Crash Course on Combinatorial Game Theory

for the Mathcamp 2019 Qualifying Quiz

The purpose of this brief primer on combinatorial game theory is to introduce you to the techniques needed to solve Problem #6 on the Mathcamp 2019 Qualifying Quiz. If you already know about Nim addition, the MEX rule, and the Sprague-Grundy Theorem, you can skip this crash course and proceed directly to the problem.

We created this primer to make your life easier, but you are also welcome to learn this material from any other sources: it is quite standard and can be found in many books and online references. Just make sure to cite all the sources that you use (other than this primer) in your solutions.

The suggested exercises in this primer are there to help you learn. Please do not submit solutions to these exercises as part of your solution to Problem #6: we will not read them.

You may ask other people for help with understanding the general concepts in this document, as long as you are not asking about Problem #6 itself. Conversely, **please do not ask clarification questions on this primer, or general questions on combinatorial game theory, through the Qualifying Quiz hotline (`quiz19@mathcamp.org`)**. That email address is reserved only for questions about the Quiz itself.

Good luck, and enjoy!

1 A Mathematical Introduction to Games

The games we are interested in

Definition 1.1. A **combinatorial game** is a game in which two players take turns making moves; both of them have complete information about what has happened in the game so far and what each player's options are from each position.

Common examples of combinatorial games include chess, go, and tic-tac-toe. Rock-Paper-Scissors is not a combinatorial game, since the two players move simultaneously. Most card games are also not combinatorial, since usually a player does not know what cards her opponent has or what cards she herself will draw on the next turn.

Combinatorial Game Theory is a broad field of mathematical research. In this crash course, we will only deal with a particularly simple type of combinatorial games, satisfying the following conditions:

- **Standard play**. This means that the winner and loser are determined based on who runs out of moves first. In standard play, *the first player who cannot make a valid move on their turn is the loser.* In particular, there are no ties.
- **Finiteness**. This means the game is guaranteed to end: no matter what the players do, at some point one of them will be unable to move and will lose the game.

- **Impartiality**. This means that the moves available to each player from each position are exactly the same. For example, chess is *not* an impartial game, since one player can only move the white pieces, while the other can only move the black pieces. On the other hand, Candy Split, as described in Problem #6 of the Mathcamp 2019 Qualifying Quiz, is impartial.

We will refer to combinatorial games that satisfy these three conditions as **FISP games** (FISP = finite, impartial, standard play).[1] So far, Candy Split is our only example of a FISP game; we will see more examples soon.

The game graph

Let us denote the position of Candy Split with piles of size X and Y (where $X \leq Y$) by CANDY(X, Y).

To any position of a FISP game, we can associate a *game graph* that shows all the ways the game can proceed starting from that position. The vertices of the game graph are positions of the game, and the edges are valid moves. For instance, Figure 1 shows the game graph for CANDY$(3, 4)$.

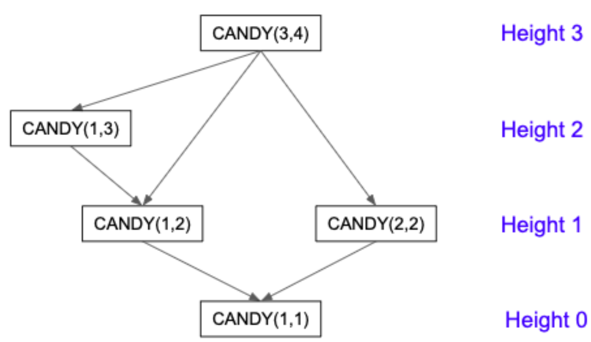

Figure 1: The game graph for CANDY$(3, 4)$

Since FISP games are finite, a game graph will always contain only a finite number of vertices and no loops.

Definition 1.2.

- A **terminal position** in a game is a position from which there are no valid moves. Every path through the game graph ends in a terminal position.
- The **height** of a position in a FISP game is the length of the longest path in its game graph. (Terminal positions have height 0.)

Every FISP game has a terminal position; some games have more than one. In Candy Split, the only terminal position is CANDY$(1, 1)$. From Figure 1, we see that the height of CANDY$(3, 4)$ is 3.

[1]This is not standard terminology. Most of the time when one talks about FISP games, finiteness and standard play are taken for granted, so the games are called simply "impartial".

N-positions and P-positions

The most important question you can ask about a position in a FISP game is: "Is it better to go first or second?" We'd like to be able to classify all positions into two categories:

Definition 1.3.

- A position is an **N-position** (or "has status N") if the player whose turn it is to move (the "***N***ext" player) has a winning strategy. In other words, assuming that the Next player plays correctly, she can win no matter what her opponent does.
- A position is a **P-position** (or "has status P") if the player whose turn just ended (the **P**revious player) has a winning strategy. (If it's the beginning of the game and there was no previous move, then the Previous player is the player who goes second.)

Intuitively, N-positions are ones from which you want to go first; P-positions are ones from which you want to go second. Here are some simple examples:

- CANDY$(1, 2)$ is an N-position: the Next player can immediately move to the terminal position CANDY$(1, 1)$ and win the game.
- CANDY$(3, 3)$ is a P-position: the only thing the Next player can do is go to CANDY$(1, 2)$, from which the Previous player (who will now be in the role of Next) can win.

Obviously, a position can't be both N and P, since it's not possible for *both* players to have a winning strategy. But how can we be sure that every position is either N or P? Perhaps there are positions from which neither player has a fail-safe strategy?

Lemma 1.1. *In any FISP game,*

(a) A position from which you can move to a P-position is an N-position. (There may also be moves to other positions, but they don't matter.)

(b) A position from which the only moves are to N-positions is a P-position.

Proof:

(a) When you move to a P position, it is now your opponent's turn, and you become the Previous player. By definition of P, you have a winning strategy from that P-position.

Now consider your original position: if you are the Next player and you can move to a P-position, then you have a winning strategy, so your current position is an N-position.

(b) Say you are the Next player. If you can only move to N-positions, this means that whatever you do, you opponent (who becomes Next after your turn) can win the game. This means your current position is a P-position.

Theorem 1.1. *Every position in a FISP game is either N or P.*

Proof: We use strong induction on height h.

