

Game Theory and Algorithms*

Lecture 19: Nim & Impartial Combinatorial Games

May 17, 2011

Summary: We give a winning strategy for the counter-taking game called Nim; surprisingly, it involves computations with binary numbers. We then proceed to explain how a slight generalization gives the optimal strategy for the larger class of all *impartial combinatorial games*. (Nim is thus “complete” for these games.)

1 Nim

The game of Nim works as follows:

- The initial configuration has several piles; each pile has several counters. (The prototypical example is one pile of size 3, one pile of size 5, and one pile of size 7.)
- Player one moves first, then player two, then player one, etc.
- On your turn, you pick a pile and delete one or more counters from it. (So as the game proceeds, the piles get smaller, and some piles become empty/totally removed.)
- When all counters are gone, the next player to move loses. (I.e. you win when you delete the last counter(s).)

The winning strategy for Nim will generalize to the following class of *impartial combinatorial games*: such a game has a finite set P of positions (states), one of which is the *initial* position $p_0 \in P$. From every position $p \in P$ a certain set of actions $A(p)$ are possible. Gameplay is as follows:

- The current position p of the game is initially the initial position p_0 .
- Player one moves first, then player two, then player one, etc.
- On your turn, pick a new position $p' \in A(p)$ replacing the current position, so $p := p'$.
- Once $A(p) = \emptyset$ (no moves are possible) the next player to move loses.

* Lecture Notes for a course given by David Pritchard at EPFL, Lausanne.

Assumption: in the directed graph with vertex set P and edges from p to p' for all p and all p' , there are no directed cycles. This forces the game to end in a finite amount of time, since at most each state will occur once.

Clearly Nim fits in this framework: e.g. in the variant with initial piles 3-5-7, we can model the game using the set of 192 positions

$$\{0, 1, 2, 3\} \times \{0, 1, 2, 3, 4, 5\} \times \{0, 1, 2, 3, 4, 5, 6, 7\}.$$

In the rest of the lecture:

- We give an easy general method to play optimally in impartial combinatorial games, which takes time proportional to the number of positions.
- We show that in the special case of Nim, there is a much simpler method: there is a non-obvious formula which tells you how to play optimally.
- We show that the Nim formula generalizes: when games are split into multiple independent parts, the formula gives a much more efficient means for optimal play.

2 \mathcal{P} - and \mathcal{N} -Positions

The most straightforward way to analyze an impartial combinatorial game uses a sort of backwards induction. We will give a proof by example, which introduces useful terminology and ideas.

In the *10-coin game*, we have 10 coins in a pile. On each player's turn, they can delete 1, 2, or 3 coins. Player 1 moves first and then alternates with player 2. The player to take the last coin wins.

Note that this is an impartial combinatorial game with $P = \{0, 1, \dots, 10\}$.

What is an optimal strategy? It is not too clear what to do at the start of the game, but at the end of the game things are clear. If it is your turn and there are no coins left, then you lose. On the other hand if it is your turn and there are 1, 2 or 3 coins left, you can win by taking them all.

Definition 1. A position p is called a \mathcal{P} -*position* if when the game reaches state p , the player who moved **p**reviously has a winning strategy. It is called an \mathcal{N} -*position* if the player who will move **n**ext has a winning strategy.

(Formally, a *strategy* specifies which action $p' \in A(p)$ to choose for each position p that could arise in the future.)

So for the 10-coin game, 0 is a \mathcal{P} -position, while 1, 2, and 3 are \mathcal{N} -positions. The backwards analysis can be continued, using the following observation.

Lemma 2 (Recursive \mathcal{N}/\mathcal{P} characterization). (i) If for all $p' \in A(p)$, p' is a \mathcal{N} -position, then p is a \mathcal{P} -position (this includes the case $A(p) = \emptyset$). (ii) If there exists a $p' \in A(p)$ such that p' is a \mathcal{P} -position, then p is an \mathcal{N} -position. (iii) Both (i) and (ii) hold with if-and-only-if.

