An Introduction to Hyperbolic Geometry

Prof. Abhijit Pal

Department of Mathematics & Statistics

Indian Institute of Technology-Kanpur

Lecture – 41

Hyperbolic Groups and the Properties of Quasi-Convex Subgroups

(Refer Slide Time: 00:13)

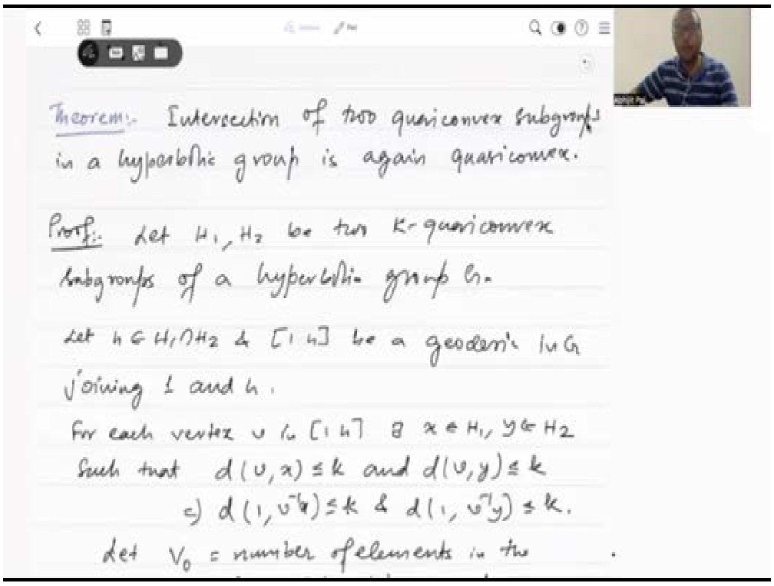

Hello! In this lecture, we will demonstrate that the intersection of two quasi-convex subgroups within a hyperbolic group is itself a quasi-convex subgroup. Additionally, we will use this result to show that a hyperbolic group cannot contain an abelian subgroup of rank greater than or equal to 2.

Let's begin by proving this theorem. Consider H_1 and H_2 as two k-quasi-convex subgroups within a hyperbolic group G. Now, let's take an element h that lies in the intersection of these two groups, H_1 and H_2. We will join 1 and h by a geodesic in the group G.

Our goal is to show that for any vertex v on this geodesic from 1 to h, there exists a vertex v' in the intersection $H_1 \cap H_2$ such that the distance between v and v' is less than or equal to a function of k and the hyperbolicity constant of the group G.

For each vertex v on the geodesic connecting 1 to h, we can find vertices $x \in H_1$ and

$y \in H_2$ such that the distance from v to x is less than or equal to k, and the distance from v to y is also less than or equal to k. This property holds because both H_1 and H_2 are k-quasi-convex.

Consequently, this implies that the distance from 1 to v^{-1} x is less than or equal to k, and the distance from 1 to v^{-1} y is likewise less than or equal to k.

(Refer Slide Time: 02:21)

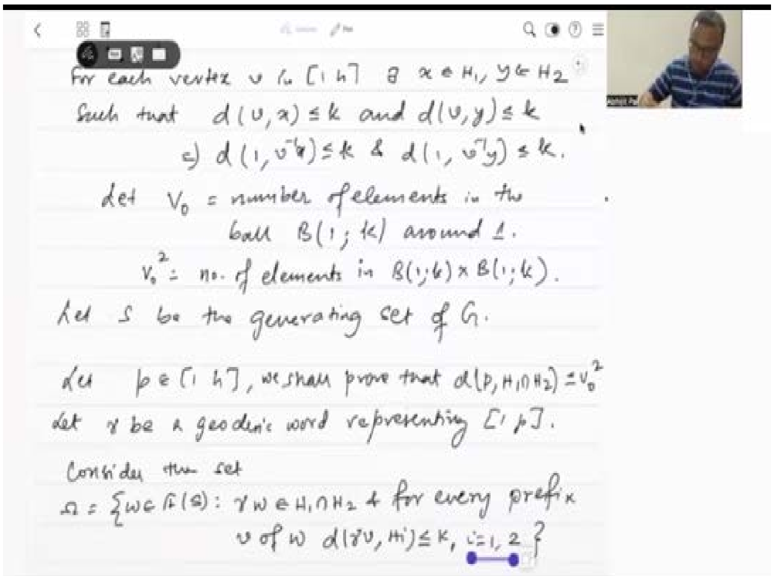

Now, let us define v_0 as the number of elements in the ball of radius k centered around 1. Consequently, v_0^2 will represent the number of elements in the Cartesian product of these two balls $B_1(k)$. We also take S to be the generating set of the group G.

Our aim is to prove that for any vertex p on the geodesic connecting 1 to h, the distance from p to the intersection $H_1 \cap H_2$ is less than or equal to v_0^2. This will demonstrate that $H_1 \cap H_2$ is a quasi-convex subgroup.

Let γ be a geodesic word representing the path from 1 to p. We will now consider the following set. We define the set **Ω** as the collection of all words w that belong to the free group generated by S such that the concatenation of γ and w is an element of the intersection $H_1 \cap H_2$. Furthermore, for every prefix v of w, the distance between γ v and H_i is less than or equal to k, where i is either 1 or 2.

Let me illustrate this concept with a diagram. We have the set **Ω** in focus. First, we take an element h that belongs to the intersection $H_1 \cap H_2$. We connect 1 to h via a geodesic.

Next, consider an element p that lies on this geodesic. The geodesic path from 1 to p is represented by the word γ.