- *Base Case: h=0.* A position of height 0 is a terminal position. The Next player has just lost, so it's a P-position.
- *Inductive step:* Suppose all positions of height less than h have been labeled N or P. Given a position of height h, all moves from it must be to positions of height less that h. We can therefore determine its N/P status according to the criteria in Lemma 1.1. QED

The proof of Theorem 1.1 suggests an algorithm for determining the N/P status of any position g in a FISP game:

- Draw the game graph for g and label all positions of height 0 as P.
- Once you've labeled all positions up to height h, label all positions of height h according to the rules of Lemma 1.1.
- Keep incrementing h until you get to your starting position g.

Suggested exercise 1.1. Use this algorithm and the game graph in Figure 1 to determine the N/P status of CANDY$(3,4)$. Then check your answer below.

Answer: To find the N/P status of CANDY$(3,4)$, we label the positions in its game graph in the following order (Figure 2).

- *Height 0:* CANDY$(1,1)$ is P, since it's terminal.
- *Height 1:* CANDY$(1,2)$ and CANDY$(2,2)$ are N, since you can move from them to CANDY$(1,1)$.
- *Height 2:* CANDY$(1,3)$ is P, since the only move from it is to CANDY$(1,2)$, which is N.
- *Height 3:* CANDY$(3,4)$ is N, since you can move from it to CANDY$(1,3)$, which is P. The fact that you can also move to two other positions, both of which are N, is irrelevant. (Those are bad moves that a player who knows how to play the game would never make.)

Solving a FISP game

Solving a game means figuring out how to play it optimally from any position. For a FISP game, knowing the N- and P-positions is all you need. The optimal strategy is:

(a) If you're in an N-position, there must be a P-position you can move to. Do that!

(b) If you're in a P-position, nothing you can do can help you win unless your opponent makes a mistake. So make whatever moves you want, but keep watching: if your opponent messes up and moves to an N-position, go to (a).[2]

[2] Actually, the advice in part (b) is outside the realm of mathematics: when you study combinatorial game theory, you always assume that both players play perfectly. But in practice, when you're playing with an actual human, watching for errors is essential!

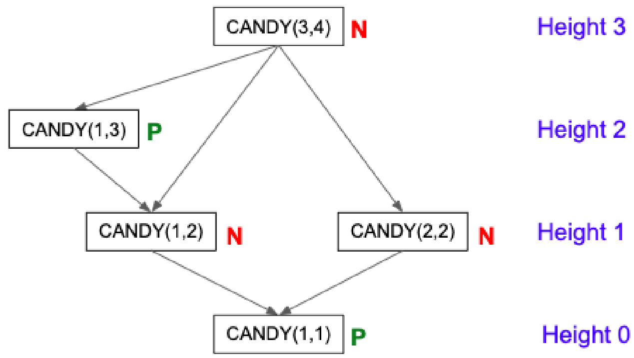

Figure 2: Figure 2: N and P positions for CANDY$(3,4)$

Here is one more example of a simple FISP game for you to practice on. In this game, you start with a pile of X stones. On her turn, a player must remove 2, 3, or 5 stones from the pile. Whoever can't move loses.

This game is part of a large class of games called "take-away games", so we will call it $\mathbf{T}$. We will denote the position of $\mathbf{T}$ with X stones by $\mathbf{T}_X$. This time, instead of drawing the full game graph for each position, let's just make a table of values of X and not put in the edges. It is easy to keep track of what the valid moves are from each position; putting in all the arrows would only make the picture messier.

X	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
status of $\mathbf{T}_X$	P	P	N	N	N	N	N	P	P	N	$\dots$									

The first few entries in the table above were computed as follows:

- The terminal positions are $\mathbf{T}_0$ and $\mathbf{T}_1$; both are P-positions.
- $\mathbf{T}_2$, $\mathbf{T}_3$, and $\mathbf{T}_5$ are N-positions, since you can move from them to the P-position $\mathbf{T}_0$.
- $\mathbf{T}_4$ and $\mathbf{T}_6$ are N-positions, since you can move from them to the P-position $\mathbf{T}_1$.
- From $\mathbf{T}_7$, you can only move to $\mathbf{T}_2, \mathbf{T}_4$, or $\mathbf{T}_5$. Since these are all N-positions, $\mathbf{T}_7$ must be a P-position.
- From $\mathbf{T}_8$, you can only move to $\mathbf{T}_3, \mathbf{T}_5$, or $\mathbf{T}_6$. Since these are all N positions, $\mathbf{T}_8$ must be a P-position.
- From $\mathbf{T}_9$ you can move to $\mathbf{T}_7$, which is a P-position, so $\mathbf{T}_9$ is an N-position.

Note that we did not have to fill in the table strictly in order of height. To find the N/P status of position g, you don't actually need to know the status of *all* positions of height less than g; you just need to know the status of all the positions you can move to from g. The concept of height is useful for proofs, but rarely referred to in actual calculations.

Suggested exercise 1.2. Fill in the missing values in the table before you read on.

The table looks periodic, and it's not hard to prove that the pattern does indeed continue. Thus we have found a strategy for this game: always move to a number that is 0 or 1 mod 7. If you can't, then you've lost.

Suggested exercise 1.3. Use the method developed so far to solve Candy Split, i.e. to figure out which positions are N and which are P. Instead of a one-dimensional table, you will need a two-dimensional table where the (X, Y) entry is the N/P status of CANDY(X, Y). As you fill out the table, look for patterns; see if you can spot a pattern that makes the strategy for the game easy to summarize in a few sentences. Then see if you can prove it.

Addition of games

The tools that we currently have are enough to solve Candy Split. But Problem #6 asks you to play a much more complicated game: three games of Candy Split at the same time! This too is a FISP game, so in theory you could approach it as above. However, it will get very ugly very fast; even just recording the N/P status of each position of this game would require a six-dimensional table. Fortunately, there is a better way.

Definition 1.4. Given two FISP games, G_1 and G_2, we define their sum, $G_1 + G_2$, to be a FISP game with the following rules:

- A position of $G_1 + G_2$ consists of a position g_1 of the game G_1 and a position g_2 of the game G_2. We refer to this position as $g_1 + g_2$.
- A valid move from $g_1 + g_2$ is *either* a valid move in G_1 from g_1 to some other position g_1' or a valid move in G_2 from g_2 to some other position g_2'. The resulting position of $G_1 + G_2$ is then either $g_1' + g_2$ or $g_1 + g_2'$.
- A player loses the game $G_1 + G_2$ if it is their turn and they are unable to move. In other words, they have no valid moves in either G_1 or G_2.

