Game Simulation and Analysis

Sarah Eichhorn, Jason Wilkinson, and Jacquelyn Rische
Department of Mathematics
University of California, Irvine

June 11, 2013

Abstract

In the following notes, we present an introduction to game theory with adaptive learning computational analysis. Only a minimal mathematical and computer programming background is assumed. These notes are meant to accompany the UCI Math/ICS 77D course in "Game Simulation and Analysis". The goal of this course is to provide students an introduction to the fields of game theory and computer adaptive learning and prepare students for research in these fields.

1 Introduction to Game Theory

To get a overview of what the study of game theory entails, let us begin with an example of game.

1.1 Poison Game

Poison is a two player game in which there are 10 objects on the table. Players alternate turns and on each turn a player is required to take 1 or 2 of the objects from the table. The person to take the last object loses the game.

After playing this game several times, you may come up with questions about this game. For example you might ask: Is this a fair game? Does one player have an advantage? What is the best strategy to play by? How would this game change if there were more or less objects on the table initially?

It turns out that with just a little bit of game analysis, most of these questions can be answered relatively easily. Let us start by reducing this game to a simpler version. Begin by assuming instead of having 10 objects at the beginning, there is only 1 object. Who loses? The answer is obvious, the player who goes first will lose since they have to take the one and only object.

Now assume there are 2 objects and determine who should win. If the first player takes 2, they would lose since they would have taken the last object. But that is a silly decision, since they have the choice of only taking 1 object and would not lose if they made that move. If Player 1 did only take 1 object, they would leave Player 2 with only 1 object, and we are the situation we just analyzed. Player 2 would have to lose. Player 1 has a choice of two different **strategies**, or choices of moves, that he/she could employ. The move to take 1 object at the beginning is called a **winning move or strategy** because if Player 1 makes this move they are guaranteed to win regardless of what the other player does. To recap, with only 2 objects, Player 1 should take 1, and would force Player 2 to lose.

Let us continue this analysis with 3 initial objects. In this case, if Player 1 takes 1 object, they would leave the board with 2 objects. However, we just showed whoever had the first move, or next move, with 2 objects

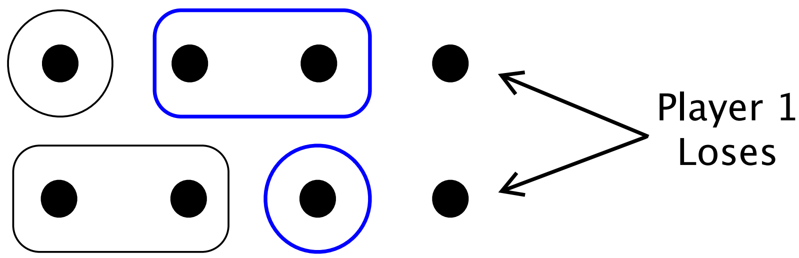

Figure 1: Poison with 4 Objects

should always win. In this cases, the winner would be Player 2. Thus it seems that Player 1's only chance for winning would be if they took 2 objects. This leads to success, since Player 2 is left with only 1 object, which results in a loss. So in the case of 3 objects remaining, the first player to act has a **winning strategy**, in other words a plan of future moves which will always result in them winning, assuming it is played perfectly.

Again, let us continue this process but with 4 initial objects. If Player 1 takes 1 object, they leave Player 2 with 3 objects remaining. If we can look at our analysis for three objects we see what should happen. We showed that the first person to act with 3 objects remaining should win with the appropriate winning strategy. In this case, Player 2 should win. Thus, it seems Player 1 choosing 1 object is a bad decision. Instead, assume Player 1 takes 2. Now they leave the board with 2 objects for Player 2. But, again, we refer to our 2 object analysis, and see the player who goes first with 2 remaining should win. This is also Player 2. So, if Player 2 plays perfectly, they should always win there are 4 initial objects. Regardless of what opening move Player 1 makes, Player 2 can make a move such that Player 1 is immediately forced to take the last object and lose the game. The execution of the winning strategy performed by Player 2 is diagramed in 1.

In our analysis, we began referring to previous cases to help determine who would be the winner. This is called **recursive analysis** because we are referring back to a previous case. The discussion of the recursive analysis may have led to some confusion over who was Player 1 and Player 2. Instead of determining which player is first or second, we will instead discuss the next player to move and the previous player to move. (Note: Once the game is underway, it does not matter who went first only how many objects are left and whose turn it is to move.) In the case of 1 object remaining, the next player to move is always going to lose. This is called a **P position**, meaning the previous player should win assuming they employ a winning strategy. However, with 2 objects remaining, the next player to move should always win because they could take just 1 object and force the other player to take the last one. This is an example of an **N position**. How about with 3 objects remaining? Again, this is an N position, since the next player can force the previous player into a losing move by taking 2 objects. With 4 objects remaining, we are now in a P position, since the next player to move doesn't have a winning strategy. One can show that this pattern of P and N positions repeats itself from this point onward for larger and larger board sizes (i.e. number of objects).

Table of Number of Objects n vs Type of Position for Poison

n	Position
1	P
2	N
3	N
4	P
5	N
6	N
7	P
8	N
9	N
10	P

If the number of objects is 1 mod 3 (the remainder of the number of pieces remaining is 1 when divided by 3), then we are always in a P position, and otherwise we are in an N position. Let us assume we have 1 mod 3 pieces remaining. This should be a P position, so the next player to move should always lose. Why? The next player will have to take one or two objects, and the previous player should choose the opposite number of objects (i.e. if the previous player takes 1 you take 2 and if the previous player takes 2 you take 1). After this pair of moves, we will again be at 1 mod 3 pieces remaining and therefore the one player can ensure they always remain in a winning position when it is their turn. Eventually, we should be at 1 object remaining, and the next player would lose.

However, if the remaining number of pieces is 0 or 2 mod 3, we are in an N position. In these cases, the next player could take the appropriate number of pieces to leave their opponent with 1 mod 3 objects remaining. With 0 mod 3, this number would be 2 and with 2 mod 3, it would be 1. Then we are in the case we finished, and previous player has a known winning strategy.

Poison is an example of an **unfair game**: a game in which there exists a winning strategy for one of the players. Games for which there is no winning strategy are called **fair**. It is often very difficult to determine whether a game is fair or unfair because this relies upon studying all of the possible combination of moves for the game. In the case of the Poison game, there are relatively few moves available and a small number of different game boards one could encounter. Figure 2 illustrates a collection of the game boards with 10 objects where Player 2 is playing the winning strategy.

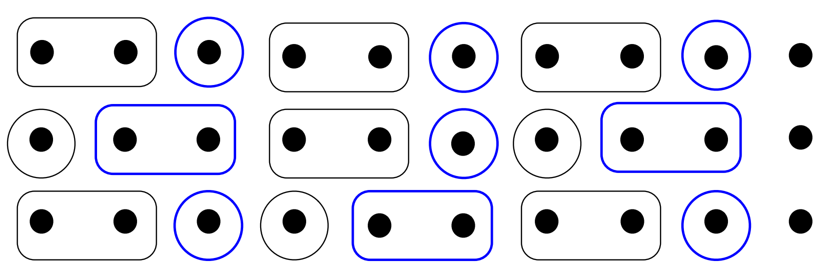

Figure 2: Poison with 10 Objects in which Player 2 Plays Optimally

1.2 Cupcake Game

Consider a slight variation to the rules of the Poison game. Cupcake is a two player game in which there are 10 objects on the table. Players alternate turns and on each turn a player is required to take 1 or 2 of the objects from the table. The person to take the last object wins the game.

The rules to Cupcake are identical to Poison, except that the person to take the last object wins instead of loses the game. A similar analysis can be done to show the possible winning strategies as was done for Poison. In the case of Cupcake, whenever there are 1 or 2 mod 3 objects remaining, the board is in an N position. However, if there are 0 mod 3 remaining objects, the board is in a P position.

