Computational Arithmetic - Geometry for Algebraic Curves

Prof Nitin Saxena

Dept of Computer Science and Engineering

IIT Kanpur

Week - 12

Lecture – 26

Proof of RH for curves II: Multilinear algebra

We have a picture in which the projective line is covered by the given curve through a point P_0 , and the Galois cover C' provides all the conjugates of P_0 that cover it. So that was one thing; the other is that it essentially reduces to the Galois properties of field extensions that are of finite degree. The degree of K_c' / K_c is equal to the order of the Galois group of this finite Galois extension. So, this has been made possible because we have moved from the curve to its function field. So, in terms of geometry, it does not look very simple, but in terms of field extensions, it is quite straightforward. This is the starting point of Galois theory.

And yes, one important thing to remember is that the Frobenius action on the function field is a bit different from what I had defined; it is not f^q. I have to use variables for this. So, actually, on, let us say, x1 and x2, only the variables are raised to the power of q. Yes, the important thing is that the Frobenius map, due to the characteristic p and q being a power of p , is always present because of the characteristic p .

At the level of functions, what it does is raise the variables to the power of q . Functions may also have constants; however, the constants are not exponentiated—only the variables are. Constants, yes. Right now, I think it is Fq, but in the future, we may move to Fq bar. I do not want this to be tied to the base field; I want it to be tied only to Q.

So, this is the Frobenius Q with Frobenius. So q will always be a power of p ; p is the characteristic, and the q -th Frobenius will operate independently of the function field and the constants. It does not affect the constants; it only changes the variables to the q -th power. What does it have to do with the issue? Sir, can there be a point on Kc' from F_q^m ? No, no, so yes, the image is a q-th power, but not of f. No, no, C ‘ has points from the algebraic closure.

You want α beta to be— I mean, in general, you want it to be anything that comes from

fp bar. So, yes, this map is actually intended for the q-th Frobenius to be general, regardless of the base field. So there is the q -th Frobenius and the p -th Frobenius. Their definition is independent of what small k represents. So, do you want this important definition to remain? For example, at x1 - α, it changes only x1.

So, it is x_1 ^{q} - α , and in this case, α can be anything. So, let us check the order of the items. So, $\sqrt{}$ x1 - α = α, correct? So, what is the Frobenius action on α that is α to the power of q inverse? And you see that α to the power of q inverse is a root of the image, which is x1 $^{q - \alpha}$, correct? The action of f on the point and on the polynomial is consistent, and you can see that the order is simply multiplied by q . Valuation is simply multiplied; this is clear because, in the image, whatever α is, this entire polynomial is a qth power. Again, it is something linear.

The order is actually multiplied by Q. This is an important example to remember because, in the proof of the Riemann Hypothesis, all these elements will become implicit. Yes, and then we defined this important generalization of point counting. It is with respect to the Galois automorφsm, σ. So N_1 C', σ consists of those points that are fixed by σ^{-1}F.

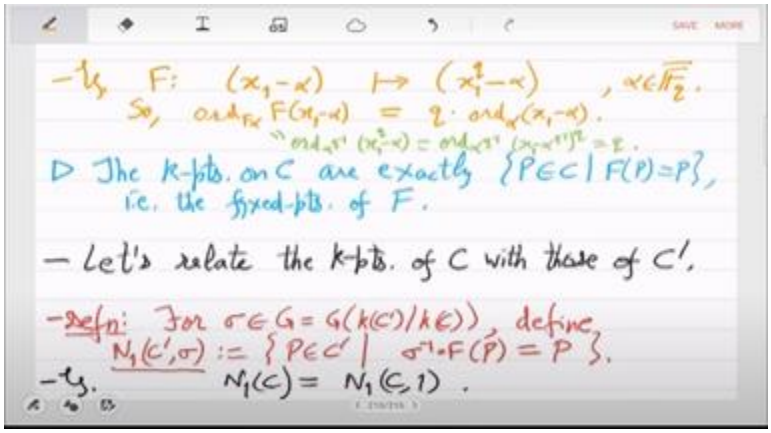

In other words, when you apply Frobenius, you obtain a conjugate under the action of σ , and we will prove this averaging theorem. This averaging theorem states that the count we are interested in, N_1 , is the average of all Galois actions. Of N1c ‘, so N1c ‘, σ is divided by the sum, and it is averaged by the Galois group. It is this, and we have seen the proof of it; this is not difficult. So, what will this allow? It will allow us to connect N1C to N1C ‘ σ tightly; in fact, it will connect N1P1 tightly.