For example, suppose G_1 is Candy Split and G_2 is the **T** game. One possible position of $G_1 + G_2$ is CANDY$(2, 3) + T_8$. Possible moves from this position include CANDY$(1, 2) + T_8$ and CANDY$(2, 3) + T_6$. However, since you can only move in one of the games, you cannot move from CANDY$(2, 3) + T_8$ to CANDY$(1, 2) + T_6$.

The order of addition does not matter: there is no difference between $G_1 + G_2$ and $G_2 + G_1$. Similarly, $(G_1 + G_2) + G_3$ clearly has the same set of positions and moves as $G_1 + (G_2 + G_3)$. In other words, addition of games is commutative and associative. These properties will come in very handy later on.[3]

How does addition of games interact with N/P status? Here are two basic and important theorems:

Theorem 1.2. *For any position g in any FISP game, $g + g$ is always a P-position.*

[3] Technically, we haven't defined what it means for two games to be "the same". For example, if you and I play Candy Split with chocolates and our friends play it with marshmallows, are we playing the same game? If we wanted to be really precise, we would define isomorphism of games via isomophism of their game graphs – but that would take us too far afield. Hopefully the fact that $G_1 + G_2$ and $G_2 + G_1$ are the same game for all intents and purposes is clear enough that you can take it on faith.

Theorem 1.3. *If g is a P-position, then for any other position h in a FISP game H, $h+g$ has the same N/P status as h.*

Suggested exercise 1.4. Try to prove these two theorems yourself before reading on.

Proof of Theorem 1.2:

We need to show that the second ("Previous") player has a winning strategy from $g+g$. Indeed, she can always win by imitating her opponent's moves. For instance, if her opponent plays in the first copy of g and moves to $g'+g$, she can make the same move in the second copy of g and move to $g'+g'$. With this strategy, whatever moves her opponent makes, the second player is guaranteed to be able to make a valid move in response, by copying. Thus it will have to be her opponent who eventually runs out of moves and loses the game. QED

Proof of Theorem 1.3: First, suppose h is a P-position. Then the second (Previous) player has a winning strategy from each of h and g individually. She can combine these into a strategy for $h+g$: if her opponent moves in h, she responds according to her strategy for h; if her opponent moves in g, she responds according to her strategy for g. Thus, for any move her opponent makes, the second player is guaranteed to have a response. Her opponent will run out of moves first and lose. Thus $h+g$ is a P-position.

Now suppose h is an N-position. Then, by Lemma 1.1, there is a move from h to some P-position h'. Since h' and g are both P-positions, so is $h'+g$ (as we just proved above). Since you can move from $h+g$ to the P-position $h'+g$, Lemma 1.1 implies that $h+g$ is an N-position. QED

One way to summarize Theorem 1.3 is:

$$N+P=N, \qquad P+P=P.$$

But what about $N+N$? The answer is: it could be either! Here are two examples of sums of N-positions, with different results:

- $\text{CANDY}(2,2)+\text{CANDY}(1,2)$: This is a P-position, since there is only one valid move from each of the summands. The first player will move in one of the games, the second player will move in the other, and that's it. Second player wins.
- $\text{CANDY}(3,4)+\text{CANDY}(1,2)$: This is an N-position: you can move from it to $\text{CANDY}(2,2)+\text{CANDY}(1,2)$, which we just showed is P.

In other words, P and N are good enough if you are playing just one game. But if we want to understand how games behave under addition, we need a finer classification.

2 Nim

The key to understanding addition of FISP games turns out to reside in one particular game – the most famous FISP game of all.

The rules of Nim

The most common version of Nim is 3-pile Nim. Two players start with three piles of stones. (These piles are often called *nimheaps.*) On her turn, a player must remove a nonzero number of stones from any one pile. (In particular, she is allowed to remove all the stones from a pile.) As usual, whoever can't move loses, i.e. whoever takes the last stone wins.

Suggested exercise 2.1. Satisfy yourself that Nim is indeed a FISP game. You may also want to play a few games (with small nimheaps) to get the hang of it.

Notice a key difference between Nim and Candy Split: in Candy Split, a single move affects both piles of candy. In Nim, a move affects only one pile at a time. The rule that a player must choose one pile and move in that pile should remind you of something. In fact, 3-pile Nim is precisely the *sum* of three separate games of 1-pile Nim!

Of course, 1-pile Nim on its own is an extremely boring game. A single nimheap of size X, which we denote by X^*, is very easy to analyze:

- X^* is an N-position if $X > 0$. (The Next player can win in one move by removing all the stones in the heap.)
- X^* is a P-position if $X = 0$.

And yet, when you add three games of 1-pile Nim together, you get something quite interesting and complicated!

A few words on notation:

- We will denote a Nim position with heaps of size X, Y, and Z by $X^* + Y^* + Z^*$. As usual with addition of games, order does not matter.
- If, on her turn, a player removes an entire heap (say Z^*), we will skip the resulting 0^* (the empty heap) and record what is left as just $X^* + Y^*$. Obviously, leaving out the empty heap makes no difference to the game.
- We will only use 0^* when there are no other piles. This is the terminal position of every game of Nim, regardless of how many piles you start with.

Suggested exercise 2.2. Figure 3 shows the full game graph of the position $1^* + 2^* + 3^*$. Label all the positions in the graph with N and P to deduce that $1^* + 2^* + 3^*$ is a P-position.

