THE GAME OF SPROUTS

ISABEL VOGT

1. Brussels Sprouts

We begin with the game of Brussels Sprouts. The rules are as follows:

- Draw some number of crosses/spotss. We'll generally refer to this number with the variable n.
- Players alternate drawing a line between two "free" ends of crosses and making a new cross in the middle. This line must not cross any other lines already drawn.
- The game ends when there are no new lines that can be drawn.

An obvious question regarding the game of Brussels Sprouts is:

Is there a formula for the total number of moves in a game of sprouts and what does it depend upon?

If this depends upon only n, the beginning number of spots, then the winner of the game is completely determined by who begins the game, irrespective of their strategy in the game.

We will prove the following theorem, which answers this question:

Theorem 1.1. *A game of Brussels Sprouts beginning with n spots has exactly $5n - 2$ moves. Thus player 1 wins if n is odd, and player 2 wins if n is even.*

To prove this, we need to recall information about the Euler Characteristic of a graph.

1.1. The Euler Characteristic.

Let G be a connected planer graph. That is, G is a collection of vertices, edges, and faces, such that it can be made to lie in the plane without any edges crossing eachother.

Then we can define a quantity as follows: let $\chi_E(G) = V - E + F$ for V the number of vertices, E the number of edges, and F the number of faces. We call $\chi_E(G)$ the Euler Characteristic of the graph G. Then we have the following Theorem:

Theorem 1.2. *For G a connected planer graph,*

$$\chi_E(G) = V - E + F = 2.$$

Recall that χ_E is a topological invariant. The fact that $\chi_E = 2$ for all connected planer graphs is a reflection of the fact that they can lie in a plane (which is equivalent to lying on a sphere).

We can prove this theorem by reducing to the case of a certain class of graphs known as trees. A tree is a graph that contains no cycles: that is it contains no nonzero paths beginning and ending at the same vertex. The following shows examples:

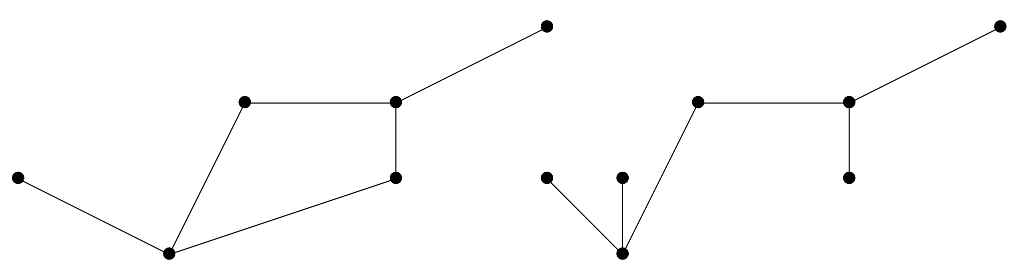

the first is note a tree, but the second is a tree. You can see that another definition is that it only has 1 face – the outside one – as cycles create extra faces. This leads to another key point: if we were to add in any one edge between vertices to a tree, then it would be no longer be a tree!

Another way to think about a tree is that it is built out of some number of "vertex-edge" units plus one extra vertex.

Now using this we can prove the Euler Characteristic formula for planer graphs:

Proof. First we prove this for trees T. Let v denote the number of vertices of the tree. Then by the previous discussion we know that there are $v-1$ edges and 1 face. Hence the Euler characteristic of T is

$$\chi_E(T) = v - (v-1) + 1 = 2.$$

Now for any planer connected graph G, we may take a tree spanning G – that is just delete the extraneous edges so that it is a tree. Call such a tree T_G. Then

$$\chi_E(T_G) = 2$$

by the previous paragraph. Write

$$\chi_E(G) = \chi_E(T_G) + \Delta V - \Delta E + \Delta F,$$

where ΔV is the difference in the number of vertices between T_G and G, ΔE is the difference in the number of edges, and ΔF is the difference in the number of faces.

By the fact that T_G is a spanning tree for G, they have the same vertices, hence $\Delta V = 0$. Now, say that we add back in an edge at a time, then this edge *must* add in a face to the graph. Similarly for every edge we add back in.

So $\Delta E = \Delta F$ and hence

$$\chi_E(G) = \chi_E(T_G) + 0 - \Delta E + \Delta E = \chi_E(T_G) = 2. \qquad \square$$

For more on Euler Characteristics, see the section at the end.

