Wythoff's variant of NIM

Matt Roughan

August 17, 2017

1 Problem

There are two piles of pebbles on a table. There are two players, A and B, who move alternatively. Player A moves first. The rules of the game are the same for both players: at each move they can remove any number of pebbles provided they come from the same pile or the same number of pebbles from both piles. The winner is the player who takes the last pebble. What is the winning strategy for player A, if one exists?

To be more specific, imagine there are two piles of pebbles, say 31 and 48. What is the winning move of player A?

2 Notes

Notation is as follows. We denote the state of the game by the number of stones in each pile (N, M), so we will consider potential points on the lattice $Z^+ \times Z^+$. The following moves are possible:

- $(N, M) \rightarrow (N-k, M)$, for $k \leq N$, which corresponds to moving along a row (towards the y-axis).
- $(N, M) \rightarrow (N, M-k)$, for $k \leq M$, which corresponds to moving along a column (towards the x-axis).
- $(N, M) \rightarrow (N-k, M-k)$, for $k \leq N$ and $k \leq M$, which corresponds to moving along a diagonal (towards the origin).

We illustrate a set of possible moves in Figure 1.

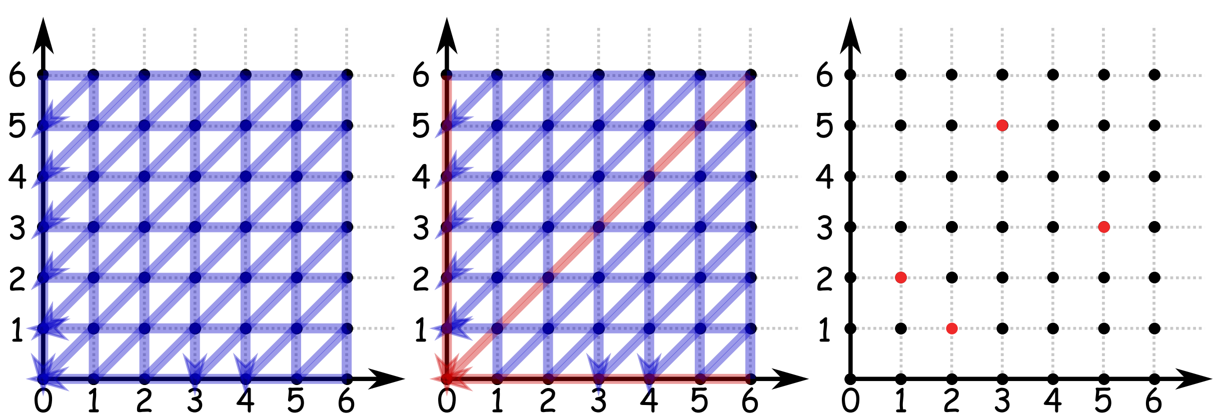

Figure 1: Possible moves, noting that we do not illustrate all, but rather illustrate sets of moves along horizontal, vertical, and diagonals by a single arrow (in most cases). Figure (a) illustrates moves, Figure (b) positions that can reach (0,0) in one hop, and (c) "safe" positions defined recursively as positions that cannot reach another safe position or (0,0) in one move.

These are a few notes on a strategy, leading to some general comments about the solution.

1. The winner is the one who moves the game to the state $(0, 0)$. A winning strategy involves making a move either directly to $(0, 0)$, or to a another state such that the other player does not have a winning strategy.

   There are positions from which player A cannot win with a single move. Therefore, if A has a winning strategy from one of these points, there must also be other points where there is no winning strategy, otherwise B can win.

2. Everything is symmetric, so anything said of (N, M) also applies to (M, N).

3. There are a set of trivial positions from which A can win in one move, *e.g.,* $(N, 0)$, $(0, N)$, (N, N).

4. The position $(2, 1)$ has no winning strategy, *e.g.,* if B is in this position it can make the following moves:

   - $(2, 1) \to (0, 1)$
   - $(2, 1) \to (2, 0)$
   - $(2, 1) \to (1, 1)$

   Each final position has a winning strategy for A as noted above, so one won't win from $(2, 1)$.

