The Angel wins

Péter Gács

Computer Science Department
Boston University

October 19, 2007

The game

The board is an infinite "chessboard" (represented by the integer lattice $\mathbb{Z}^2$), with black and white squares, but initially, all squares are white.
The Angel (a star in the picture) is always located on some square. In each step, first the Angel moves, then the Devil.
In her move, the Angel makes $\leqslant p$ unit steps, (p is constant, say $p = 1000$). The steps can be both horizontal and vertical, so we are using the distance

$$d(\langle x_1, y_1 \rangle, \langle x_2, y_2 \rangle) = \max(|x_2 - x_1|, |y_2 - y_1|).$$

She has to land on a white square.
The Devil can turn a square black (eat it). This square stays black from now on.

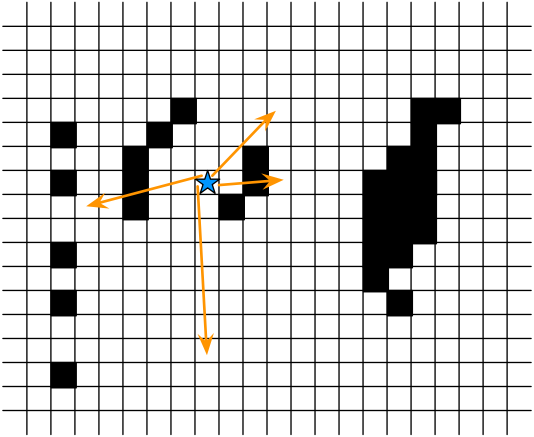

Here, the Angel can move at most 6 steps (horizontal or vertical).

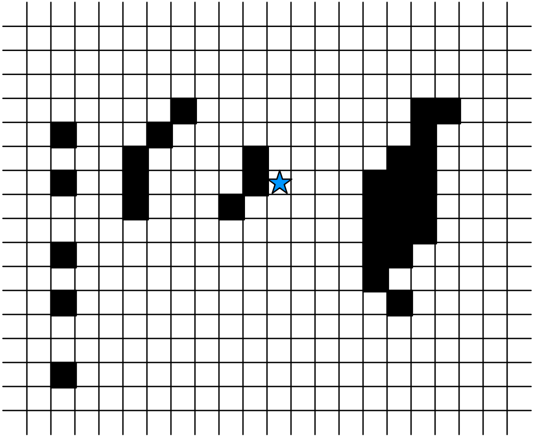

Here, the Angel can move at most 6 steps (horizontal or vertical).

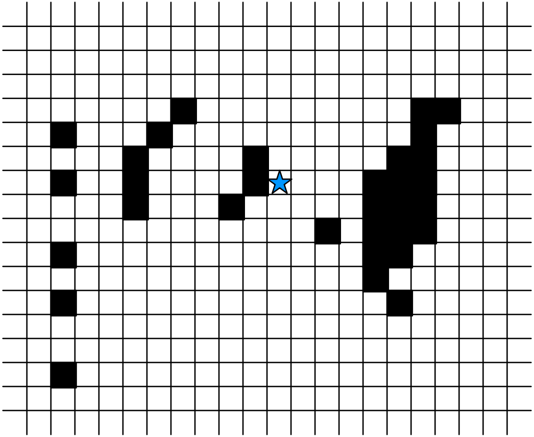

Here, the Angel can move at most 6 steps (horizontal or vertical).

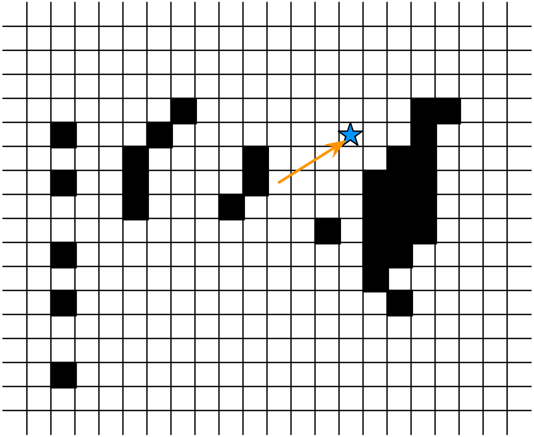

Here, the Angel can move at most 6 steps (horizontal or vertical).

Problem

Can the Devil always capture the Angel?

How could he? For example, by buiding a wall of width p around her.
Berlekamp: for $p = 1$ the Devil wins.
Of course, with p large, the Angel seems to have a large advantage.

Theorem

For sufficiently large p, the Angel has a strategy in which she will never run out of places to land on.

Four independent solutions, in order of increasing merit:

- I have proved the theorem for large p.
- Bowditch: $p = 4$.
- Máthé: $p = 2$ (optimal).
- Kloster: $p = 2$, with a simple algorithm.

Most of the talk is adapted from Kloster's website (and paper).

Why has it been a challenge to prove this theorem? Because it has been found difficult to translate the Angel's advantage into a strategy that the Devil cannot outwit. Here is an example:

Theorem

Suppose that the Angel is not allowed to ever decrease her x coordinate. Then the Devil can capture her.

How? Let me show this only for the case when the Angel has to *increase* the x coordinate in every step. Then from every position, her future steps are confined to a cone of angle with tangent $(p-1)$.

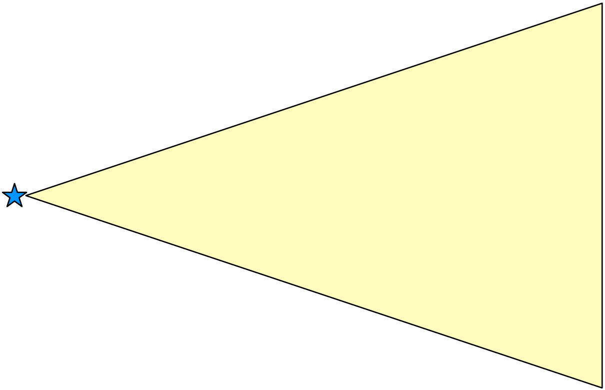

The Devil thickens a smaller and smaller segment of a far-away vertical wall. While the Angel is at distance $2^{n-1} \leqslant d < 2^n$, he increases the thickness of a segment of size $M/2^n$ by an additive constant.

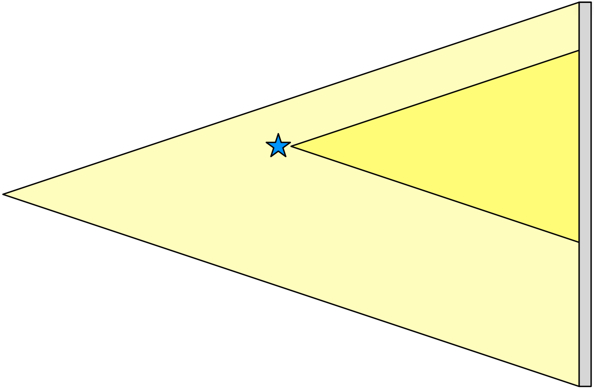

The Devil thickens a smaller and smaller segment of a far-away vertical wall. While the Angel is at distance $2^{n-1} \leqslant d < 2^n$, he increases the thickness of a segment of size $M/2^n$ by an additive constant.

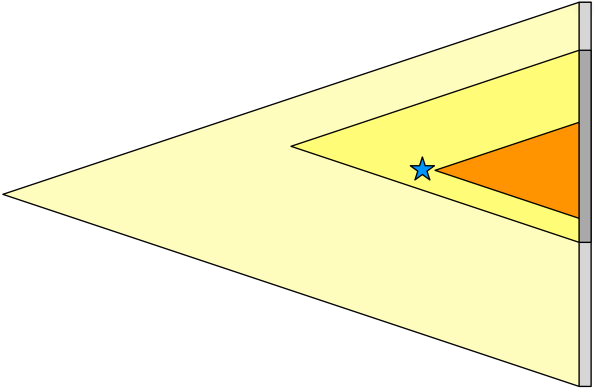

The Devil thickens a smaller and smaller segment of a far-away vertical wall. While the Angel is at distance $2^{n-1} \leqslant d < 2^n$, he increases the thickness of a segment of size $M/2^n$ by an additive constant.

