20200107

The Josephus Problem

The original Josephus Problem is based on a story told by Flavius Josephus, a Jewish/Roman historian of the first century of the Common Era[1]. He described a time when he was leading a group of trapped warriors who preferred death over capture by Roman soldiers. They decided to stand in a circle and kill every third man until only two were left. These two would then kill each other. According to the story (told by Josephus himself, without any corroborating evidence) he quickly placed himself and a friend so that they would be the last survivors, after which they surrendered to the Romans. Regardless of whether the story is true, it makes an interesting mathematical problem: how do you choose the correct position in the circle so that you survive the elimination process?

We looked at a simplified version: instead of eliminating every third person, we eliminate every second person, and we are only interested in the final survivor, not the final two survivors. We number the seats at the table 1, 2, 3 ... n, where n is the number of people in the group. We start the counting process with 1, so the first few eliminations are 2, 4, 6 etc. **We will use J(n) to represent the seat number of the person who survives.** (In class I used F(n) but I think J(n) is a better choice ... J for Josephus ... so I'm using J(n) here.)

Now this is not really a very practical problem – it's a bit unlikely that you will ever find yourself in a game of this sort. But it's a wonderful exploration of how we can mathematically attack a problem – we will encounter several old friends from discrete math along the way, so we can think of this as a bit of a review.

Example: Suppose there are 10 people at the table. The elimination sequence starts 2, 4, 6, 8, 10 ... at which point we are back to 1, with only persons 1, 3, 5, 7, 9 remaining. We continue the process, eliminating 3, 7, 1, 9 ... leaving 5 as the sole survivor. Thus J(10) = 5

At this point we might observe that $5 = \frac{10}{2}$, and we might guess that $J(n) = \frac{n}{2}$ ... at least for even values of n.

We can test this by trying another value. Let's compute J(8). With 8 people at the table the elimination sequence is 2, 4, 6, 8 (leaving 1, 3, 5, 7) followed by 3, 7, 5 and then 1 is the final survivor ... so J(8) = 1. This refutes the guess that $J(n) = \frac{n}{2}$ but it offers some new ideas. We might notice that J(10) and J(8) are both odd. A moment's thought reveals that J(n) **must** be odd, since all the even numbered people are eliminated on the first time around the table. We might also think about the fact that 8 is a power of 2, meaning that when we eliminate the first half of the people we are left with a smaller power of 2, and this is true each time we go around the table: the

1 https://en.wikipedia.org/wiki/Josephus

number of players starts at 8, then reduces to 4, then to 2, then to 1. Is it a coincidence that we ended up on player 1 as the survivor?

We can look at a few more powers of 2:

$2^0 = 1$ ... with only 1 player, they are automatically the winner, so $J(1) = 1$

$2^1 = 2$ ... with 2 players, we eliminate player 2 and player 1 wins, so $J(2) = 1$

$2^2 = 4$ ... with 4 players, the elimination sequence is 2, 4, 3 and player 1 survives so $J(4) = 1$

$2^3 = 8$ ... we already know $J(8) = 1$

$2^4 = 16$ ... let's think about this. After we go around the table once, we will have eliminated 8 people so 8 are left, and we are back to person 1. We already know that with 8 players, whoever we start with is the final winner ... so $J(16) = 1$

You should be able to see how we can construct a proof by induction that

$$J(2^i) = 1 \quad \forall i \geq 0$$

But what about situations where n is not a power of 2? We can work out some small cases by hand ... see the table on the next page

n	J(n)
1	1
2	1
3	3
4	1
5	3
6	5
7	7
8	1
9	3
10	5

(I think I got some of these wrong in class!)

We can go back to J(10) = 5 and see if we can dig any more insight out of that case.

With 10 people at the table, after we go around the table once, we have eliminated 5 players and we are back to player 1 ... so it's almost like we are starting over with just 5 people. The game has no memory. It doesn't matter at this point how many people we started with. And – **here is the key point** – the person who wins this reduced 5-player game is the winner the original 10-player game, since this 5-player game is part of the original game.