The Nim table

As you can see in Figure 3, the full game graph for 3-pile Nim can get very complicated, even when all the heaps are small. Fortunately, the task of figuring out the P- and N-positions of 3-pile Nim is greatly simplified by the following theorem:

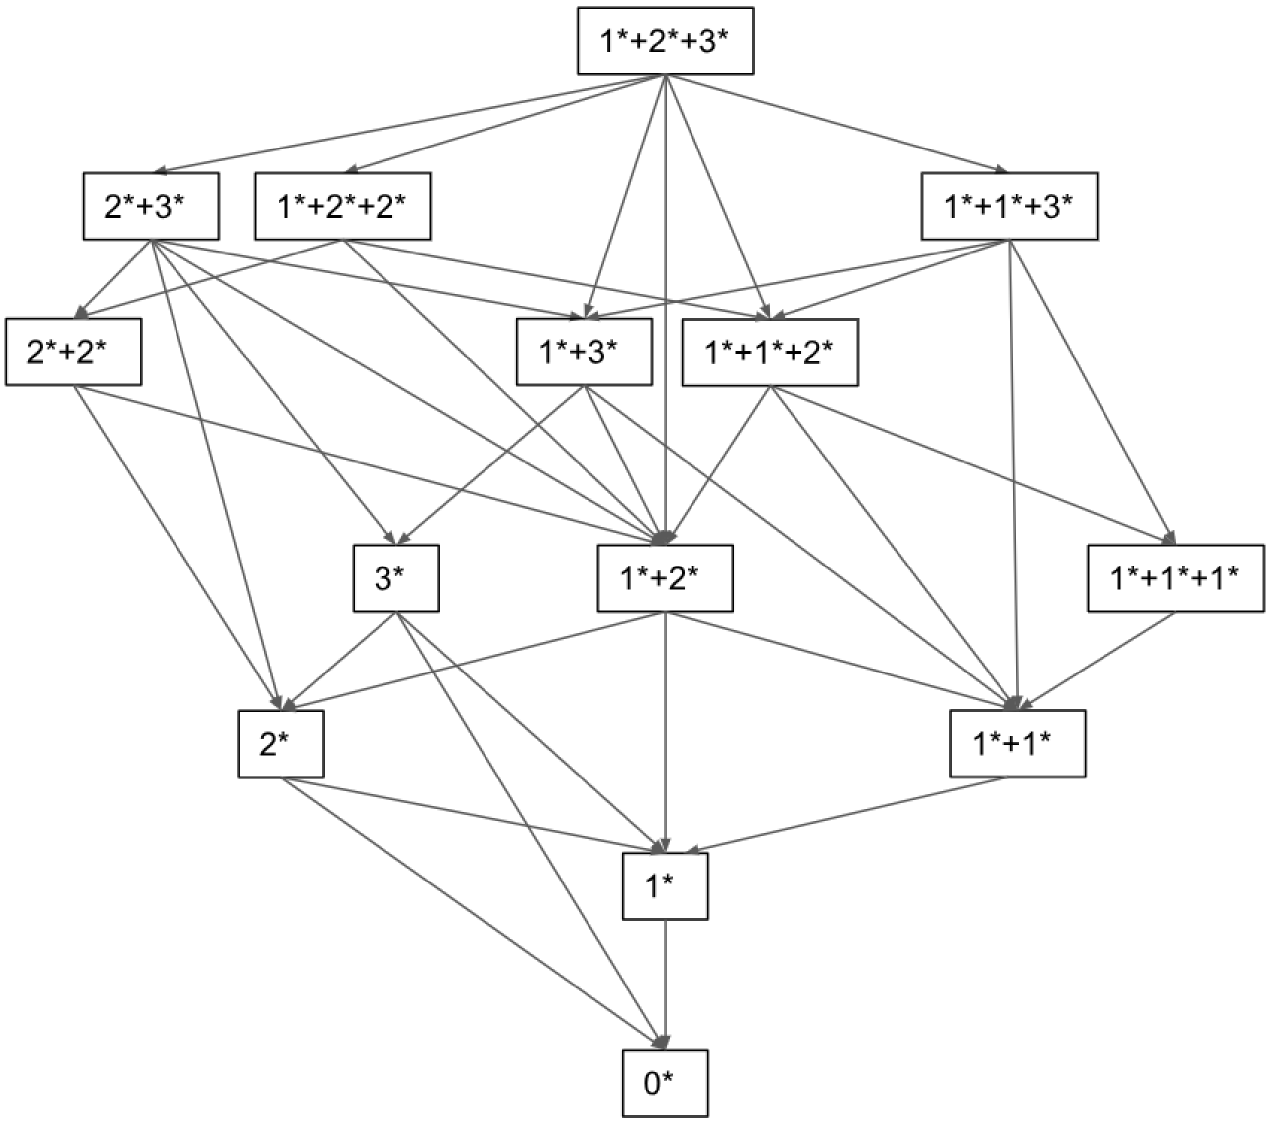

Figure 3: The game graph of $1^* + 2^* + 3^*$

Theorem 2.1. *For any two numbers X and Y, there exists at most one number Z such that $X^*+Y^*+Z^*$ is a P-position.*

Proof: Suppose $X^* + Y^* + Z_1^*$ and $X^* + Y^* + Z_2^*$ are both P-positions with $Z_1 > Z_2$. Note that you can move from the former to the latter by taking $Z_1 - Z_2$ stones from Z_1^*. But there cannot be a move from a P-position to a P-position, by Lemma 1.1. Thus we obtain a contradiction. QED.

Note that we have not shown that a number Z satisfying the condition of the theorem actually exists; we only know that, if it exists, it must be unique. Still, we can try constructing a table whose entry in row X and column Y is the unique value of Z (if one exists) such that $X^* + Y^* + Z^*$ is a P-position. If we can construct such a table, we will know all the P-positions in Nim, which is all we need for a winning strategy.

We denote the entry in row X and column Y by $[X, Y]$. Here are some entries that we can fill in right away:

- By Theorem 1.2, $X^* + X^*$ is a P-position for all X. We can think of this as $X^* + X^* + 0^*$. Thus, for all X, we have
$$[X, X] = 0, \quad [X, 0] = [0, X] = X$$
- By analyzing Figure 3, we can show directly that $1^* + 2^* + 3^*$ is a P-position. Thus:
$$[1, 2] = [2, 1] = 3, \quad [2, 3] = [3, 2] = 1, \quad [1, 3] = [3, 1] = 2.$$

Here is the table with what we know so far:

	0	1	2	3	4	5	6	7	8	...
0	0	1	2	3	4	5	6	7	8	...
1	1	0	3	2						...
2	2	3	0	1						...
3	3	2	1	0						...
4	4				0					...
5	5					0				...
6	6						0			...
7	7							0		...
8	8								0	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Let's try to fill in some more values. For instance, what can we say about $[1, 4]$?