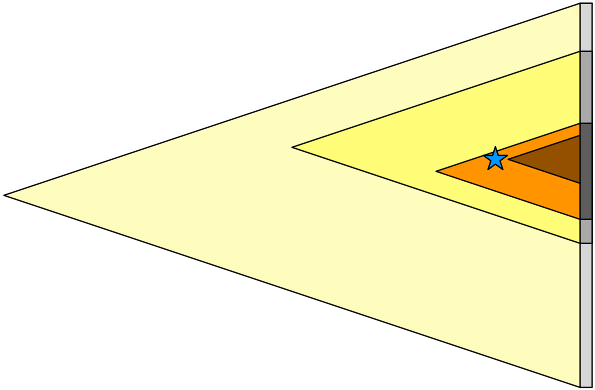

The Devil thickens a smaller and smaller segment of a far-away vertical wall. While the Angel is at distance $2^{n-1} \leqslant d < 2^n$, he increases the thickness of a segment of size $M/2^n$ by an additive constant.

Why interesting?

To most mathematicians, the problem does not need justification: it is sufficient that it is simply stated and still nontrivial. But I had two other reasons to be interested.

The Angel seems to need a "multi-level" strategy (my specialty). She must be aware of possible threats posed by the Devil on many "scales": She should not walk into a trap of size 10, 100, 1000, ...

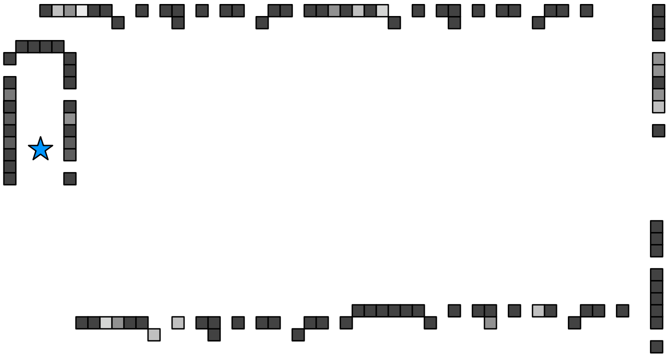

I indeed gave a hierarchical solution, but everybody else's solution is simpler, and is not hierarchical.

Another reason:
The problem fits interestingly into a general project of developing an "approximate topology". It measures a sort of connectivity of the lattice in terms of an adversary's cutoff abilities.

Kloster's solution

The algorithm

1. The Angel declares part of the board red. At start, this is the left half-plane, and will always be "homeomorphic" to it. Its perimeter (initially directed upward) is called the path.
2. At all times, the Angel stays next to the red area, next to a segment on the path called the perch.
3. On the Angel's turn, she advances two units along the path, keeping the red area on the left.
4. Every time the Devil has eaten a square, the Angel may paint additional squares red, while satisfying the following conditions.
   - The red area remains connected.
   - The path behind and including the perch is unchanged.
   - The path length increases by no more than two units for each eaten square that is now painted red.

   When she can do this, she also must, for the maximal number of eaten squares.

Walking along the path.

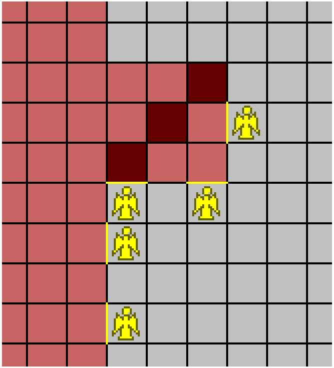

So, whenever the Devil eats a threatening amount of squares too close to the Angel's path, the Angel will paint the area red and walk around instead. Examples:

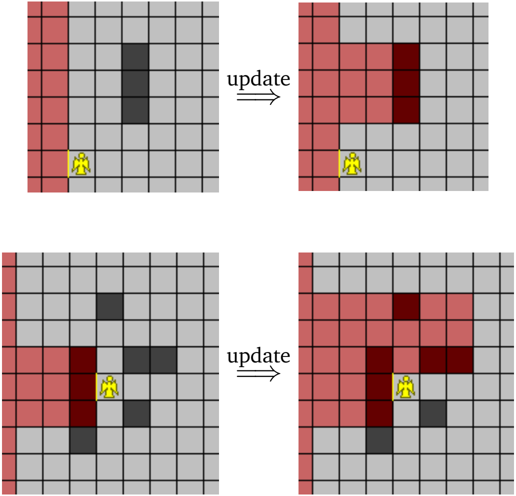

The Devil can also eat the square the Angel is standing on, and this might result in a—slightly—degenerate path:

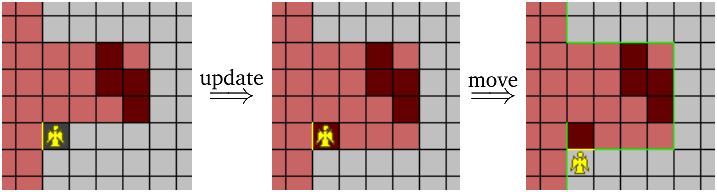

A situation that would capture the Angel but will never occur:

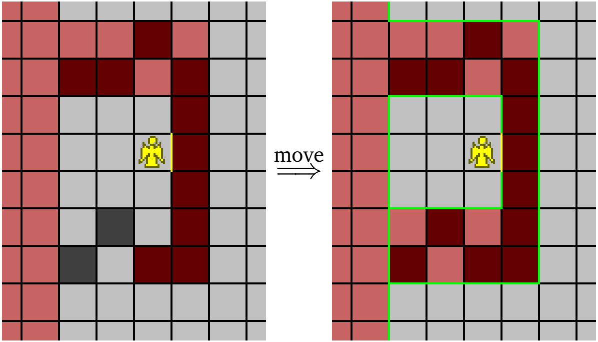

Proof that it works

Easy case: the Angel never steps on an eaten cell that is not red.

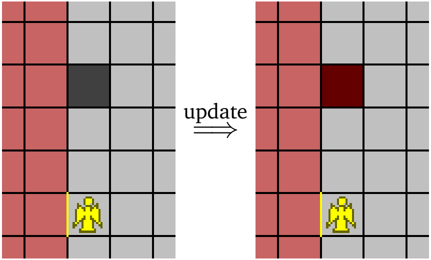

Harder case: the forbidden area will never surround the Angel in a way that she cannot get out.

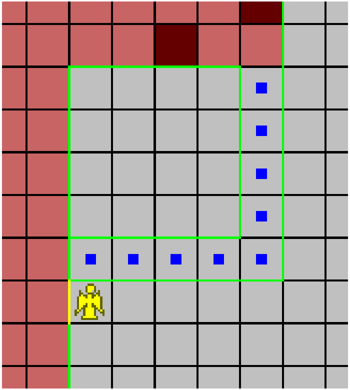

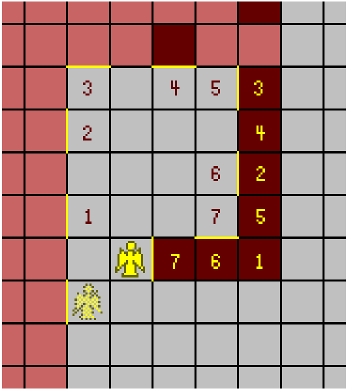

In this example, the Devil can start building the vertical wall of the trap only in places where he is below the Angel. This gives her enough time to escape. *(demo on surround-history)*

- For a path λ, let $|\lambda|$ be its length increase from the initial one (after ignoring the identical infinite parts).
- λ_t = path before step t.
- $r_t(\lambda)$ = number of eaten squares painted red by time t for λ.

Lemma (Potential)

$|\lambda_t| - 2r_t(\lambda_t)$ is nondecreasing in t.

This is easy to check.

Assume that λ_j (black looped curve here) is the first curve for which both sides of some segment get red (stripes here). It is easy to see that this can happen only with the Angel being in the loop.

- s = the first such segment.
- κ = the green curve outside of λ_j, after cutting out the loop including s.
- i = the first time the Angel gets past s.