Proof. (i) Clearly if $A(p) = \emptyset$, p is a \mathcal{P} -position: the next player cannot move & loses and the previous wins. Otherwise, no matter what the upcoming player X does, they move to an \mathcal{N} -position, which by definition allows the opponent Y of X to have a winning strategy. Altogether this is a winning strategy for the player Y who moved previous to the arrival at p .

(ii) This is simpler: the next player X to move from p can move to p' . Since relative to p' player X was previous, they then have a winning strategy.

(iii) This amounts to seeing that the contrapositive of (i) is the converse of (ii) and vice-versa (and using that every position is \mathcal{N} or \mathcal{P} but not both). To say it more constructively, when we test all $p' \in A(p)$, note that either they are all \mathcal{N} (so p is \mathcal{P}) or one is \mathcal{P} (so p is \mathcal{N}). \square

Again in our example, this means that 4 is a \mathcal{P} -position: a player forced to move from a pile of size 4 doesn't have a winning strategy, since their opponent can just empty the pile. Then 5, 6, 7 are \mathcal{N} -positions, 8 is a \mathcal{P} -position, and 9 and 10 are \mathcal{N} -positions. I.e. in the 10-coin game, the first player has a winning strategy.

3 Winning in Nim

The notation about \mathcal{N} and \mathcal{P} positions lets us give a coherent proof of Bouton's Theorem. The result uses some notation. The¹ *binary representation* of a whole number n is a string $s_t s_{t-1} s_{t-2} \cdots s_2 s_1 s_0$ with each $s_i \in \{0, 1\}$ and $n = \sum_i s_i 2^i$. I.e., the binary representation of 23 is 10111 since

$$23 = 1 \cdot 2^4 + 0 \cdot 2^3 + 1 \cdot 2^2 + 1 \cdot 2^1 + 1 \cdot 2^0 = 16 + 4 + 2 + 1.$$

The *exclusive-or* of two or more binary strings is obtained by right-aligning them, adding each column modulo 2, and interpreting the column sums as a new binary number. We denote the exclusive-or by \oplus . For example,

$$\begin{array}{r} 23 \qquad 10111 \\ 37 \qquad 100101 \\ \hline 23 \oplus 37 \quad 110010 \end{array}$$

(note the 3rd column from the right showed the "column addition modulo 2") and since 110010 in binary is 50 in decimal, $23 \oplus 37 = 50$. The exclusive-or satisfies some easy properties like $x \oplus y \oplus z = (x \oplus y) \oplus z$, $x \oplus y = y \oplus x$, $x \oplus x = 0$, $x \oplus 0 = x$, and when $x \neq y$, $x \oplus y \neq 0$. We write $\bigoplus_i m_i$ as short for $m_1 \oplus m_2 \oplus \cdots \oplus m_k$.

Theorem 3 (Bouton, 1901). *In the game Nim, denote an arbitrary position by $p = (p_1, p_2, \dots, p_k)$ where there are k piles, and the i th has p_i counters. Then p is a \mathcal{P} -position if and only if $\bigoplus_i p_i = 0$.*

¹The representation is unique if we assume there are no leading zeroes, i.e. if $t > 0 \Rightarrow s_t = 1$.

Proof. We prove the theorem by induction on the number $p_1+p_2+\dots+p_k$ of counters. Clearly if there are no counters, the position is a \mathcal{P} -position, and as claimed we have $\bigoplus_i p_i = 0$. So by induction we may assume there is a positive number of counters, and that Bouton's theorem holds for all $p' \in A(p)$, since all such p' have a smaller number of counters.

In both cases, we combine the induction with the recursive \mathcal{N}/\mathcal{P} characterization (Lemma 2).