Who wins a 5-player game? The elimination sequence is 2, 4, 1, 5 so the winner is 3. How can we relate that to the winner of the 10-player game? We have to ask ourselves what player in the original 10-player game is in seat 3 when the game is reduced to 5 players. Consider this diagram:

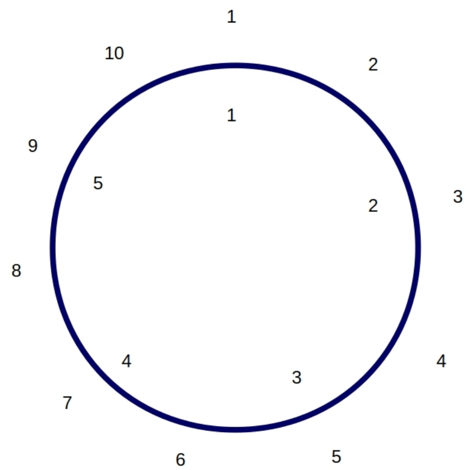

Outside the circle we have the 10 original players, numbered 1 to 10. Inside the circle we have the 5 players remaining after the first time around the table, re-numbered 1 to 5 since there are now only 5 people in the game. As we expect, the player numbered "3" in the 5-player game (the winner of this game) is the same player who is numbered "5" in the 10-player game.

Now we can figure out the relationship between the seat numbers in the original and the reduced game. It is a function that gives these values, where i is the seat number in the reduced 5-player game and $f(i)$ is the seat number in the original 10-player game

i	$f(i)$
1	1
2	3
3	5
4	7
5	9

It only takes a little experimentation to come up with $f(i) = 2 * i - 1$. So if the winner of the 5-player game is person x (that is $J(5) = x$), that person is in the chair numbered $2 * x - 1$ in the original 10-player game. So $J(10) = 2 * J(5) - 1$

Note: the figure shown above and our discussion is NOT a proof that this formula is correct. As an exercise, prove it!

Once we have proved that this analysis is correct for any situation where we start with an even number of players, we can write

$$J(n) = 2 * J\left(\frac{n}{2}\right) - 1 \qquad \forall \text{ even values of } n$$

It's useful to think a moment about what exactly is being said here. It means that when n is even, the winning seat number ($J(n)$) is found by taking the winning seat number for a game of size $\frac{n}{2}$, (that is, $J\left(\frac{n}{2}\right)$), doubling it and subtracting 1 .

<Quick check – does this fit with our earlier observation that $J(2^k) = 1 \quad \forall k \geq 0$? >

But what about when n is odd?

Example: Suppose there are 7 people at the table. The elimination sequence is 2, 4, 6, 1, 5, 3 ... leaving 7 as the sole survivor. Thus J(7) = 7

This is different from the case where n is even because when we go around the table and eliminate all the even numbers, the ***next*** number eliminated is 1 – so we can't say that we are reducing to a smaller game and starting the smaller game in the same location as we started the larger game (this was what made the analysis easy for even values of n). But we **are** reducing to a smaller game, and we do know where that game starts: it starts with the person who is sitting in chair 3. (When we eliminate 1, 2 is already gone so we will skip over 3. We can think of position 3 as the start of the smaller game, in which the chairs have original numbers 3, 5, 7, and "reduced" numbers 1, 2, 3. This figure shows what I mean:

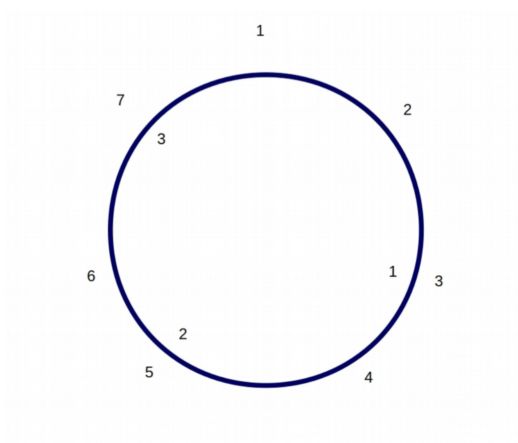

We need a function that provides the following mapping from the seat numbers in the reduced game (inside the circle) to the seat numbers in the original game (outside the circle)

i	$f(i)$
1	3
2	5
3	7

Once again it is not hard to come up with a function that works: $f(i) = 2 * i + 1$ But here again we must remember that there are infinitely many functions that fit these three points. For completeness, we would need to ***prove*** we have the right one.