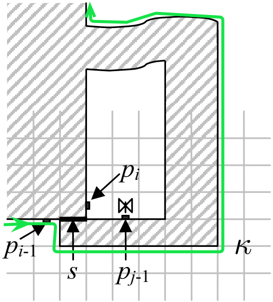

$$|\lambda_j| - |\kappa| \geqslant 2(j-i)$$

because the Angel moves two segments per turn

$$\geqslant 2(r_j(\kappa) - r_i(\kappa))$$

because the Devil cannot eat more

$$\geqslant 2(r_j(\lambda_j) - r_i(\kappa))$$

since κ surrounds more than λ_j, hence

$$|\lambda_j| - 2r_j(\lambda_j) \geqslant |\kappa| - 2r_i(\kappa) \geqslant |\lambda_i| - 2r_i(\lambda_i)$$

by optimality of λ_i

$$\geqslant |\lambda_j| - 2r_j(\lambda_j)$$

by the Potential Lemma.

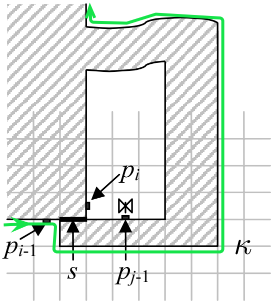

Hence equality throughout, showing $|\lambda_j| - |\kappa| = j - i$, so p_{j-1} is immediately before s, and the Angel will jump out.

Complexity

Recall the Angel's algorithm:

> *Every time the Devil has eaten a square, the Angel may paint additional squares red, while satisfying the following conditions.*
>
> - *The red area remains connected.*
> - *The path behind and including the perch is unchanged.*
> - *The path length increases by no more than two units for each eaten square that is now painted red.*
>
> *When she can do this, she also must, for the maximal number of eaten squares.*

This is an optimization problem for a path, by its form it could have a complexity exponential in the (relevant) length.

In practice, the Angel never needs to do anything this complex. She can always keep the future part of the path in the following simple form:

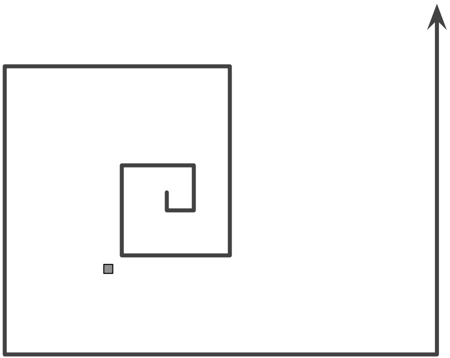

This spiral has a logarithmic number of sides. After the Devil eats a square, unfortunately you may still have to modify all of the sides, so the search may still be of the order of $n^{\log n}$. The possiblity of a polynomial algorithm still seems open.

Amusingly, the past part can be quite complicated (in a fractal way, maybe not with the proportions shown).

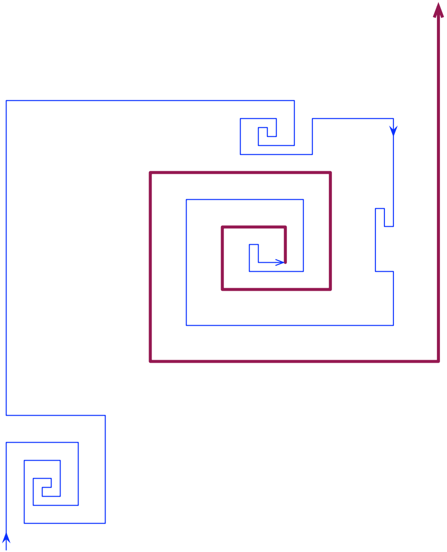

Máthé's solution

The Nice Devil

An interesting observation:

Theorem (Conway)

If the Angel has a winning strategy then she also has a winning strategy in which she will never go to a site that she has visited, nor to any site that she could have passed to in an earlier move but did not.

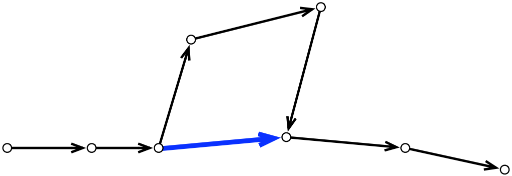

(The figure gives a very rough idea of the proof.)

Máthé found an important recasting of this theorem.

Definition

A Nice Devil is a Devil that never blocks a square on which the Angel has visited, nor to any site that she could have passed to in an earlier move but did not.

Theorem (Nice Devil)

If the Devil catches the Angel then the Nice Devil can entrap her in some finite domain.

This theorem almost solves the problem. It frees the Angel from worrying about walking into most kinds of trap: she can walk back out, the Devil cannot stop her!

Proof of the Nice Devil theorem

To each journey $\boldsymbol{v} = \langle v_0, \dots, v_n \rangle$ of Angel, reduced journey

$$\rho(\boldsymbol{v}) = \boldsymbol{u} = \langle u_0, \dots, u_k \rangle.$$

For this,

1. Draw an arrow from each v_i to the earliest v_j within distance p.
2. Take the path formed by these backward arrows starting from v_0, and number it forward as $\langle u_0, \dots, u_k \rangle$.

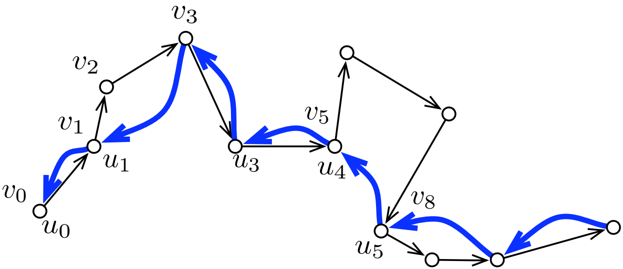

Let $\Phi(\boldsymbol{v}) = \Phi(v_0, \dots, v_n)$ be the position where the Devil would put something after seeing journey $\boldsymbol{v}$ (he has also the option of doing nothing). We define the Nice Devil's strategy Ψ as follows.

$$\Psi(\boldsymbol{v}) = \Phi(\rho(\boldsymbol{v})) = \Phi(\boldsymbol{u})$$

if this move is permitted to the Nice Devil, and nothing otherwise. Suppose Φ captures $\boldsymbol{u} = \langle u_0, \dots, u_k \rangle$, we will show that Ψ captures $\boldsymbol{v} = \langle v_0, \dots, v_n \rangle$.

There are $s < t$ with $u_t = \Phi(u_0, \dots, u_s)$. Let $s' = \min_{v_i = u_s} i$, similarly for t'. It is easy to see

$$\langle u_0, \dots, u_s \rangle = \rho(v_0, \dots, v_{s'}).$$

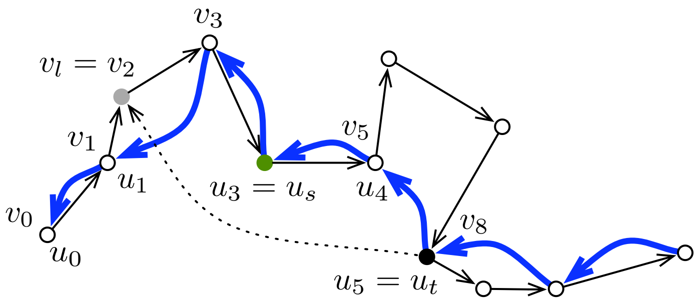

We have $\Psi(v_0, \dots, v_{s'}) = \Phi(u_0, \dots, u_s) = u_t = v_{t'}$ or nothing. If it is $v_{t'}$ then the Nice Devil captures $\boldsymbol{v}$.

Let us show that it cannot be nothing. Indeed, this could only be if the Nice Devil could not eat $v_{t'}$, which assumes an $l < s'$ with distance $d(v_l, v_{t'}) \leqslant p$. But then the construction of the backward path $u_t, u_{t-1}, \dots,$ would have bypassed the node $u_s = v_{s'}$.

Constructivity

Formally, the transformation done for the Nice Devil makes Máthé's solution non-constructive. Bowdich's solution uses almost the same transformation. It is claimed that the proof of Conway's theorem turns them constructive.

A non-constructive solution would be quite interesting if nothing else (at least nothing else so simple) was available. But Kloster's solution is simple, constructive and even efficient (as shown here).

Example (Game with only non-constructive solution)

It is known that there is a finite set of square tiles (with various marks on their edges) such that the plane can be tiled with copies of them (touching edges must have matching marks), but only in a non-recursive way. So let our game be: in each step, the Angel puts down a tile, adjacent to the others, in a circular order. The Devil does nothing, but still wins if the Angel cannot continue.

