The Mathematical Vision of Maryam Mirzakhani

Professor Holly Krieger

24th May 2023

We're here today to discuss the work of one of mathematics' brightest modern stars, Maryam Mirzakhani. Maryam was an Iranian mathematician whose deep and creative mathematical work received the highest accolades in mathematics, and a warm and enthusiastic woman whose life and work were cut short at just age 40, when she succumbed to breast cancer. What I'd like to do today is to communicate to you some of Maryam's work, not just the mathematical content but also a sense of why mathematicians find it so impressive and compelling. Really, I'd like to share with you all the appreciation I feel when I think about her mathematics, and to give you a sense of the excitement and enthusiasm that she felt while working - to pass on to you some part of her legacy, as best I can.

Maryam's work concerned the connection between two fields. The first, geometry, is familiar to use from the very early school years - even in nursery, children learn "circle", "triangle", "bigger", and so on. To a mathematician, geometry is not just the study of shape, which we more commonly call "topology", but the study of how shapes interact with measurements of size, distance, or curvature. Geometry has innumerable applications outside mathematics, to physics, art, architecture, biology, and so on, but it also is deeply connected to many other fields of mathematics, like number theory and algebra.

The second specialist field that Maryam's work involved is the field of dynamics, particularly ergodic theory. A dynamical system is one that evolves according to some fixed rules, usually pretty simple rules - think of the basic equations of motion governing the movement of planets in our solar system, for example. In this talk there will be some very simple dynamical system rules, like "walk in a straight line", and some more complicated ones, like "walk to another mathematical object that looks like where you started".

Maryam's work focused on a family of mathematical objects known as Riemann surfaces and used dynamics to understand the surfaces themselves as well as the some more meta-data things like the collection of all such surfaces. We'll take a little tour of two of her most important contributions in that direction, as well as a warmup problem from her teenage years. But beyond the mathematical content, what I really want is to help you understand *why* her work was recognized by the highest honours in mathematics, and to see how - though her mathematical career was cut so unfairly short - her contributions will resonate through mathematics in the future.

Maryam was born in Tehran, Iran in 1977, just before the Cultural Revolution and Islamic reformation of Iranian academia. There was enough educational stability by the time she entered middle school that she was able to attend a school for talented girls. By 7th grade it was clear she had a standout talent in mathematics. With the support of her school principal, Maryam and her friend Roya were the first women permitted to participate in the preparatory national competition and join the national team for the International Mathematical Olympiad.

The IMO is a competition in problem solving for pre-university students that draws immense mathematical talent from all over the world. Maryam scored 1 point from a perfect score in 1994, and a perfect score in 1995, a rare accomplishment. Maryam was one of two women on the 1994 Iranian team, and the only woman on the 1995 team, but it is worth noting that being a woman accomplished in mathematics and science in Iran at the time was not in itself so unusual - in a 2008 interview, she mentioned that she often had to explain in the US that in Iran, a talented woman had access to high quality education.

As an example of that, I'd like to discuss as a mathematical warm-up a problem that Maryam thought about as a teenager - indeed, the results became her first published paper. To zoom a bit on this quote from Ramin

Takloo-Bighash: Maryam was attending a summer school for gifted children at the university, and one of the topics he had talked about was decomposing graphs into disjoint unions of cycles, including the rather curious example of decomposing a tripartite graph into a union of 5-cycles. This was considered the first difficult case of the general problem. Mahmoodian had asked the students to find examples of tripartite graphs that were decomposable as unions of 5-cycles, offering one dollar for each new example. By the time of the next lecture, Maryam had found an infinite family of examples. She had also found a number of necessary and sufficient conditions for the decomposability.

Let's discuss this problem a little bit, because it gives a wonderful glimpse into the sense of exploration and persistence that Maryam, indeed any mathematician, must use to approach a new problem. Graphs are models for a huge array of relations and processes - network analysis and neural networks, biological networks like protein-protein interactions, computational neuroscience and social science, etc. But all these fields with many syllables are using a simple idea: a visual, mathematical encoding of relationships.