But let's skip over the proof for now, and accept that this function is correct. This means we can use it to write J(7) in terms of J(3): $J(7) = 2 * J(3) + 1$

And what is the relation between 7 and 3? Well, 3 is the number of players in the reduced game, which we arrived at by going around the table once and then eliminating one more player. In terms of the number of people eliminated, this is the same as eliminating one player and then eliminating half of the remaining players ... so the number of players left in the reduced game when we start with an odd number n is exactly $\dfrac{n-1}{2}$ .... and yes, $3 = \dfrac{7-1}{2}$

So we generalize from n = 7 and arrive at

$$J(n) = 2 * J\left(\frac{n-1}{2}\right) + 1 \qquad \forall \text{ odd values of } n$$

Now we can put these together to get a recurrence relation:

$$J(1) = 1$$

$$J(n) = 2 * J(\frac{n}{2}) - 1 \qquad \forall \text{ even } n > 1$$

$$J(n) = 2 * J(\frac{n-1}{2}) + 1 \qquad \forall \text{ odd } n > 1$$

So, do we have a solution? NO! We have a hypothesis that matches all the examples we have looked at. To claim it is a solution, we have to prove it. The most natural way to prove this is by induction – **I encourage you to do this proof as an exercise**. You will discover that this recurrence is in fact correct.

Like most recurrence relations, this one "looks back" in the sense that it tells us how to compute J(n) for any value of n. We can turn it around and make it "look forward" so that we can see how to use a known value of J(n) to compute more values.

Turning the first recursive line around is easy. If we introduce a new variable $x = \frac{n}{2}$, then the line becomes $J(2 * x) = 2 * J(x) - 1 \qquad \forall \text{ even } 2 * x > 1$

and since $2 * x$ is even and > 1 if and only if $x \geq 1$, the constraint just becomes $\forall x \geq 1$

Another way to write this is

$$J(x) = k \quad \Rightarrow \quad J(2 * x) = 2 * k - 1 \quad \forall x \geq 1$$

and that is what I mean by "looking forward"

Now for the second recursive line. We introduce a new variable $x = \frac{n-1}{2}$, from which we get $n = 2 * x + 1$. The recursive line becomes

$$J(2 * x + 1) = 2 * J(x) + 1 \qquad \forall \text{ odd } 2 * x + 1 > 1$$

and since $2 * x + 1$ is odd and > 1 if and only if $x \geq 1$, the constraint just becomes $\forall x \geq 1$

We can rewrite this as $J(x) = k \quad \Rightarrow \quad J(2 * x + 1) = 2 * k + 1 \quad \forall x \geq 1$

Now the two lines have exactly the same restriction ($x \geq 1$) and the same antecedent ($J(x) = k$) so we can combine them:

$$J(x) = k \quad \Rightarrow \begin{cases} J(2 * x) = 2 * k - 1 \\ J(2 * x + 1) = 2 * k + 1 \end{cases} \qquad \forall\, x \geq 1$$

Starting with $J(1) = 1$, we can use this to compute J(x) for increasing values of x. J(1) gives us J(2) and J(3), J(2) gives us J(4) and J(5), etc.

n	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
$J(n)$	1	1	3	1	3	5	7	1	3	5	7	9	11	13	15	1

And now, after all this work, we see one of the most beautiful and unexpected patterns in elementary mathematics. The sequence of Josephus Numbers consists of all the positive odd integers less than 2, then all the positive odd integers less than 4, then all the positive odd integers less than 8, and so on forever.

That's all well and good, but if you're trapped in a cave with 1242 desperadoes who are bent on playing a death-game (incidentally, the Josephus elimination process is sometimes called "Roman Roulette") you may not have time to extend the table out to n = 1243 to figure out which chair you should choose. Is there an easier way to compute J(1243) ?