- $[1, 4]$ cannot be any of the values that already appear in Row 1. For example, since $[1, 2] = 3$, we can't have $[1, 4] = 3$, since this would mean that both $1^* + 2^* + 3^*$ and $1^* + 4^* + 3^*$ are P-positions, contradicting Theorem 2.1.
- For the same reason, $[1, 4]$ cannot be any of the values that already appear in Column 4.
- We conclude that $[1, 4]$ cannot be any of the numbers 0, 1, 2, 3, or 4. So maybe it's 5?

Suggested exercise 2.3. Convince yourself that $[1, 4]$ is indeed 5, by checking that all possible moves from $1^* + 4^* + 5^*$ lead to N-positions. Don't draw the game graph: everything you need should be deducible from what's already in the table!

If you don't understand how to do this exercise, don't worry: it just a special case of the proof of Theorem 2.2, which is explained below.

The MEX rule

From the previous example, you might be tempted to guess that $[X, Y]$ has to be greater than all the numbers that previously appeared in row X or column Y. Perhaps it is the first number that is greater than all of them?

But then you notice that $[1, 5] = 4$, even though 5 has already appeared in Row 1. So it's not quite that simple ... but close.

Definition 2.1. Let S be a set of nonnegative integers. MEX(S) is defined to be the smallest nonnegative integer that does not appear in S. (**MEX** stands for "minimal excluded".)

For instance:

$$\begin{aligned}\mathrm{MEX}(0,1,2,3,4) &= 5\\ \mathrm{MEX}(0,1,2,3,5) &= 4\\ \mathrm{MEX}(1,2,3,4,5) &= 0\end{aligned}$$

Theorem 2.2 ("the MEX rule"). *For all integers $X, Y \geq 0$,*

$$[X,Y] \,=\, \mathrm{MEX}\,(\,\{[X',Y] : X' < X\} \,\cup\, \{[X,Y'] : Y' < Y\}\,)$$

In other words, $[X,Y]$ is the smallest integer that does not appear directly above or to the left of its position in the table.

Suggested exercise 2.4. To make sure you understand the statement of the theorem, go back and finish filling out the table on the previous page. Then compare your answer to the full table as shown below. You can do this either before or after you read the proof.

Proof of Theorem 2.2: Given $X, Y \geq 0$, let

$$M = \mathrm{MEX}\,(\,\{[X',Y] : X' < X\} \,\cup\, \{[X,Y'] : Y' < Y\}\,)\,,$$

as in the statement of the theorem. We want to show that $X^* + Y^* + M^*$ is a P-position.

Consider the three ways that the Next player can move from $X^* + Y^* + M^*$:

- Suppose she moves to $X_1^* + Y^* + M^*$ with $X_1 < X$. By definition of MEX, $M \neq [X_1, Y]$. Thus $X_1^* + Y^* + M^*$ is not a P-position, so it is an N- position.
- Suppose she moves to $X^* + Y_1^* + M^*$ with $Y_1 < Y$. The same argument as above shows that this is an N-position.
- Finally, suppose she moves to $X^* + Y^* + M_1^*$ with $M_1 < M$. Then, by definition of MEX, M_1 must be an element of the set

  $$\{[X',Y] : X' < X\} \,\cup\, \{[X,Y'] : Y' < Y\}.$$

  Without loss of generality, suppose $M_1 = [X_1, Y]$ with $X_1 < X$. Then $X_1^* + Y^* + M_1^*$ is a P-position, which means that $X_1^* + Y^* + M^*$ must be an N-position by Theorem 2.1.

We have shown that all possible moves from $X^* + Y^* + M^*$ lead to N-positions. Thus $X^* + Y^* + M^*$ itself must be a P-position, so $[X,Y] = M$. QED

Here is the complete Nim strategy table for $X, Y \leq 8$, constructed recursively using Theorem 2.2:

	0	1	2	3	4	5	6	7	8
0	0	1	2	3	4	5	6	7	8
1	1	0	3	2	5	4	7	6	9
2	2	3	0	1	6	7	4	5	10
3	3	2	1	0	7	6	5	4	11
4	4	5	6	7	0	1	2	3	12
5	5	4	7	6	1	0	3	2	13
6	6	7	4	5	2	3	0	1	14
7	7	6	5	4	3	2	1	0	15
8	8	9	10	11	12	13	14	15	0

Using the Nim table

The Nim table tells us everything we need to know about how to play 3-pile Nim. For instance, suppose you and I start playing from the position $3^* + 4^* + 6^*$. I generously let you choose whether to go first or second.

- *Should you go first or second?* From the table, you see that $[3, 4] \neq 6$, so $3^* + 4^* + 6^*$ is an N-position. You should go first.
- *How should you play?* Recall that the winning strategy is always to move to a P-position. Since you're at an N-position, there must be at least one such move available, but it might take you a few tries to find it.

  For instance, since $[3, 4] = 7$, you might try to move to $3^* + 4^* + 7^*$. But you can't go from 6^* to 7^*: you can only remove stones from piles, not add them. Here the Nim table tells you that removing *any* number of stones from the 6^* heap is a bad move.

  Similarly, since $[3, 6] = 5$, you have no good moves from the 4^* heap (since $5 > 4$). Fortunately, $[4, 7] = 2$, and 2 is less than 3. Thus the *only* good move from this position is to remove one stone from the 3^* heap, moving from $3^* + 4^* + 6^*$ to $2^* + 4^* + 6^*$.

Suggested exercise 2.5. Teach a friend to play Nim without teaching them the strategy. Play a bunch of games with up to 8 stones in each pile, and impress them by beating them every time! (You can even let them choose whether to go first or second. They might get lucky and make the right choice initially, but if they don't know how to play, they'll almost certainly make a mistake at some point by moving to an N-position. From there, you can win.)

Nimsum

Suppose we want to compute $[2019, 2020]$. Right now, the only way we know how to do this is by constructing the whole 2019×2020 Nim table using the MEX rule. This seems really wasteful! Wouldn't it be nice to have a closed formula for $[X, Y]$ that did not require knowing the entire table up to that point?