Here is an example you might be familiar with: a competition that is structured as a round robin. For example, the group stage of the World Cup. I have here the groups from 2022, with a zoom on Group E: Spain, Costa Rica, Germany, and Japan. You might think Group B a better choice for this talk, but as an American I didn't want to raise any emotions to distract from the mathematics. A graph can encode the relationship of "playing a match" - here, I think of each team as a point of the graph, and I connect the points if the teams need to play a match. In round robin, every one of the teams in a given group plays each other, so each pair of teams is connected by an edge.

This property - that every pair of points is connected by an edge - is known as being a *complete* graph. One of the key tools in graph theory is the idea of graph decomposition. Just like we study the whole numbers by decomposing them into pieces - that is, their factors, particularly the prime numbers - we can study graphs by decomposing them into pieces to make them easy to understand.

The group stage round robin graph decomposes into 8 identical graphs: each of these has four points, with every pair connected by an edge. This decomposition helps, for example, if we want to figure out how to schedule these matches - with no team playing twice in the same day, a total of 48 matches with 12 days to play them, and so on.

This can be encoded in the language of graph theory via an edge colouring, where the colour of an edge indicates the day of the match.

The decomposition allows us to easily schedule according to these rules, propagating the colouring on the small graph to the larger graph from which it is composed. This basic example gives us a hint of the utility of graph decomposition.

Graph theory slide 4: Mahmoodian asked a group of talented high schoolers in a summer graph theory course - can you come up with examples of complete, tripartite graphs which can be decomposed into 5 cycles? We have already seen how what "complete" means - "complete tripartite" is a little different. A graph is complete tripartite if its vertices can be split into three groups, and we have an edge between any two points in *different* groups. Here I have given the groups three colours - red group, green group, blue group - to help visualize how these graphs are complete tripartite. Any pair with vertices of different colour have an edge between them, and no edges exist between vertices of the same colour. A 5-cycle is just what it sounds like - a cycle, that is, a closed path on the graph - with 5 points in it.

Now, how would you approach such a question? First you'd try to construct some complete tripartite graphs, and look for 5-cycles inside them, ideally a decomposition into 5 cycles. I've got for you here some of the complete tripartite graphs that aren't too big - notice, this type of graph is completely determined by the number of red, green, and blue points it contains, since "complete tripartite" tells you precisely how to join them up. Right away we see - maybe even without doing any experiments - that a 5-cycle decomposition can only happen if the total number of edges is divisible by 5, since a decomposition must include each edge exactly once. This is what is called "necessary" condition - it is satisfied by every complete tripartite graph which has a 5-cycle decomposition.

So here I've got the two graphs in my list which have edge number divisible by 5 - K1,1,2 has 5 edges, and K1,3,3 has 15 edges. But there's an impediment to trying to decompose K1,1,2 into a 5-cycle - if you start at a vertex, say maybe the green one, then you've got to return to it at the end of the cycle. But the green vertex has three edges coming out of it, not two - how could we ever return to the green vertex at the end of the 5 cycle? Either we visit it once again and there's no edges left to use, or we don't, and we have one leftover edge.

So this is a complete tripartite graph with edge number divisible by 5 which nonetheless cannot be decomposed into 5-cycles. It turns out that with a little effort, the next one K1,3,3 *can* be decomposed into 5-cycles. But we don't want to go one-by-one here! Can we figure out what property prevented us from decomposing K1,1,2 that can be generalized to other complete tripartite graphs? The issue was the vertex (green or red) with *three* edges - since a cycle always includes "entering" and "leaving" any given vertex, any one vertex must have an *even* number of edges connected to it.