It turns out there is, and it ties together two ideas we have seen so far:

1. The game has no memory – after each person is eliminated, it is as if the game starts over again with the person next to the one just eliminated.
2. When the number of players is a power of 2, the first player wins.

So here's a way to think about the game: eliminate people until the number of remaining players is a power of 2. The next player is effectively Player 1 in a game with 2^m players for some value of m – which means they will win!

Let n be the number of people in the group. We can write n uniquely as

$n = 2^m + k$ where $0 \leq k < m$

Basically, k is the number by which n exceeds the largest power of 2 that is $\leq$ n

If we write n in binary notation, it will look something like this:

$$1001110101$$

Written this way it is easy to see that the first 1 represents 2^m and the rest of the bits represent k.

When the game starts, after k people are eliminated, 2^m players remain, and the very next player is the winner. The people who are eliminated are in seats $2, 4, 6, ..., 2 * k$, so the next player is in seat $2 * k + 1$. (Note: we can be sure that $2 * k + 1 \leq n$ .... why?)

Therefore $J(n) = 2 * k + 1$

It's just that easy. In a group of 1243? Easy – the largest power of 2 that is $\leq$ 1243 is 1024. Thus

$k = 1243 - 1024 = 219$

The formula gives

$2 * 219 + 1 = 439$

Therefore $J(1243) = 439$ Make your way quickly to Chair 439 and hope there are no other discrete mathematicians in the group.

And now at last we are happy. We have found a simple and easy formula for computing J(n) for any positive integer m.

But are we finished? Of course not! We left quite a few bits unproved along the way – they all need to be properly proved. And remember we did not solve the original problem, which involved eliminating every ***third*** person, not every second person. So here's the completely general problem for your consideration:

Let n and k be positive integers. Suppose we have n people sitting around the table, and we eliminate them by skipping over k-1 people and eliminating the next one, continuing until there is only one person left. Let J(n,k) be the seat number of the person who survives. How can we compute J(n,k)?

What is the take-away from this? Again, it is all about looking for patterns and building solutions out of the patterns we find. Plus, now you know where to sit next time someone suggests playing Roman Roulette to decide who gets the last slice of pizza.

Addendum:

When I first learned about the Josephus problem (not first-hand, I am not quite that old) a different approach was taken to defining a recurrence relation. I think the approach outlined above is superior, but I'll give you this one to ponder at your leisure.

Let's focus on the simple form of the problem where we eliminate every second person. Instead of numbering the chairs 1, 2, ..., n we will number them 0, 1, 2, ..., n-1

So we start by skipping person 0 and eliminating person 1. As we have seen, after each elimination the game basically starts again, just with n-1 people at the table. If we renumber the people at this point, everyone has their original number reduced by 2: person 2 becomes person 0, person 3 becomes person 1, etc. The sole exception to this is person 0, who becomes person n-2. Now suppose we know which number wins when there are n-1 people at the table. The original number for that person is just their "reduced" number + 2. The only exception is person n-2, whose original number is 0. But we can take care of this by applying a "mod n" operation.

This gives the recurrence:

$J(1) = 0$ (remember, now the first seat is seat 0)

$$J(n) = (J(n-1) + 2) \mod n \quad \forall n > 1$$

Check it out:

$$J(2) = (J(1) + 2) \mod 2 = (0 + 2) \mod 2 = 0$$

$$J(3) = (J(2) + 2) \mod 3 = (0 + 2) \mod 3 = 2$$

$$J(4) = (J(3) + 2) \mod 4 = (2 + 2) \mod 4 = 0$$

$$J(5) = (J(4) + 2) \mod 5 = (0 + 2) \mod 5 = 2$$

$$J(6) = (J(5) + 2) \mod 6 = (2 + 2) \mod 6 = 4$$

$$J(7) = (J(6) + 2) \mod 7 = (4 + 2) \mod 7 = 6$$

$$J(8) = (J(7) + 2) \mod 8 = (6 + 2) \mod 8 = 0$$

It looks correct – can you prove it? And does it lead to a different closed form solution?