My solution
Reformulation

1. The Angel can make only one (horizontal or vertical) step at at time.
2. The Devil manages a weight distribution μ. At time t, the weight at site $\boldsymbol{x}$ is

$$\mu_t(\boldsymbol{x}).$$

   The Angel cannot land on a site $\boldsymbol{x}$ with weight $\mu(\boldsymbol{x}) \geqslant 1$.
3. The weight of a set S of sites is $\mu_t(S) = \sum_{\boldsymbol{x} \in S} \mu_t(\boldsymbol{x})$. There is a small constant $\sigma > 0$ bounding the total weight increase per step:

$$\mu_{t+1}(\mathbb{Z}^2) - \mu_t(\mathbb{Z}^2) \leqslant \sigma.$$

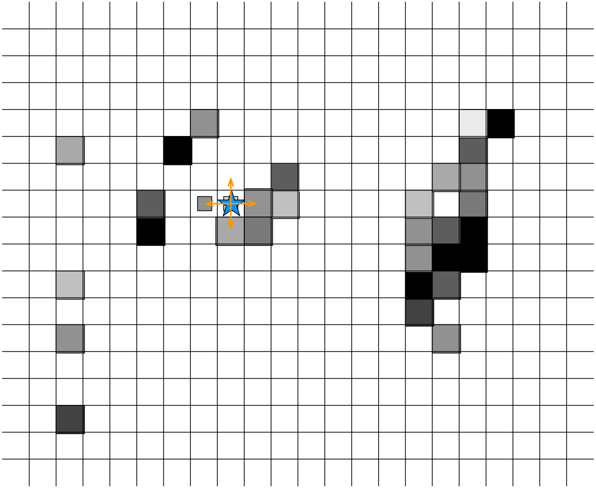

Now the theorem says that the Angel wins for small enough σ.

Multi-level terminology

Fix a (large) integer constant $Q > 1$. A k-colony is a square whose corners have coordinates that are multiples of Q^k.
When looking only at level k and $(k+1)$, the $(k+1)$-colony is called a colony, the k-colony is called a cell.

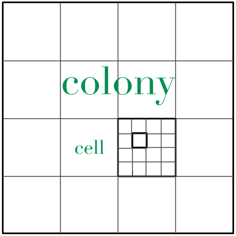

The side length of a square U is denoted $|U|$. A square U is bad (for the current measure μ), if

$$\mu(U) \geqslant |U|.$$

Otherwise it is good. Note that U becomes bad as soon as its weight is as large as its side length. The Devil need not "fill" a colony to spoil it, it is sufficient (for example) to "draw a line" through it.

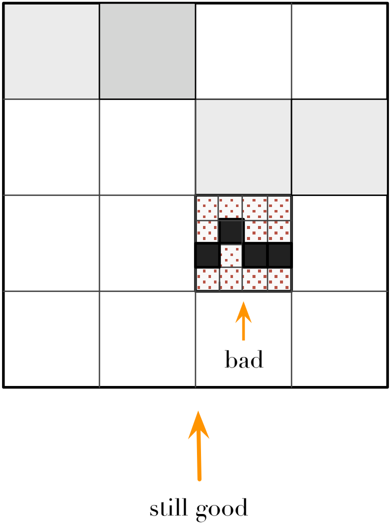
bad
still good

Failure of planning

Suppose we want to pass a colony. We cannot plan our path all in advance based on what we see at our start, even on two levels.

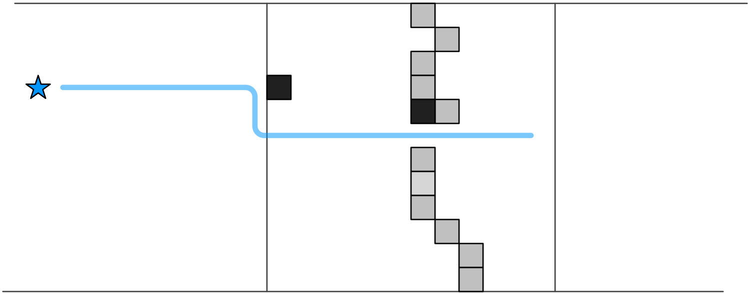

Failure of planning

Suppose we want to pass a colony. We cannot plan our path all in advance based on what we see at our start, even on two levels.

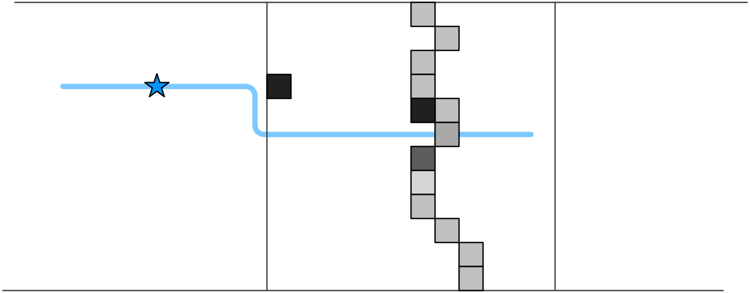

Failure of planning

Suppose we want to pass a colony. We cannot plan our path all in advance based on what we see at our start, even on two levels.

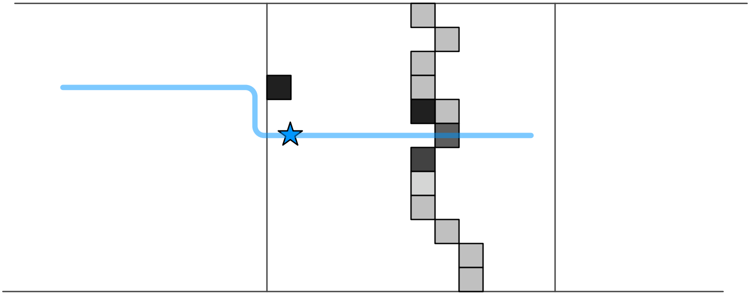

Failure of planning

Suppose we want to pass a colony. We cannot plan our path all in advance based on what we see at our start, even on two levels.

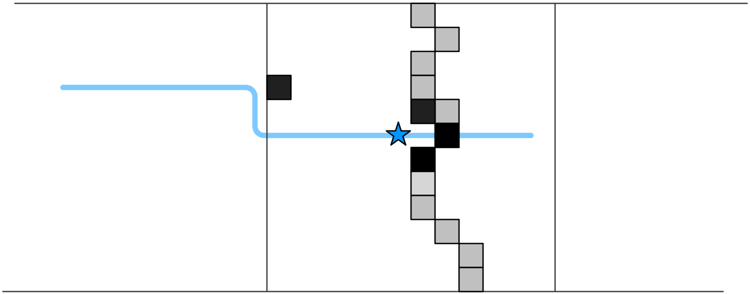

Failure of planning

Suppose we want to pass a colony. We cannot plan our path all in advance based on what we see at our start, even on two levels.

