The Svarc-Milnor Lemma

Georgi Kocharyan

Abstract

We present the entire path to the Svarc-Milnor lemma from scratch. After defining groups and metric spaces, we define actions of groups on metric spaces, a concept of metric on a group and show how they are related. This is done via the *Svarc-Milnor lemma*, a beautiful result giving an exact relationship under some well-behavedness conditions. The primary use of this document is the 240-minute course at Maths Beyond Limits 2023.

We will prove a *non-topological* version of the Svarc-Milnor lemma, which might look considerably different to what you will find if you look it up online. This is good because people who don't know any topology will be able to understand it. But this is also bad because our version of the theorem will sound considerably more long-winded. But translating between the two versions isn't hard and we will offer a derivation of the more familiar topological version in an optional final chapter. In this document, an exercise is a fairly easy question that builds up theory and you should try and solve for completeness. In contrast, a *problem* is at the end of a chapter and is typically harder and optional for understanding the theory. We will assume the axiom of choice. Starred sections are optional for the proof of Svarc-Milnor and we will only discuss them if we have enough time.

The problems in this text draw heavily from [2].

1. What are groups and metric spaces?

We draw on Chapters 1,2 and 3 from Evan Chen's Napkin [1] for a concise introduction. If you have time you can read these chapters in advance, but we will go through the important parts in the course. Especially relevant is the beginning of Chapter 3 on generating sets.

2. Group actions

Lots of mathematical structures have some way of rearranging their own elements in a meaningful way. It turns out that we can sometimes discover groups 'inside' of these rearrangements, which give us a lot of information about both the group, and the structure. We start with the simplest example of this, which is the case of a set (i.e. just the elements with no additional structure specified).

2.1. Acting on sets

Let X be a set in the remainder of this section, unless specificed otherwise.

Definition 2.1. A *group action* of a group G on X, written $G \curvearrowright X$, is a function $\alpha: G \times X \to X$ with the properties that

(i) $\forall x \in X : \alpha(1_G, x) = x$.

(ii) $\forall g,h \in G, x \in X : \alpha(gh,x) = \alpha(g,\alpha(h,x))$.

If α is understood, we will write $g.x$ instead of $\alpha(g,x)$.

With this, each group element corresponds to a bijection from X to itself. Further, the composition of group elements corresponds to the composition of the bijections, and the identity element to the identity map. We can translate this idea into the language of homomorphisms as follows.

THEOREM 2.2. Every group action of G on X corresponds to a homomorphism from G to $Sym(X)$.

PROOF. The correspondence is that given α, the image of $g \mapsto \alpha(g, \_)$ is a bijection from X to X for any g. Check yourself that this correspondence is a homomorphism, and that the reverse also works. $\square$

EXAMPLE 2.3. Any group acts on any set via the *trivial group action,* namely the one corresponding to the trivial homomorphism.

EXAMPLE 2.4. The dihedral group D_{2n} acts on the $2n-$gon by rotations and reflections.

EXAMPLE 2.5. We can consider the group of ways to rotate a cube, which acts on the cube. How many elements does this group have?

2.2. *ORBITS AND STABILISERS*

It can be useful to consider the two following objects given a group action of a group G on a set X.

DEFINITION 2.6. Given a fixed $x \in X$, its *orbit*, written $\mathrm{Orb}(x)$ is the subset of X of elements that can be reached by applications of the group action. Written more precisely,

$$\mathrm{Orb}(x) := \{y \in X \mid \exists g \in G : g.x = y\}.$$

Exercise 2.1. Prove that for any group action on a set, the orbits partition the set. Equivalently, being in the same orbit is an equivalence relation.

DEFINITION 2.7. Given a fixed $x \in X$, its *stabiliser*, written $\mathrm{Stab}(x)$ or G_x is the subset of G of elements that don't change x under the group action, that is,

$$\mathrm{Stab}(x) := \{g \in G \mid g.x = x\}.$$

Exercise 2.2. Prove that for any group action of G on a set X and any $x \in X$, the stabiliser of x is a subgroup of G.

The next theorem is an extremely useful statement that you should test on some examples of group actions, such as Example 2.5.