So before that, let me apologize. P1—no, no, here the base field continues to be Fq, Q, Q + 1. Yes, so the picture is as follows: you have P1 covered by C, covered by C ‘, and

what is Galois is this; this is Galois, and this is also Galois. So, what we will do is prove the Riemann Hypothesis for Galois extensions first. Once we have proven this, we will first demonstrate the Riemann hypothesis for C' over P_1 , and once we establish that, we will essentially use these two Kalwa extensions, C' / C and C' / P_1 , to deduce the Riemann hypothesis for C .

This is what will happen. So, let us first state the Riemann Hypothesis. We have to finish the proof by today, let's say. So, Weil proved this in the 1930s, and the version we are studying in this course is Bombieri–Stepanov, which is much later. I think the spelling is incorrect.

So let C' be a Galois cover defined over F_q , and let me continue with C here. Let q be the power of p , where p is prime and is raised to α , with α being sufficiently large and even. Thus, G is the arithmetic genus of the curve over this base field, and we need to ensure that the field is sufficiently large compared to the genus. That is okay because we have seen that we can prove the Riemann Hypothesis in any field, and it will then hold in every other field. Then, with this setting for all Galois automorφsms, the count n_1 has an upper bound given by this expression.

Okay, this is a significant weakening of what we wanted to prove because it only states that N_1 is upper-bounded by this, right? So, for all you know, N1 could be zero. This seems to be much weaker than the Riemann hypothesis, but through this Galois diagram, we will soon see that this upper bound will actually lead to the Riemann hypothesis. Okay, we only need to provide an upper bound. That was mainly, I think, Bombieri's idea. This upper bound, as shown in the Galois cover diagram, implies everything.

So, why is that? Yes, the projective line has q + 1 points. I understand what you are saying. Yes, so basically what I am saying in the second line is, "without loss of generality." So your curve is defined over a field F_q , the genus is defined there, and then you move to a sufficiently large q' . Yeah, so maybe we can call this q ‘ for now.

Although I do not want to do that because I want to continue with Q, I can add a comment here. Think of C / FQ as being defined over a smaller field, and then let FQ be large enough. So, the curve is defined over some small field; for example, it can even be F_p , but the proof may not work over F_p ; we need a sufficiently large field. So, from there, we will go to a sufficiently large field F_Q that is larger than the genus, like this. It is always possible because wherever you go and whatever field you are in, the genus is fixed; the genus depends only on the degree of the given curve, which is, in fact, the polynomial that defines it.

So that will not depend on the value of k . For example, if the genus is around δ^2, then you just need to go to a field Q that is larger than δ to the eighth power; that's all. Yeah, it seems to be a cycle, but you can break it with this comment. You are given the curve at a specific point, and then.