(Refer Slide Time: 04:08)

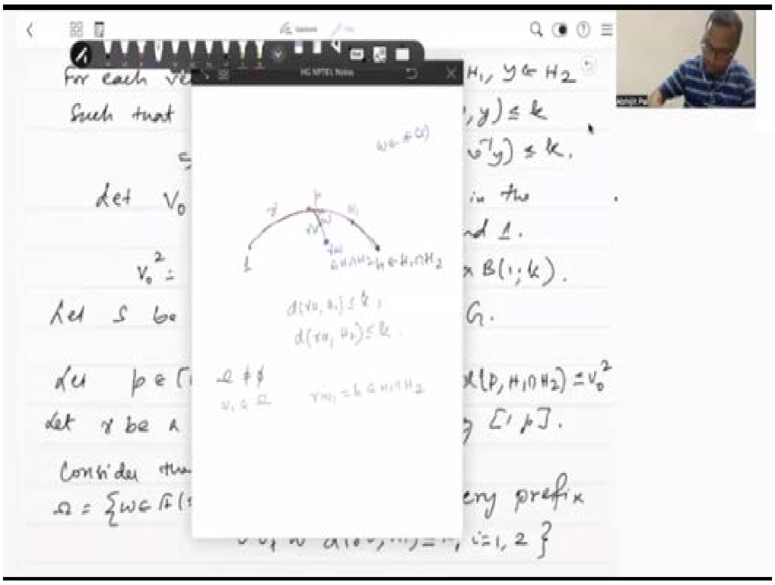

Now, what exactly is the set Ω? It consists of all those words w in the free group generated by S such that when you concatenate γ with w, denoted as γ w, you reach a vertex γ w that lies within $H_1 \cap H_2$. If I take any vertex along this path, labeled by a subword u, that vertex corresponds to γ u. The distances from γ u to H_1 and H_2 are both less than or equal to k.

(Refer Slide Time: 07:37)

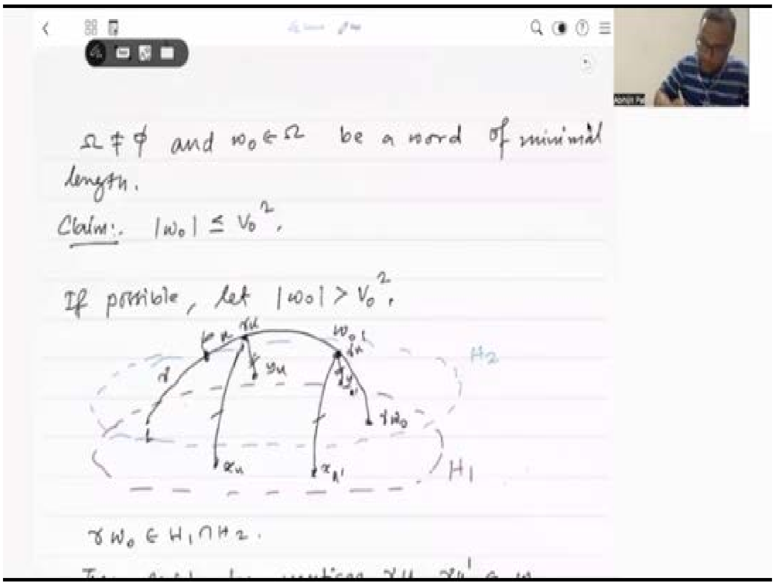

Having defined the set Ω, we can assert that it is non-empty. This is because, if we consider the geodesic path connecting v to h, let's denote this geodesic word as w_1. The word w_1 will belong to Ω since γ w_1 is equal to h, which indeed lies in the intersection $H_1 \cap H_2$. Moreover, since both H_1 and H_2 are quasi-convex, any vertex on this path will have a distance to H_1 or H_2 that is less than or equal to k. Therefore, we conclude that the set Ω is indeed non-empty.

Thus, we can conclude that the set Ω is non-empty. Now, let's denote w_0 as an element of Ω that has minimal length. It is indeed possible to identify such a word within Ω that exhibits this minimal length property. The claim we wish to establish is that the length of w_0 is less than or equal to v_0^2. If we can demonstrate this, then our proof will be complete.

However, let's assume for the sake of contradiction that the length of w_0 is strictly greater than v_0^2. This assumption will lead us to explore the implications further.

(Refer Slide Time: 08:41)

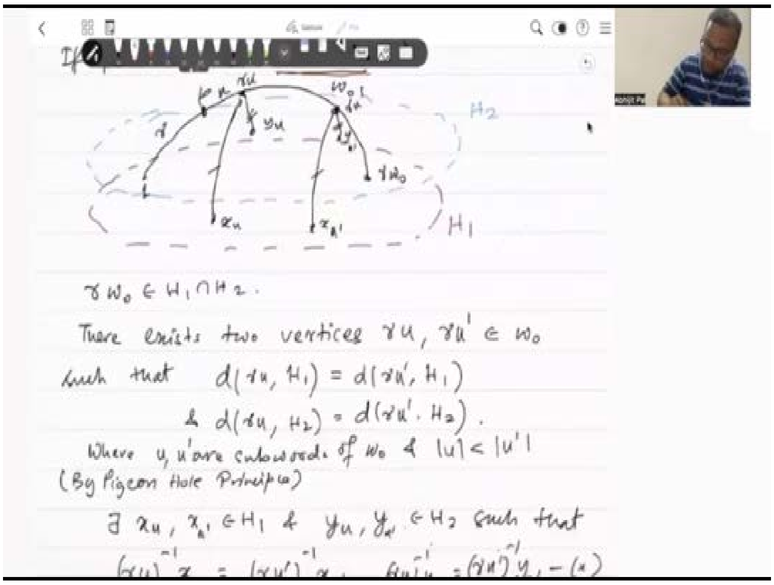

Given our setup, we have that γ w_0 belongs to the intersection $H_1 \cap H_2$, as defined by our set Ω. Now, by applying the Pigeonhole Principle, we can conclude that there exist two vertices, γ u and γ u', within the word w_0 such that the distance from γ u to H_1 is equal to the distance from γ u' to H_1. Similarly, the distance from γ u to H_2 is equal to the distance from γ u' to H_2.

But why does this hold true? The reasoning stems from the fact that w_0 is constructed from a finite number of vertices along the geodesic. When we explore the distances to the quasi-convex subgroups H_1 and H_2, the properties of quasi-convexity dictate that there cannot be too many distinct distance values. Hence, by the Pigeonhole Principle, some vertices must share the same distances to both H_1 and H_2.