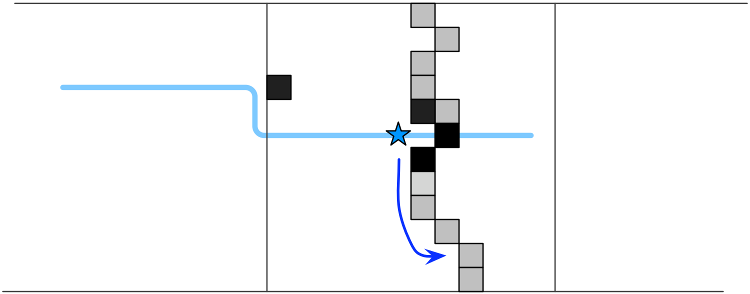

Time bound

Too many digressions

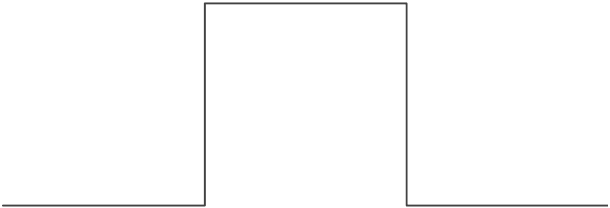

Time bound

Too many digressions

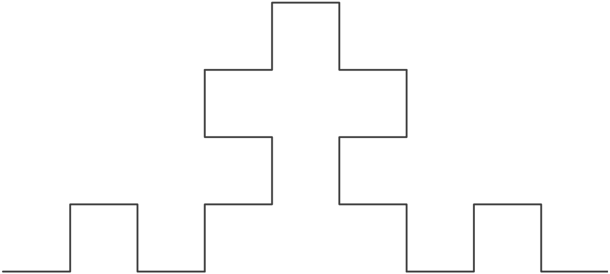

Time bound

Too many digressions

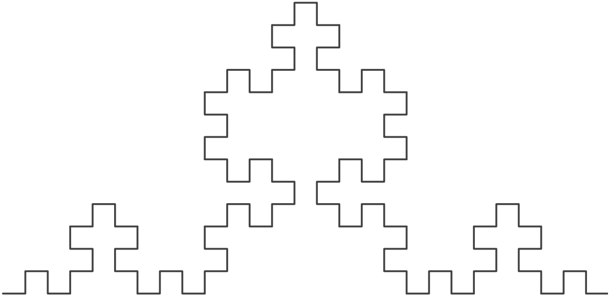

Time bound

Too many digressions

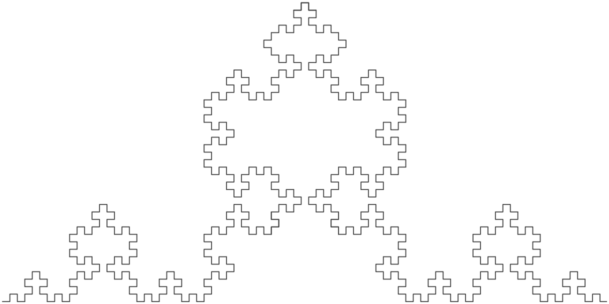

Time bound

However: the Devil must "pay" for every "digression" we are forced into. This is formally expressed using a time upper bound

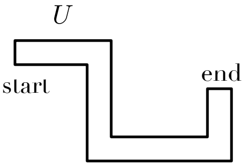

$$\tau_{gc}(U) + \rho\mu(U).$$

Here $\tau_{gc}(U)$ is the geometric cost of moving from the start to the end of region U: it is just the sum of the lengths of straight runs making up the region.
The part exceeding the geometric cost is charged to the Devil via $\rho\mu(U)$.

Scaling the time bound

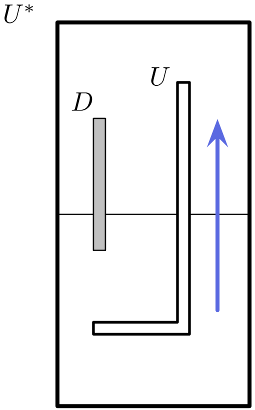

The lower-scale time bound will imply a similar higher-scale time bound:

$$
\begin{aligned}
&\tau_{\mathrm{gc}}(U) + \rho\mu(U) \\
&\leqslant \tau_{\mathrm{gc}}(U^*) + \rho\mu(D) + \rho\mu(U) \\
&\leqslant \tau_{\mathrm{gc}}(U^*) + \rho\mu(U^*).
\end{aligned}
$$

The geometric cost of U^* is its length. The geometric cost of U is larger by the horizontal “digression”, but this will be estimated via the charge $\rho\mu(D)$, where D is the bad area outside U causing the digression.

The simple cases

Let us set up a structure helping us to pass through areas that are not too beleagered by the Devil.
A (horizontal or vertical) union of adjacent colonies is a run.

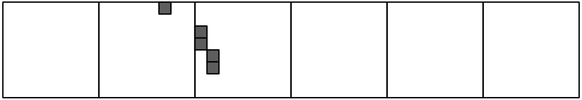

We will have constants

$$\delta > 0, \quad 1/2 < \xi < 1 - \delta,$$

where δ is appropriately small. A run U whose squares have width B is 1-good (for μ) if $\mu(U) < (1 - \delta)B$. It is safe if $\mu(U) < \xi B$.

Observation

With sufficiently large Q, the following holds:

- *If a run (row or column) is good then at most one of its cells is unsafe (we used $\xi > 1/2$).*
- *If a colony is safe then at least 4 of its columns (and its rows) are 1-good.*

The colony of largest weight in a run will be called its obstacle. In a good run, (hopefully) we can move around rather freely, *except for jumping somehow over the obstacle*. Let us postpone the problem of the obstacle.

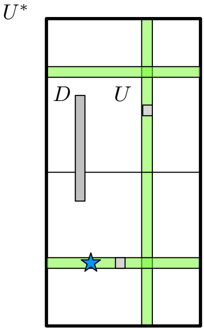

In this safe pair of colonies, there will be several good columns along which we can pass from the bottom to the top. There are also good rows. Assuming we start from a good row, we can pass to the column, and in the column we can pass to a good row in the destination colony, running above the "scapegoat" area (see earlier).

If we move fast, then the situation does not change too much while we execute all this.

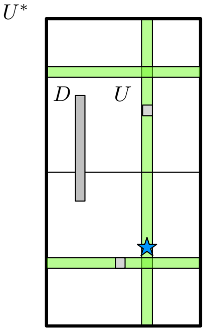

In this safe pair of colonies, there will be several good columns along which we can pass from the bottom to the top. There are also good rows. Assuming we start from a good row, we can pass to the column, and in the column we can pass to a good row in the destination colony, running above the “scapegoat” area (see earlier).

If we move fast, then the situation does not change too much while we execute all this.

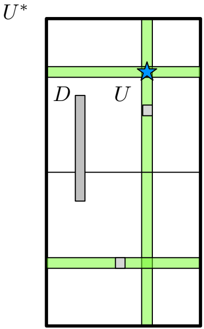

In this safe pair of colonies, there will be several good columns along which we can pass from the bottom to the top. There are also good rows. Assuming we start from a good row, we can pass to the column, and in the column we can pass to a good row in the destination colony, running above the “scapegoat” area (see earlier).

If we move fast, then the situation does not change too much while we execute all this.

An attack

How to pass an obstacle? Note that here $\mu(U) < |U|$ but we cannot plan in advance, where to pass through square U.

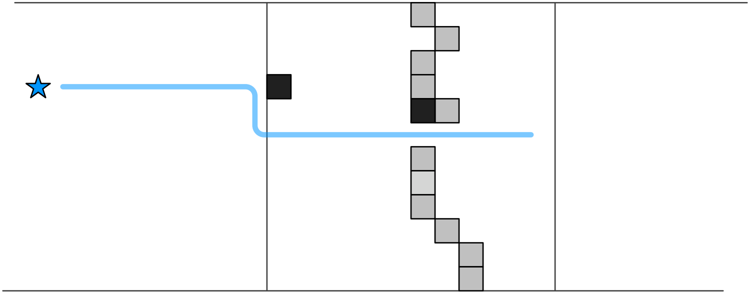

An attack

How to pass an obstacle? Note that here $\mu(U) < |U|$ but we cannot plan in advance, where to pass through square U.

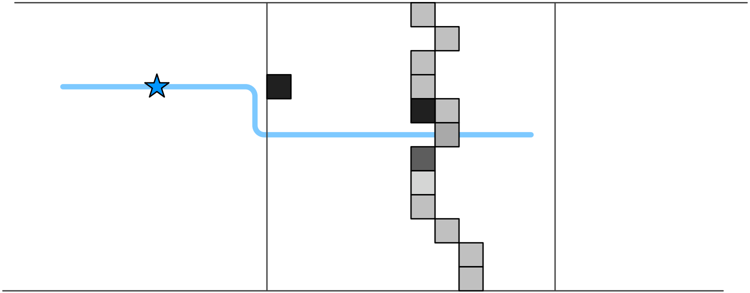

An attack

How to pass an obstacle? Note that here $\mu(U) < |U|$ but we cannot plan in advance, where to pass through square U.

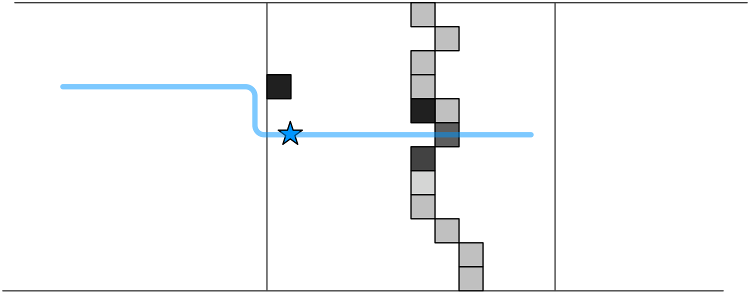

An attack

How to pass an obstacle? Note that here $\mu(U) < |U|$ but we cannot plan in advance, where to pass through square U.