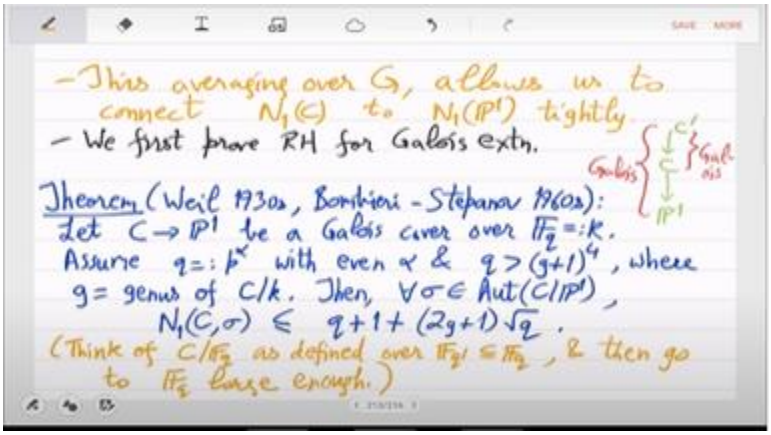

You are jumping to a larger field, which is your k , and that is where the Weil-Bombieri-Stepanov bound holds; that is the theorem. No, the arithmetic genus for a given field was not defined; this was the error in Riemann's theorem. The difference between L and the degree defines the genus. The L was essentially both the L dimension and L space, and the degree to which they were sensitive to k depended on k. However, the way we proved Riemann's theorem indicates that the genus can never be more than the degree squared.

Therefore, even if it changes, it cannot continue to change indefinitely. No, no, that k -\bark was simply meant to indicate that x is a transcendental function; that is all. No, no, no. As I said, this dimension is sensitive to the base field because the functions originate from it. Similarly, the degree is sensitive to the base field because the degree of the point—meaning the point itself—may become ambiguous.

That is sensitive to the base field, isn't it? So this is the tower: Fp contained in Fq ', contained in Fq. So, these are potentially the three fields, but we will only be interested in the last one, which is K. Let us see why it implies the Riemann Hypothesis. So let C_0 over K be a curve that may or may not be Galois; it is a general curve with a function field K_0 . Let K be; obtain all the information.

I mean, basically, also obtain FQ bar. It is a composite of the two fields, so you have the constant FQ at your disposal. So, let this be the function field over F_q , which is, of course, the same as F_p' . Yes, so now let C be the curve corresponding to the field K .

Okay, so C is essentially C_0. I mean, when we do point counting, not much will change because the point counting we are interested in is only in F_x. So, N1C0 and N1C are the same, but C has a better geometric structure; therefore, we are currently working with C. The point is that C is still not a Galois cover in general, so we need to take a Galois cover here; let us call it C '. Let K' be the smallest field and the smallest Galois extension that defines the Galois group C' over C . Now, we can apply the theorem to C' .

At this point, let us draw that once more. So, C' → C → P1, where this is Galois and that is also Galois. There is no loop; it is just a K. What is unclear about this? Yes, by the averaging proposition, we can connect C ' with C, the count. What you will obtain is: n_1 c is the average over this Galois group H , along with some error that depends only on the degree of the definition of this curve c , which I will call δ .

We have connected N1C with N1C ' H by going over all H; similarly, we can connect C ' with the line P1. This is done by the Galois group K ' / K bar X. Yes, I would like the algebraic closure. Yes, indeed. The function field K already contains constants from F_q , and the same applies to the line, as well as to C' and K' .

I can now use K' over whatever the function field of P^1 is, which I am calling K(X) . The Galois group is G , and G is a supergroup of H from the diagram. Each is a subgroup of G, and again, by the averaging proposition, what you will obtain is as follows: How does it depend on the cover? It depends on the average value.

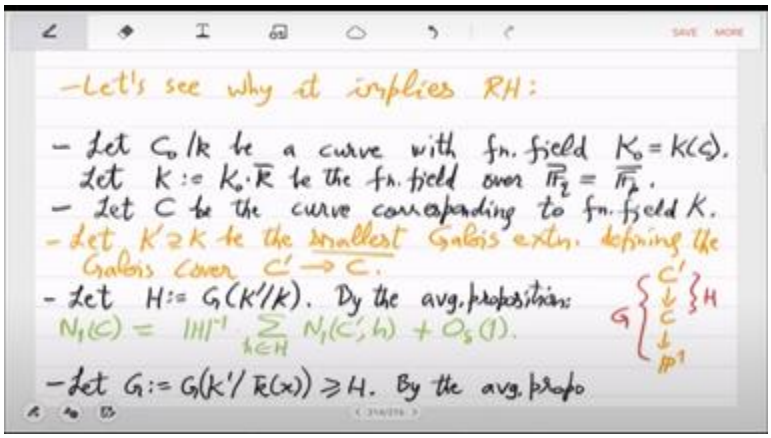

Okay, this is great because it invokes the weak Riemann hypothesis, which is the main theorem we need to prove in the future. What it says is that n_1'(σ) is something like q + 1 ; I mean, it is at most q + 1 + √q .