Table of Number of Objects n vs Type of Position for Cupcake

n	Position
1	N
2	N
3	P
4	N
5	N
6	P
7	N
8	N
9	P
10	N

1.3 Exercises

1. Consider the game of Poison in which each player is allowed to take 1, 2 or 3 objects. Create a table for n versus N or P position for n=1 to 10.
2. Consider the game of Cupcake in which player is allowed to take 1, 2, 3 or 4 objects. Create a table for n versus N or P position for n=1 to 12.

2 Introduction to Adaptive Learning

In the previous section, the game of Poison was introduced and a winning strategy was found. In the case of the Poison game, there are a relatively small number of move choices to analyze and it was relatively straightforward to find this winning strategy. However, for most games there are a large number of possible move combinations and tracking all the possible choice combinations by hand is prohibitive. In this section, we introduce the idea of computer adaptive learning.

In adaptive learning, the computer will start with no knowledge of how to win the game and will pick what move to make completely at random. The computer will then "learn" how to play the game optimally by playing a large number of games either against itself or a human player. After each game, if the computer lost it will learn not to make the last move which lead to its defeat. To illustrate this, let us discuss a simple adaptive learning scheme for the Poison game.

Each student in the class was given a role in the game or in the computer. One student was the human player, who would play the game as any regular person would. The rest of the students were given different tasks which would illustrate the different tasks the computer would have to do. There was a main program, who would run the game. There was a game simulator would set up the game and alternate turns, asking the current player their move and adjusting the game board appropriately. There was a game over checker that would determine if the game was over. There scorekeeper that would keep the numbers of wins by each player. There were ten students that covered the ten different number of objects remaining. Each one of these students would start with two markers, in our cases two different colored poker chips. When called upon, they would randomly produce one, which would be the computers move at that time. There was also a punisher, who, upon a computer loss, would remove the marker the from the student who last produced one. This would remove those plays which led to a loss.

The game was then played between the human player and the computer, with the human player going first. At first, the choices by the computer were random, and a win came out of pure chance. However eventually the computer would lose. The first loss would result in the punisher removing the markers from the student would stood for 1 object remaining, since it is impossible to win from that position. Also, the student in the spot with 2 objects remaining would lose there take 2 marker, since that would lead to a loss. This would continue until over time the computer would eventually would have the winning strategy for going second in a game with 10 objects, which was shown in the previous day's analysis to exist.

Afterwards, the class began to compile a Pseudo Code, an outline for the program written in simple words, for this computer program. Each part of the program was written, first simply, then fleshed out as the class determined the different concerns that arouse.

At the end of class, a homework assignment was given. In it, each student was asked to give a list of 10 different games and write out a definition for what a game is.

3 Game Theory Terminology

Stop for a moment and jot down ten examples of games you enjoy. Now think about how one would define what properties makes something a **game**. If you look over the ten games on your list, do all of them satisfy your definition of game?

3.1 Terminology Definitions

Def: **Games** are characterized by a number of players or decision makers who interact and take actions under certain conditions and receive some benefit or reward or possibly punishment.

Def: **Game theory** consists of ways of analyzing conflict of interest in game play, to seek information about a "solution", or find a best way of playing various games. For some games, we can only rule out certain types of decisions. **Game theory** is a set of ideas and techniques for analyzing mathematical models for conflict of interest.

5 Elements of a Game:

- Players
  - How many?
  - Does chance play a role?
- Set of all possible actions
  - What players can do
- Information available to players when choosing an action
- Payoff consequences
  - description for all possible combinations of actions by players
- Player's preferences over payoffs
  - utility

Classifying Games:

Determinate vs Random	In a random game there is some amount of chance involved, i.e. rolling a die.
Zero Sum vs Non-Zero Sum	A zero sum game is one in which the amount one player takes away from the game, must be equal to the amount lost by the other players. Games like Checkers and Chess are zero sum games. In a non-zero sum game, the aggregate gains and losses is not zero. An example of this is the prisoner's dilemma.
Symmetric vs Asymmetric	A symmetric game is one where the payoff depends only the strategy, not on who plays them. For an asymmetric game, there are differing strategies for each player.

Perfect vs Imperfect Information	In a perfect information game, the player has available all of the information needed to determine all possible games. Chess is an example of such a game. Games in which it is not possible to determine all possible games are imperfect information games. Card games where each player has hidden hands are examples of this type.
Sequential vs Simultaneous	Sequential games are ones in which the players alternate turns, unlike in simultaneous games where both players perform their turns at the same moment. Checkers and Chess are sequential games, while most sports are simultaneous.
Misére vs Normal	A misére game is one in which the last player to move loses, while in a normal game the last player to move is the winner. Poison was a misére game and Cupcake was normal.
Fair vs Unfair	In an unfair game, there exists a winning strategy for one of the players. In contrast, a fair game has no such winning strategy.

Game Analysis:

Tree Graph	-	Each vertex is a game state in which a move must be made, each line represents a possible move, called "extensive form"' of game.
Payoff Matrix	-	Gives payoff for each player for each combination of strategies possible in the game.

3.2 Nim

Nim is another example of a two player game in which the players alternate turns. In this game, there are chosen number of piles with each pile having some number of objects which can vary from pile to pile. On their turn, a player can remove any amount of objects from a single pile. Nim can be played as either Normal or Misére.

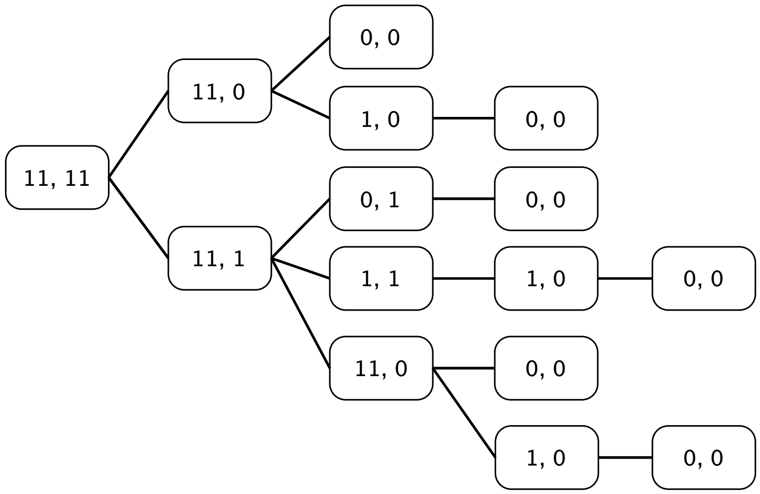

Figure 3: Tree Graph of Nim with 2 Piles of 2 Objects

For the tree graph in Figure 3, a list of every possible choice a player is constructed. For example, Player 1 can choose to take 1 object from either pile their first turn. Player 1 will not have any impact on Player 2's choice, but potentially they could be left with a (II,0) board. Since this is a possibility, they need to be prepared to make a decision on what to do in this case. Thus, if Player 1 does take only 1 on their first turn, there are two possible strategies they could implement, take 1 if (II,0) is presented to them, or take 2 in the same situation. Also, on their first turn, Player 1 could take 2 from any pile. In this case, Player 1 will never have a turn after this in which they need to make a decision. This gives three possible strategies for Player 1. They are listed below, with the 6 possible strategies for Player 2.

Player 1 Strategies:

- Take 1 in (II,II) case and 1 in (II,0) case
- Take 1 in (II,II) case and 2 in (II,0) case
- Take 2 in (II,II) case

Player 2 Strategies:

- Take 1 from P2 in (II,I) case and 1 from P1 in (II,0) case
- Take 1 from P1 in (II,I) case and 1 from P1 in (II,0) case
- Take 2 from P1 in (II,I) case and 1 from P1 in (II,0) case
- Take 1 from P2 in (II,I) case and 2 from P1 in (II,0) case
- Take 1 from P1 in (II,I) case and 2 from P1 in (II,0) case

- Take 2 from P1 in (II,I) case and 2 from P1 in (II,0) case

Payoff Matrix:

After determining all possible strategies, the payoff matrix can be constructed. In it, we pit every Player 1 strategy against every Player 2 strategy and determine the winner. If Player 1 is the winner, a 1 is inserted into the matrix and if Player 2 wins, a -1 is inserted. For the example of Nim with 2 piles of 2 objects each, we would have a 3 x 6 matrix with Player 1's strategies as the rows and Player 2's as the columns.