These kind of experiments and deductions allowed Maryam to make progress - as a teenager! - on the question of 5-cycle decompositions for this graph, and even more impressively, construct not just one or two but an infinite family of examples that allow a decomposition. While this problem is elementary and different in topic than her more sophisticated mathematical work, notice how different this is than the type of mathematics you encounter in school. There's very little memorization, no technique that you already know to solve the problem. Mathematics is not often portrayed as an experimental science, but in this playful sense, it really is. As a professional mathematician, Maryam had impressive technical expertise thanks to tenacious work and talent, but you can already see in her youth her ability to experiment and visualize mathematics, the key tools necessary for truly novel vision and progress.

OK, now that you're warmed up with mathematical experimentation and visualization, let's move on to discuss the more technical accomplishments of her academic mathematical career.

Maryam's work focused on Riemann surfaces. These are geometric objects which, if you look closely, resemble the basic plane of complex numbers. I really mean "complex" numbers here, not just the Euclidean plane - the difference is that I require a notion of angles between curves that is consistent no matter where I'm zooming in. You're very familiar with one Riemann surface, which we're all standing (or sitting) on right now - idealized as a sphere, the surface of the Earth (or any hollow sphere) is a Riemann surface, and we're familiar with this idea of "zooming" to see a plane, since we often use flat 2-dimensional maps to describe small areas of the surface of the Earth.

We can describe the possible shapes of Riemann surfaces very easily, so long as they don't have any boundaries - they are completely characterized by the number of handles that they have, known as the genus of the curve. If you've ever heard the joke that a topologist can't tell a coffee cup from a donut, this is - if all you care about it is the shape, if you allow yourself to stretch and squish a Riemann surface, then you can always squish it into one of these, just like a coffee mug could be squished down to a donut shape. So for each number 0, 1, 2, 3, 4, and so on, we have a surface of that genus - a surface with that many handles.

Once we start to care about finer properties like size, shape, and curvature, it turns out that Riemann surfaces admit some notion of geometry, albeit with restrictions. When I say "geometry", I mean a way to describe size, angles, and curvature, in a way that is intrinsic to the surface. Now here is where it starts to get tricky. For the sphere, the surface as we see it here in 3-dimensional space is a good representation of the geometry - we see something which is curved outwards, we could use a string to measure distances or angles, and so on. But in genus larger than 0 this is not the case, and the geometries we can put on the surface don't have the same features as what we see when we look at it in 3-dimensional space.

Let's take curvature as an example. The sphere has positive curvature - it is known as a spherical or elliptic geometry. If you ever learned about non-Euclidean geometry you might recall that a feature of spherical geometry is that triangles have angles which add up to more than 180 degrees.

The donut here - known as a torus - looks like it is positively curved on the outside, and negatively curved on the inside. But that is just how it looks in our 3-dimensional world! In reality the geometry carried by a genus 1 surface is a "flat" geometry, which means that it looks like the usual Euclidean plane, and the angles of a triangle add up to 180 just as we were taught in our school geometry classes.

Maryam's work focuses on the surfaces of genus at least 2 - everything except the sphere and the torus. Just as with the torus, looks are deceiving here - although we can't see it when we look at the shape with Euclidean eyes, the geometry on a surface of genus 2 or higher is hyperbolic, meaning negatively curved, everywhere. You should think of the inside of the handles as doing a pretty good job of visualizing the actual curvature for the geometry on the hyperbolic surface - triangles have angles which add up to less than 180.

If this all makes your head hurt, rest assured that the same is true even for professionals in the field! For this reason, mathematicians often try to work with other visualizations of hyperbolic Riemann surfaces - in fact, coming up with and studying alternative, useful visualizations for hyperbolic Riemann surfaces is an important piece of geometry on its own. While these pictures don't do a good job helping us visualize curvature, they don't do too badly with distance - each type of visualization has its strengths and weaknesses.

I need to make an admission - I will sometimes accidentally use the word “curve” when I’m talking about a Riemann surface. This is the difference between thinking of them like a mathematician in a world of the real numbers (where they are surfaces) vs a world of complex numbers (where they are curves). I know it is confusing!