It could have been much smaller than q + 1 , but that is not possible because, for all σ , when you average, you must get a result of q + 1 . Right? This is q + 1 . If you want to achieve an average of q + 1 , then each of the values must be around q + 1 . According

to the theorem, for all σ in G_n1 (the ‘ σ), the main term must be q + 1 , along with the upper bound, which is 2g + 1 √q . So, this is δ 2 + g.√q, where big O has absolute constants.

We are obtaining big O of δ 2 due to this O of δ 1 error, which is essentially bounded by the resultant argument we discussed last time—the points that ramify. The original √ q is obtained from the theorem. The theorem had that much error, and I mean it was on the positive side. On the negative side, the error cannot be worse than this because if it were, then the average could not be greater than q + 1. Oh, that is based on the resultant argument that you have to count how many points ramify.

Yes, the O δ 1 is O δ 2. Yes, this part comes from the theorem. Yes, with this beautiful averaging, you average twice. You are now able to connect C' to the line. So, the benefit now is that, from an upper bound, you are obtaining bounds on both sides: a lower bound and an upper bound. This is actually what the strong Riemann hypothesis entails.

Wait, so we get c' , and then you use the other one—yeah, this one, the second averaging. You substitute that here, and you will obtain a value for n_1c . So, if you are interested in n1c, we go from n1p1 to n1c ‘, and then from n1c ‘ to n1c; that is the route. Let me just number this; it starts with 1. Okay, so what this implies is that the Riemann Hypothesis for C_0 holds over all fields.

Is that clear? Oh, this C_0 refers to the specified curve. So, what you have is n_1c ; n_1c is equal to n_1c0 because the only difference between c and c0 is the introduction of these f \barq constants in the functional field. So, you have the same curve c_0 , but you are viewing it as a curve over k . That does not change: N1, N1C, and N1C0 are equal because N1 is one of the fixed points of Q with respect to the Frobenius.

Yes, but that will not change the n_1 value. You can go anywhere, but the n_1 value remains the same. We have this old result that if you demonstrate it for n_1 , you will obtain the Riemann Hypothesis. Why did we need to go to K Bar? Why not just go to a large field instead? Why did we want to go to K-Bar? Yes, that probably isn't necessary. The theorem does not require it. I am not sure, but perhaps when we convert the function field to a curve, we may need algebraic constants from F_q .

But anyway, this is just for the sake of the proof. It ultimately demonstrates that establishing the upper bound with this setting will imply the Riemann Hypothesis; this is what we have shown. So, yes, we just have to prove this theorem. Let us mark this as the statement for the Riemann hypothesis; this is the version. The idea is to use the L-sheaf

extensively, employing this concept in multiple ways and applying the various Frobenius maps: the Q-th Frobenius and the P-th Frobenius.

Also, the automorφsm σ is included. So, look at the actions of all these on the L-sheaf; this is what will ultimately lead to the Riemann Hypothesis. So, use LAP; P is a point. For a single point, you examine various values of a, where a represents integers. You examine the vector space LAP and the actions of σ. In the theorem, the setting is a Galois cover, specifically G of C over P1, along with the actions of σ and Q with Frobenius.

The Qth Frobenius F and the Pth Frobenius, which we will refer to as F absolute, are associated with this sheaf. So, we will use this L to essentially define the location of the pole. And what the multiplicity can be is the LAP sheaf, and what happens when we apply the automorφsm σ, the Q-th Frobenius F, or the P-th Frobenius F is absolute with respect to these functions. We will essentially set up a commutative diagram connecting various L-sheets and comparing their dimensions, among other things.