Ex.: Payoff Matrix for Nim with 2 piles of 2 objects

	II_1	II_2	II_3	II_4	II_5	II_6
I_1	1	1	-1	1	1	-1
I_2	-1	1	-1	-1	1	-1
I_3	-1	-1	-1	1	1	1

Nim falls into a class of impartial combinatorial games where the game is determined by a set of positions, including the initial position, and the player whose turn it is to move. Players alternate turns until a terminal position is reached.

Algorithm to solve a game: (Normal)

- Label every terminal position as a P-position
- Label every position that can reach a P-position in one move as an N-position
- Label positions whose only moves are to a labeled N-position as P-positions
- Repat until no new P-positions

Characteristic Property:

For impartial combinatorial games under normal play rule, P and N-positions are defined recursively by:

- All terminal positions are P-positions
- From every N-position, $\exists$ at least one move to a P-position
- From every P-position, every move is to an N-position

Ex.: Normal Nim with piles (5,7,9)

(0,0,0) is the terminal position, so it is a P-position
(X,0,0) is an N-position because ∃ a move to a P-position
(1,1,2), (1,1,3), (1,2,2) are N-positions
(1,2,3) is a P-position

We would like to find a simple criteria to check if a game position is a P or N position without recursively labelling all the game positions.

Def: The **nim-sum** of $(x_m ... x_0)_2$ and $(y_m ... y_0)_2$ is $(z_m ... z_0)_2$ and we write

$$(x_m ... x_0)_2 \oplus (y_m ... y_0)_2 = (z_m ... z_0)_2$$

where for all k, $z_k = x_k + y_k (mod 2)$. This is just base 2 addition without carry.

Ex.: $22 \oplus 51 = (10110)_2 \oplus (110011)_2 =$

$$\begin{array}{rll} & 010110_2 & \\ \oplus & 110011_2 & \\ \hline & 100101_2 & = 37 \end{array}$$

The nim-sum is associative and commutative.

Theorem: (Bouton 1902)
A position $(x_1, x_2, x_3)_2$ in normal Nim is a P-position iff the nim-sum of the components is zero.
i.e. $x_1 \oplus x_2 \oplus x_3 = 0$

Ex.: (5, 7, 9) Normal Nim Game

$$\begin{array}{rll} & 0101 & \\ & 0111 & \\ \oplus & 1001 & \\ \hline & 1011 & = 11 \neq 0 \end{array}$$

This shows that the starting positions in a (5, 7, 9) is an N-position by Bouton's Theorem.

Proof: (Generalized Bouton's Theorem)

Let $\mathcal{P}$ denote set of Nim positions with zero nim-sum.
Let $\mathcal{N}$ denote the complement of $\mathcal{P}$.

- Terminal position (0, 0, ..., 0) is $\mathcal{P}$ position because $0 \oplus 0 \oplus ... \oplus 0 = 0$

- Form nim-sum as column addition
  Look at the left most column with an odd number of 1's and change any number in that column which is 1 to a 0.
  Repeat this process for the columns to the right which have an odd number of 1's, but instead for the already altered row change the entry to its opposite

This is a legal move to a position in $\mathcal{P}$.

- If $(x_1, x_2, ...) \in \mathcal{P}$ and if $x_1 \to x_1^{'}$
  then we cannot have $x_1 \oplus x_2 \oplus ... = 0 = x_1^{'} \oplus x_2 \oplus ...$
  because the cancellation law gives $x_1 = x_1^{'}$, so $(x_1^{'}, x_2, ...)$ is in $\mathcal{N}$

This shows $\mathcal{P}$ is set of all $\mathcal{P}$-positions.

3.3 Exercises - Normal Nim

1. (3, 5, 7) Normal Nim - Which player has winning strategy?

2. (8, 12, 25) Normal Nim - Player 1 has winning strategy. Why? What is their first move?

3. Find a 5 pile game where Player 2 has winning strategy.

4 Graph Games

The two games that have been presented so far, Cupcake/Poison and Nim, are examples of graph games.

4.1 Graph Definitions

Def: A **directed graph**, G, is a pair (X, F) where X is a nonempty set of vertices (positions) and F is a function that gives for each $x \in X$ a subset of X, $F(x) \subset X$. If $F(x)$ is empty, x is a terminal position.

Def: A **path** is a sequence x_0, x_1, ...,x_m, such that $x_i \in F(x_{i-1})$ $\forall i = 1, ..., m$, where m is the path length.

Def: A graph is **progressively bounded** if $\exists n$ such that every path from x_0 has length $\leq n$.

Figure 4 shows a directed graph of Poison/Cupcake with 5 objects. It can been seen that the longest path from x_0 to x_4 is 5, thus this graph is progressivly bounded.

Def: The **Sprague-Grundy function** of a graph (X, F) is a function function g defined X and taking non-negative integer values such that

$$g(x) = min\{n \geq 0 : n \neq g(y) for y \in F(x)\}.$$

i.e. - $g(x)$ is recursively defined to be the smallest, non-negative integer not found among the Sprague-Grundy values of the followers of x.

For terminal vertices, $g(x) = 0$.

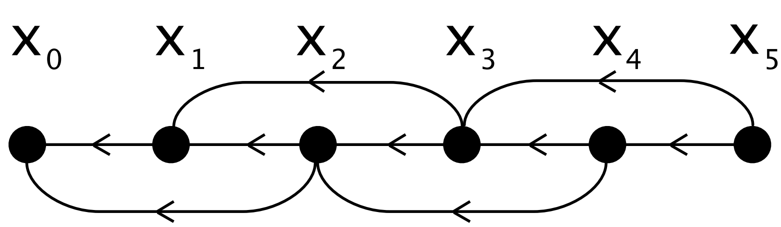

Figure 4: Poison/Cupcake with 5 Objects

Ex. Find the Sprague-Grundy function value for each vertex.

Figure 5: Exercise 1

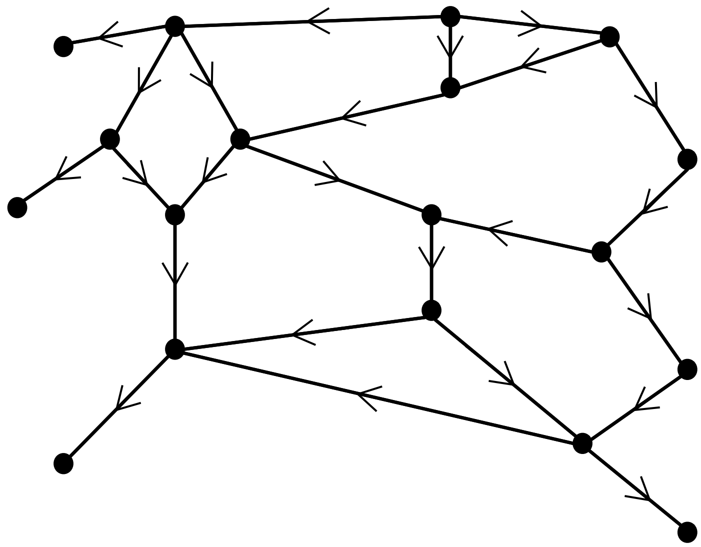

Figure 6: Exercise 2

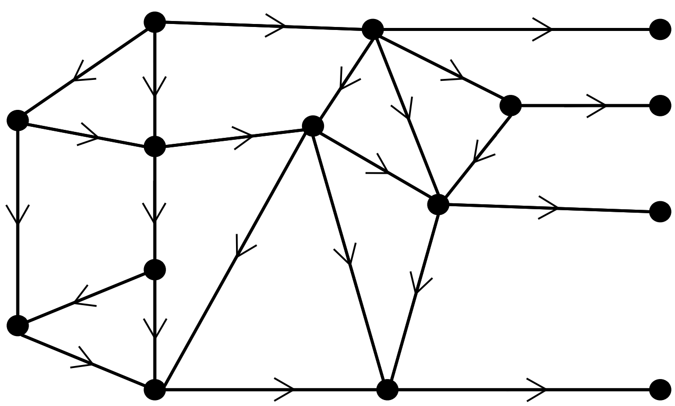

Sprague-Grundy Theorem: The Sprague-Grundy function of a sum of graph games is the nim-sum of the Sprague-Grundy functions of its component games.