1.2. **Analysis of Brussels Sprouts.** Ideally we'd like to figure out who wins in the game of Brussels Sprouts. We could determine this if we knew how many moves there were in a game, we'd be done. The first player always wins with an odd number of moves and the second player always wins with an even number of moves – by definition of winning! In order to know who will win a Brussels Sprouts game, we just need to determine if the number of moves is even or odd.

Theorem 1.3. *The number of moves in a game of Brussels Sprouts starting with n spots is:*

$$m = 5n - 2.$$

Proof. By using the Euler Characteristic formula $2 = v - e + f$, we can then break it down in order to create a formula for Brussels Sprouts, to figure out who wins a game. First v turns out to be $n + m$ because we start with the number of spots n, every move we create a new spot, so m in total.

Second, e breaks down into $2(m)$ because after every move, two new edges were added to the line drawn.

Third f breaks down into $4n$ because the number of free ends of the cross is constant throughout the game and each face has a free "end" at the end of the game. The number of these free ends at the end of the game is thus $4n$.

$$\begin{aligned} 2 &= v - e + f \\ 2 &= (n+m) - (2m) + (4n) \\ 2 &= 5n - m \\ m &= 5n - 2. \end{aligned}$$

To conclude we have the new formula of $m = 5n - 2$. □

This shows that player 1 wins if n is odd and player 2 wins if n is even. So the game of Brussels Sprouts is not very interesting to play, since we already know who wins!

2. Various Questions About Sprouts

2.1. The game of sprouts.

The game of sprouts is very similar to the game of Brussels Sprouts, except that instead of being allowed to have 4 lines meeting at a vertex, you're only allowed to have 3. So the rules can be summarized as:

- Draw some number of spots. We'll generally refer to this number with the variable n.
- Players alternate drawing a line between two spots and making a spot in the middle. This line must not cross any other lines already drawn, and **only three lines can meet any given spot**.
- The game ends when there are no new lines that can be drawn.

2.2. Questions.

Here are some natural (to me) questions regarding the game of Sprouts that might have interesting answers.

(1) Does the game of Sprouts necessarily ever end? Is there some maximum number of possible moves in a game? If so, what does it depend upon, can it depend upon only n?
(2) Must a game last a certain number of moves before it finishes?
(3) Is there a formula for the number of moves in a game of Sprouts? What does it depend upon?
(4) Who has a winning strategy for 1-spot sprouts? For 2-spot sprouts? For 3? Is there an explicit formula for n-spot sprouts?
(5) Looking at a finished game of Sprouts, can you tell how many spots it began with? Can you tell who won?

Some of these questions have good answers. In particular, let's look at the first one first.

3. Bounding the Number of Moves from Above

Note that if we can find a formula for the maximum number of moves, then we will have proven that (with finite starting conditions) the game must terminate in a winner. We'll define terms as we go, but here are a couple of useful ones:

- n denotes the number of spots at the beginning of the game of Sprouts. We call the game an n-spot game.
- m denotes the total number of moves.
- $\mathcal{L}$ denotes the total number of "liberties" in the game. As every spot can have a maximum of 3 lines meeting it, it is "free" to use if it have less than 3 lines meeting it at any time. Thus we say the liberties of a spot (denoted by ℓ) are 3 minus the number of lines currently meeting the spot. The total liberties $\mathcal{L}$ is the sum of the liberties for all spots. As an example:

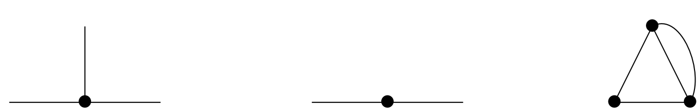

$\ell = 0$ $\qquad$ $\ell = 1$ $\qquad$ $\mathcal{L} = 1$

- s denotes the number of survivors at the end of the game. This is the number of spots that have exactly one liberty when the game finishes.

Now, back to the question of a bound on the number of moves m in terms of the number of starting spots n. We'll prove this in the following manner: first we'll show that $\mathcal{L} = 3n - m$. That is, at any point in time, the number of total liberties is equal to 3 times the original number of spots, minus the number of moves that have already occured.