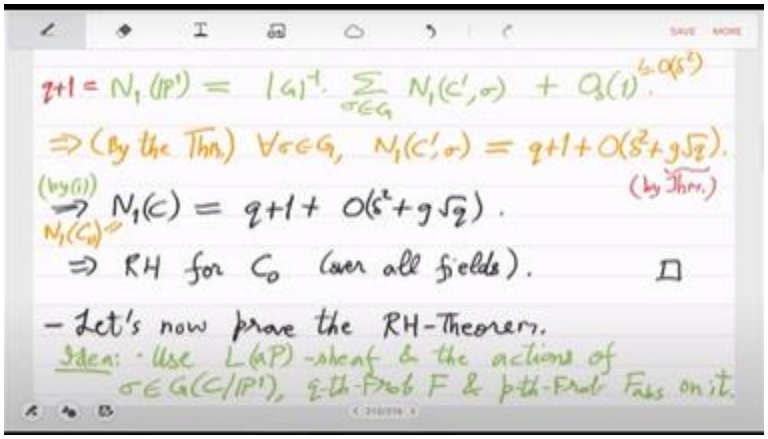

So, this is, of course, a very clever proof. Let us delve into the details. I am not sure whether this algebraic closure will actually be needed, but let us continue with it just to be safe, as we want to prove something about n_1 c , n_1 , c , and σ . That is the statement of Theorem 1: c and σ are upper bounded. So, changing the base field will not change anything. I mean n_1 is defined only by the fixed points of the q -th Frobenius.

Changing the base field will not alter this number, and that is the upper bound we wish to demonstrate. Yes, if the count is 0, then we are done; otherwise, there is a point that is fixed by the inverse of σ f. Such that $\sigma^{-1}(f)$ at p is equal to p. No, no, no. The curve C is defined by some equation, right? I mean, we have seen many examples.

Let's say it is y^2 = x^3 + x . There may not be any x such that x^3 + x is a square

because it is a finite field. Therefore, x^3 + x may never be a square unless you move to a larger extension. So, a priori, anything is possible. But the thing is, if N1 is 0, then it is okay because we are not stating anything about the lower bound; we only want an upper bound. In the lower bound, you can go all the way down to 0, which was not the case with the Riemann Hypothesis; however, we will not concern ourselves with that here.

That is why we are actually proving a weak version of the Riemann hypothesis; however, through Galois covers, it becomes more precise. It becomes stronger So anyway, let us pick a point where it exists, and its degree is 1. Yes, I think that is actually where we are using the definition of algebraic closure: every point has a degree of 1. Actually, that's also the reason we included K-bar in the connection. Here, k bar is also needed because, in the proof, I am not using anything.

But I guess, okay, let me not get into this again. Let's just move forward. So, the degree of this point is 1, and yes, we should choose a large enough number. We are not fixing it yet; actually, no, I can tell you what it should be: we will take this to be more than 2g - 2. So that Riemann Rock functions properly. So, basically, what will happen is that we will define L_a as the vector space L of a_p .

So now Riemann Rock tells you that since the degree is well A, which is more than 2g - 2, you have this exact formula: L equals a + 1 - g. Okay, so that is our first L-sheaf. I have a question: Why do you say DVR when this random variable n has a σ of 0? The reason is that.

.. Yes. The point simply means that there exists an x in Fq such that x^q + x is a square in Fq. Therefore, it is an even simpler statement. Sir, this variable Fq is in a polynomial, right? We can see every point, so it is a DVR, correct? Okay, I mean you can do that, but it won't give you any insight. Yes, N1 is counting the Frobenius points fixed by the Frobenius, so they are in FQ. N2 will count points in FQ^2, but we will only focus on N1, which represents the actual points.

When we define a class group, we say we are working with points that actually form a cloud of points or ‘ ideals, but they will not be counted in N_1 ; that is the main point. N1 actually provides you with rational points in Fq on the curve; however, it is not entirely accurate to say that it is N1c, σ. So, it gives you points for which the Frobenius moves to a conjugate element. Yes, there is an additional σ here, but suppose you are talking about N1c; N1c provides you with real, actual points.