Implication - Every impartial game is equivalent to some nim pile.

$\star$ Positions x for which $g(x) = 0$ are P-positions and all other positions are N-positions.

Check:

1. If x is a terminal position, $g(x) = 0$.
2. At positions x for which $g(x) = 0$, every follower y of x is such that $g(y) \neq 0$.
3. At postions x for which $g(x) \neq 0$, there is at least one follower y such that $g(y) = 0$.

4.2 Green Hackenbush

Green Hackenbush is a two player game in which the board is a collection of rooted graphs or graphs which are connected to the ground. On their turn, a player can choose one edge of a rooted graph where that edge and all others not connected to the ground are removed. The last player to make a move wins, thus the game is Normal.

Ex. Bamboo Stalks

$\rightarrow$ equivalent to Nim!

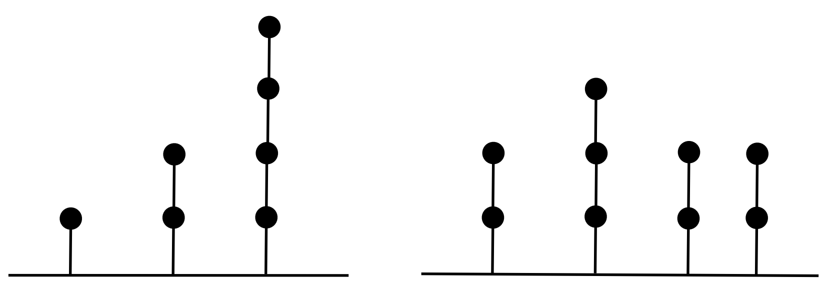

Figure 7: Bamboo Stalks Board Examples

Colon Principle: When branches come together at a vertex, one may replace the branches by a non-branching stalk of length equal to their nim-sum.

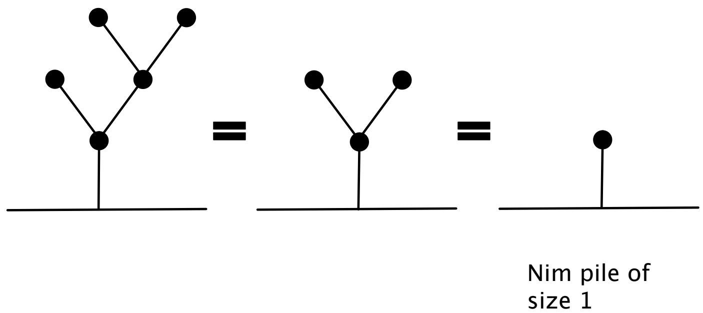

Figure 8: Colon Principle Example 1

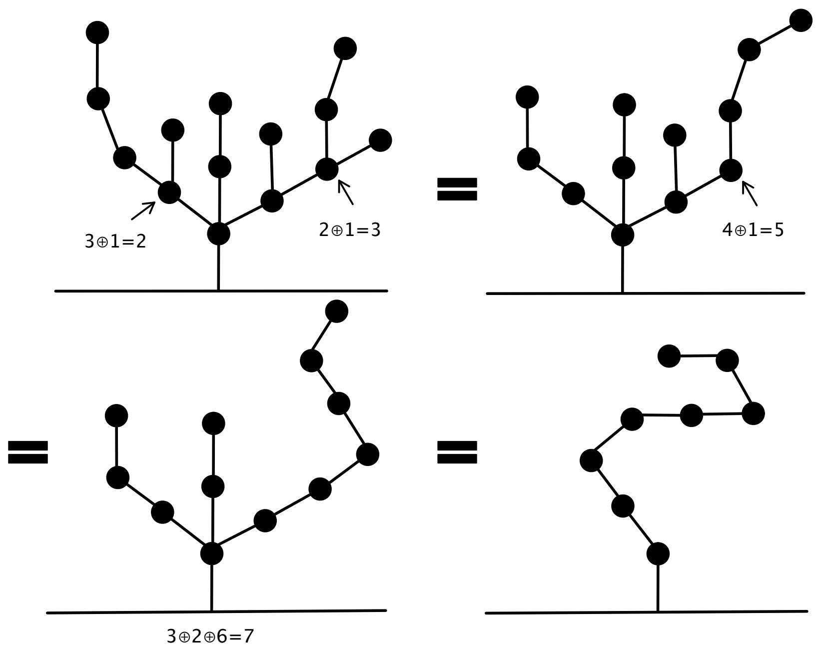

Figure 9: Colon Principle Example 2

Fusion Principle: The vertices on any circuit may be fused without changing the Sprague-Grundy value of the graph.

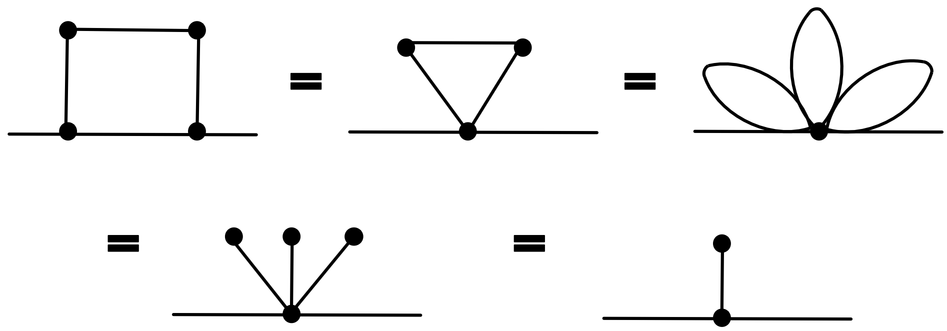

Figure 10: Fusion Principle Example 1

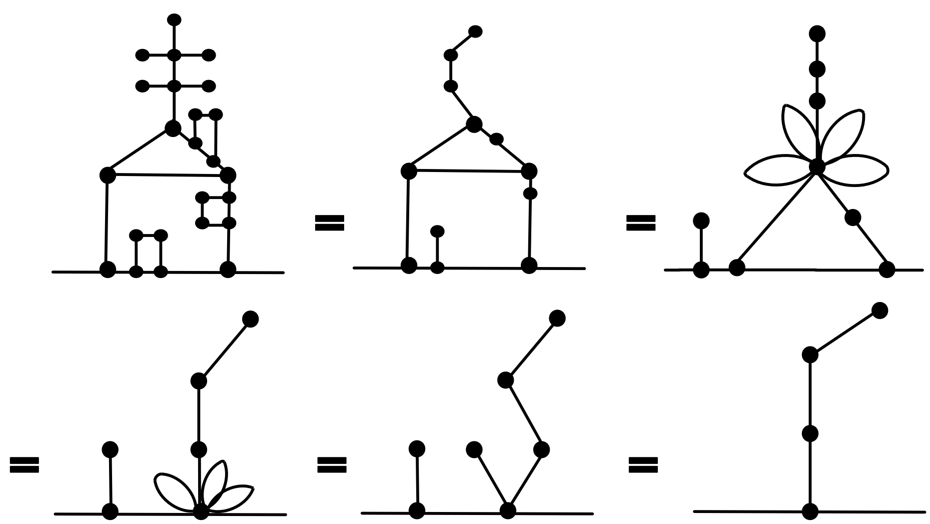

Figure 11: Fusion Principle Example 2

Note:

- Circuit with an odd number of edges = one edge
- Circuit with and even number of edges = single vertex

Exercise: Show that the Sprague-Grundy value is 4.