In any case, one of these hyperbolic curves admits many different geometries - perhaps the handles are thin in one geometry and thick in another, for example. How can we understand these geometries? The answer is the same as it was for understanding whole numbers, or for understanding graphs - we try to decompose them into simpler versions. Here you might try to take your two-handled surface and chop it into two tori, for example.

But if we do that, we introduce a boundary. Can we instead study Riemann surfaces with boundaries allowed?

The answer is yes - there is a similar “stretch and squish” result that tells us that the shape of a Riemann surface with a finite collection of boundary curves can just be described by the number of boundary curves. Here we have a famous surface with boundary known as a “pair of pants” - if you imagine inflating one leg and deflating the other, this is the same as a sphere with three little missing circles. It turns out that these are the building blocks for all Riemann surfaces with boundaries, but you have to be really careful - if you want to decompose a Riemann surface into pairs of pants in a way that respects the *geometry*, then you’d better make your cuts along straight lines. A straight line cut on a surface is what is known as a closed geodesic.

A closed geodesic is a walk on a surface that has no acceleration - no turns, no speeding up, no slowing down. A key point is that this is all relative to the surface’s geometry! If I asked you to go take a straight walk right now, you wouldn’t walk off the edge of the earth into space - your straight-line walk would be relative to the surface of the earth, and eventually (if you were persistent, and had a really good boat) form a great circle on the earth. That great circle would be a geodesic on the earth.

Other Riemann surfaces have geodesics too, but they get a lot more interesting as we add more handles to our surface. Here is an example of a lady bird walking a geodesic on a torus. To get a better glimpse, let’s see a movie.

You see that a geodesic from the torus’ perspective has nothing to do with a straight line in this 3-d world - sometimes the lady bird walks towards us, sometimes away, etc. But from the lady bird’s perspective, it is always walking straight ahead at a constant speed. This is a geodesic. Often geodesics are described as “minimal length” paths on a surface, which is true in a slightly tricky technical sense.

Thanks to the pair of pants decompositions (and more, of course), the study of geodesics on Riemann surfaces is a fundamental part of hyperbolic geometry, and it’s the part of Maryam’s PhD work that I can most easily describe to you. Maryam completed her PhD at Harvard in 2004 under the supervision of Curt McMullen, with what most any mathematician would consider a much more notable accomplishment than her stellar IMO performance - a thesis composed of three papers that appeared in the top three mathematical journals.

So what did Maryam discover about these geodesics - straight line paths - on hyperbolic Riemann surfaces? I claimed that pairs of pants decompositions, and correspondingly, geodesics, form a basic class for the study of these surfaces in the same way that prime numbers form a basic class for the study of the integers. It turns out that there is a beautiful analogy between the two that comes from *counting*. One of the reasons prime numbers are useful is because there are infinitely many of them, but nonetheless if you are clever enough, you can count them by imposing a length bound. For example, you might ask: how many prime numbers have L digits, where L is some number you choose? If L is small, we can just write them down to compute the answer. More interestingly, it is possible to figure out approximately what the answer should be for any* value of the length! I’ll modify the question a little bit to make the numbers nicer, and ask how many primes natural logarithm have less than L. A mathematical result known as the Prime Number Theorem tells us how the answer should grow as L gets large: very quickly, like this function e^L/L. Perhaps the most interesting interpretation here is probabilistic - since there are e^L total numbers satisfying the length bound, the probability that a number of length L is prime is about 1/L - that is, prime numbers get harder to find as numbers get longer.

Even though a hyperbolic surface has infinitely many closed geodesics, they can still be counted if we impose a length bound: how many geodesics of length less than L are on the surface?

It turns out that geodesics on a hyperbolic surface satisfy a kind of Prime Geodesic Theorem - so long as you don’t overcount (“primitive” geodesics only), the total number of geodesics of length less than L grows like e^L/L, this exponential growth…

which is precisely the same as the growth rate for prime numbers of length less than L. The amazing thing about this theorem for geodesics is that it doesn't care which hyperbolic surface you started with! Even the genus isn't important.