This is probably more clearly stated as follows: $\mathcal{L}(m) = 3n - m$. I am using this "functional" notation as the liberties is a function of the number of moves, in particular, it changes as the game progresses.

Next we'll prove that $\mathcal{L} \geq 1$ during the game. In fact, we'll show that $\mathcal{L}$ decreases steadily until settling upon a terminal values. This "terminal value" must be greater than or equal to 1, and is the number of "survivor spots" s left at the game. At this point, the game finishes.

Finally, we'll put these two facts together to understand and prove the following Theorem:

Theorem 3.1. *A game of Sprouts starting with n spots, has a finite maximal number of moves. In particular, this is*

$$m \leq 3n - 1.$$

We'll begin with the followig lemma:

Lemma 3.2. *After m moves, the liberties can be expressed as:*

$$\mathcal{L}(m) = 3n - m.$$

The proof of this relies upon the technique of induction, so we'll review it here. It is a "staricase argument". The "steps" of the staircase are the "rounds" of the game; thus they are indexed by m. Our goal is to be able climb to the top of the staircase. Being on any one step of the staircase means the the statement $\mathcal{L}(m) = 3n - m$ is true for that particular value of m.

So to prove that we can get to any step, we do two things

(1) We show that we can get onto the bottom step $m = 0$ (this is called the base case).
(2) Then we show that if we happened to be able to get onto the kth step, we can get from there to the $(k+1)$st step.

Then for $k = 0$ the statement is true by (1). *AND* by (2) we can get to $k = 1$. Then we plug in $k = 1$ and get $k = 2$, and then $k = 3$, $k = 4$, etc. So we are done!

Now, let's use this to prove our Lemma:

Proof. We will prove this by induction on the number of moves m. We need to show that this is true if the game has just begun, i.e. if $m = 0$. Then there are no lines drawn, so by definition, the number of liberties is 3 times the number of the spots at the beginning, thus $\mathcal{L}(o) = 3n$ and the base case is true!

Now assume that $\mathcal{L}(k) = 3n - k$ for some k. We want to compute $\mathcal{L}(k+1)$ and see if it is $3n - k - 1$. Well, let's compute it in terms of $\mathcal{L}(k) + \Delta\mathcal{L}$ where $\Delta\mathcal{L}$ is the change in the liberties in one move. Well, one move always connects two spots and adds a new one in the middle. So the two original spots each have a decrease in ℓ of 1, so together a decrease of 2. But the new spot we added has only 1 incident line, so it *adds* one liberty to the total. Thus $\Delta\mathcal{L} = -1$. Thus we have

$$\begin{aligned} \mathcal{L}(k+1) &= \mathcal{L}(k) + \Delta\mathcal{L} \\ &= \mathcal{L}(k) + \Delta\mathcal{L} \\ &= \mathcal{L}(k) - 1 \\ &\stackrel{\checkmark}{=} 3n - k - 1. \end{aligned}$$

So by our previous description of proof by induction, we are done!

□

Now, on to the second Lemma, that is:

Lemma 3.3. *The number of liberties $\mathcal{L}$, at any stage in the game, is always bounded from below by* 1. *That is*

$$\mathcal{L}(m) \geq 1 \qquad \text{for all } m.$$

Further, it decreases monotonically from $\mathcal{L}(0) = 3n$ to $\mathcal{L} = s$, the number of survivors at the end of the game.

Proof. At the beginning of the game, $\mathcal{L} = 3n$ as noted above. Since $n \geq 1$, $3n \geq 3$, which satisfies the hypotheses. At each stage, $\Delta\mathcal{L} = -1$, so it decreases steadily at each step. Can $\mathcal{L} = 0$? This would mean that right before $\mathcal{L} = 0$, $\mathcal{L} = 1$. In order to continue one more move, we first need to connect two dots by a line. We can *only* do this if both dots have a liberty, but then $\mathcal{L} \geq 2$. Thus the game ends no later than when $\mathcal{L} = 1$. So at all points of the game $\mathcal{L} \geq 1$. □

Finally, we put these two together to prove the theorem above:

Proof. We know that $\mathcal{L}(m) = 3n - m$, that is $m = 3n - \mathcal{L}(m)$. And we know that $\mathcal{L}(m) \geq 1$. So putting this goether, m is at most $3n-1$ when $\mathcal{L}(m) = 1$ (hypothetically). Thus we get the desired result:

$$m \leq 3n - 1.$$

□

As a corollary, we get one expression for the number of moves in a game of sprouts. First we note the following proposition:

Proposition 3.4. *If we let $\mathcal{L}_f$ be the number of liberties at the end of a game of Sprouts, then*

$$\mathcal{L}_f = s.$$

Proof. It suffices to prove that at the end of the game, every survivor spot has exactly 1 liberty and every other spot has no liberties. But this is just the definition of a survivor and the fact that if a spot had 2 liberties, then we could draw a line from that spot to itself and complete another move! □

So we get the following exact formula for the total number of moves as $m = 3n - s$.

From this, we can figure out who wins the game, as an even number of moves is a win for Player 2, and an odd number of moves is a win for Player 1. So, putting this together into one important theorem, we see:

Theorem 3.5. *Let m be the total number of moves, n the number of starting spots, and s the number of survivors at the end of the game. Then,*

(i) The game ends after m moves, where $m = 3n - s$.
(ii) Player 1 *wins if $n + s$ is odd. Otherwise player* 2 *wins.*

4. Bounding the Number of Moves from Below

Now we'll turn our attention ot bounding the number of moves from below. That is, the game must be *at least* this long. This will also give us some more clues about how to play well, even if you can't figure out an optimal strategy from the beginning because there are too many spots.

First we need to define some more terminology:

(1) g denotes the number of "guard spots". A guard spot is one of the two closest spots to a survivor spot at the end of the game. This can be in one of two ways:

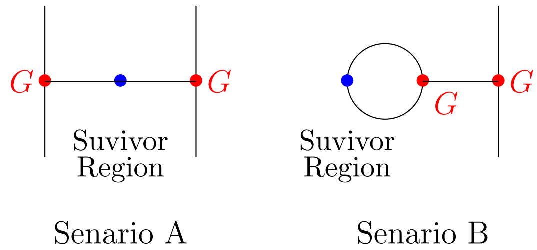

In both cases, though, by virtue of being the nearest neighbors, the guards are dead (or else we could connect a guard to a survivor and do another move).

(2) p denotes the number of "pharisees". A pharisee is a spot that is dead but not a guard spot.

Using this we will investigate a minimal bound for the number of moves. As in the maximal bound, this end up giving us another explicit formula for the number of moves m in terms of different parameters of the game: this time n the number of starting spots, and p the number of pharisees. In particular, first we prove the following theorem:

Theorem 4.1. *Let n denote the original number of spots, m the number of moves at the end of the game, and p the number of pharisees. Then a game of n-spot sprouts lasts exactly*

$$m = 2n + p/4$$

moves.

Proof. First let's get to know the guards a little better. We know that the guard spots are dead. But further, the survivor must have "access" to at least two of the "regions/sites" of each guard. Each dead spot has three regions:

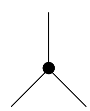

and each survivor spot has 2 regions. It must be that at least two of the survivor's regions are the same as the guard's region.

But, this means that each guard must guard only 1 survivor spot, as two survivor spots cannot have the same region, or we could connect them!

So, the number of guards at the end of the game is equal to 2 times the number of survivors, i.e. $2s$. Now let's list the facts that we know:

(1) We know that the total number of spots at the end of the game is $n+m$, because we started with n and then we add one every move of the m total moves.
(2) We know that every spot is either a survivor, a guard, or a pharisee
(3) We know that there are $2s$ guards at the end of the game
(4) We know that the number of survivors is given by $s = 3n - m$, from the theorem above.

So let's put all of this together to get a formula for the number of pharisees at the end of the game:

$$\text{number of pharisees} = \text{total number of spots} - \text{number of survivors} - \text{number of guards}$$

$$\begin{aligned} p &= n + m - s - 2s \\ &= n + m - 3s \\ &= n + m - 3(3n - m) \\ &= n + m - 9n + 3m \\ &= -8n + 4m \end{aligned}$$

Now we just rearrange a bit:

$$p = 4m - 8n \Rightarrow 4m = 8n + p$$

$$\Rightarrow m = 2n + p/4,$$

which is exactly what we wanted to prove! □

Now, we just not that there cannot be *negative* pharisees, thus the total number of move m must be bounded by the possibility of $p = 0$, i.e.