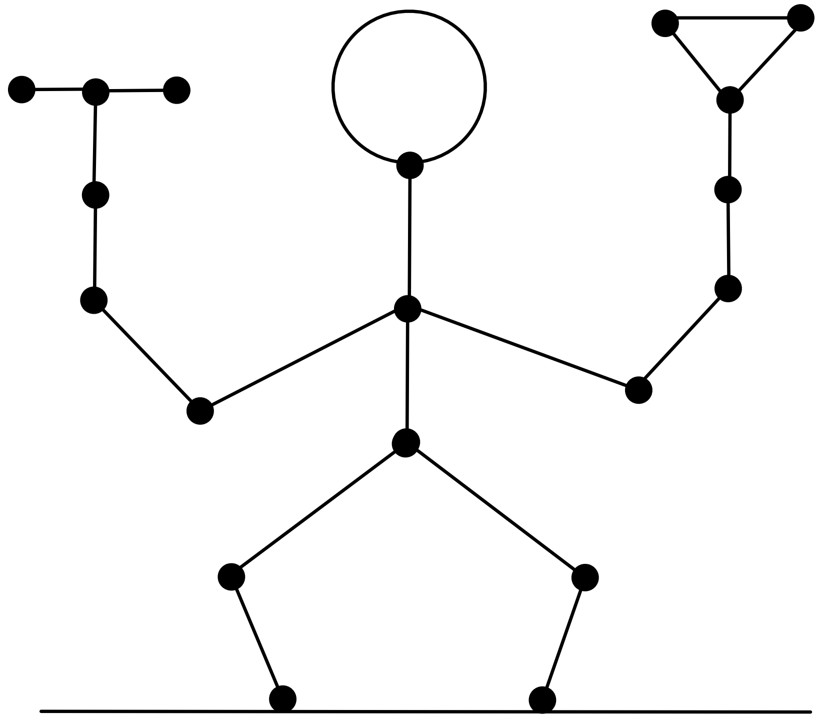

5 Hex and Chomp

5.1 Hex

Hex is a two player game played on a diamond-shaped board made of hexagons. The players, White and Black, each own two opposite edges of the board. The players alternate turns coloring any open hex their color. The object of the game is to create an unbroken chain of their color connecting their two edges of the board. The first player to create such a chain is the winner.

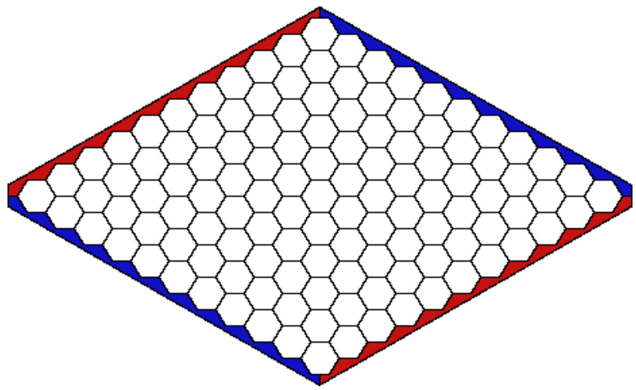

Figure 12: 7 × 7 Hex Game Board

Hex was created in 1942 by Piet Hein and in 1948 by John Nash. John Nash invented it to play on the bathroom tiles at Princeton. The game was called "John" in honor of Nash and since it was played in the restroom. Parker Brothers marketed the game in 1952.

There are two things to note about Hex. First, the game cannot end in a tie. And, secondly, for any size board n × n, the first player has a winning strategy. The proof that first player can is sadly nonconstructive, so the winning strategy for n larger than 7 is still unknown.

Proof: (First Player Can Win)

Hex is a two player, perfect information, symmetric game. As was noted previously, the game cannot end in a tie. Now, for the purpose of contradiction, assume that the second player the winning strategy. Another item of note is that in this game no "extra play" can be a disadvantage. Thus, on their first turn, Player 1 could anywhere and lose nothing. At this point, Player 1 would become Player 2 and could play the second player's strategy. Therefore, Player 1 would have the winning strategy. This is a contradiction of our assumption, thus Player 2 cannot have the winning strategy. Hence, since the game cannot end in a tie, the first player must have the winning strategy. This is type of proof is called a **strategy stealing argument**.

5.2 Chomp

Chomp is another two player game; this time played on a rectangular "chocolate bar" made of smaller rectangles. On a player's turn, can select any block and remove it from the board along with any other blocks to the right and below the chosen piece. The top-left block is considered to be "poisoned" and whomever eats it loses.

Chomp, like Hex, is a two player, perfect information, symmetric game. Also, the game cannot end in a tie, since one person must eat the ”poison” block. And lastly, the first player has the winning strategy. This can be proven with the same strategy stealing argument that was made for Hex. On Player 1’s first move, they could take the most bottom right block. Player 2 should have a response which would lead to victory. But then Player 1 could have played that move first, since any move would remove the bottom right block. Thus the first player had the winning strategy to start with.

6 Hexapawn

Hexapawn is a two player game played on a 3x3 board. On a player's turn, they can move forward one square to an empty location or they can move diagonally to capture an opponent's piece (like pawns in chess). The game ends in one of three ways:

- A player captures all of the opponent's pieces
- A player advances one of their pieces to the opposite side
- A player has no more legal moves.

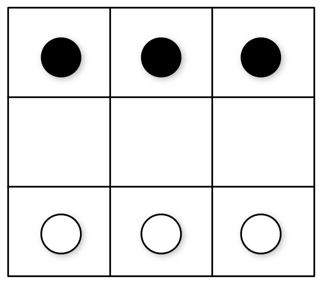

Figure 13: Starting board for Hexapawn

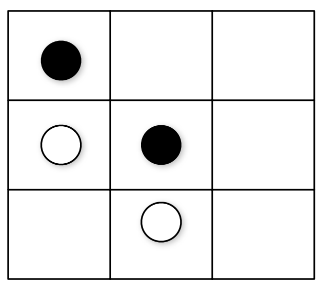

Figure 14: A board with no more legal moves

Hexapawn is a perfect information combinatorial game, with a maximum game length, a finite number of combination, and it cannot end in a tie. Therefore, it is an unfair game. In fact, player 2 has the winning strategy. Hexapawn was invented by Martin Gardener. He presented it in an article in *Scientific America* in 1968.

Theorem: All perfect information combinatorial games that are progressively bounded and cannot end in a tie are unfair.

7 Super TTT

Super TTT is a version of the game Tic-Tac-Toe. The board is a 3x3 grid, where each spot of the grid has a 3x3 tic-tac-toe board in it (see figure 15). Tic-tac-toe is played normally in each of the small grids. However, where you move tells your opponent where they have to move next (see figure 16). When a player wins a small grid, they get that square on the big grid (for example, if X wins a small grid, then that box gets a big X in it). The first player to get three boxes in a row on the big grid wins. It a small board ends in a cat's game, then that spot on the big grid gets an X and an O. Finally, if I make a move that directs the next player to a board that is already finished, then that player can pick any of the small grids to play in.

Super TTT is a combinatorially, perfect informations, progressively bounded (longest game is 81 moves) game with a guaranteed win or tie for one player. Therefore, it is unfair with the possibility of a tie.

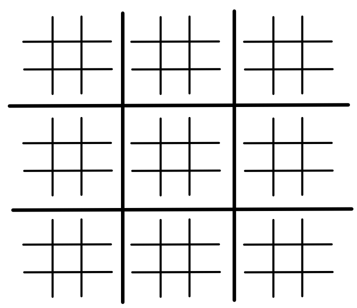

Figure 15: Super TTT Board

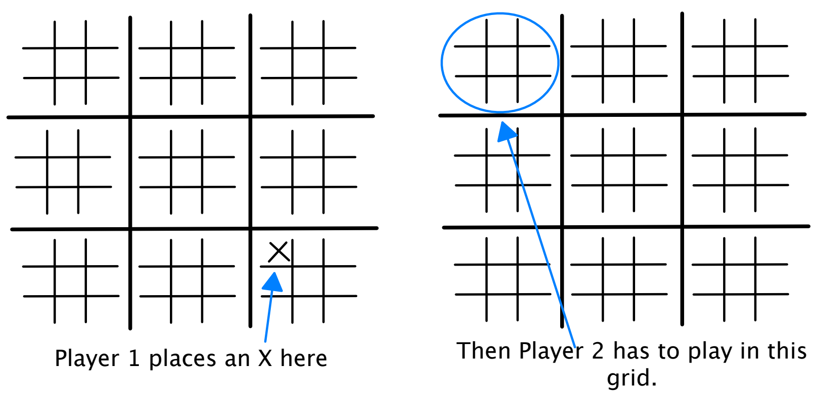

Figure 16: Super TTT Moves

8 Simultaneous Games (Non-combinatorial)

8.1 Nash Equilibria

Definition: **Nash equilibrium** is a solution concept in which each player is assumed to know the strategies of the other players and no player has anything to gain by changing only his own strategy unilaterally.