Maryam considered a refinement of this counting which is much trickier. We call a geodesic "simple" if it never crosses itself. In a sense, these are even more basic building blocks, and they are much harder to find on hyperbolic surfaces.

Let me explain a bit why this is a tougher problem. Even though a torus is not hyperbolic, it's much easier to visualize what's going on there, so let's look at that example. I mentioned that mathematicians use different models to try and visualize the geometry on a Riemann surface; for a torus, one very useful such model is that of a "video game" square, a square with parallel sides identified. Some experiments with paper and tape - or a good squint at this diagram - should convince you that if you glue the top to the bottom, the right to the left, then a torus can be constructed by this kind of square. In fact, that square is a good model for visualizing geometry on the torus - it is where the one of the possible geometries on the torus comes from!

Just as in the hyperbolic case, we have geodesics on tori, which might be short and simple as the torus on the left, or long and not simple as the torus on the right.

In the square model, a geodesic is really a straight-line path, but we need to respect the gluing of the sides. So we could, for example, count geodesics through a point by counting lines in the video game square, which are determined by their slope.

In fact, we can remove the confusion of the identified sides entirely if we consider not just one but many copies of our square. Repeating the square over and over we see that our geodesic on the torus is the same thing as this blue line that connects two red dots in our "many copies" model.

If we want to think about all the possible geodesics at once, we need to consider infinitely many copies of the square - that is, a plane of repeating squares, with vertices at points with integer coordinates. In this infinite-copies model, our geodesic corresponds to a line segment from the origin to some other point with integer coordinates.

Our infinite squares model allows us to turn our geodesic-counting problem into a point- counting problem! Counting closed geodesics on the torus of length < L is akin to counting integer points in a disk of radius L (for the sake of the analogy, I'm forgetting the "primitive" requirement). Counting *simple* closed geodesics on the torus of length < L is equivalent to counting so-called primitive (confusing notation, ack!) points - points whose two integers are coprime i.e. have no common factors. This is not just analogy - this is really the counting argument for this geometry on the torus, and only a slight modification is needed for the other torus geometries.

Counting integer points in a disk is not too hard, since it's almost the same as computing the area of the disk - imagine identifying each unit square in the disk with its upper right corner, which should give us the right answer up to some error that grows like the perimeter. Counting coprime integer points in a disk is a more delicate question, as you can imagine! It is equivalent to knowing the probability that two numbers up to a given length are coprime, which turns out to be (independently of the length!) 6 divided by pi squared, or about 61%. But notice the key feature here - counting *simple* geodesics on the torus required an additional layer of understanding.

Now we're in a position to understand one of the significant results of Maryam's thesis work. She provided the first precise counts of the number of simple closed geodesics of length at most L on a hyperbolic surface, in terms of mathematical constants that we understand. Notice (as was already known at the time) that the growth rate is much smaller with the assumption of simplicity - the likelihood that a "random" long geodesic on a hyperbolic surface is simple is vanishingly small.

More striking than this result, though, were the means by which Maryam approached this problem, which was a technique that ran through her three early papers which comprised her thesis. The key idea is to consider not just a single geometry on your surface to count its geodesics, but rather to consider the counting problem across all possible geometries at once. This is very meta, so I am going to try to explain it with an analogy to counting integer points first.

Imagine again we want to count integer points, but perhaps inside shapes other than disks. I'm going to break my integer points into some (arbitrary) groups - I'll colour them green if one coordinate is odd and the other even, blue if both are odd, and red if both are even. The probability that an integer is even or odd doesn't depend on its size, so the density of green (or red, or blue) points is the same no matter where I look.

This fact means that if we change our disk to another shape *by a transformation which doesn't change point colour*, the density of green (or red, or blue) points inside the new shape will be the same as inside the disk. So all we really need to know to count green points in the ellipse is the area of the new shape.

Maryam proceeds along these lines, but instead of a parity grouping she divides geodesics into groups based on their so-called mapping class orbits. This means that she splits the geodesics into groups depending on what shape is obtained when we cut along the geodesic (just like in the pair of pants decomposition procedure!). Understanding the densities themselves and relating them to the geometry of the surface relates to her other impressive thesis work but is a bit too technical for this talk.