$\bigoplus_i p_i = 0 \Rightarrow p$ is a \mathcal{P} -position: We need to show that every $p' \in A(p)$ is an \mathcal{N} -position. By the rules of Nim, p' is the same as p except one pile x was decreased: $p' = (p'_x, p_{-x})$ with $p'_x < p_x$. From the basic properties of exclusive-or,

$$\bigoplus_i p'_i = \bigoplus_{i \neq x} p_i \oplus p'_x = \left(\bigoplus_i p_i \right) \oplus (p_x \oplus p'_x) = 0 \oplus \text{nonzero}.$$

By induction this p' is an \mathcal{N} -position, as needed.

$\bigoplus_i p_i \neq 0 \Rightarrow p$ is an \mathcal{N} -position: We need to find a winning move for the next player. Let $\bigoplus_i p_i = V \neq 0$. Along the lines of the previous argument, it is enough to find a pile x such that $p'_x := p_x \oplus V$ satisfies $p'_x < p_x$, because $p' := (p'_x, p_{-i})$ would be a valid move ($p' \in A(p)$) as well as a winning move (p' is \mathcal{P}):

$$\bigoplus_i p'_i = \bigoplus_{i \neq x} p_i \oplus p'_x = \left(\bigoplus_i p_i \right) \oplus (p_x \oplus p'_x) = V \oplus V = 0.$$

This x exists since (1) $p_x \oplus V < p_x$ iff the leftmost column with a 1 in V also has a 1 in p_x and (2) there must be an odd number of such x 's. \square

So in the standard 3-5-7 Nim, how should we play optimally?

$$\begin{array}{r} 3 \qquad 11 \\ 5 \qquad 101 \\ 7 \qquad 111 \\ \hline 3 \oplus 5 \oplus 7 \quad 001 \end{array}$$

One of the three winning moves is to remove 1 counter from the pile of 7, leaving your opponent facing a game with sum $3 \oplus 5 \oplus 6 = 0$. Suppose for example they proceed to take 4 counters from the pile of 6, leaving

$$\begin{array}{r} 3 \qquad 11 \\ 5 \qquad 101 \\ 2 \qquad 10 \\ \hline 3 \oplus 5 \oplus 2 \quad 100 \end{array}$$

Then your only winning reply is to take 4 counters from the pile of 5, again changing the game sum to $3 \oplus 1 \oplus 2 = 0$. You proceed in this way, always bringing the sum back to 0 after your move.

4 Sprague-Grundy Numbers

The broad generalization of Bouton’s theorem, due to Roland Sprague (1936) and Patrick Grundy (1939), says that for every impartial combinatorial game G , there is exactly one integer m so that G is equivalent (in a way made precise below) to a Nim pile with exactly m counters. The theory introduces the very natural idea of game sums, which leads to an efficient algorithm for playing optimally in sums of games.

The *sum* of two impartial combinatorial games is obtained by placing the games side by side, with each player choosing on their turn to make a move in one or the other, and the game terminating once no moves are possible in either game. Formally,

Definition 4. For $G = (P, p_0, A)$ and $G' = (P', p'_0, A')$, the sum $G + G'$ has positions $P \times P'$, initial position (p_0, p'_0) , and from position (p, p') valid moves are $\{p\} \times A'(p') \cup A(p) \times \{p'\}$.

The game sum operator naturally extends to $(G_1 + G_2) + G_3 = G_1 + (G_2 + G_3) =: G_1 + G_2 + G_3$ which corresponds to players having their choice of 3 games to move in. Notice that $\text{Nim}(p_1, p_2, \dots, p_k)$ is the same as the game sum $\text{Nim}(p_1) + \dots + \text{Nim}(p_k)$.

Definition 5. The *type* of G is \mathcal{P} (resp. \mathcal{N}) if its p_0 is a \mathcal{P} -position (resp. \mathcal{N} -position).

This definition just makes wording easier. Note that type \mathcal{P} (resp. \mathcal{N}) means player 2 (resp. 1) has a winning strategy.

