and the order of devil moves. However, only the angel's current position and the set of vertices blocked by the devil at this stage should be relevant. If the angel has forgotten the additional information, it is sufficient to reconstruct some hypothetical starting point and order of play, and to base its next move on these. Again only devil plays in $\sigma(i) + W(3i)$ are relevant here. (It won't matter if this entails superimposing σ_- on top of devil moves outside this square, since this does not influence our strategy anyway.)

To be more formal, the inductive hypothesis in the above set-up states that at a given time i, there exists a point $o = o(i)$ (the hypothetical origin) and a total order on the set $\Delta(i)$ of blocked positions (the hypothetical order of devil plays) such that if the angel and phantom had played to the above strategy, they would both have arrived at the actual current location of the angel, and, moreover, the earlier inductive hypotheses as laid out above, are satisfied in this hypothetical situation (interpreting in terms of the derived Game 1). The angel now plays, pretending this is really what happened. Since an angel move is always non-trivial, by the above strategy, the angel is never trapped. Again, in the above hypothetical reconstruction, only the square $\sigma(i) + W(3i)$ is relevant, and the angel can choose that reconstruction that is minimal in some a-priori enumeration of possibilities. This gives an explicit algorithm. We know inductively that a hypothetical reconstruction exists at time $i + 1$, even if we have forgotten once again what it is.

This proves Theorem 0, when 4-angel is replaced by 5-angel and $W(4n)$ is replaced by $W(3n)$. For a 4-angel, we use crosses instead of 3×3 arrays, we take $w = 5$, and then (L) becomes "$l(\gamma_i) \leq 4i$". This time, the radii of the squares is determined by have the radius bound of the loops, $l(\gamma_i)/2 = 2i$ (rather than the length of the angel's initial segment as before). Only the multiplicative factors in the argument change.

5. Remarks.

While the above gives a resolution to the angel problem in the euclidean plane, it is in some ways less than satisfactory. Since it involves a systematic search, it provides little insight into the nature of the evasion procedure. Moreover, it makes essential use of planar topology, in particular, variants of the rectilinear Jordan curve theorem. This means it may not be very adaptable. It would be nice to have a broader perspective on the issue.

It seems natural to formulate the problem more generally. We have already referred to the recent work on the 3-dimensional integer game by Kutz and Bollobás and Leader. One might make the following, more general, definition.

Let Γ be a connected uniformly locally finite graph (i.e. with a bound on the degree of its vertices).

Definition : We say that Γ is *diabolical* if the devil can trap an arbitrarily powerful angel regardless of its starting position.

The definition is quite robust. Clearly any subgraph of a diabolical graph is diabolical. More generally, the following statement is easily verified. Suppose that Γ and Γ' are connected and uniformly locally finite, and that there exists a map $f : V(\Gamma) \longrightarrow V(\Gamma')$ such that the preimage of every point has bounded cardinality, and the images of any two adjacent points are a bounded distance apart. If Γ' is diabolical, then so is Γ.

From this, we may conclude easily that diabolicity is quasi-isometry invariant among uniformly locally finite graphs. Indeed the above statement is equivalent to subgraph closure together with quasi-isometry invariance. (One can easily find a graph Γ'' and a quasi-isometry $\Gamma'' \longrightarrow \Gamma'$ such that this quasi-isometry precomposed with an embedding $\Gamma \longrightarrow \Gamma''$ and restricted to $V(\Gamma)$ equals f.)

One can ask the general question of when a graph is diabolical. We have seen that $\mathbf{Z}^2$ is not. Neither is the 3-regular tree: it is easily seen that a 1-angel can escape.

To make this more tractable, we could restrict the question to finitely generated groups. If view of quasi-isometry invariance, we can define a finitely generated group to be diabolical if some (hence any) Cayley graph is. Recall that a group is *virtually cyclic* if it has and infinite cyclic subgroup of finite index. It seems natural to ask:

Question : Is every diabolical group finite or virtually cyclic?

From the main result of this paper, a diabolical group cannot contain any non-cyclic free abelian subgroup or free subsemigroup. An attack on this question may take us into the realm of exotic group theory.

References.

[BeCG] E.R.Berlekamp, J.H.Conway, R.K.Guy, *Winning Ways, Volume 2* : Academic Press (1982).

[Bo] B.Bollobás, *Modern graph theory* : Springer-Verlag (1998).

[BoL] B.Bollobás, I.Leader, *The angel and devil in three dimensions* : J. Combin. Theory Ser. A **113** (2006) 176–184.

[C] J.H.Conway, *The angel problem* : in "Games of no chance" (ed R.J.Nowakowsi), Cambridge (1996) 3–12.

[E] R.A.Epstein, *Theory of gambling and statistical logic* : Academic Press (1967).

[G] M.Gardner, *Cram, crosscram and quadraphage: new games having elusive winning strategies* : Sci. Amer. **230** (1974) 106–108.

[Kl] O.Kloster, *A solution to the angel problem* : Preprint, SINTEF Oslo.

[Ku] M.Kutz, *Conway's angel in three dimensions* : Theoret. Comput. Sci. **349** (2005) 443–451.

[KuP] M.Kutz, A.Pór, *Angel, devil and king* : preprint (2006).

[M] A.Máthé, *The angel of power 2 wins* : to appear in Combin. Probab. Comput.

[S] D.L.Silverman, *Your move* : McGraw Hill (1971).