5. As a result of $(2, 1)$ being a "safe point" for A to move to, then A has a winning strategy from all points $(2, N)$, $(N + 1, N)$, and the symmetric counterparts (for $N > 2$).

6. If a point has *no* winning strategy, then any point you can reach from that point must have a winning strategy.

7. We can continue this process of finding a "safe point", *i.e.,* point where there is no winning strategy. The points have properties:

   - There can be no more than one in any row, column or diagonal of $Z^+ \times Z^+$, because if there were, the (larger) safe point would have a winning strategy (go to the other one in the same column/row/diagonal).
   - There is at least one safe point that can be reached in one move, from every point that is not safe (unless you can win directly from that point), because, by definition, if a point is not safe, it must have a winning strategy.
   - From symmetry, if there is a safe point in a row, there is also one in a column.

8. *Hypothesis: there is exactly one safe point in every row, column and diagonal?* The only part of this that we haven't proved is that there is at least one safe point in each diagonal.

The net result of these observations is that we can derive an algorithm for finding safe points, namely

```
(1) z = {1,2,3,... }
(2) for d=1,2,3,...
        (a) find the smallest pair of integers (i,j), i>j in z such that i-j=d
        (b) add (i,j) and (j,i) to the list of safe points
        (b) remove i and j from z
    end for
```

The algorithm finds points that have a unique diagonal, row, and column, and finds exactly one per row, column and diagonal. There are other ways to find these points, as we shall see, but the fact that we have an algorithm that finds a safe point for every column, row and diagonal proves the hypothesis above. By definition, the set of safe points must be unique, and the above list satisfies the requirements on the safe points imposed by the solution. The fact that it has one safe point in each diagonal is therefore proof that the safe points satisfy this property.

Figure 2 illustrates the results. The green circles indicate safe points (on two different scales), and all other points have a winning strategy. A list of points is also included in Table 1.