A group of players is in Nash equilibrium if each one is making the best decision they can, taking into account the decisions of the others.

Nash equilibrium does not necessarily mean the best cumulative payoff for all the players involved.

8.1.1 Examples

Ex. Chicken

In the game of chicken, two cars drive head on until either one or both cars swerve out of the way or they collide. The winner is the driver who does not swerve when his oppenent does. Below is the payoff matrix for chicken.

	Swerve	Straight
Swerve	$(0,0)$	$(-1,1)$
Straight	$(1,-1)$	$(-10,-10)$

The entry $(1,-1)$ is a Nash equilibrium. This because if Player 1 were to change his mind, he would go from a payoff of 1 to 0, thus no longer making a better decision. The same is true for Player 2 since he would go from a payoff of -1 to -10. For a similar reason, $(-1,1)$ is also a Nash equilibrium. Neither $(0,0)$ nor $(-10,-10)$ is a Nash equilibrium, because if either player changed their strategy, they would increase their payoff.

Ex. Prisoner's Dilemma

The Prisoner's Dilemma involves two prisoners being interrogated about a crime they both committed. Each prisoner has the option of staying silent or betraying their partner. Below is the payoff matrix for the Prisoner's Dilemma.

	Silent	Betray
Silent	$(-1,-1)$	$(-10,0)$
Betray	$(0,-10)$	$(-3,-3)$

The entry $(-3,-3)$ is the only Nash equilibrium. If either player is to change their mind in this case, they would go from a payoff of -3 to -10, thus not improving their prospects. $(-1,-1)$ is not a Nash equilibrium because if either player where to change their mind, they would improve their payoff. For the cases $(0,-10)$ and $(-10,0)$, if the player who is silent changes their strategy, they would improve their payoff. Therefore, these cases are not Nash equilibria either.

If we consider the iterative prisoner's dilemma (where the game is played multiple times), the best strategy is tit-for-tat (your move is what your opponent did the last time).

Ex. Presentations

During a presentation, there is the presenter and the audience. Both have the option of giving effort in their half of the presentation; for the presenter this is in preparing for and giving a good presentation, and for the audience this is in listening intently. Below is the payoff matrix.

	Effort	No Effort
Effort	$(4,4)$	$(-16,-14)$
No Effort	$(0,-2)$	$(0,0)$

For this payoff matrix, both $(4,4)$ and $(0,0)$ are Nash equilibria. For both points, if either the presenter or the audience were to change their strategy their payoff would decrease. The other two points are not Nash equilibria because in both cases changing one's strategy lowers their payoff.

Exercise: What is the equilibrium of the following game?

	X	Y
X	$(10,10)$	$(15,5)$
Y	$(5,15)$	$(12,12)$

8.2 Strategic Form

Definition: The **strategic form** or **normal form** of a two person zero-sum game is given by a triplet (X, Y, A) where

1. X is a nonempty set of strategies of Player I
2. Y is a nonempty set of strategies of Player II
3. A is a real-valued function defined on $X \times Y$.

An interpretation of this is if Player I chooses $x \in X$ and Player II chooses $y \in Y$ simultaneously, then Player I gets payout $A(x, y)$.

8.3 Minimax algorithm

Definition: A **minimax algorithm** is a recursive algorithm for choosing the next move in a game.

The idea is to look down the game tree and rank each board, then to pick your path based on the "best" options. In practice, we pick a search depth, rank those boards, and choose a move based on the best collective options. We also want to be sure that if there is an immediate win available, then we pick it. In general, the board rating algorithm can be tough, but we can use machine learning to help.

Minimax Theorem For every finite two-person zero-sum game:

1. There is a value V of the game
2. There is a mixed strategy of Player I such that Player I gains at least V no matter what Player II does
3. There is a mixed strategy of Player II such that Player II gains at most V no matter what Player I does.

For simultanous games, we can just switch the I and II. This gives us the **maximin theorem**.

8.3.1 Examples

Ex. Odd or Even

The game odd or even has two players. On each turn the players simultaneously pick the number 1 or 2. Player I wants the sum of the two number to be odd and Player II wants the sum to be even. The winner gets the sum. Below is the payoff matrix for Odd or Even.

	One	Two
One	$(-2,2)$	$(3,-3)$
Two	$(3,-3)$	$(-4,4)$

This game does not have any Nash equilibria.

Let's figure out the expected payout for Player I. Let p equal the proportion of the time Player I shows one. If Player II shows one, then Player I's payout is $-2p+3(1-p)$. If Player II shows one, then Player I's payout is $3p-4(1-p)$. We will set these equal to each other because we want to know when the payouts will be the same (no matter what Player II picks). We get the equation:

$$-2p+3(1-p)=3p-4(1-p).$$

Solving this for p gives $p=7/12$. Therefore, Player I's **equalizing mixed strategy** is is to show one 7/12 of the time and to show two 5/12 of the time.

To find Player I's **expected payout** we take the value for p we found above and substitute it back into one of the equations so that:

$$E=-2p+3(1-p)=-2(7/12)+3(5/12)\approx 8 \text{ cents}.$$

We can perform the same calculations for Player II and find that he has the same optimal strategy, and his payout is -8 cents.

Ex. Soccer

Below is a table for soccer giving the percent of the time that the kicker makes a goal given whether they and the goalie go left or right. We would like to figure out what percent of the time each player should go left or right. Also, we would like to see the percent of the time that the goal kicks should be successful.

	Goalie: Left	Goalie: Right
Kicker: Left	.6	.9
Kicker: RIght	.95	.75

Let's look at the kicker first. Let p be the probability that the kicker kicks left and $(1-p)$ be the probability that he kicks right. Then, if the goalie goes left, we get $.6p+.95(1-p)$ and if the goalie goes right, we get $.9p+.75(1-p)$. Setting these equal to each other and solving for p we find that $p=.4$. Therefore, the kicker should kick left 40% of the time and he should kick right 60% of the time.

For the goalie, let q be the probability that he goes left and let $(1-q)$ be the probability that he goes right. Then if the kicker kicks left, we get $.6q+.9(1-q)$ and if the kicker kicks right we get $.95q+.75(1-q)$. Setting these equal to each other and solving for q we get $q=.3$. Therefore, the goalie should go left 30% of the time and go right 70% of the time.

The total percent of goals should be

$$.6p+.95(1-p)=.6(.4)+.95*(.6)=.81.$$

Therefore, in soccer we expect the kicker to score a goal 81% of the time.

Ex. Consider the following game tree (or **Kuhn tree**) diagram (see figure 17).

We can represent this as a matrix of strategies and payoffs. This is called the **strategic form** or **extensive form** of the game. This is called the **bimatrix form** if the game is non-zero sum strategic form.

$$\begin{bmatrix} (1,4) & (2,0) & (-1,1) & (0,0) \\ (3,1) & (5,3) & (3,-2) & (4,4) \\ (0,5) & (-2,3) & (4,1) & (2,2) \end{bmatrix}$$

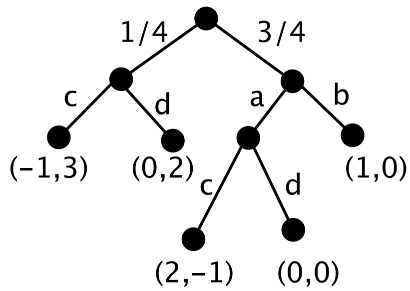

Figure 17: A Game Tree

Ex. Consider the following strategic from of a game:

$$\begin{bmatrix} (2,0) & (1,3) \\ (0,1) & (3,2) \end{bmatrix}$$

We can split this into the matrices for each player:

$$A = \begin{bmatrix} 2 & 1 \\ 0 & 3 \end{bmatrix} \qquad B = \begin{bmatrix} 0 & 3 \\ 1 & 2 \end{bmatrix}$$

Let's find Player I's maximin strategy. Suppose that Player I makes their first move with probability p. If Player 2 picks move α, we get $2(p) + 0(1-p)$. If Player 2 picks move β, we get $1(p) + 3(1-p)$. Setting these equal to each other, we find that $p = 3/4$. Therefore, Player I's strategy is to pick the fist move 3/4 of the tome and the second move 1/4 of the time.