N1c, σ, gives you what? Right, okay. No, we have proved those things, but what is the thing to which you want to apply them? Because n1c σ can still be 0 despite all those

connections. There is still a finite number of fields.

Yes. Yes. It tells you about the x-coordinate, not the y-coordinate. The Galois cover is doing something with the y-coordinates. The cover of the line is always present, but that doesn't mean you have both coordinates in the FQ. That's only one side of the coin. Yes, but we don't think in terms of those explicit projections.

Yeah, but those examples can also be considered quite explicitly. Yes, that is the very first L-sheaf. Now, we will do more with it. For example, we'll see how this map, σ, and f relate to each other, so let us define φ as the inverse of σ applied to f. Shall we? So, this is a map from a curve to curve ; it is a morφsm. To make sense of this, I guess we need to go to FQ Bar because Frobenius is fine, but there is also this σ, which is actually sending conjugates.

So, conjugates will not be in FQ; they will be higher up. However, you are correct that going to a sufficiently large degree will still be adequate. So, C is an algebraic closure here, and φ is a morφsm from C to itself. Its an automorφsm of its curve So, what does it do in terms of its functions at the functional level? So, define its pullback. So, remember that when you have a morφsm from one curve to another, there is an opposite arrow when studying the functions or the actions on those functions.

So, this L A φ represents the functions. So, this is basically F composed with φ. F is in L.A., okay? So, what you are doing is that when you want the action at a point, you should apply φ to go to the image and then use the function f there. Therefore, the composition of f with φ will be your function on the left-hand side. And f is the function on the right-hand side, is that correct? So, this is a pullback, but you do not need the geometric picture; you can also think of LFI as just another vector space of functions because both are defined on \mathbbC , and the curves are the same.

On that point, yes, it is somewhat inverted. So, it's Q in the Q-inverse. No, there is also this σ. So, are you only talking about the Frobenius action? Yes. So, if your function is x^1-α , then Frobenius will only affect the variables, changing the zeros through Q^-1 , which is the opposite of how it affects the variable. And now we will use this multiplicity connection, which is the order connection.

This is the right place. So let us now talk about the divisor of F\varφ . For any point Q in C on the curve and F in LA, what do you know about the divisor of F composed with φ? So, basically, what is happening is that the order of the point Q in F composed with φ is the same as the order of Q in essentially some other polynomial raised to Q. Because the characteristic is p , q is a p -power. So, when you apply this φ map, even though φ has

σ, the Frobenius action is such that f of φ is the q -th power of some function f' . Again, we are using the algebraic closure here Wherever the constants are, Q inverse will give you another constant that essentially provides you with the multiplier Q. So, when f is composed with φ , the pullback is of the q -th power, and what essentially happens is that the multiplicity increases by q .

Therefore, overall, the divisor also increases by q . Is that correct Okay Yes, this is a more abstract notation than that example, but it basically derives from it. I wrote this now to make the point that when you look at the degree of f composed with φ, it has been multiplied by q times. Yes, but what do functions do? Functions basically evaluate points on the curve, so it ultimately boils down to actions on points on the curve. It belongs to them; it acts on the curve.