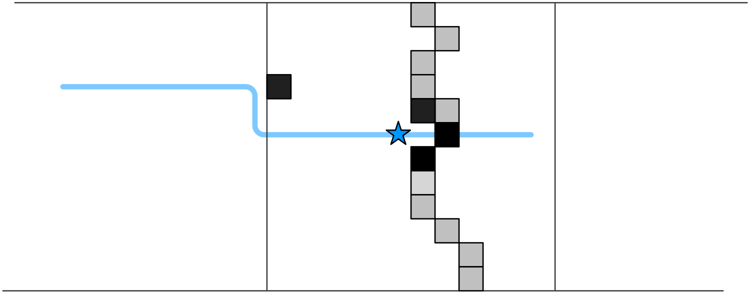

An attack

How to pass an obstacle? Note that here $\mu(U) < |U|$ but we cannot plan in advance, where to pass through square U.

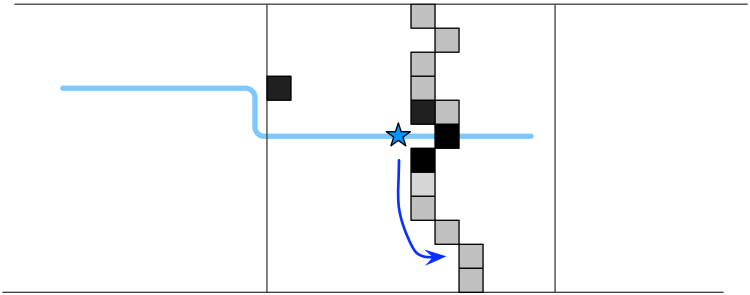

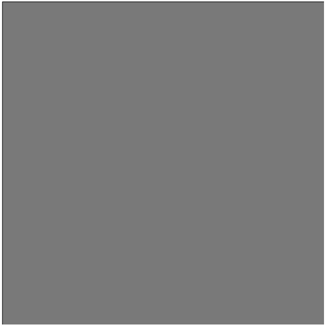

We cannot avoid U. Indeed, if we avoid all squares U with $\mu(U) < \lambda|U|$ where $\lambda < 1$ then on level k the Devil will scare us even with a square U having $\mu(U) < \lambda^k|U|$!

Moral: attempt to pass, allowing for possible failure. (You cannot win a war "from the air", avoiding all dangerous situations.)

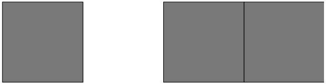

We cannot avoid U. Indeed, if we avoid all squares U with $\mu(U) < \lambda|U|$ where $\lambda < 1$ then on level k the Devil will scare us even with a square U having $\mu(U) < \lambda^k|U|$!

Moral: attempt to pass, allowing for possible failure. (You cannot win a war “from the air”, avoiding all dangerous situations.)

We cannot avoid U. Indeed, if we avoid all squares U with $\mu(U) < \lambda|U|$ where $\lambda < 1$ then on level k the Devil will scare us even with a square U having $\mu(U) < \lambda^k|U|$!

Moral: attempt to pass, allowing for possible failure. (You cannot win a war "from the air", avoiding all dangerous situations.)

We cannot avoid U. Indeed, if we avoid all squares U with $\mu(U) < \lambda|U|$ where $\lambda < 1$ then on level k the Devil will scare us even with a square U having $\mu(U) < \lambda^k|U|$!

Moral: attempt to pass, allowing for possible failure. (You cannot win a war "from the air", avoiding all dangerous situations.)

We cannot avoid U. Indeed, if we avoid all squares U with $\mu(U) < \lambda|U|$ where $\lambda < 1$ then on level k the Devil will scare us even with a square U having $\mu(U) < \lambda^k|U|$!

Moral: attempt to pass, allowing for possible failure. (You cannot win a war "from the air", avoiding all dangerous situations.)

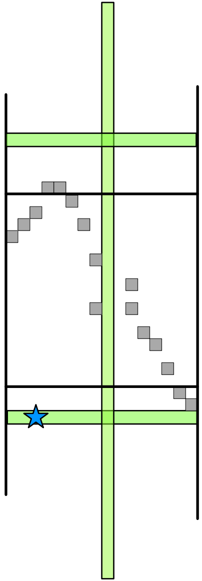

If there is a colum without obstacle, pass in it. Otherwise each column contains at most one obstacle.

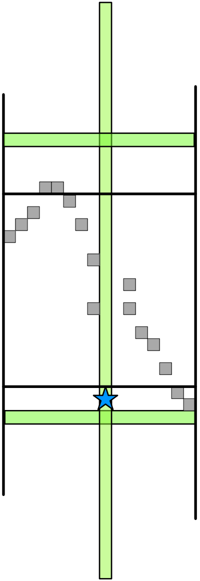

If there is a colum without obstacle, pass in it. Otherwise each column contains at most one obstacle.

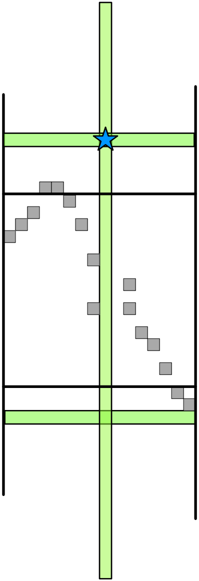

If there is a colum without obstacle, pass in it. Otherwise each column contains at most one obstacle.

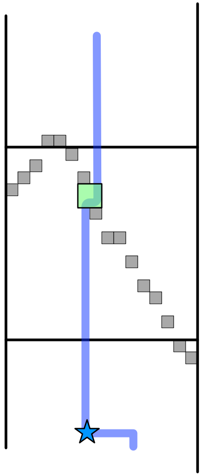

If obstacles in neigboring columns are not close, pass between them.

The case remains (marginal case) when the obstacles form a “chain”.

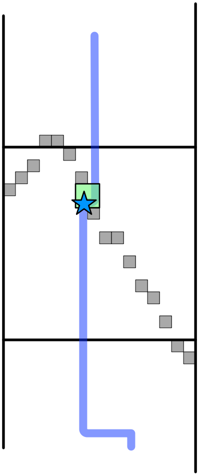

If obstacles in neigboring columns are not close, pass between them.

The case remains (marginal case) when the obstacles form a “chain”.

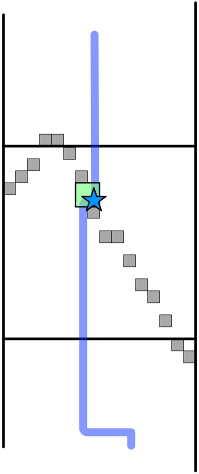

If obstacles in neigboring columns are not close, pass between them.

The case remains (marginal case) when the obstacles form a "chain".

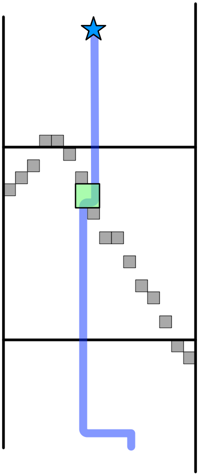

If obstacles in neigboring columns are not close, pass between them.

The case remains (marginal case) when the obstacles form a “chain”.

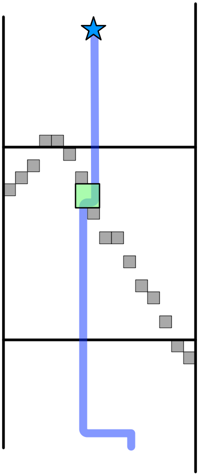

If obstacles in neigboring columns are not close, pass between them.

The case remains (marginal case) when the obstacles form a “chain”.

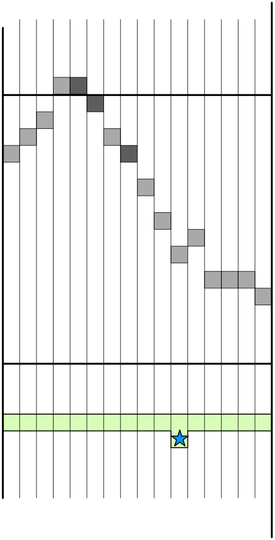

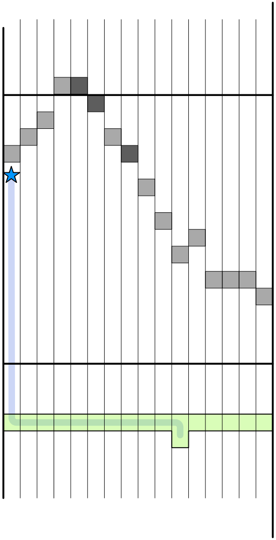

Preparing to attack in column 1.