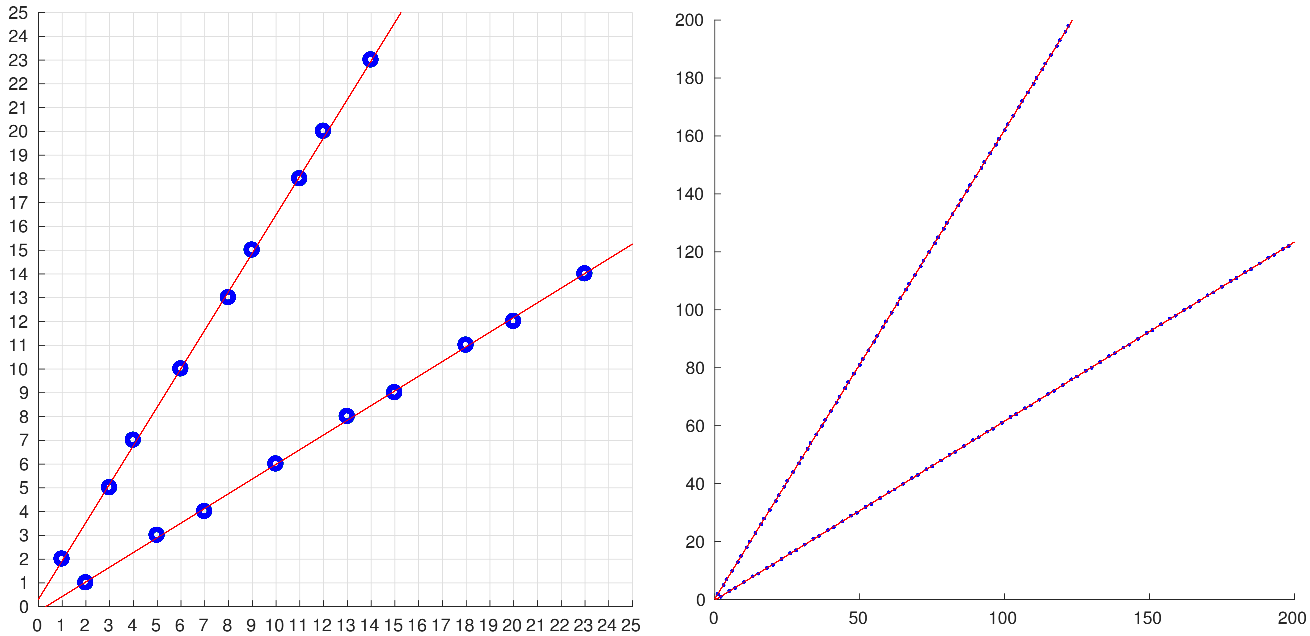

(a) (b)

Figure 2: Safe points and lines fitted to these points..

N	M	$M-N$	$N_i - N_{i-1}$
1	2	1	N/A
3	5	2	2
4	7	3	1
6	10	4	2
8	13	5	2
9	15	6	1
11	18	7	2
12	20	8	1
14	23	9	2
16	26	10	2
17	28	11	1
19	31	12	2
21	34	13	2
22	36	14	1
24	39	15	2
25	41	16	1
27	44	17	2
29	47	18	2
30	49	19	1
32	52	20	2

Table 1: The first 20 safe points (not including symmetric points (M, N)). Note that $M - N$ tells us which diagonal the point is in.

3 Winning Strategy

A's winning strategy is now clear. If A starts at a safe point, she/he is stuffed: there is no winning strategy unless B is stupid. Otherwise, A should find a move to a safe point (there must be one unless you start on such a point) and then move to it. B must then move to a non-safe point, whereupon A can repeat the same strategy until he/she wins.

For example: $(N, M) = (31, 48)$ then A could take 4 off each pile to move to $(27, 44)$ (there are two other possible moves). The next move (by A) will depend on B's move, but is easy enough to calculate.

The trick is, if you start with $N, M \gg 1$ then how do you (easily) calculate the safe points. The safe points found above have several interesting properties. Firstly, there is no (immediately) obvious pattern that shows which points will be safe points. However, we learn something by considering Figure 2. The graph also includes two lines plotted through the set of points above and below the line y=x. The lines were obtained by linear regression on the points. The lines have slope approximately equal to ϕ and $1/\phi$, *i.e.,* 1.6180 and 0.6180, the golden mean and golden ratio.

The two quantities have some interesting properties, in particular, they are closely related to the Fibonacci sequence. When we look closer at the safe points, in particular, the steps between them, this is very closely related to the Fibonacci string. The string can be generated in several different ways (see Appendices below), but the way of most relevance here is to start with a string 2, and replace any value of 2 in the string by 21, and any value of 1 by 2. See below for details. The first 20 terms of the Golden string = `21221212212212122121`, which we can see matches $N_i - N_{i-1}$ in Table 1.

All of the safe points all seem to lie near the two lines noticed above, but there are other points close to the line that are not safe points. So it would be nice to exploit the Fibonacci numbers to find safe points easily.

Why do we see this interesting structure, and is it real, or just a temporary coincidence that stops at some points? The answer lies in the relationship between the generation process for the safe points, and the generation of the Fibonacci string.

3.1 Finding the safe points

Start by partitioning the safe points into those above (U), and those below (L) the diagonal line $x = y$. For some cutoff value k, denote

$$\begin{aligned} L_k &= \{(N_i, M_i) | N_i > M_i \text{ and } M_i \leq k\}, \\ U_k &= \{(M_i, N_i) | N_i < M_i \text{ and } N_i \leq k\}, \end{aligned}$$

where (N_i, M_i) are the safe points sorted in order so that $N_{i-1} < N_i$. $U = \lim_{k\to\infty} U_k$ and $L = \lim_{k\to\infty} L_k$.

Our proposition is that the differences $d_i = N_i - N_{i-1}$ for L follow the Fibonacci string (with characters 3 and 2, instead of 1 and 0). We have already shown it is true for small values of k. We will proceed to demonstrate it for larger values using induction.