Theorem 4.2. *In a game of n-spot Sprouts, the minimum number of moves m in a game is given by*

$$m \geq 2n.$$

5. The Euler Characteristic of a Polyhedron

In this talk we will work almost exclusively with the Euler characteristic of a polyhedron. A polyhedron is a solid in three dimensions that consists of straight edges and faces. Two distinct faces intersect in an edge and two distinct edges intersect in a vertex.

Definition 5.1 (The Euler Characteristic)**.** Let P be a polyhedron with V vertices, E edges, and F faces. Then we define the Euler characteristic to be

$$\chi_E(P) = V - E + F.$$

Notice that the Euler characteristic is the alternating sum of the "building blocks" of our polyhedron in each dimension: the vertices are dimension 0 and get a $+$ sign, the edges are dimension 1 and get a $-$ sign, and the faces are dimension 2 and get a $+$ sign.

The Euler characteristic is best understood by example.

Example 5.2 (Cube)**.** Let C be the cube in 3-dimesions:

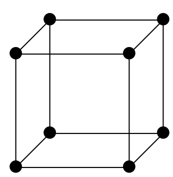

There are 8 dots in the above diagram of the cube, so $V = 8$. Simiarly, $E = 12$ and $F = 6$. So

$$\begin{aligned}\chi_E(C) &= V - E + F \\ &= 8 - 12 + 6 \\ &= 2.\end{aligned}$$

Example 5.3 (Tetrahedron)**.** Let T be the tetrahedron in 3-dimensions:

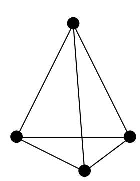

Here $V = 4, E = 6, F = 4$, so

$$\chi_E(T) = 4 - 6 + 4 = 2.$$

From these examples it seems that the Euler characteristic of polyhedra also satisfy the seemingly mysterious relation the $\chi_E = 2$ seen for connected planer graphs. However, the above polyhedra have a very special property: they can be drawn *on the surface of a sphere*. Imagine that one of the above polyhedra were made of an elastic rubber material like a balloon. Then you could image "blowing up" the polyhedron through an imaginary hole at one of the vertices, and the figure would lie on the face of a sphere!

From here it is not hard to see why the Euler characteristic of these examples must be 2, given the theorem for planer graphs. Given a graph on the circle, we could puncture a tiny hole in one of the faces of the graph, grab the sphere through that whole, and "splay" it out so that it lies flat, so that the boundary of our tiny puncture becomes the surrounding boundary of the... **connected planer graph**!

There are, however, polyhedra that cannot be drawn on the surface of a sphere. Consider the following examples.

Example 5.4. Let D be the cube with a central shaft removed:

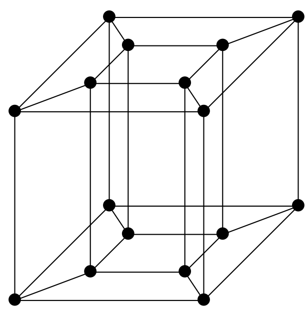

In this case, $V = 16, E = 32, F = 16$. So in total

$$\chi_E(D) = 16 - 32 + 16 = \mathbf{0} \neq 2!$$

If we imagine doing the same procedure of "blowing up" to the polyhedron D above, we would end up with something that looked like a donut, called a *torus*:

(insert image of a torus)

Clearly, for this polyhedron D lying on the surface of a torus, the Euler characteristic is 0. But is this true of all polyhedra that can be drawn on the surface of a torus?

6. A Topological Invariant

In order to understand if it should be true that the Euler characteristic "reads" the information of what kind of surface a certain polyhedron can be drawn on, we must understand what kinds of changes to a graph leave the Euler characteristic unchanged.