To find Player II maximin strategy, we look at matrix B. We can see that Player II has a **dominant strategy**, which is to always make the second move (we can see this be looking at the second column of matrix B. This strategy is clearly better if I only care about what happens to me, but it I want bad things to happen for Player I too, this might not be the best strategy.

Definition: The **safety levels** are the payout of a minimax strategy.

Player I's safety level is $2 * p = 2 * (3/4) = 6/4$.

This game is not zero sum, so does the the minimax theorem tell us anything? If does if we combine together what we learn from it with the safety levels.

9 Combinatorics

9.1 Addition and Multiplication Principles

Addition Principle: If there are r_1 different objects in the first set, r_2 objects in the second set, ... , r_m objects in the m^{th} set, and if the different sets are disjoint, then the number of ways to select an object from each of the m sets is $r_1 + r_2 + \cdots + r_m$.

Multiplication Principle: Suppose a procedure can be broken down into m successive stages with r_1 outcomes in the first stage, ... , r_m outcomes in the m^{th} stage. If the composite outcomes are all distinct, then the total procedure has $r_1 \cdot r_2 \cdots r_m$ different composite outcomes.

Ex. Suppose there are 40 students in an algebra class and another 40 students in a geometry class. If there are exactly 10 students in both classes, how many total students are there?

Solution: Assuming there were no overlap in the two class, the max number of students there could be is 80. However, since 10 are in both, there are only 30 students in each class who only take that one course. Thus, the 30 students only in algebra plus the 30 students in geometry plus the 10 taking both courses gives 70 students total by the addition rule.

Ex. Given two dice, one of which is green and the other is red, how many total outcomes are there? How many outcomes are there with no doubles?

Solution: Since each die is independent of the other, and each die has 6 possible outcomes, by the multiplication rule we would have $6 \cdot 6 = 36$ total outcomes. For the second question, we know there are 6 outcome that involve doubles, thus the total with no doubles must be $36 - 6 = 30$.

Ex. Suppose there is a stack of 5 Spanish, 6 French, and 8 Russian books. How many ways are there to pick 2 books not in the same language?

Solution: When picking two books of different languages, the possible outcomes are Spanish and French books, Spanish and Russian books, and French and Russian books. In the case of ending up with a pair of Spanish and French books, there are 5 possible choices for the Spanish book and 6 for the French book. Since the number of books in each language is independent from the other, by the multiplication rule, there are $5 \cdot 6 = 30$ possible outcomes for the pair of Spanish and French books. Similarly, there are 40 possible outcomes for Spanish and Russian and 68 possible outcomes for French and Russian. Thus the total number of outcomes is $30 + 40 + 48 = 118$.

Ex. How many was to form 3-letter sequence from a, b, c, d, e, f?

(a) With repetition allowed.

(b) Without repetition allowed.

(c) Without repetition, containing an e.

(d) With repetition, containing an e.

Solution:

(a) Since each letter in the sequence is independent, there are $6 \cdot 6 \cdot 6 = 216$ outcomes.

(b) The first letter in the code could be anything, thus there are 6 possible choices for the first letter. The second letter can be anything besides the letter already chosen, leaving 5 choices. The last letter can be any besides the first two, thus in total there are $6 \cdot 5 \cdot 4 = 120$ outcomes.

(c) Since an e must be contained, there are 3 possible locations for that e. For the other two letters, the first slot can be filled by anything besides an e, thus there are 5 options and the other slot can be filled by anything besides an e and whatever was in the first slot, leaving 4 choices. Thus for the two choices not involving 3, there are $5 \cdot 4 = 20$ outcomes. Thus in total there must be $3 \cdot 20 = 60$ outcomes.

(d) The main concern with the question is to avoid potential double counting of some outcomes. From a), there are 212 total outcomes with repetition. Since the code must contain an e, let us look at how many codes there are with repetition that do not contain an e. This is easy to solve since each option has 5 choices, there are $5 \cdot 5 \cdot 5 = 125$ outcomes without an e. Thus the number of outcomes containg an e must be $216 - 125 = 91$.

9.2 Arrangements and Selections

Def. A **permutation** of n objects is an arrangement, or ordering, of the n objects.

Def. An r**-permutation** of n distinct objects is an arrangement of r of the n objects. The number of r-permutations of n objects is given by the formula $P(n,r) = \dfrac{n!}{(n-r)!}$.

Def. An r**-combination** of n objects is an unordered selection of r out of the n objects. The number of r-combinations of n objects is given by the formula $C(n,r) = \dfrac{n!}{r! \cdot (n-r)!}$.

Ex. Suppose there are n candidates for Chief Wizard. How many ways are there to order these candidates? What is the probability that Gandalf is 2nd?

Solution: The total number of orderings of n wizards out of n is given by $P(n,n) = \dfrac{n!}{(n-n)!} = n!$. Now given that Gandalf is 2nd, the number of arrangements of the other n-1 wizards is $P(n-1, n-1) = (n-1)!$. Thus the probability that Gandalf is 2nd is $\dfrac{(n-1)!}{n!} = \dfrac{1}{n}$.

Ex. How many ways are there to position 12 elves at a circular table?

Solution: First let one elf sit anywhere. Then one of the other 11 can sit on his left, followed by one of the remaining elves on his left, and so on. Thus there are 11! ways.

Ex. How many ways to arrange seven letters of SYSTEMS?

Ex. Find the probability of being dealt a flush in a 5 card poker hand.

Soultion: The probability is given by

$$P = \frac{\text{Number of flushes}}{\text{Total number of hands}}.$$

The total number of hands is given by

$$C(52,5) = \frac{52!}{5!(52-5)!} = 2,598,960.$$

The number of flush hands is given by

$$4 \cdot C(13,5) = 4 \cdot \frac{13!}{5!8!} = 5148.$$

Therefore, the probability of getting a flush is

$$P = \frac{5148}{2,598,960} \approx 0.2\%.$$

Exercise: In a race containing 10 runners, what is the number of ways of getting 1st, 2nd, and 3rd place?

9.3 Pigeonhole Principle

Def. The **pigeonhole principle** states that if there are more pigeons than pigeonholes, then some pigeonhole must contain 2 or more pigeons. In other words, if there are n objects and m slots and $n > m$, then in order to place all n objects in all m slots, some slot must contain at least 2 objects.

Ex. Suppose an area has 20 towns. How many people from the area would be needed to insure there are 3 from the same town?

Solution: It is possible that the first 40 persons could be such that there is a pair from every single town. But when a 41st person is added, they must be from one of the 20 towns also. Since every town already has two people from it, one now must have a third.

Ex. There is a group of 20 couples. How many people do we need to talk to to ensure that we have talked to a husband and wife.

Solution: We must talk to 21 people. If we talk to 20 people, we might have been really unlucky and talked to only one person from each couple, but once we talk to the 21st person, they must be spouse of someone we have already spoken to.

Ex. A professor only tells 3 jokes each year. How many jokes would the professor need to never repeat a triple of jokes over 12 years?

Solution: The professor needs 6 jokes. $C(6,3) = 20$ so we actually have enough triples for 20 years. We can see that 5 jokes is not enough since $C(5,3) = 10$.

10 Instant Insanity

Instant insanity is a puzzle game played with four cubes. The cubes are colored as shown below in figure 18. The goal of the puzzle is to stack the cubes into a tower such that each side of the tower has each of the four colors on it.