Here is a fun application of Maryam's result! Let's think about a genus 2 surface and call a geodesic separating if it splits the surface into two genus 1 (with boundary) surfaces - something we might want to do to construct a pair of pants decomposition! The question is: if I pick a geodesic at random, what is the likelihood that I chose one that is separating?

It follows from Maryam's theorem (which actually counts those groups) that the answer should be independent of the surface's geometry, asymptotically! In fact, no matter which geometry we put on a curve of genus 2, a randomly chosen long simple closed geodesic is about 48 times more likely to be non-separating than separating. The real mathematical machinery behind what Maryam did was her understanding of the so-called *moduli space* - a geometric object which is a map of all geometries we can put on aa surface of genus g with n punctures. Just like a map, each point represents something - but rather than a place, it represents a particular choice of geometry.

We study moduli spaces because the help us understand the objects they contain, they are often geometrically interesting themselves, and because they help us understand how to "deform" objects - what it means, for example, when two geometries are "close to" each other.

After finishing her PhD, Maryam accepted a prestigious Clay postdoctoral fellowship at Princeton University, and shortly after was appointed a Professor at Stanford University. She began a long-term and celebrated collaboration with Alex Eskin studying a different kind of moduli space which is a map of *translation surfaces*, a type of mathematical object we can describe just using polygons which is related to the study of billiard dynamics.

Here we see some translation surfaces related to the pentagon - I'll explain how these surfaces are described.

Suppose I want to make a shot in billiards. When the ball hits the side of the table, the angle it leaves the side is equal to the reflection of the incoming angle - just like if we pointed a laser at a mirror, the light would be reflected off the mirror at the reflected angle. If we want to study the trajectory of the ball, it all gets complicated rather quickly - here, we have two edge bounces, and we have to keep track of two reflection angles.

We can simplify the description of the ball's trajectory by *unfolding* the table - that is, placing a reflected copy of the table along the edge. Here we've unfolded in the direction of our first bounce. If we unfold the trajectory too, we see that our ball trajectory at the bounce point unfolds to a straight line.

Unfolding again by reflecting across the second bound, the entire trajectory unfolds to a straight line! So we can simplify the trajectory by allowing ourselves to unfold the table. Notice that we don't have to keep doing this process forever - if I reflected the lower left table to the left, I'd get another copy of my original table. So I might as well instead identify the lower left edge with the lower right edge, just like we did in the torus world.

Translation surfaces generalize this idea broadly. Roughly speaking, they are surfaces which are defined by gluing parallel edges of polygons with rational angles. Here's an example where we glue parallel edges of an octagon - and get a surface of genus 2!

So we have seen these surfaces before…in shape. The torus, for example, we've already seen is a gluing of the square along parallel edges, and since the four angles add up to 360 degrees, this flat geometry actually works on the entire torus. But now for a general translation surface, we will consider them with the flat geometry away from the corner points - a type of "pseudo-geometry" because it is not defined everywhere. Really, it's precisely the geometry coming from the flat polygon itself, together with information about how the corner angles glued up. Just like in the Riemann surface setting, there is a moduli space which is a map for all translation surface geometries that a surface of genus g can carry.

Since polygons are composed of line segments and parallel lines are glued to make a translation surface, any linear transformation of a polygon can give a new translation surface. The two translation surfaces will share some properties - their topological type, for example. But some things will change, like (for example) the length of the shortest or longest diagonal on the surface. The powerful thing is that these types of transformations give us a way to group together translation surfaces in moduli space, just like previously we had mapping class orbits to group together Riemann surfaces in *their* moduli space.

So, a basic question then is: what do these orbits look like? It's not a discrete set, in this setting, we can stretch or skew a polygon by as small an amount as we want. In general, describing orbits of dynamical systems is very hard! You might have heard of a "strange attractor" before, here's a picture, and if we zoom to the red box cross section…

…it turns out that the cross section is a complicated fractal set called a Cantor set. This is a feature of strange attractors, which are "strange" because of their features and not "strange" because they are rare.