Assume we have the found the set of all safe points L_m for some m, and that differences d_i follow the Fibonacci string. Now, let us try to find L_n for $n > m$, in particular, let us find it for $n = \max\{N_i | (N_i, M_i) \in L_m\}$. We know that L_n will be the reflection of U_n through the line $x = y$. We can find U_n directly from L_m by noting that there is exactly one safe point in each column. Hence, the sets L_m and U_n partition the columns $\{1, 2, \ldots, n\}$.

For example take $m = 4$. From Table 1 we get $L_4 = \{(2, 1), (5, 3), (7, 4)\}$. Consequently $n = 7$, and $U_7 = \{(1, 2), (3, 5), (4, 7), (6, 10)\}$. We can see that there is one safe point in each of columns 1,2,. . . ,7.

Hence, we can construct U_n by filling in the gaps in the columns of L_m, *e.g.,* for $m = 4$, there are gaps in columns $1, 3, 4$ and 6. The differences $d_i = N_i - N_{i-1}$ for L are important. When such a difference is 2, there is one gap to be filled in, and when the difference is 3, there are two gaps to fill.

Now, adding the requirement obtained by our construction algorithm that the distance from the diagonal $x = y$ of each successive safe point in L (ordered by N_i) must increase by one, we get $N_i - M_i = i$ for points in L, and $M_i - N_i = i$ for points in U. In the example above, we see that this means when we fill in the gap between columns 5 and 7 (in L_4) we add the 4th point of U, and so the co-ordinates must be $(6, 10)$.

In general, this means that where we have a single column to fill, there will be gap before its previous member of U, and so its differences (in co-ordinate) from the previous member of U will be $(2, 3)$, where the second term is dictated by $M_i - N_i = i$. Where there are two adjacent columns to fill in, the first will be preceeded by a gap, and the next will be directly adjacent, so there will be two new points in U_n, and these will have differences (in co-ordinate) of $(2, 3)$, $(1, 2)$ respectively. When we reflect this to obtain L_n, the gaps will become, respectively $(3, 2)$ and $(3, 2), (2, 1)$.

To summarize, where L_m had a gap of 2 (or more precisely $(2, 1)$) we will find a gap of $(3, 2)$ in L_n. However, where L_m had a gap of $(3, 2)$, this will be replaced by two differences $(3, 2), (2, 1)$. This is exactly the construction rule for the Fibonicci string! So we have demonstrated the inductive step from L_m to L_n, and hence in the limit it should be true for L (and consquently for U in reflected terms).

Let's look at the upper set U. The differences between points on this line take the form $(1, 2)$ or $(2, 3)$ with the number of each jump being determined by the Fibonacci string. We know that the number of points of each type will follow the form h_n and g_n respectively, and that g_n follows a Fibonacci sequence rule, and $h_n = g_{n-1}$, so the limiting value of such points will be

$$\lim_{n\to\infty} \frac{2h_n + 3g_n}{h_n + 2g_n} = \lim_{n\to\infty} \frac{2[h_n + g_n] + g_n}{g_{n+1} + g_n} = \lim_{n\to\infty} \frac{2g_{n+1} + g_n}{g_{n+2}} = \lim_{n\to\infty} \frac{g_{n+1} + g_{n+2}}{g_{n+2}} = \lim_{n\to\infty} \frac{g_{n+3}}{g_{n+2}} = \phi.$$

Hence the slope of the upper line approaches ϕ (as we include more points), and the slope of the lower line will be (by symmetry) $1/\phi = \phi - 1$.

Appendices:

A The Fibonacci sequence

The Fibonacci sequence is $1, 1, 2, 3, 5, 8, 13, 21, \ldots$. It is the sequence defined by

$$f_{n+1} = f_n + f_{n-1},$$

for $n \geq 1$, and $f_0 = f_1 = 1$. It has some remarkable properties, for instance

$$f_n = \frac{\phi^n - (1-\phi)^n}{\sqrt{5}} = \frac{\phi^n - (-\phi)^{-n}}{\sqrt{5}} = \left\lfloor \frac{\phi^n}{\sqrt{5}} + \frac{1}{2} \right\rfloor,$$

where $\phi = 1.6180...$ is the Golden ratio, and $\lfloor \cdot \rfloor$ is the floor function.