(Refer Slide Time: 09:36)

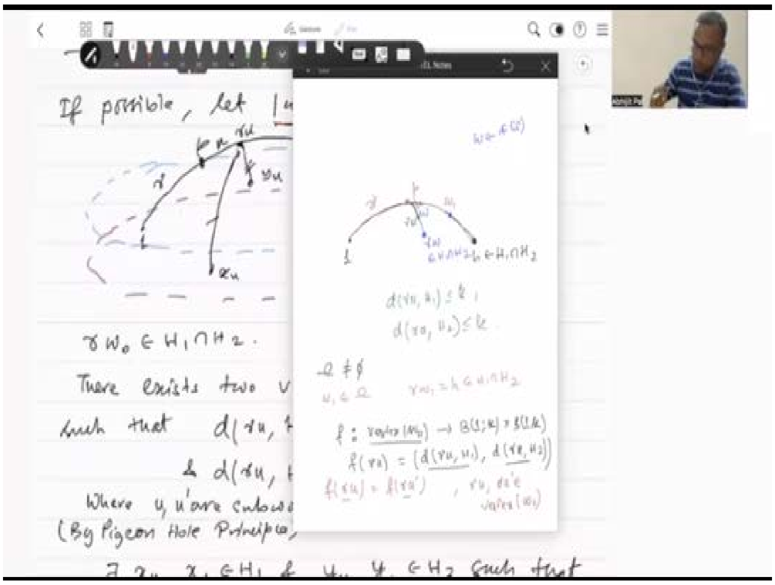

Let us consider a function f defined from the vertex set of w_0 to the Cartesian product of two balls. Specifically, for a vertex γ u, we define f(γ u) as the ordered pair consisting of the distance from γ u to H_1 and the distance from γ u to H_2.

The first coordinate of this ordered pair will belong to the ball of radius k around the identity because the distance from γ u to H_1 is less than or equal to k. Similarly, the second coordinate will also be within this ball, as the distance from γ u to H_2 is likewise bounded by k.

Now, let’s analyze the cardinality of the domain of this function f. The domain's cardinality is equal to the length of w_0. Given our earlier assumption, we took the length of w_0 to be greater than v_0^2.

Thus, according to the Pigeonhole Principle, since the cardinality of the domain (the length of w_0) exceeds that of the codomain (which is limited by the product of two balls of radius k), there must exist two elements, say γ u and γ u', in the vertex set of w_0 such that f(γ u) = f(γ u').

Consequently, we find that there exist two vertices, γ u and γ u', both belonging to w_0, such that the distance from γ u to H_1 is the same as the distance from γ u' to H_1, and the distance from γ u to H_2 is equal to the distance from γ u' to H_2. Moreover, we can select these subwords u and u' from w_0 such that the length of u is strictly less than the length of u'..

(Refer Slide Time: 13:44)

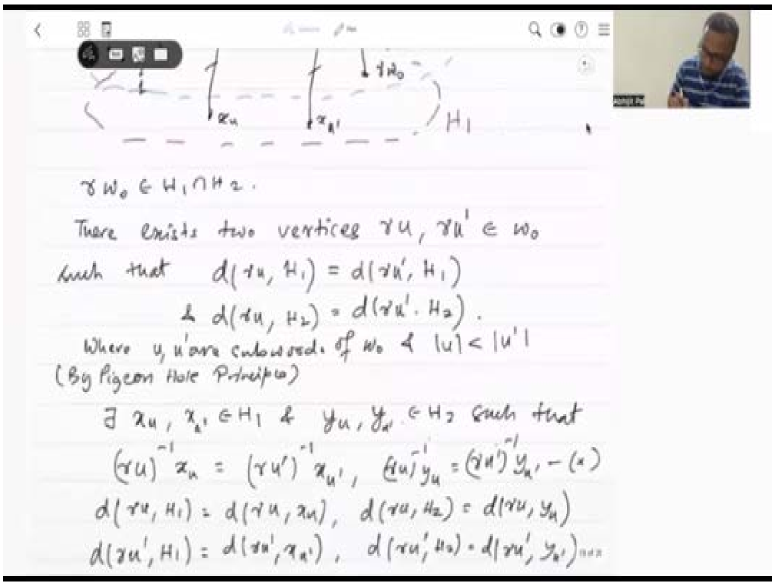

Therefore, we can conclude that there exist points x_u and $x_{u'}$ that belong to H_1, as well as points y_u and $y_{u'}$ that belong to H_2, such that the distance from γ u to x_u is the same as the distance from γ u' to $x_{u'}$. This means that these two geodesics are labeled by the same word, which we can express as $\gamma u^{-1} x_u = \gamma u'^{-1} x_{u'}$.

Similarly, the distance from γ u to y_u is identical to the distance from γ u' to $y_{u'}$. Here, both x_u and $x_{u'}$ lie within H_1, while y_u and $y_{u'}$ are elements of H_2. We can also state that $\gamma u^{-1} y_u = \gamma u'^{-1} y_{u'}$.

Now, let's visualize this situation. The word w_0, which belongs to our set ω, represents the minimal length path. We have established two vertices, γ u and γ u', along this path, such that the distance from γ u to x_u equals the distance from γ u' to $x_{u'}$. Similarly, the distance from γ u to y_u matches the distance from γ u' to $y_{u'}$. Notably, x_u and $x_{u'}$ are both in H_1, while y_u and $y_{u'}$ belong to H_2.

(Refer Slide Time: 16:07)

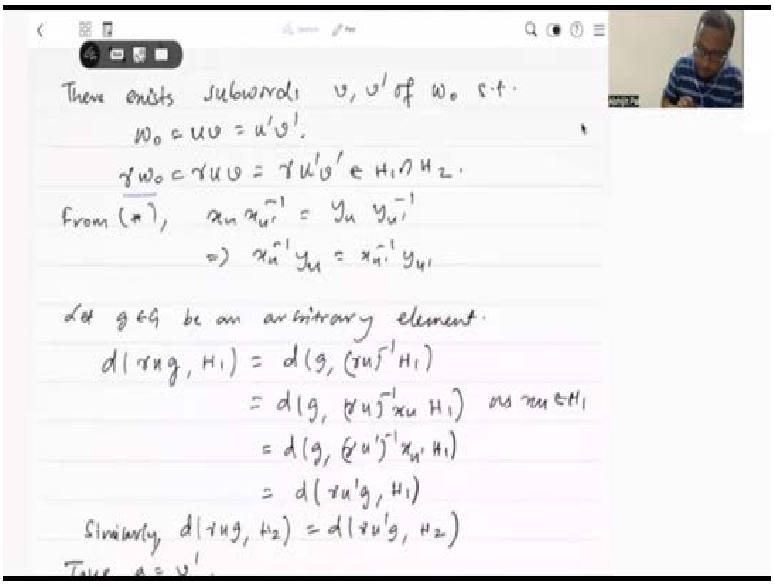

Now, we can identify subwords v and v' of w_0 such that w_0 can be expressed as u v and also as u' v'. If we refer back to our previous diagram, we can see that this path is labeled by w_0. Specifically, w_0 is represented as the concatenation of u followed by the geodesic labeled by v. Similarly, it can also be expressed as u' followed by the segment labeled v'.