Lemma 6 (Copycat lemma). *For any impartial combinatorial game G , $G + G$ has type \mathcal{P} .*

Proof. We need to establish that the second player has a winning strategy. Here it is: copy what the first player just did, but in the other pile. \square

Lemma 7 (\mathcal{P} -ignorance). *If P has type \mathcal{P} , the type of $P + Q$ is the same as the type of Q .*

Proof sketch. The main idea is to establish a winning strategy for the appropriate player in $P + Q$. It is the same as the winning strategy in Q , except that when their opponent moves in P , they “cancel” the move in P by bringing it back from an \mathcal{N} -position to a \mathcal{P} -position. \square

Corollary 8. *If $G + G'$ has type \mathcal{P} , then for any game H , the types of $H + G$ and $H + G'$ are the same. (I.e. G and G' are equivalent in determining winners of game sums.)*

This motivates the following, which you can prove is an equivalence relation:

Definition 9. Games G and G' are *equivalent* if $G + G'$ has type \mathcal{P} .

Theorem 10 (Sprague-Grundy). *Every impartial combinatorial game G has a unique integer g (its Grundy value $g(G)$) such that G is equivalent to $\text{Nim}(g)$. These g also satisfy*

Refines \mathcal{N}/\mathcal{P} : $g(G) = 0 \Leftrightarrow G$ has type \mathcal{P}

Xor rule: $g(G + H) = g(G) \oplus g(H)$

Mex rule: $g(G) = \min(\mathbb{Z}_{\geq 0} \setminus \{g(G') \mid G' \in A(G)\})$.

In the theorem statement, $G' \in A(G)$ is being used for a shorthand: where $G = (P, p_0, A)$, each $p' \in A(p_0)$ gives rise to the game $(P, p', A) \in A(G)$. I.e. $A(G)$ is the set of all games which are “ G and then one move has happened.”

The name *mex* rule comes from *minimal excluded*: the Grundy value of G equals the minimum nonnegative integer which does *not* appear in the Grundy values of games which are valid moves from G .

Here is an example of using the theorem and its rules, before proving it.

Example 11 (Grundy). Consider the following counter/pile game: on your turn, take any pile of counters and split it into two piles of unequal size. Let X_i denote this game, starting with a single pile of i counters.

A pile must have size 3 or more to be split. So the Grundy value $g(X_i) = 0$ for $i = 0, 1, 2$ as these games are \mathcal{P} -positions (they admit no valid moves).

Now consider X_3 . It has exactly one valid move, which results in $X_1 + X_2$. By the xor rule,

$$g(X_1 + X_2) = g(X_1) \oplus g(X_2) = 0 \oplus 0 = 0$$

and by the mex rule, $g(X_3) = \min(\mathbb{Z}_{\geq 0} \setminus \{0\}) = 1$. We continue similarly:

- $g(X_4) = \min(\mathbb{Z}_{\geq 0} \setminus \{g(X_1 + X_3)\}) = \min(\mathbb{Z}_{\geq 0} \setminus \{1\}) = 0$
- $g(X_5) = \min(\mathbb{Z}_{\geq 0} \setminus \{g(X_1 + X_4), g(X_2 + X_3)\}) = \min(\mathbb{Z}_{\geq 0} \setminus \{0, 1\}) = 2$

etc.

Having computed this sequence of values, we can play the game optimally. On your turn, look at the sizes of the existing piles, and compute each pile’s Grundy value using the table generated as sketched above. If the xor of these numbers is 0, you will lose if the opponent plays perfectly. Otherwise, you can win by making a move such that the xor of the new piles’ Grundy values becomes 0.

Exercise. *Treblecross* is a degenerate version of tic-tac-toe: the board is 1-dimensional and only Xs are allowed. On your turn, you mark an X anywhere you like. The first player to cause three Xs in a row is the winner. Describe how to use the Sprague-Grundy method to play this game optimally. (Hint: in what sense does a new X cut the existing game into a sum of two smaller games?)