Of importance here is the a result that can be easily derived form those above, *i.e.,*

$$\lim_{n\to\infty} \frac{f_{n+1}}{f_n} = \phi.$$

There are many other interesting results, but they are not of relevance here. See, for instance, `http://en.wikipedia.org/wiki/Fibonacci_number`, or `http://www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/fib.html` for more information.

B The Fibonacci string

The Fibonacci string can be defined as the limit of a series of string defined by starting with string

$$s_0 = 1$$

and repeatedly applying the replacement rules

$$\begin{array}{rcl} 1 & \to & 10 \\ 0 & \to & 1 \end{array}$$

So the series of strings is

$$\begin{array}{rcl} s_0 & = & 1 \\ s_1 & = & 10 \\ s_2 & = & 101 \\ s_3 & = & 10110 \\ s_4 & = & 10110101 \\ s_5 & = & 1011010110110 \\ s_6 & = & 101101011011010110101 \end{array}$$

The Golden string is the limit of this process.

Note that the 1 and 0 in the string are just characters, *i.e.,* their numerical value has no meaning, and so we could equally use x and y, or any other character. In the above work we use 1 and 2.

The string has certain properties. Firstly, prefixes are preserved, i.e, the strings are identical up to their length. Secondly the number of terms in the sequence follows a Fibonacci sequence in itself, but even more importantly, if we count the numbers of 1's and number of 0's these also follow a Fibonacci sequence. We can see this as follows. Denote

$$\begin{array}{rcl} g_n & = & \text{the number of 1's in } s_n, \\ h_n & = & \text{the number of 0's in } s_n. \end{array}$$

i	i*Phi	trunc(i*Phi)	diff	Fibonacci string
1	1.618034...	1		
2	3.223606...	3	2	1
3	4.854101...	4	1	0
4	6.472135...	6	2	1
5	8.090169...	8	2	1
6	9.708203...	9	1	0
7	11.326237...	11	2	1

Table 2: The Fibonacci string

Then, using the replacement rules above we can see that

$$\begin{aligned} g_{n+1} &= g_n + h_n, \\ h_{n+1} &= g_n. \end{aligned}$$

Now, substituting the second equation in the first we get

$$g_{n+1} = g_n + g_{n-1}.$$

Likewise, we can get the same defining equation for h_n, but note that it is delayed by one step, *i.e.,* $h_1 = h_2 = 1$, as opposed the the normal initial conditions of the Fibonacci process. This leads to the result giving the total length of s_n to be f_{n+1}. As a result

$$\lim_{n\to\infty} \frac{g_n}{h_n} = \phi.$$

Interestingly, $\lfloor i \times \phi \rfloor$ gives a series of numbers determined by the sum of the sequence (if we use 2's and 1's instead of 1's and 0's), *e.g.,* see Table 2. Another related result is that if we plot a line with slope ϕ and write a string with 0 where it crosses a horizontal grid line (at an integer), and 1 where it crosses a vertical grid line, then the result will be a Fibonacci string.

See `http://www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/fibrab.html` for more details. We use `goldstr` to generate the string, *e.g.,* see `http://www.ibiblio.org/pub/linux/apps/math/calc/!INDEX.short.html`.

C The Golden ratio

The Golden ratio is also related as we have seen above. It is an irrational number

$$\phi = \frac{1+\sqrt{5}}{2} \simeq 1.6180339887...$$

It satisfies the quadratic equation

$$\phi^2 - \phi - 1 = 0$$

and as a result has several interesting properties, *e.g.,*

$$\phi - 1 = 1/\phi,$$

which is sometimes called the Golden section.

See `http://en.wikipedia.org/wiki/Golden_ratio`, `http://www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/phi.html` and `http://www.mcs.surrey.ac.uk/Personal/R.Knott/Fibonacci/fibnat2.html` for more details.