Lemma 6.1. *The Euler characteristic of a polyhedron does not change when we subdivide a face or an edge.*

Proof. As an example, consider the following triangle that may be part of the underlying graph of a polyhedron:

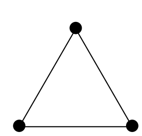

Now, if we subdivide the face of triangle, this involves adding a single vertex, and then 3 edges to connect this vertex to the other 3 vertices of the trianglular face:

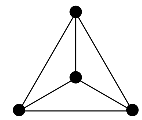

In general, a face might be an n-gon. We still add a single vertex, so the number of vertices V changes by $\Delta V = 1$. This vertex must connect to all outer vertices on the n-gon, of which there are n. This means that E increases by n, so $\Delta E = n$. These new edges divide the face into n compartments (subfaces) where there used to just be 1. Thus F has increased by $\Delta F = n + 1$. So in total the Eular characteristic changes by

$$\begin{aligned}\Delta\chi_E &= \Delta V - \Delta E + \Delta F \\ &= 1 - n + n - 1 \\ &= 0,\end{aligned}$$

so the process of subdividing a face does not change the global Euler characteristic.

Similarly, if we "subdivide" an edge we introduce one new vertex in the middle of the edge, so $\Delta V = 1$. We also have divided one edge into two, so $\Delta E = 1$. Thus in total

$$\Delta\chi_E = \Delta V - \Delta E + \Delta F = 1 - 1 + 0 = 0,$$

and again this process does not affect the global Euler characteristic.

□

Lemma 6.2. *The Euler characteristic of a polyhedron does not change when we delete an edge bordering distinct faces, or a vertex bordering 2 distinct edges.*

Remark. It is necessary to have these qualifications on to which edges and vertices the lemma applies, in order to disqualify the following examples:

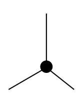

vertex meets more than 2 edges

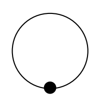

loop around a single vertex

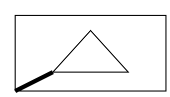

edge bounds same face

for which the vertex (respectively edge) is essential for the "integrity" and information inherent in the graph.

Proof. As in the previous lemma, we will argue by looking at the local changes to the Euler characteristic caused by our deletion. Upon deleting an edge, clearly $\Delta E = -1$. But, additionally, as long as the edge bordered 2 distinct faces, the number of faces decreases by 1 as well, so $\Delta F = -1$. So in total:

$$\Delta\chi_E = \Delta V - \Delta E + \Delta F = 0 + 1 - 1 = 0.$$

Similarly, if we delete a vertex bordering 2 distinct edges then $\Delta V = -1$, $\Delta E = -1$ and $\Delta F = 0$, so

$$\Delta\chi_E = \Delta V - \Delta E + \Delta F = -1 + 1 + 0 = 0.$$

□

This last lemma suggests that the underlying graph of a polyhedron drawn on some surface S can be "stripped down" to its essentials - that is vertices that do not border distinct edges and edges that do not border distinct faces - without changing the Euler characteristic.

Using both of these lemmas, we can now answer our question posed originally,

Do all polyhedra that can be drawn on the same surface have the same Euler characteristic?

with the following theorem.

Theorem 6.3. *The Euler characteristic of a polyhedron depends only on the surface on which it can be drawn; in other words, it is constant for polyhedra that can be drawn on the same surface.*

Sketch of Proof. Any graph can be transformed into any other graph by subdividing faces and edges and deleting edges and vertices. By Lemmas 6.1 and 6.2 these do not affect the Euler characteristic.

We could use this theorem to give a definition of the Euler characteristic of a surface as the Euler characteristic of the class of polyhedra that can be drawn on its surface. There is a more sophisticated way to define the Euler characteristic of a topological space, but for the surfaces we are considering, these definitions agree.

Thus the Euler characteristic is something we might call a crude *topological invariant*. Knowing the Euler characteristic of a surface, or polyhedron tells you something about its shape, its topology, irrespective of stretching, skewing, "blowing up", or similar nondestructive deformations.

We have already seen that the Euler characteristic of a sphere is 2 and of a torus is 0. What about other surfaces? It turns out that for a donut with g holes, what we call a g-hole torus, the Euler characteristic has a very nice form.

Theorem 6.4. *Let M_g be a g-hole torus. Then we have*

$$\chi_E(M_g) = 2 - 2g.$$

Remark. Note that in the cases we already know, this theorem agrees with what we have calculated. When $g = 0$, the surface is just the sphere, and so

$$\chi_E(M_0) = 2 - 2 \cdot 0 = 2.$$

And when $g = 1$, the surface is the torus so

$$\chi_E(M_1) = 2 - 2 \cdot 1 = 0.$$

Department of Mathematics, MIT
E-mail address: ivogt@mit.edu