Therefore, we have γ w_0 equal to γ u v and also equal to γ u' v', which means that this expression lies within the intersection $H_1 \cap H_2$.

Now, let's analyze the implications of this equation. From the relationships established, we find that $x_u x_{u'}^{-1} = y_u y_{u'}^{-1}$. This implies that $x_u^{-1} y_u = x_{u'}^{-1} y_{u'}$.

What this tells us is that if we consider the geodesic connecting x_u and y_u, its labeling will be identical to that of the geodesic connecting $x_{u'}$ and $y_{u'}$. This suggests a strong structural similarity between the paths defined by these vertices, reinforcing the coherence of our findings in this context.

Now, let’s take note of our observations and proceed by considering g as an arbitrary element of the group G. We need to examine the distance between γ u g and H_1. This distance is equal to the distance between g γ u^{-1} and H_1, because x_u belongs to H_1.

From this relationship, we can express the distance between g γ u'^{-1} $x_{u'}$ in a similar fashion. This distance is equivalent to the distance between g γ u'^{-1} $x_{u'}$ and H_1 due to the equality we established earlier. Thus, we can also say that this distance is the same as that between γ u' g and H_1.

(Refer Slide Time: 18:23)

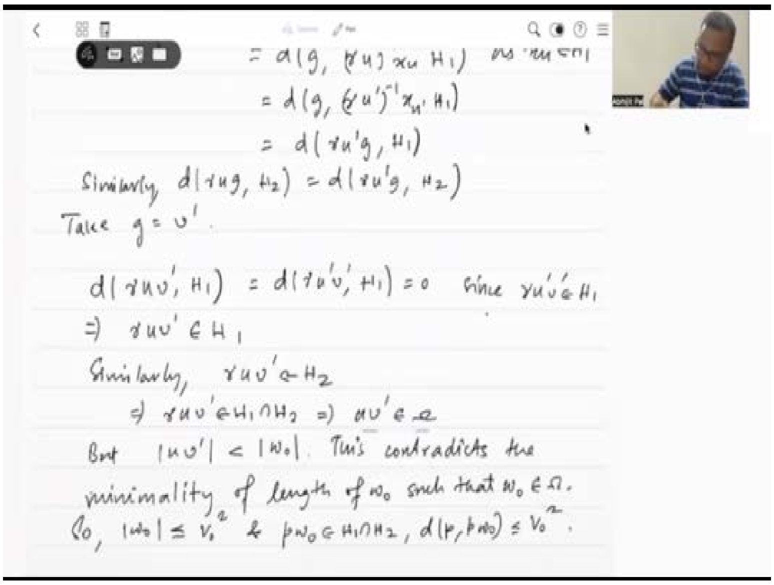

Furthermore, the distance between γ u g and H_2 will also match the distance between γ u' g and H_2. Now, let’s specifically set g to be v'. While we’ve considered g as an arbitrary element, we can indeed take g to be v'. Consequently, the distance between γ u v' and H_1 is equal to the distance between γ u' v' and H_1. Since both of these elements lie in H_1, this distance will be zero.

So, when we substitute g with v' in our equation, we find that the distance equals zero, which implies that γ u v' must belong to H_1. Similarly, we can conclude that γ u v' also belongs to H_2. Therefore, it follows that γ u v' is an element of both H_1 and H_2, meaning that γ u v' belongs to the intersection $H_1 \cap H_2$.

This establishes that u v' is indeed an element of the set **Ω**.

Now, if we look back at our earlier illustration, we can see that u followed by v' belongs to **Ω**. The length of u v' is strictly less than the length of w_0, which directly contradicts our assumption regarding the minimality of the length of w_0.

Thus, we conclude that the length of w_0 must be less than or equal to v_0^2. Additionally, since p w_0 lies within $H_1 \cap H_2$ and the distance between p and p w_0 is at most v_0^2, we have effectively proven that the intersection $H_1 \cap H_2$ is a quasi-convex subgroup.

(Refer Slide Time: 22:17)

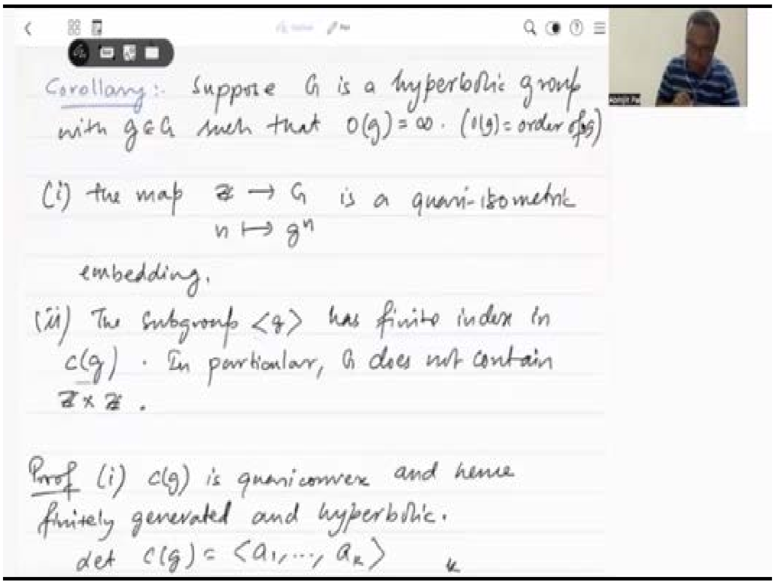

Now, as a corollary, we can draw some significant conclusions. Let's consider G to be a hyperbolic group with an element g such that the order of g is infinite. Our first observation is that if we define a map from Z to G given by $n \mapsto g^n$, this map serves as a quasi-isometric embedding.