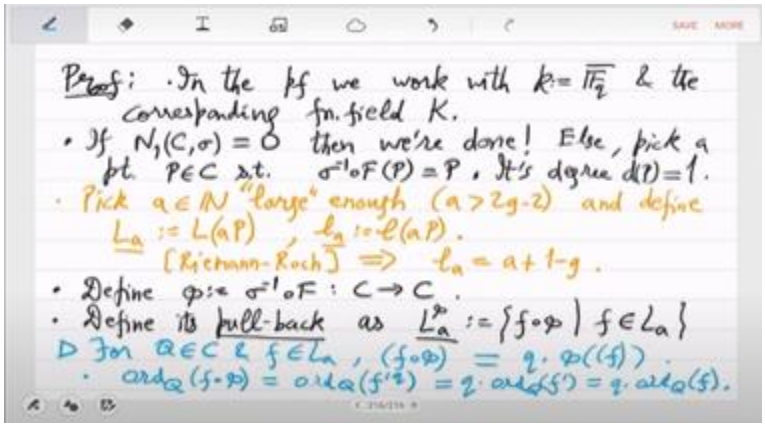

The curve remains the same; I mean, φ is an automorφsm of the curve. Well, if you want to see it that way, then it is essentially this: the map L_φ is a pullback of the sheaf \mathcalL. Φ maps the first curve to the second curve. Over the second curve, the La sheaf is defined, and it pulls back to the La φ sheaf over the first curve. The arrow is reversed, and I have just made this explicit.

This is the pullback, denoted like this, or perhaps I should place it on the right side. So, φ sends the first curve to the second curve, while φ* sends the sheaf on the second curve to the sheaf on the first curve. It is always like that. Okay. Whenever we mention x1 or x2, x1 is essentially a value in the field of constants given a point x1 at p. This is why we say that these functions are defined over a variety; they do not exist on their own.

Instead, they evaluate points in the variety, mapping them to constants. As I mentioned, you can think of LFI as something distinct. Vector space LA is a space of functions, and LA φ is another space of functions. Kind of, yeah, but I don't want to write that because

there are some variables in the function that are raised to the power of q, while the constants are not.

So, what happens with the constants is actually the inverse of q . But anyway, the function is a qth power. I do not want to discuss f ‘, but I know that it is some f ‘ raised to the power of q. No, as I said, think about this case; that is all there is to this example. However, if you're looking for a slight refinement for clarity, you could say: "It is x_1 - α q^-1 raised to the power of q ; thus, the image is always a q -th power. Yes, that’s correct. You can— I mean, just think of P inverse. So, P^-1 on F_p is—yes, it's Q^-1 . I mean, it brings Q inside; you get x1Q - α. So, yes, that example should help you understand everything; otherwise, it might be a bit confusing.

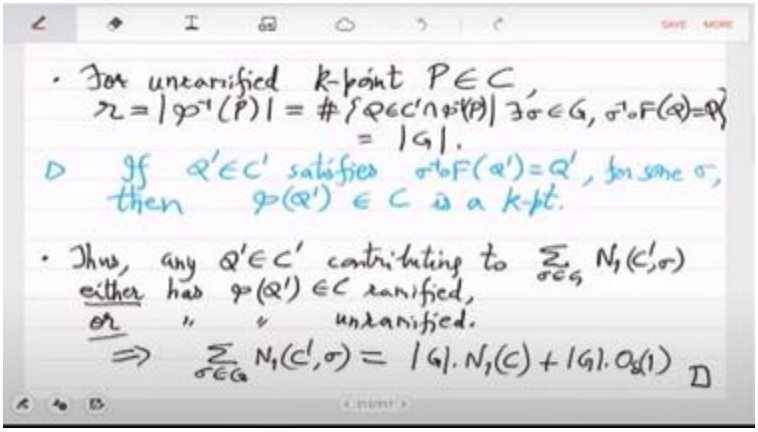

This one. But here is a slightly refined version for clarity: "So, it is written in green here. The root has changed from α to the inverse of q. What is the order of that root? Well, it is the same as what you started with, multiplied by q.

Where is it? Oh, it's in the final stage. Yes, in the last step, I should apply something. What should I apply for? Yes, you are right; something needs to be done. That is also the final step. So, yes, in the last step, you just apply the map—oh, that will be F inverse, I suppose. So, apply φ inverse. That's what you said, right? Correct.

Yes, when you become familiar with this notation, it can actually be somewhat complicated. Okay, but it essentially comes from the qth power; everything originates from the qth power. So, should I also include φ inverse here in the divisor? This is a valid version as well because it’s the same thing. No, no, this is the divisor—the divisor of F.

So, you obtain the zeros. Aren't we saying that the inverse of φ of Q is √ F, right? You should place φ inverse of Q in the divisor in the correct position. Okay, yes, we can't

finish the proof today. We need to complete it quickly next time, but let me provide the maps that will appear multiple times in this proof.