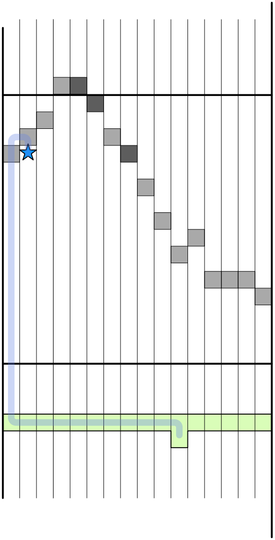

Attack in column 1 and its continuation in column 2 failed.

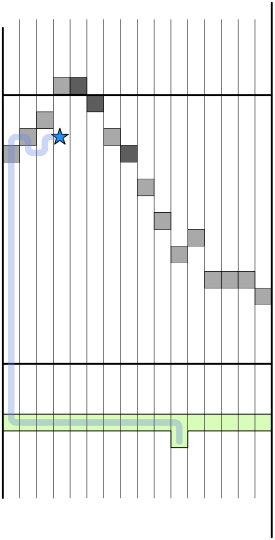

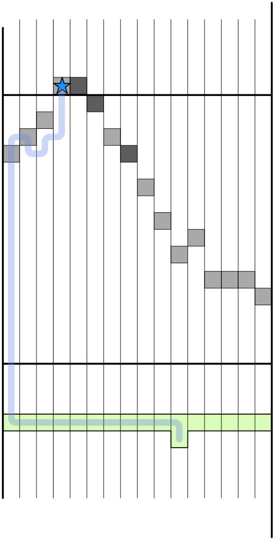

Attack in columns 3, 4 failed, too. Columns 5,6 are not good, we will not even attack: but we have to evade them.

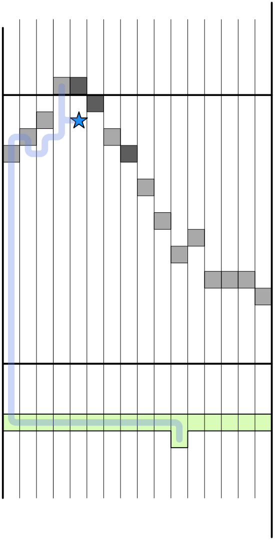

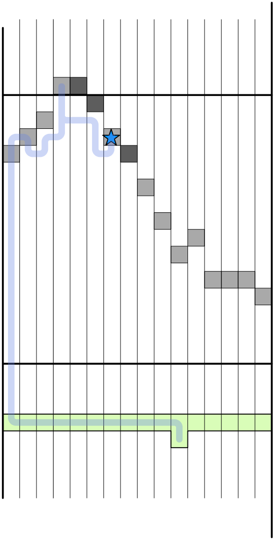

Attack in column 7 fails again.

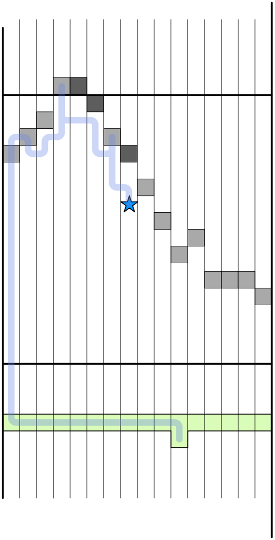

Evading in column 8.

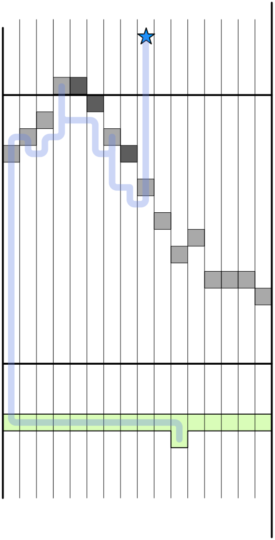

Escape after successful attack in column 9.

Paying for an attack series

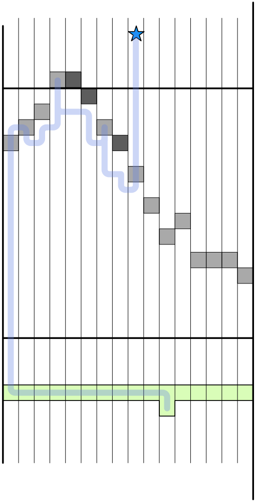

New problems

- The path shown here intersects itself (even if only slightly, in the retreats). But the time bound $\tau_{\mathrm{gc}}(U) + \rho\tau(U)$ only works for simple (self-nonintersecting) paths.
- Who pays for the initial digression?

Paying for an attack series

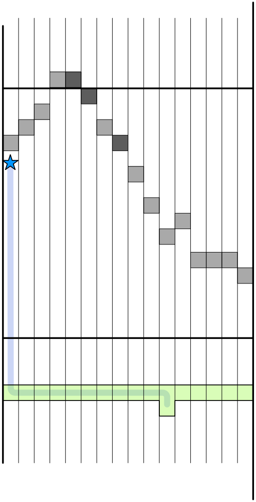

New problems

- The path shown here intersects itself (even if only slightly, in the retreats). But the time bound $\tau_{gc}(U) + \rho\tau(U)$ only works for simple (self-nonintersecting) paths.
- Who pays for the initial digression?

Solution for the retreats: many types of move. The body of an attack move will (essentially) include the possible retreat.

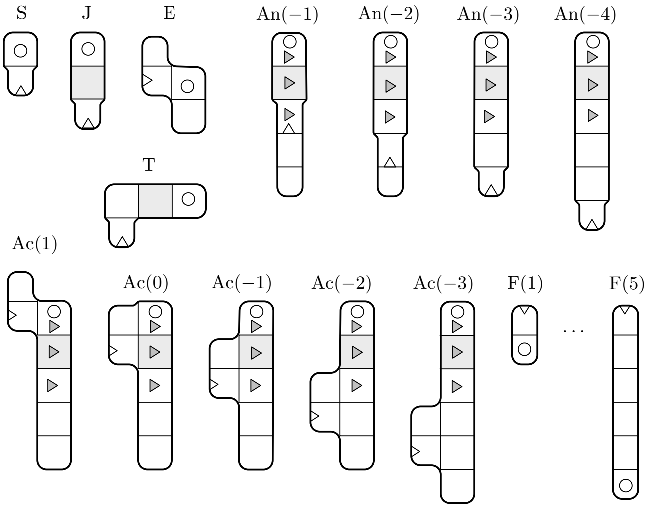

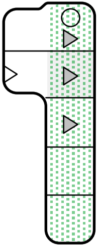

A continuing attack move.

Jump. Step. Turn.

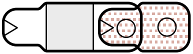

Jump. Step. Turn.

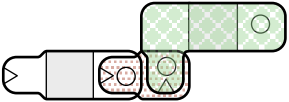

Jump. Step. Turn.

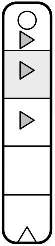

New attack. Continuing. Continuing.

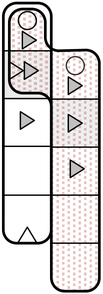

New attack. Continuing. Continuing.

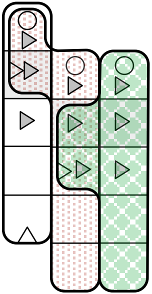

New attack. Continuing. Continuing.

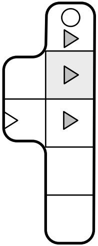

Continue. Finish. Step.

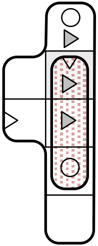

Continue. Finish. Step.

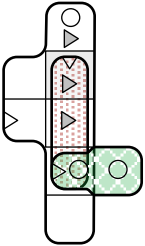

Continue. Finish. Step.