The second point requires proof. Specifically, the subgroup generated by g forms an infinite cyclic subgroup, which has a finite index within its centralizer. This leads to a crucial implication: a hyperbolic group cannot contain $Z \times Z$.

To understand why, suppose for the sake of argument that G does contain Z × Z. If we take the generator of the Z component, it will represent an element of infinite order. Consequently, this would imply that the centralizer must contain Z × Z as well.

However, if the centralizer indeed contains Z × Z, then the index of the cyclic subgroup generated by g would not be finite within the centralizer. This situation leads to a contradiction, indicating that our initial assumption, that the hyperbolic group G contains Z × Z, must be false.

Therefore, we conclude that the hyperbolic group G cannot contain Z × Z. Now, let us first focus on proving the assertion that the map from Z to G, defined by $n \mapsto g^n$, is indeed a quasi-isometric embedding.

We have already established that the centralizer of an element is quasi-convex, which implies that it is finitely generated and, consequently, hyperbolic. Let's denote the centralizer of g as C_G, generated by a_1, a_2, …, a_k. Our first claim is that the center of this centralizer is the intersection of the centralizers of each a_i, where a_i forms the generating set of C_G.

(Refer Slide Time: 24:18)

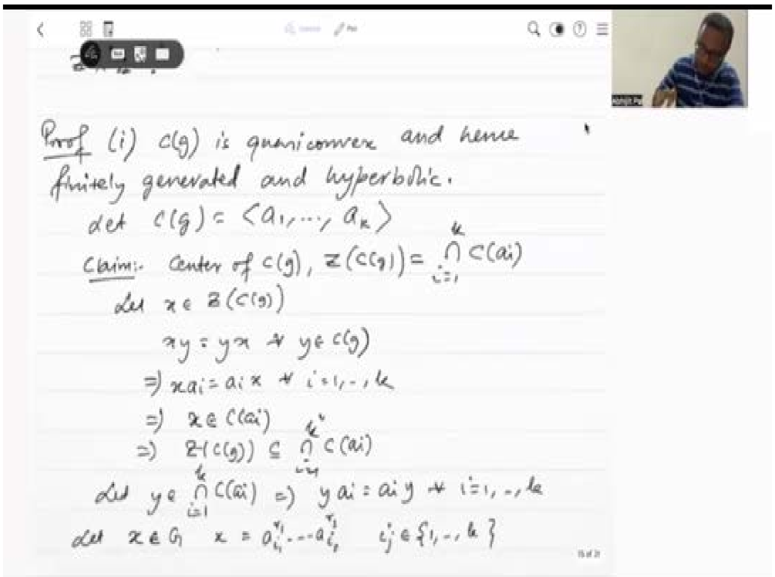

So, how do we prove this claim?

First, let's take an element x that resides in the center of the centralizer of g. This means that x commutes with every element in the centralizer. In particular, for each a_i, we have $xa_i = a_ix$. Therefore, we can conclude that x belongs to C_{ai} for all i ranging from 1 to k. This shows that the center of the centralizer of g is indeed contained within this intersection of the centralizers.

Now, let's examine the converse. Suppose we take any element y that belongs to this intersection. This means that $ya_i = a_iy$ for all i from 1 to k.

Next, consider any element x in g. We can express x as a product of the generators a_i. This establishes a connection between elements in the centralizer and their behavior in relation to x.

Now, if we consider the element xy and express it in relation to the generators a_i, we find that y commutes with each of these generators. Consequently, we can equate this element to another expression, ultimately concluding that it is equal to yx. This leads us to the conclusion that x indeed belongs to the center of the centralizer.

(Refer Slide Time: 25:50)

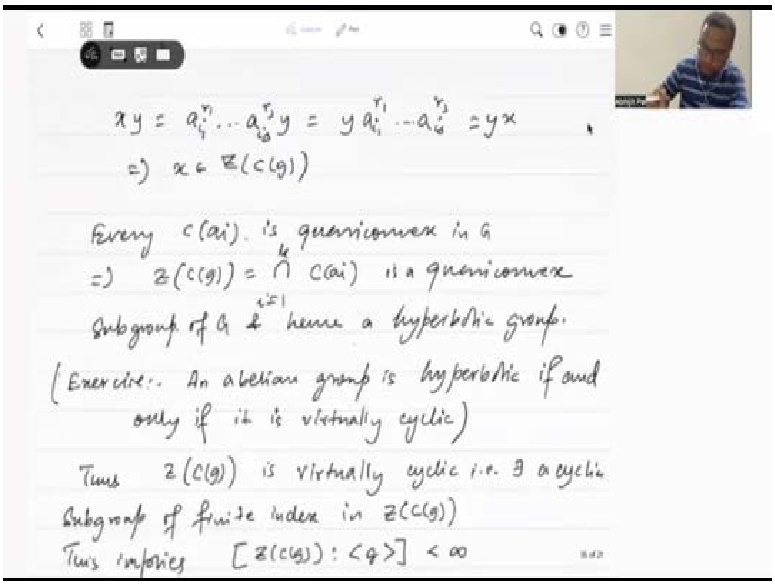

Next, it is important to note that each centralizer of the generators a_i forms a quasi-convex subgroup within the hyperbolic group. The center of the centralizer, being the finite intersection of these quasi-convex subgroups, will itself be a quasi-convex subgroup of G. Therefore, we can assert that the center of the centralizer is also a hyperbolic group.