You have probably already noticed some interesting patterns in the Nim table, involving 2×2 squares, 4×4 squares, 8×8 squares, etc. The challenge is to describe this pattern precisely, so that we can derive a formula for $[X, Y]$ and then try to prove it.

Here is an extremely useful problem-solving strategy: when you see a pattern involving powers of 2, rewrite all your numbers in binary (base 2). This usually makes the pattern much easier to describe and to understand.

In binary, our Nim table looks like this:

	0	1	10	11	100	101	110	111	1000
0	0	1	10	11	100	101	110	111	1000
1	1	0	11	10	101	100	111	110	1001
10	10	11	0	1	110	111	100	101	1010
11	11	10	1	0	111	110	101	100	1011
100	100	101	110	111	0	1	10	11	1100
101	101	100	111	110	1	0	11	10	1101
110	110	111	100	101	10	11	0	1	1110
111	111	110	101	100	11	10	1	0	1111
1000	1000	1001	1010	1011	1100	1101	1110	1111	0

Suggested exercise 2.6. If you want to figure out the formula for yourself, this is your chance! Stare at the table until you see the pattern. Or you can just read on.

The pattern that emerges from the table is as follows: to compute $[X, Y]$, first write X and Y in binary. If one of them has fewer digits, pad it with 0's until the two numbers are the same length. Now simply add the individual digits of X and Y mod 2. The difference between this and normal addition is that we are adding *without carry*. The resulting number is $[X, Y]$.

For instance, we can see in the table that

$$\begin{array}{ccccc} 5 & = & 101_2 & & \\ 3 & = & 011_2 & & \\ \hline [5,3] & = & 110_2 & = & 6 \end{array}$$

Assuming this pattern holds, we can now easily compute:

$$\begin{array}{ccrcc} 2019 & = & 11111100011_2 & & \\ 2020 & = & 11111100100_2 & & \\ \hline [2019, 2020] & = & 111_2 & = & 7 \end{array}$$

This operation of adding binary digits mod 2 without carry is called *nimsum* (although it also occurs in many mathematical contexts other than Nim). To distinguish it from regular addition, we denote it by $\oplus$. Note that $X \oplus X = 0$ for all X. It is also easy to check that $\oplus$ is commutative and associative.

Suggested exercise 2.7. Convince yourself that $X \oplus Y = Z$ if and only if $X \oplus Y \oplus Z = 0$.

We are now ready to state and prove our formula for $[X, Y]$:

Theorem 2.3. *$X^*+Y^*+Z^*$ is a P-position if and only if $X \oplus Y \oplus Z = 0$. In other words, $[X, Y] = X \oplus Y$.*

Proof: The proof is by strong induction on $X + Y + Z$. (Note that this is the usual sum of integers, not nimsum.)

The theorem is clearly true when $X + Y + Z = 0$. Now suppose the theorem holds for all triples such that $X + Y + Z < n$, and consider a position $X^* + Y^* + Z^*$ where $X + Y + Z = n$.

Suppose first that $X \oplus Y \oplus Z = 0$. We need to show that in this case, $X^* + Y^* + Z^*$ is a P-position. A move from this position involves reducing the value of one of the numbers – say going from X to $\tilde{X} < X$. But changing even one digit of X will result in the nimsum no longer being 0 at that digit, so $\tilde{X} \oplus Y \oplus Z \neq 0$. Since $\tilde{X} + Y + Z < n$, we can apply the inductive hypothesis to conclude that $\tilde{X}^* + Y^* + Z^*$ is an N-position. This shows that all moves from $X^* + Y^* + Z^*$ are to N-positions, so $X^* + Y^* + Z^*$ itself is a P-position.

Now suppose that $X \oplus Y \oplus Z \neq 0$. We will show that this is an N-position by finding a move to a P-position.

Denote the binary digits of X (after padding with added 0's if necessary) by $X_1, X_2, \ldots X_n$, and similarly for Y and Z. (Thus X_n is the 1's digit, X_{n-1} is the 2's digit, etc.) Find the smallest m such that $X_m + Y_m + Z_m \neq 0 \mod 2$. At least one of X_m, Y_m, and Z_m is not zero; WLOG, say $X_m = 1$.

Let $\tilde{X}$ be the integer with binary expansion

$$\tilde{X}_i = \begin{cases} X_i & \text{if } i < m \\ 0 & \text{if } i = m \\ Y_i + Z_i \mod 2 & \text{if } i > m \end{cases}$$

In other words, we are modifying all the digits $X_m, X_{m+1}, \ldots, X_n$ to make the nimsum work out correctly. By construction, $\tilde{X} \oplus Y \oplus Z = 0$ and $\tilde{X} < X$. By the inductive hypothesis, $\tilde{X}^* + Y^* + Z^*$ is a P-position, so $X^* + Y^* + Z^*$ is an N-position, as required. QED

Suggested exercise 2.8. Although we were just looking for a solution for 3-pile Nim, we get the solution for k-pile Nim for free! Generalize Theorem 2.3 to any number of nimheaps and prove it. (It is almost word for word the same proof as above.)

3 The Sprague-Grundy Theorem

At this point you may be wondering: why are we talking so much about Nim? I am trying to solve a problem about Candy Split. How is Nim going to help me?

Good question. Keep reading.

Equivalence of games

As we have seen, two game positions can both be N and yet behave quite differently in combination with other games. We need a more refined classification.

Definition 3.1. Let G_1 and G_2 be two (possibly different) games, and let g_1 and g_2 be positions of G_1 and G_2 respectively. We say that g_1 is **equivalent** to g_2 ($g_1 \simeq g_2$) if, for any position h of any game H, $g_1 + h$ has the same N/P status as $g_2 + h$.

In other words, in the context of game addition, equivalent positions will produce the same results.

Checking whether two positions are equivalent turns out to be surprisingly easy: just add them up!

Theorem 3.1. *$g_1 \simeq g_2$ if and only if $g_1 + g_2$ is a P position.*

Proof: First, suppose $g_1 \simeq g_2$. Substituting g_2 for h in the definition of equivalence, we see that $g_1 + g_2$ must have the same N/P status as $g_2 + g_2$. Since the latter is a P-position (Theorem 1.2), so is the former.