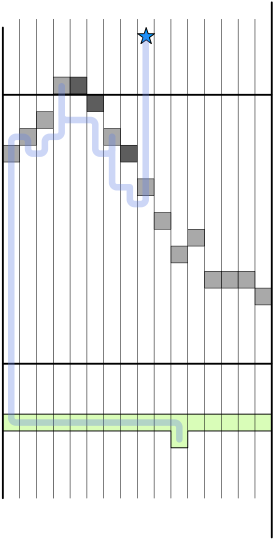

- Columns of dark squares: no attack.
- First column and columns after the dark squares: new attacks.
- All other columns: continuing attacks, all failed but the last one. They need no retreat.
- Each dark square pays for the digression around it.
- Problem: The body of a failed attack includes its obstacle: who will pay?

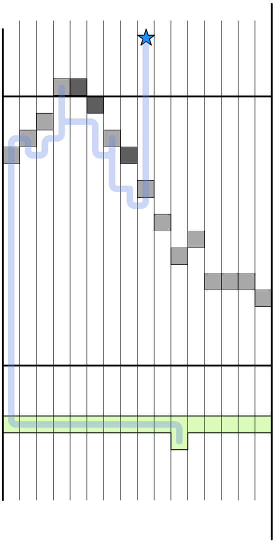

Solution: we define the problem away (push it into recursion). Let each *failed continuing attack* pay for itself (even bring profit).

New time bound:

$$\tau_{\mathrm{gc}}U + \rho_1\mu(U) - n\rho_2 B,$$

where B is the cell size and n the number of failed continuing attacks in the path.

Scale-up

OK, so what is the strategy?

Essentially, it is an implementation of each big move by a series of small moves (in case of attacks, contingent on what the Devil does).

(Of course, this is done recursively. On every sufficiently high level we are just trying to translate a big horizontal step.)

Let us recall some of the earlier examples.

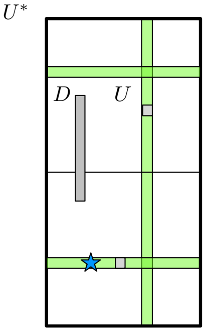

Now we understand how to pass the little grey squares: the jump moves (which are like attacks with a little more weight guarantee) will just jump over them.

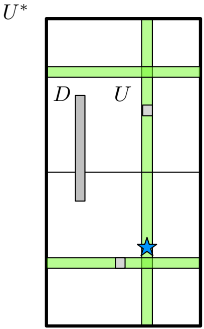

Now we understand how to pass the little grey squares: the jump moves (which are like attacks with a little more weight guarantee) will just jump over them.

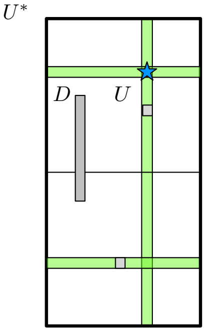

Now we understand how to pass the little grey squares: the jump moves (which are like attacks with a little more weight guarantee) will just jump over them.

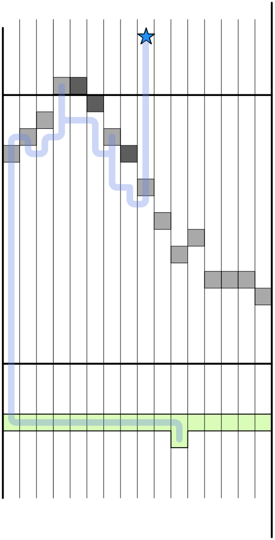

This is the implementation of a jump or attack.
Its details are contingent on what the Devil does, since by the time we are about to perform the next small attack, it may be infeasible (the little grey square may turn too dark). Then we evade instead of attacking.

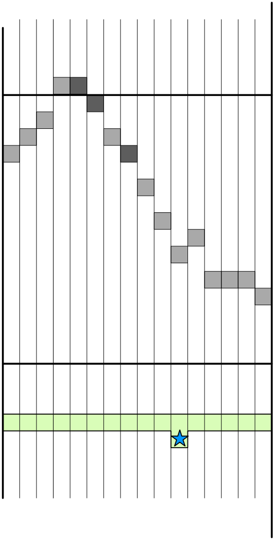

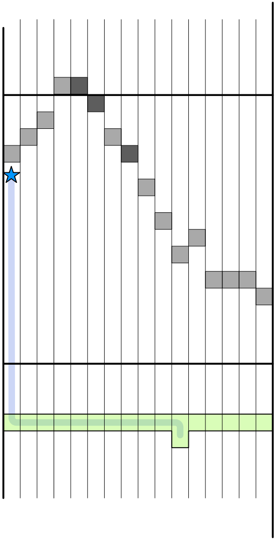

Preparing to attack in column 1.

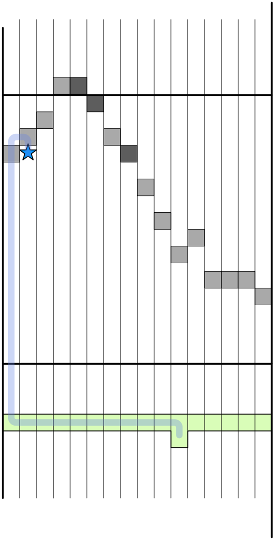

Attack in column 1 and its continuation in column 2 failed.

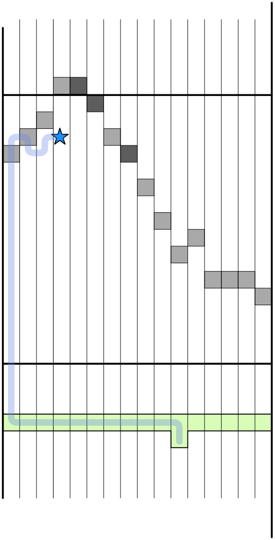

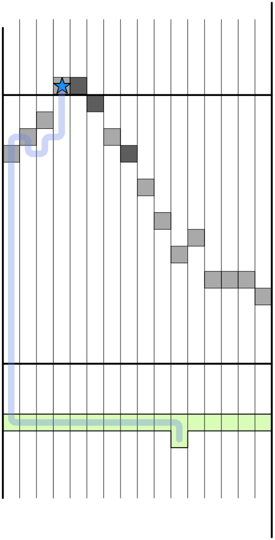

Attack in columns 3, 4 failed, too. Columns 5,6 are not good, we will not even attack: but we have to evade them.

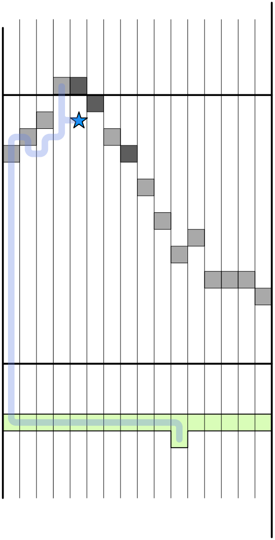

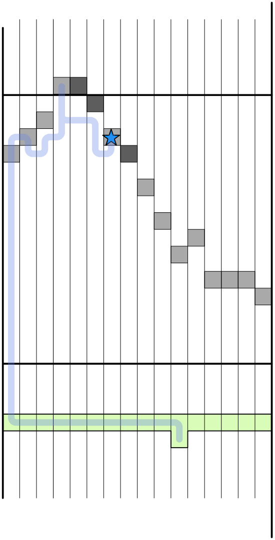

Attack in column 7 fails again.

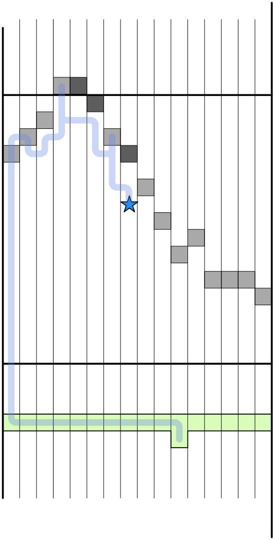

Evading in column 8.

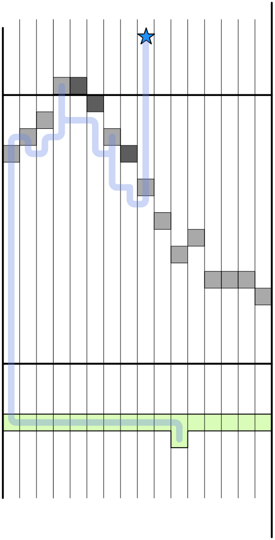

Escape after successful attack in column 9.

Of course, there are some more formal details, (and quite a bit more dirty details). But I hope that I have conveyed the spirit of the solution.