We'd like to figure out the probability of getting a solution at random. What is the total number of possible towers? There are 24 symmetries of a cube (there are six sides that can each be made the top, and for each of these, the cube can be placed 4 different ways). Therefore, the total number of towers is $24^4 = 331,776$. We can reduce this number to $41,472$ since it doesn't really matter which color faces you for the first cube. There are 2 distinct solutions.

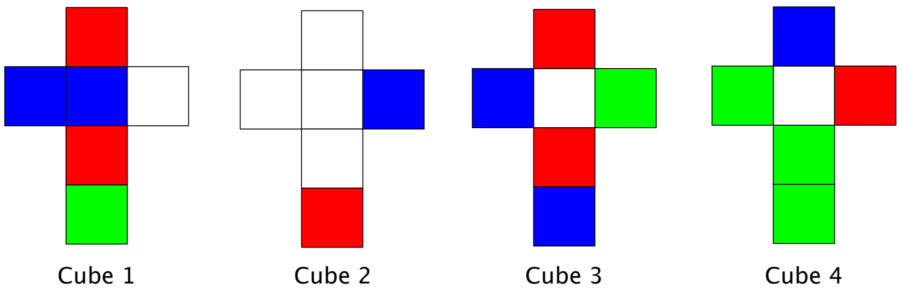

Figure 18: Instant Insanity Cubes

10.1 Mathematical Representation

To represent this game mathematically, we can draw a **multigraph:** a graph with labeled edges. For use, a numbered line between two colors will represent that those colors or on opposite faces of that numbered cube (see figure 19).

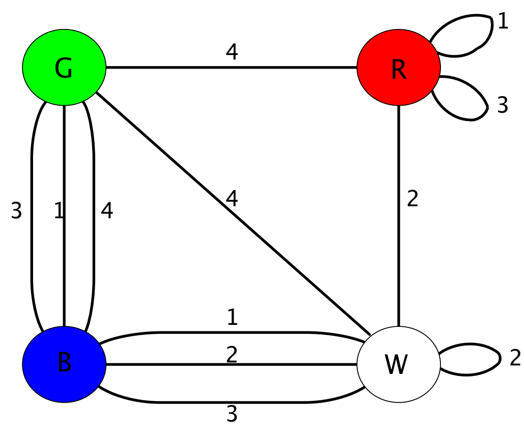

Figure 19: Multigraph for Instant Insanity Cubes

To solve the puzzle, we want to apple the decomposition principle:

Decomposition Principle

1. Pick one pair of opposite faces on each cube for the left and right sides so that the sides have one face of each color.

2. Pick a different pair of opposite faces on each cube for the front and back sides so that the sides have one of each color.

How can we use the multigraph to help us solve the puzzle? We want to pick a subgraph with exactly one edge of each label and each vertex incident with two edges (i.e. two edges touch each vertex). In graph theory terminology, we want to find a subgraph which is a **labeled factor of degree two**. Figure 20 gives some examples of labeled factors of degree two that our found in figure 19. We then want to add direction to each line to make a **digraph**. Let the tail of each line the color that goes on the left side of the tower and the head of each line be the color that goes on the right side. We now have a tower that is correct on the left and right sides, but we still need the front and back sides (see figure 21). We will pick another labeled factor a degree two that does not share any edges with the one we already used. We will use this digraph to get the front and back sides correct (see figure 21).

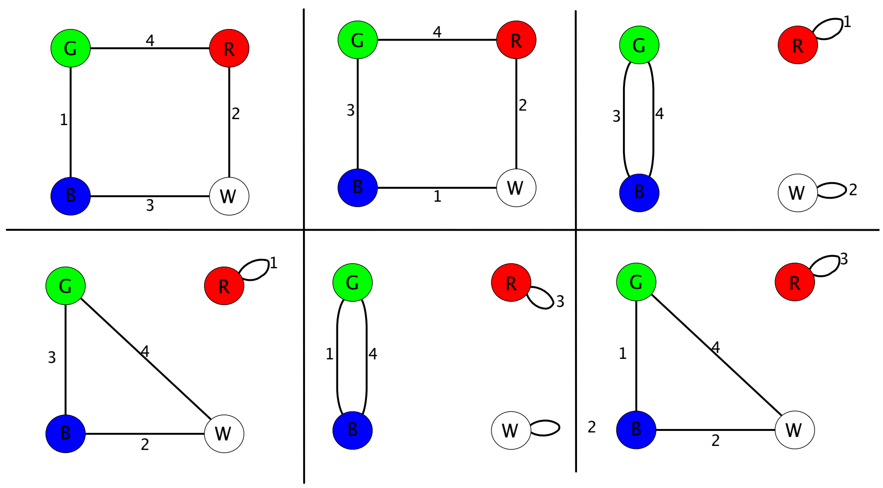

Figure 20: Labeled Factors of Degree Two

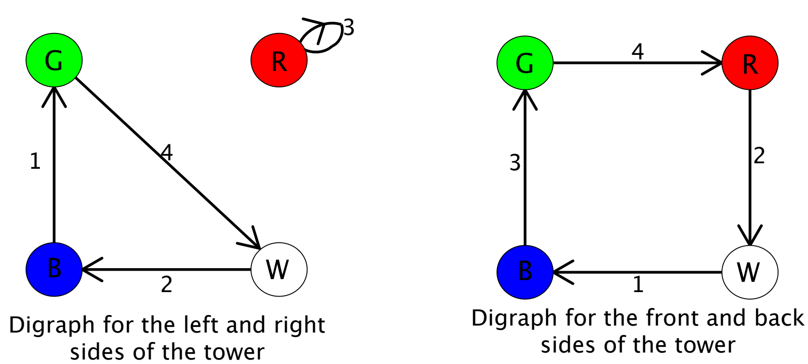

Figure 21: A solution to Instant Insanity

11 Rubik's Cube

A Rubik's Cube is a 3x3x3 cube where each side is a different color (see figure 22). The sides can be rotated, mixing up the colors, and the goal is to solve the cube by getting back to the form where each side is all one color.

Figure 22: Rubik's Cube (picture from www.rubiks.com)

11.1 Number of Configurations

What are the number of configurations of a Rubik's cube?

- There are 8! ways to organize the corners
- There are 7 corners that can orient independently, so this gives 3^7 total ways
- There are 12!/2 ways to organize the edges
- There are 11 edges that can orient independently, so this gives 2^{11} total ways

- Therefore there are a total of $\frac{8!\ 3^7\ 12!\ 2^{11}}{2} = 43,252,003,274,489,856,000$ configurations
- In other words, there are over 43 quintillion configurations!

11.2 God Number

The **God number** is the number of moves to make to solve any given cube. In 2010, it was found that the God number is 20. Therefore, any Rubik's cube can be solved in at most 20 moves.

11.3 Solving the Rubik's Cube

Definition: Singmaster-Frey Notation The Singmaster-Frey notation denotes moves by: U (up), D (down), L (left), R (right), F (front), and B (back). Each move is a quarter turn of the indicated face, clockwise (where you imagine a clock of the face to determine which way is clockwise). For example, L indicates to turn the left face 1/4 turn clockwise. L^{-1} indicates to turn the left face 1/4 turn counterclockwise.

Example: UR tells us to turn the upper face 1/4 clockwise and then turn the right face 1/4 clockwise. To undo this move we would $R^{-1}U^{-1}$.

Example: F^3=FFF tells us to move the front face 3/4 clockwise.

11.4 Commutator

We would like to do is to find move sequences that move a small number of cubies and leave the rest alone.

Definition: If g is an operation on a cube, the **support of g** (supp[g]) is the number of pieces changed by g.
Example: The maximum value for supp[g] is 20.

Example: supp[B]=8 since moving the back face 1/4 turn clockwise changes 8 of the cubies.

Example: supp[$RBLFUF^{-1}L^{-1}B^{-1}R^{-1}U^{-1}$]=3.

Definition: If g and h are two operations, then [g,h]=$ghg^{-1}h^{-1}$ is called the **commutator.** Note that [g,h]=1 if and only of gh=hg.

Example: $RBLFUF^{-1}L^{-1}B^{-1}R^{-1}U^{-1}$=[RBLF,U]