Conversely, suppose $g_1 + g_2$ is a P-position and let h be any position in any game. We know that adding a P-position has no effect on N/P status (Theorem 1.3), so $g_1 + h$ has the same status as $(g_1 + h) + (g_1 + g_2)$. Since addition of games is commutative and associative, we can rewrite this as $(g_2 + h) + (g_1 + g_1)$. Since $g_1 + g_1$ is a P-position (Theorem 1.2), adding it does not affect N/P status. Thus we have show that the following positions all have the same N/P status:

$$g_1 + h \quad \longleftrightarrow \quad (g_1 + h) + (g_1 + g_2) \quad \longleftrightarrow \quad (g_2 + h) + (g_1 + g_1) \quad \longleftrightarrow \quad (g_2 + h).$$

Thus $g_1 \simeq g_2$. QED.

Corollary 3.1. *If $g_1 \simeq g_2$ and h is any position, then $g_1 + h \simeq g_2 + h$.*

Note that Corollary 3.1 is not just repeating the definition of equivalence: it says that not only do $g_1 + h$ and $g_2 + h$ have the same N/P status, but they are actually equivalent themselves.

Proof: $(g_1 + h) + (g_2 + h) = (g_1 + g_2) + (h + h)$. This is a sum of two P-positions, so is itself a P position. By Theorem 3.1, this means that $g_1 + h \simeq g_2 + h$. QED.

Corollary 3.2. *g is a P-position if and only if $g \simeq 0^*$.*

Proof: This is almost tautological. $g + 0^*$ is just g itself, so g is a P-position iff $g + 0^*$ is a P-position.

Corollary 3.3. *Two nimheaps are equivalent if and only if they are the same size. In other words, if $X \neq Y$ then $X^* \not\simeq Y^*$.*

Proof: If $X \neq Y$, then $X^* + Y^*$ is an N-position, since you can move from it either to $X^* + X^*$ or to $Y^* + Y^*$ (depending on which of X and Y is larger). Thus $X^* \not\simeq Y^*$.

The Nim table, revisited

Theorem 3.1 sheds a new light on our Nim table. Remember that we defined $[X, Y]$ to be the unique number Z such that $X^* + Y^* + Z^*$ is a P-position. But by Theorem 3.1, this is the same as saying that $X^* + Y^* \simeq Z^*$. Thus our table turns out to be the *Nim addition table*: it tells us how to turn the sum of two nimheaps into a single nimheap, at least up to equivalence. No wonder the formula for $[X, Y]$ was a kind of sum, $X \oplus Y$.

Originally we constructed the Nim table just to play 3-pile Nim. But now that we know that it is an addition table, we can use it to play Nim with any number of piles. For example, let's analyze the position $5^* + 6^* + 7^* + 8^*$. We have:

$$\begin{array}{rcll} 5^* + 6^* + 7^* + 8^* & \simeq & 3^* + 7^* + 8^* & (\text{since } 5 \oplus 6 = 3) \\ & \simeq & 4^* + 8^* & (\text{since } 3 \oplus 7 = 4) \\ & \simeq & 12^* & (\text{since } 4 \oplus 8 = 12) \end{array}$$

Thus $5^* + 6^* + 7^* + 8^*$ is an N-position. If you add it to any additional nimheaps, or to any other position of any other game, it will behave just like a single nimheap of size 12.

To find a good move from $5^* + 6^* + 7^* + 8^*$, we use the second-to-last equivalence in our derivation above:

$$5^* + 6^* + 7^* + 8^* \simeq 4^* + 8^*.$$

Thus we can remove 4 stones from the 8^* heap to get a position equivalent to $4^* + 4^* \simeq 0^*$, i.e. a P-position. There are other good moves too, but all we need is one.

(You can also derive all of these results using the generalization of Theorem 2.3: $X_1^* + \cdots + X_k^*$ is a P-position if and only if $X_1 \oplus \cdots \oplus X_k = 0$.)

One game to rule them all

We are now ready to state and prove the main theorem of this crash course, which will finally explain why we care so much about Nim.

Theorem 3.2 (Sprague-Grundy). *Any position g in a FISP game is equivalent to a nimheap. In other words, there is some nonnegative integer X such that $g \simeq X^*$.*

By Corollary 3.3, X is unique. We call X the **nimvalue** of g and write $|g| = X$.

Proof: We use strong induction on the height k of g:

Base case: k=0. All terminal positions have nimvalue 0. (If you think about what this means, it's just saying that if you add a terminal position to the empty nimheap, you get a P-position. Well, you definitely can't make any moves from this position! So yes, it's P.)

Inductive step: Suppose the theorem holds for all positions of height less than k, and suppose g is of height k. Let $S = \{h_1, h_2, \dots h_k\}$ be the set of positions to which you can move from g. Since all the elements of S have height less than k, each $h_i \in S$ is equivalent to some nimheap X_i^*. Then we claim that $g \simeq M^*$, where

$$M = \text{MEX}(X_1, \dots, X_k).$$

We need to show that $g + M^*$ is a P-position, i.e. that all moves from $g + M^*$ are to N-positions. There are two ways to move from $g + M^*$:

- You could make a move in g to some $h_i \in S$. The resulting position is $h_i + M^*$, which is equivalent to $X_i^* + M^*$. This is an N-position unless $M = X_i$, which is impossible by the definition of MEX.

- You could remove some stones from M^*. The resulting position is $g + M_1^*$, for some $M_1 < M$. By definition of MEX, $M_1 = X_i$ for some i. Thus it is possible to move from $g + M_1^*$ to

$$h_i + M_1^* \simeq X_i^* + M_1^* = M_1^* + M_1^* \simeq 0^*.$$

This means $g + M_1^*$ is an N-position

We have shown that all moves from $g+M^*$ are to N-positions. This implies that $g+M^*$ is a P-position, so $g \simeq M^*$. QED

The Sprague-Grundy Theorem is incredibly powerful. It says that, to fully analyze a FISP game, you just need to compute the nimvalue of every position, which you do by the MEX rule. Then, when you find yourself playing a sum of games that you've already analyzed, and you're at position $g = g_1 + g_2 + \cdots + g_k$, you can just pretend that the individual g_i's are nimheaps and that g is a game of k-pile Nim. This allows you to figure out the right move from g using Nim addition, exactly as we did with $5^* + 6^* + 7^* + 8^*$ in the last section.