(Refer Slide Time: 26:45)

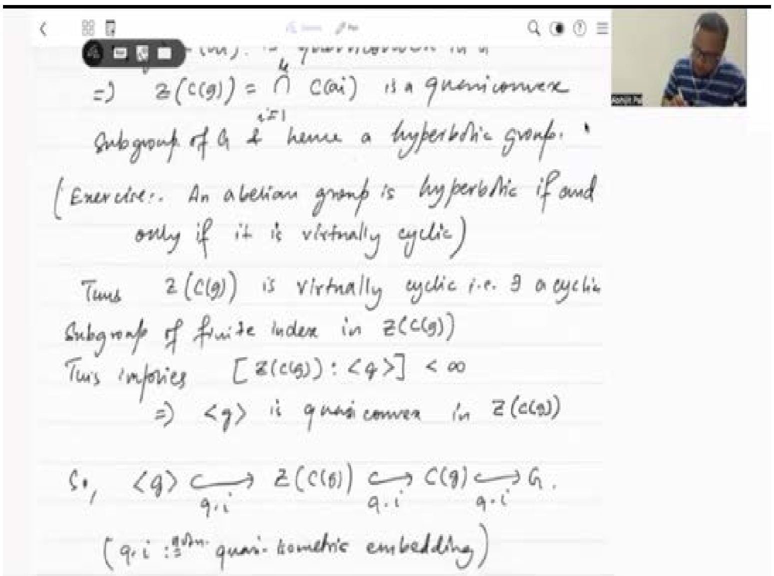

This exercise demonstrates that if we take any abelian group that is hyperbolic, it will necessarily be virtually cyclic. The term "virtually cyclic" refers to the existence of a cyclic subgroup with finite index. This is a crucial concept, and it's worth noting that the converse is also true.

Given that our group is abelian and we have established that it is a hyperbolic group, we can apply the findings from this exercise to conclude that this group must indeed be virtually cyclic. This indicates that there exists a cyclic subgroup of finite index within the center of the centralizer. As a result, we can infer that the index of this infinite cyclic subgroup within the center of the centralizer is finite. Consequently, this implies that the cyclic subgroup is quasi-convex in the center of the centralizer.

Therefore, we have a quasi-convex subgroup of the center of the centralizer. By referencing a previous theorem, we can assert that the embedding of this cyclic subgroup into the center of the centralizer will be a quasi-isometric embedding. Since the centralizer of an element is a quasi-convex subgroup, it follows that this embedding retains the property of being quasi-isometric.

Thus, we have demonstrated that this infinite cyclic subgroup is quasi-isometrically embedded in the hyperbolic group G. Therefore, we can confidently state that the map from Z to G, defined by $n \mapsto g^n$, is indeed a quasi-isometric embedding. Our next objective is to prove that this infinite subgroup has finite index within the centralizer.

(Refer Slide Time: 29:39)

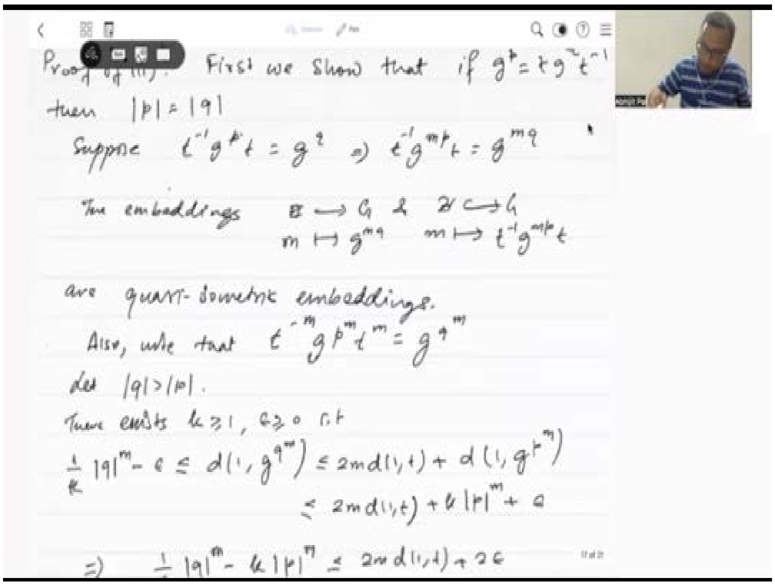

Let's begin by demonstrating the following relationship: if g^t is equal to the conjugate t g^q t^{-1}, then the absolute value of p is equal to the absolute value of q.

Now, suppose we have t^{-1} g^p t = g^q. This implies that t^{-1} g^m p t = g^m q. As we established earlier, these embeddings are indeed quasi-isometric embeddings, which we have already proven.

Additionally, consider the equation t^{-m} g^p t^m = g^{q^m}. This is a crucial observation and is presented as an exercise for further exploration.

Now, let's assume that the absolute value of q is strictly greater than the absolute value of p. This assumption will guide our subsequent analysis.

(Refer Slide Time: 30:56)

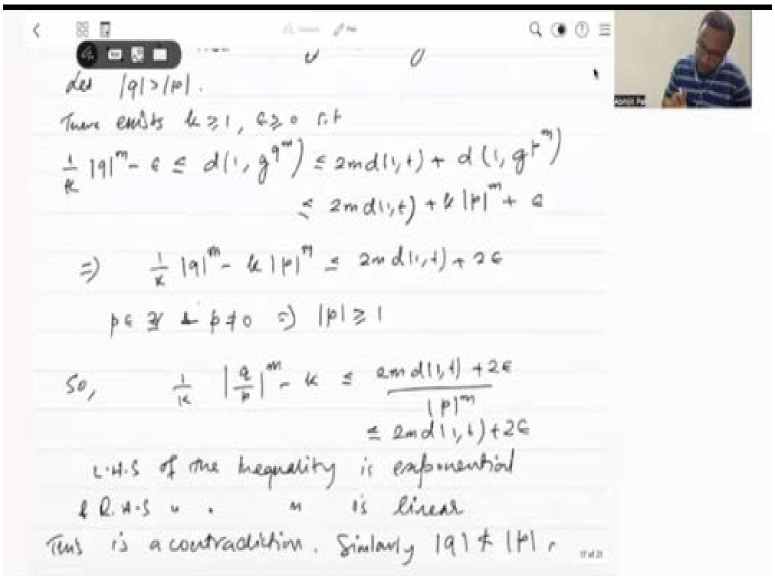

Let’s delve into our findings. We can identify constants k greater than or equal to 1 and ϵ greater than or equal to 0, such that

$$\frac{1}{k}|q^m| - \epsilon \leq d\left(1, g^{q^m}\right) \leq 2m\left(d(1,t) + d\left(1, g^{p^m}\right)\right).$$

This expression can further be refined to

$$\leq 2m(d(1,t) + k|p^m| + \epsilon).$$

This inequality arises from the fact that the embedding is a k, ϵ-quasi-isometric embedding.