So, the first map we learned about is this one. L_a to L_a φ. I guess I'll remove the star because I continue to use φ, but it is essentially the pullback at the level of functions. So, L a to L a φ is doing what? It is mapping a function to f composed with φ . How do you view this as a function in LAQ? So, let us write down the properties. So, the properties you are mentioning state that F + AP is greater than or equal to 0, and the divisor of F + AP is also greater than or equal to 0.

In the next statement, we will comment on this divisor. Now, since we know that the divisor of f composed with φ is merely a multiple of q times the width of the divisor of f , we understand that this divisor + q times p is greater than or equal to zero. Okay, which means that this is actually a function in LAQ, correct? Because of the middle property, you immediately get that G1 is in LAQ. So, that is the sequence; we will actually do this many times using different maps. Next, we will repeat this with the absolute Frobenius raised to p . So, what is the absolute Frobenius? Instead of q , you will use p , and if you apply it μ times, you are essentially applying the Frobenius raised to p , raised to μ .

We will obtain a set of sequences through this process. So, let's take it. And define the LB p μ vector space in a manner very similar to what we did before with the LA φ. So, what will happen here is that you will compose it with F absolute μ. So, syntactically, it is the same thing; you have just changed the map, and we have altered it.

We will obtain a different sequence here. So the sequence we obtain here is that LB is isomorφc to LB, p, μ. So instead of φ, the map is now p raised to the power of μ in the Frobenius. This should be based on previous experience, which indicates that it will go to b \, p^μ instead of a \, q . Now you have p times the new Frobenius, which is p^μ ; the properties are very similar to those before. So, f is mapped to this, which is then mapped to some function G .

I mean, it is the same function; it is now a function in LBP \, μ . So all that these sequences are telling you is that the multiplication in the degree is performed by the Frobenius morphism. Now, we take the tensor product of these.

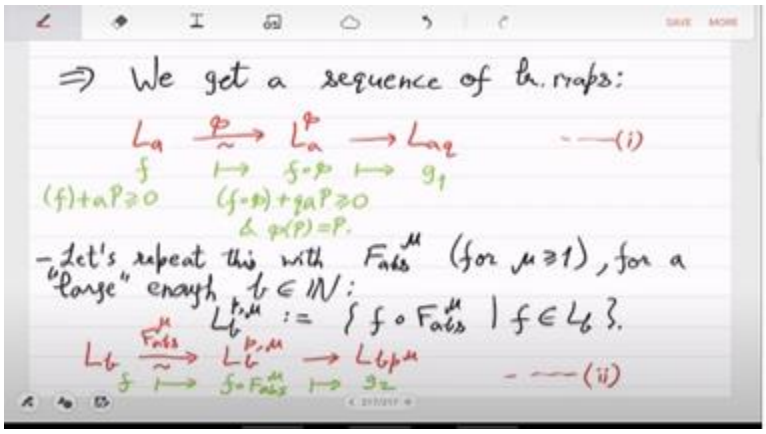

So, let us call this the multiplication map. So, what we will say here is that if μ is less than q , then... Then the multiplication map LB tensor LA φ is isomorphic to something that we need to prove next time. The content of this claim is that the tensor product can be obtained by simply taking the product of the two components. So, this also needs to be defined correctly. What is happening here is that if you have σ, the second one I wrote is G_2. So, if I have σ G2 or just a single rank-1 tensor product, it will look like this, and it will be mapped to G2 times G1, which will then be mapped to some G3.

So, you obtain a third vector space. Elements in that vector space have a natural multiplication map, where you take a rank-one tensor and multiply it, as they are functions. So you again obtain a function g_3 , which is included in that chief. The content of the claim is that there is no loss of information. Therefore, we will prove this next time. Yes, and once we have proved this, I think we will be almost done with the proof of the Riemann hypothesis, because then we will pick an element in this middle vector space, and we will show that it somehow knows about the correct count, which is N_1c σ .