4.1 Proof of the Sprague-Grundy Theorem

To prove the various parts of the Sprague-Grundy theorem, take the mex rule as the definition of g . The main part of the proof is showing that G is always equivalent to $\text{Nim}(g(G))$, i.e. that $G + \text{Nim}(g(G))$ is always a \mathcal{P} -position for the g so defined. If we can prove this, then the other parts of the theorem follow without too much work:

- There could not be two distinct g_1, g_2 for which $G + \text{Nim}(g_i)$ is \mathcal{P} , since that would imply that the following game is \mathcal{P} :

$$(G + \text{Nim}(g_1)) + (G + \text{Nim}(g_2)) = (G + G) + (\text{Nim}(g_1) + \text{Nim}(g_2)) = \text{Nim}(g_1) + \text{Nim}(g_2) = \text{Nim}(g_1, g_2)$$

but $\text{Nim}(g_1, g_2)$ is easily seen to have a first-player winning strategy when $g_1 \neq g_2$ by using the Copycat Lemma;

- The xor rule follows by applying similar manipulations to Bouton’s theorem;
- Winners rule: for one direction if $g(G) = 0$ then $G + \text{Nim}(0)$ is of type \mathcal{P} but since no moves are possible in $\text{Nim}(0)$, the game $G + \text{Nim}(0)$ is the same as G . For the other direction if $g(G) > 0$ then $G + \text{Nim}(g(G))$ is of type \mathcal{P} and so all its valid moves are \mathcal{N} , in particular $G + \text{Nim}(0) = G$ is \mathcal{N} .

So, recapping, recursively define $g(G) := \min(\mathbb{Z}_{\geq 0} \setminus \{g(G') \mid G' \in A(G)\})$ (this is well-defined, say, by using induction on the number of positions that can be reached in the game); we will show by induction that $G + \text{Nim}(g(G))$ is always a \mathcal{P} -position, essentially by giving a winning strategy for player 2. In more detail, we show that no matter what player 1 does, player 2 can take the game back to a \mathcal{P} -position by a judicious choice of move.

Case 1: Player 1 makes a “move in the Nim component,” meaning that their initial move brings the game to $G + \text{Nim}(x)$ for some $x < g(G)$. By the mex-based definition of g , there is some $G' \in A(G)$ such that $g(G') = x$. Let player 2 make the corresponding move in the left component, bringing the game from $G + \text{Nim}(x)$ to $G' + \text{Nim}(x)$. By induction on the Sprague-Grundy theorem, this is a \mathcal{P} -position, as needed.

Case 2: Player 1 makes a “move in the G component,” meaning that their initial move brings the game to $G' + \text{Nim}(g(G))$ for some $G' \in A(G)$.

Case 2a: If $g(G') < g(G)$ then player 2 can reply by making a move in the Nim component to decrease its number of counters from $g(G)$ to $g(G')$. The game $G' + \text{Nim}(g(G'))$ is a \mathcal{P} -position, again by induction.

Case 2b: Note $g(G') \neq g(G)$ by definition of g . So we need only consider the case $g(G') > g(G)$. Again by the definition of g , there is some $G'' \in A(G')$ such that $g(G'') = g(G)$. So player 2 can move to $G'' + \text{Nim}(g(G))$ which is a \mathcal{P} -position by induction.

Exercise (Moore). Find a Nim-like rule to determine winning and losing (\mathcal{P}, \mathcal{N}) positions in the following game: like Nim, there are several piles of counters, two players alternate turns and the last player to move wins; on your turn you can either remove any number of counters from one pile *and* (unlike Nim) you can, optionally, remove any number of counters from a second pile.

See “Exercises 5” for a few more exercises about \mathcal{N}/\mathcal{P} analysis and impartial combinatorial games.