Now, let's analyze the next inequality. We have $g^{q^m} = t^{-m} g^{b^m} p^m$. Therefore, the distance from 1 to g^{q^m} can be assessed as being less than or equal to the length of this word, which corresponds to the aforementioned expression.

Next, the distance from 1 to g^{p^m} satisfies

$$d\left(1, g^{p^m}\right) \leq k|p^m| + \epsilon.$$

This conclusion follows from our earlier discussion regarding the quasi-isometric embedding of the map from Z to G, where m is mapped to g^{p^m}.

Now, consolidating all of this, we find that

$$\frac{1}{k}|q^m| - k|p^m| \leq 2m(d(1, t) + 2\epsilon).$$

It's important to note that p belongs to Z and that p_0 is equal to 0. This implies that $|p| \geq 1$.

From this, we can derive that

$$\frac{1}{k}|q|^m - k \leq 2m(d(1, t) + 2\epsilon).$$

Given that $|p| \geq 1$, the left-hand side of this inequality represents an exponential growth, while the right-hand side exhibits linear growth. This leads to a contradiction, as we initially assumed $|q| > |p|$.

Consequently, we can conclude that $|p|$ cannot be less than or equal to $|q|$. Thus, the only viable outcome is that $|p| = |q|$, which we have successfully demonstrated.

This leads us to an important conclusion: the positive powers of g belong to distinct conjugacy classes. We have demonstrated that if $g^p = t\, g^q\, t^{-1}$, it necessarily follows that $|p| = |q|$. Now, if we consider p and q as positive integers, and if they are distinct, this means that g^p and g^q cannot be conjugate to one another. Hence, we can conclude that the positive powers of g are indeed contained in distinct conjugacy classes.

(Refer Slide Time: 34:45)

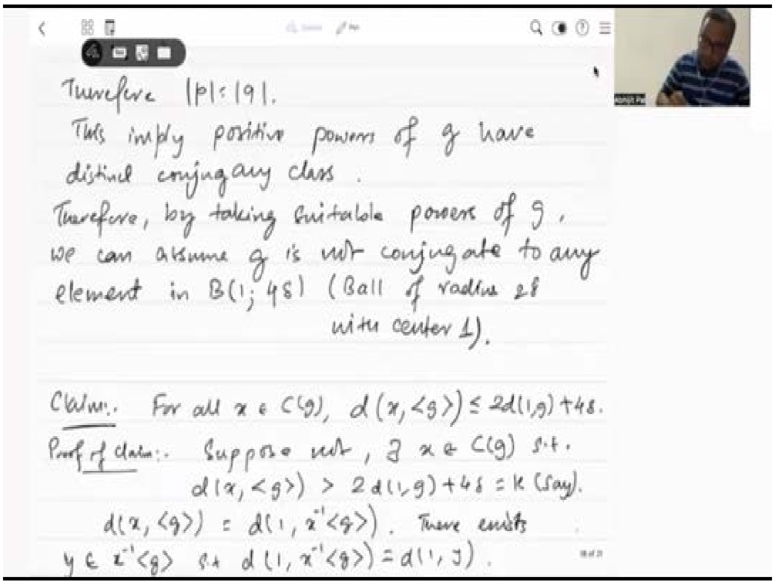

Consequently, by selecting appropriate powers of g, we can assert that g is not conjugate to any element within the ball of radius 4δ centered at 1. This gives us a clearer understanding of the relationship between the powers of g and their respective conjugacy classes, reinforcing our assertions about their distinctness.

(Refer Slide Time: 35:55)

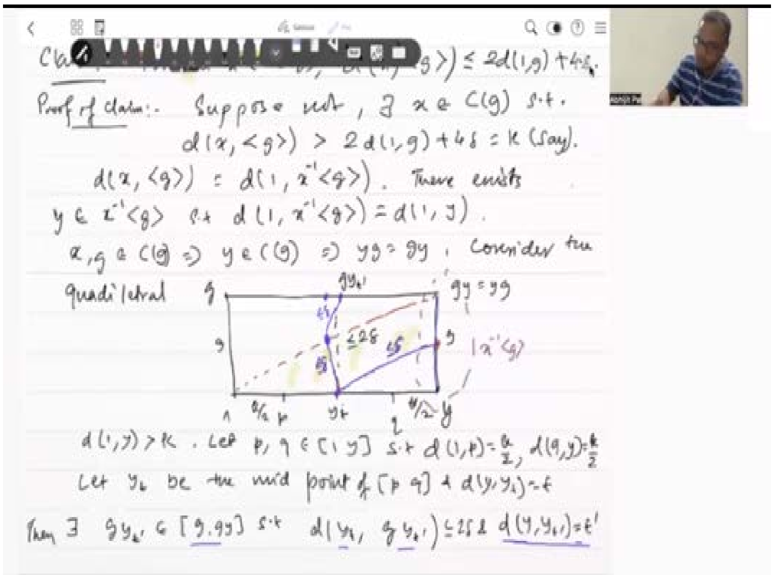

Now, let's move on to the next claim: if we take any element x that belongs to the centralizer, the distance from x to the cyclic subgroup generated by g will be less than or equal to twice the distance from 1 to g plus 4δ. Our goal here is to prove that the index of the cyclic subgroup generated by g is finite in the centralizer of the element g.

Suppose, for the sake of contradiction, that this claim is not true. This means there exists an element x in the centralizer such that the distance from x to the cyclic subgroup is strictly greater than $2 \times \text{distance}(1, g) + 4\delta$. Let's denote this quantity as k. Thus, the distance between x and the cyclic subgroup equals the distance between 1 and x^{-1} in the group generated by g.

Consequently, there must exist an element y in the subgroup generated by g such that the distance from 1 to y equals this value. Since x is in the centralizer, and g also lies in the centralizer of the element g, it follows that y must also belong to the centralizer of g. In other words, x is in the centralizer of g and y is in the subgroup generated by g.