So this question is a very difficult one. To give another analogy, suppose you have an infinite chess board, and you want to describe all possible moves a knight can make from a starting space. It's not so hard - every space can be reached. *But* suppose now that you don't know the shape of the chessboard - you take some of the squares in your infinite chessboard out, but you don't know which ones. How could you possible describe the possible sets of knight-reachable spaces for all possible chessboard shapes?

Nonetheless in a monumental work, this is what Maryam and Alex Eskin, then also with Amir Mohammadi, accomplished. They proved a strong *rigidity* result for these orbits - that they are not complicated and fractal but have relatively simple geometries. It is hard to describe in layman's terms the impact of this classification theory - the ideas in their proof open up the door to orbit classification questions in places we thought they'd never be accessible, and a number of applications have already been pursued.

Fortunately for us, one of those applications is very hands-on and beautiful, related to what's known as the illumination problem. The illumination problems asks whether light shining in a room with mirrored walls will light every point in the room - you can see on the right a beautiful installation of Yayoi Kusama that was at the Tate modern just a couple of years ago giving an idea of how such a mirrored room might fill up with illumination from interior points.

In the language of billiards, the question becomes: can a ball on the table be shot into a pocket placed anywhere on the table?

Of course for a rectangle the answer is yes - a rectangle is convex! The straight-line segment between any two points lies inside the rectangle, which gives you the shot to the pocket that you need. But what happens for more general and possibly non-convex polygons?

It turns out that there are polygons for which this is not always possible - so that there are two points which are not connected by a straight-line billiards shot.

Now that we've talked about unfolding, I can even explain why - this room is an unfolding of an isosceles triangle. One can show that there's no billiards shot that starts and ends at A (unlike C, for example, which could be shot perpendicular to the edge AB!) So in any unfolding of that triangle, there's no way for a straight-line path to connect to points that correspond to A. The trick then is to find such an unfolding so that any point corresponding to B or C is a corner point.

As a consequence of the work of Eskin-Mirzakhani-Mohammadi, a ball on a rational polygonal billiard table only has finitely many inaccessible points! If you only think about unfolding constructions, this isn't so surprising - but translation surfaces can be much more complicated than that.

In fact it is worth noting that this is really false if you allow more general shapes - on the right here is an example of a non-polygonal room where big regions fail to be illuminated.

Another related application is known as the *blocking problem*, and in what I think is one of the most charming pieces of a professional mathematics talk I've ever seen, I'd like to play a recording of Maryam describing it at a Princeton seminar in 2012.

Maryam and her husband Jan welcomed their daughter Anahita, about a year before that video was taken. A year after that seminar, she was diagnosed with breast cancer, and successfully treated it into remission. A year after *that*, she was awarded the Fields medal, the highest honour in mathematics, and the prestigious

Clay award for her research. But the cancer would soon recur, and Maryam died on the 14th of July, 2017, just 40 years old.

Maryam's loss was deepest for her family and friends but felt across the international scientific community not just for her stunning career cut short, but for the loss of a woman whose life and work inspired us with her ability to turn ambition into collegiality and collaboration instead of ego, who was willing to share her passion for mathematics with students and colleagues alike. She was brilliant and ambitious, she was kind and humble, and she was gone too soon.

There are a number of wonderful tributes to Maryam's life - international women in maths day is celebrated every 12th of May, Maryam's birthday, and scholarships, prizes, and schools in her name are supporting promoting the work of young and future mathematicians, particularly female mathematicians.

I was lucky enough today to share some of Maryam's mathematics with you. If you are interested in learning more about her personal life, hearing from her family, her former teachers and collaborators, and her friends from childhood, this lecture coincidentally shares part of the name of a beautiful documentary about Maryam's life and career from Zala films, found at http://www.zalafilms.com/secrets.