Suggested exercise 3.1. You might have noticed that the proof of the Sprague-Grundy Theorem is extremely similar to the proof of Theorem 2.2. This is because Theorem 2.2 is really just a special case of Sprague-Grundy. Can you see in what sense this is true? What is the game g in Theorem 2.2?

Example: Consider once again the game $\mathbf{T}$: a single pile of stones from which a player on her turn must remove 2, 3, or 5 stones. Previously, we computed only the N/P status of each position $\mathbf{T}_X$; now let us find the nimvalues.

X	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
$\lvert\mathbf{T}_X\rvert$	0	0	1	1	2	2	3	0	0	1	1	2	2	3	0	0	1	1	2	2

The first few entries in the table above were computed as follows:

- $\mathbf{T}_0$ and $\mathbf{T}_1$ are terminal positions, so $|\mathbf{T}_0| = |\mathbf{T}_1| = 0$.
- From $\mathbf{T}_2$ you can move only to $\mathbf{T}_0$. Thus

$$|\mathbf{T}_2| = \mathrm{MEX}(|\mathbf{T}_0|) = \mathrm{MEX}(0) = 1.$$

- From $\mathbf{T}_3$, you can move to $\mathbf{T}_0$ or $\mathbf{T}_1$, so $|\mathbf{T}_3| = \mathrm{MEX}(0) = 1$.
- From $\mathbf{T}_4$, you can move to $\mathbf{T}_1$ or $\mathbf{T}_2$, so

$$|\mathbf{T}_4| = \mathrm{MEX}(|\mathbf{T}_1|, |\mathbf{T}_2|) = \mathrm{MEX}(0, 1) = 2.$$

- From $\mathbf{T}_5$, you can move to $\mathbf{T}_0, \mathbf{T}_2$, or $\mathbf{T}_3$, so

$$|\mathbf{T}_5| = \mathrm{MEX}(|\mathbf{T}_0|, |\mathbf{T}_2|, |\mathbf{T}_3|) = \mathrm{MEX}(0, 1) = 2.$$

- From $\mathbf{T}_6$, you can move to $\mathbf{T}_1, \mathbf{T}_3$, and $\mathbf{T}_4$, so $|\mathbf{T}_6| = \mathrm{MEX}(0, 1, 2) = 3$.

- From $\mathbf{T}_7$, you can move to $\mathbf{T}_2, \mathbf{T}_4$, and $\mathbf{T}_5$, so $|\mathbf{T}_7| = \text{MEX}(1, 2, 3) = 0$. (This is consistent with our previous finding that $\mathbf{T}_7$ is a P-position.)

From here, as before, you can check that the pattern of nimvalues is periodic with period 7.

Now suppose we play a game with 3 piles of stones, in which a player, on her turn, must remove 2,3,or 5 stones from any one pile. This game is just the sum of 3 copies of $\mathbf{T}$, so we can use the Sprague-Grundy Theorem. For example, how should you play from position $g = \mathbf{T}_8 + \mathbf{T}_{12} + \mathbf{T}_{13}$?

Consulting our table of nimvalues of $\mathbf{T}$, we see that $g \simeq 2^* + 3^*$, so we can win by moving from the $\mathbf{T}$-heap equivalent to 3^* (i.e. $\mathbf{T}_{13}$) to some $\mathbf{T}$-heap equivalent to 2^*. We can achieve this by removing two stones from the $\mathbf{T}_{13}$ heap, resulting in the position $\mathbf{T}_8 + \mathbf{T}_{11} + \mathbf{T}_{12} \simeq 2^* + 2^* \simeq 0^*$. Continuing in this way, we will be able to win the game.

4 Summary

What we have learned: how to analyze a FISP game

- The first stage in analyzing a FISP game is to try to decompose it into a sum of simpler games. (This is not always possible: for example, Candy Split cannot be decomposed in this way. But if it's possible, it's very helpful.)
- Next, find the nimvalue of each position in the component games. This is done by starting at height 0 and proceeding by the MEX rule, as in the proof of the Sprague-Grundy Theorem. The hope is that once you do this for enough small positions, you will notice a pattern and prove that it holds in general.
- Finally, if you are playing a sum of games whose nimvalues you know, you can just pretend that you are playing Nim and use the Nim addition table to figure out where to move.

What we have not learned

Here are some of the many things not covered in this crash course:

- *How to solve Problem #6:* that's your job!
- *Other examples of FISP games.* There are many, including some that are actually played as games by people who are not mathematicians!
- *Non-FISP combinatorial games.* FISP games are the tip of the iceberg: most of the research in combinatorial game theory is on *partizan* games, where players may have different moves available to them. You can also change from standard play to other winning conditions, or see what happens if you allow infinity into your games.
- *Non-combinatorial games.* The mathematical field of *game theory* studies all the messy kinds of games that are banned from combinatorial game theory: games where players don't have complete

information (including games involving randomness), games with more than two players, games where players don't alternate moves, games with ties, cooperative games, etc. You might think that combinatorial game theory should be a subfield of game theory, but the techniques used to study combinatorial and non-combinatorial games are so different that they are essentially two unrelated fields. The non-combinatorial kind of game theory is sometimes called *matrix game theory*. It has a lot of application in economics, psychology, political science, etc.

Where you can learn more

You can find lots of tidbits about combinatorial game theory online. In particular, if you found this crash course unclear or confusing, you can try a different exposition. Two possible places to start are:

- the Wikipedia page for the Sprague-Grundy Theorem;
- the first 15 or so pages of the handout at `http://web.mit.edu/sp.268/www/nim.pdf`.

Don't forget to include a reference in your solutions to any outside sources that you used.

But if you really want to learn combinatorial game theory – not for the Qualifying Quiz but for the fun of it – the best place to look is the classic book *Winning Ways for Your Mathematical Plays*, by Elwyn Berlekamp, John Conway, and Richard Guy. In addition to being an excellent and thorough introduction to combinatorial game theory, it is one of the quirkiest, funniest, and best-written math books you'll probably ever read.

Good luck with the Qualifying Quiz!

The Mathcamp team