As a result, y will commute with g, meaning that yg = gy. Now, consider the following arrangement: connect the points 1 to y with a geodesic, and connect 1 to g with another geodesic. Additionally, connect the points y and gy with a geodesic. Since gy = yg, the length of this geodesic is equivalent to the length of the geodesic joining 1 and y because g commutes with y.

Furthermore, we have chosen y such that the distance from 1 to y exceeds k. Hence, we can take two points p and q that belong to the geodesic connecting 1 and y such that the distance from 1 to p is $\frac{k}{2}$ and the distance from q to y is also $\frac{k}{2}$.

Now, let's consider the point y_t to be the midpoint connecting p and q. We will also define the distance between y and y_t as t. Now, observe that there exists a point g y_t' on the geodesic connecting g and g y such that the distance from y_t to g y_t' is less than or equal to 2δ.

How can we demonstrate this? We can connect the points 1 and g y with a geodesic. If we take the point y_t, we will find either a point here or there such that the distance from y to this point is less than or equal to δ. We are considering the triangle formed by the points 1, y, and g y, which we know to be δ-slim. Therefore, either this distance will be less than or equal to δ, or the other distance will be less than or equal to δ.

Now, suppose the distance is indeed less than or equal to δ. In that case, the distance between y_t and y will be less than δ plus the distance between y and g. Thus, the distance between y and y_t will be less than k/2, where k is defined as twice the distance from 1 to g plus 4δ. When we take k/2, we find that k/2 equals the distance from 1 to g plus 2δ. This leads us to a contradiction.

Therefore, we conclude that if the distance from y_t to the side is less than or equal to δ, we will still find a point such that the distance is less than or equal to δ. By similar reasoning, there will also exist a point here, which we will refer to as g y_t', and the distance will again be less than or equal to δ.

Hence, we find that there exists a point g y_t' on the geodesic joining g and g y such that the distance between y_t and g y_t' is less than or equal to 2δ. Additionally, observe that the distance between y and y_t' will equal t'. Consequently, the distance between g and g y_t' is equal to t'. Thus, we can conclude that the distance between y_t and y_t' is equal to t'.

(Refer Slide Time: 43:30)

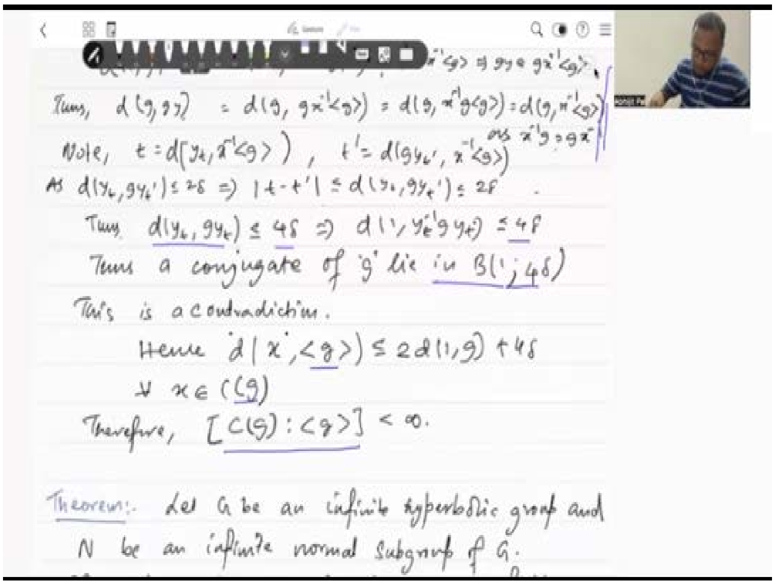

Now, let’s revisit our findings. From our earlier discussion, we can demonstrate that the distance between y_t and g y_t is less than or equal to 4δ. I won’t go through the calculations again, but I have indicated the reasoning here. This implies that the distance between y_t^{-1} g y_t is also less than or equal to 4δ. Consequently, this means that a conjugate of g lies within the ball of radius 4δ centered at 1.

However, this presents a contradiction. We initially assumed that any conjugate of g does not belong to this ball. Therefore, we conclude that the distance from x to this cyclic subgroup must be less than or equal to twice the distance from 1 to g plus 4δ. This relationship holds true for all x that belongs to the centralizer of the element g. As a result, we can infer that the index of this subgroup is indeed finite.

(Refer Slide Time: 45:03)

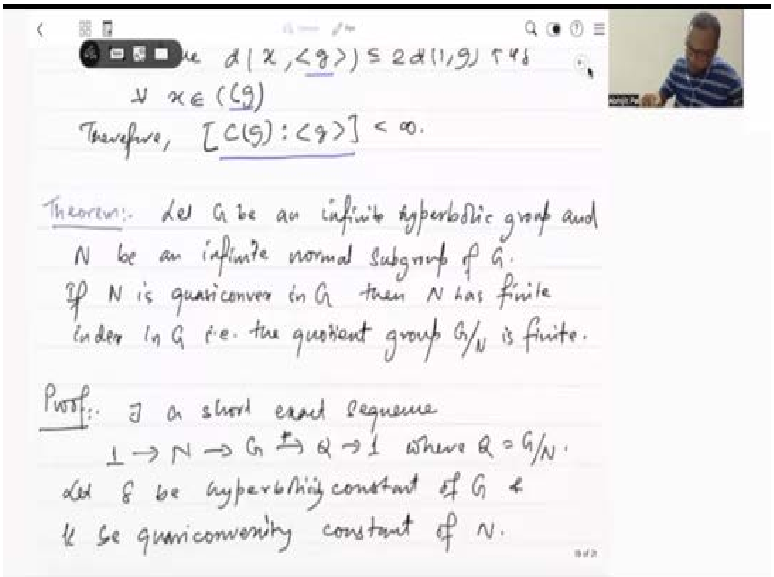

The next theorem states an important result, and while I won't go through the proof in detail, I encourage you to read it as it is provided here. The theorem asserts that if we take an infinite hyperbolic group G and an infinite normal subgroup N of G, and if N is quasi-convex in G, then N must have finite index in G. In simpler terms, this means that the quotient group G/N is a finite group. The proof of this theorem is documented here for your review. Okay, so I will stop here.