THEOREM 2.8 (Orbit-Stabiliser)**.** Let a finite group G act on a set X. Then for a fixed $x \in X$, we have

$$|\mathrm{Orb}(x)| \cdot |\mathrm{Stab}(x)| = |G|.$$

PROOF. Every element of the group maps x somewhere. Given a fixed $g.x = y$, let us think about how many elements of G also map x to y. If $h \in G$ with $h.x = y$, then

$$h.x = g.x \Rightarrow g^{-1}.(h.x) = x \Rightarrow g^{-1}h \in \mathrm{Stab}(x) \Rightarrow h \in g \cdot \mathrm{Stab}(x).$$

This is the set of all elements that we get after left-multiplying those in the stabiliser of x with g. It is not hard to see that exactly this multiplication gives a bijection between the two sets, so there are the same amount of elements here as in the original stabiliser! This means that there are as many h that map x to y as there are elements of G that stabilise x. This is true for any y in the orbit of x, so we can associate $|\mathrm{Stab}(x)|$ many elements of G with every element in the orbit of x, which proves the claim. □

2.3. *THE BURNSIDE LEMMA*

In how many ways can we paint the faces of a cube in three colours? If two paintings can be rotated to become the same, we don't count them seperately. We will see that the solution to this riddle is naturally translated into the language of group actions.

DEFINITION 2.9. Call R_C the group of rotations of a cube (that map it to a cube again). This is a group because composition of rotations gives a rotation, and rotations as all functions fulfill associativity, and are invertible. Also we count „doing nothing"as a rotation.

THEOREM 2.10. R_C has 24 elements.

PROOF. We simply count all of the elements. First, notice that the identity is one element. Next, consider those with a rotation axis through two opposite faces that fix them, and rotate the others. For each pair of opposite faces, there are 3 non-identity rotations, giving $3 \cdot 3 = 9$ such transformations, which we call *face rotations.* Now count those with an axis through two opposite corners. These have 2 non-identity rotations for each of the 3 pairs, giving $4 \cdot 2 = 8$ such rotations, called *corner rotations.* Finally, note that we can also pick a rotation axis through opposite edges, which each only give one possible rotation, so $6 \cdot 1 = 6$ possible *edge rotations.* In total we thus have $1 + 9 + 8 + 6 = 24$. □

Exercise 2.3. Prove Theorem 2.10 directly by Orbit-Stabiliser.

Now we can state our question from the beginning of this section in a much more succint way. As we identify two paintings as the same if they are related by a relation, we can take the set of all possible paintings (which has 6^3 elements), let R_C act on it, and now ask how many orbits there are! And here is the result that lets us compute the amount of orbits:

THEOREM 2.11 (Burnside)**.** Let G be a finite group and $G \curvearrowright X$. The amount of orbits of the action is given by the average amount of elements that a $g \in G$ fixes, i.e.

$$\frac{1}{|G|} \sum_{g \in G} |\mathrm{Fix}(g)|,$$

where $\mathrm{Fix}(g) = \{x \in X \mid g.x = x\}$.

PROOF. Exercise. Try rewriting the sum, and then apply Orbit-Stabiliser. □

Now let's find out how many ways there is to paint a cube in three colours. $|X|=6^3$, and $|R_C|=24$ by above. Now we classified the elements of R_C in three different groups plus the identity, it will make sense to split the face rotations in 3 that rotate by π and 6 that rotate by $\frac{\pi}{2}$.

- The identity fixes every face and thus every painting.
- The face rotations by $\frac{\pi}{2}$ fix two faces and no others, so fix those paintings which have the same colour in all 4 not fixed faces, of which there are 3^3.
- The face rotations by π fix two faces and swap the remaining ones in pairs, so fix those paintings which have the same colours in the opposite non-fixed faces, of which there are 3^4.
- The edge rotations swap all faces in pairs, so we have free choice of 3 colours, giving 3^3 fixed colourings.
- The corner rotations only fix paintings that have the same colour in all faces bordering the corners that the axis passes through. There are 3 faces for such a corner, so we can pick two colours, giving 3^2 fixed colourings.

Now applying Burnside we see that the amount of unique colourings is

$$\frac{3^6+3\cdot 3^4+6\cdot 3^3+6\cdot 3^3+8\cdot 3^2}{24}=57.$$

2.4. Isometries and actions on metric spaces

Let (X,d_x) and (Y,d_Y) be metric spaces from now on.

Definition 2.12. $f:X\to Y$ is called an *isometric embedding* if for any $x,y\in X$, the equation

$$d_Y(f(x),f(y))=d_X(x,y).$$

If f is also bijective, it is called an *isometry.* If such an isometry exists, X and Y are *isometric.*

This definition is very easily motivated – two isometric metric spaces are for all intents and purposes the same from the viewpoint of their metric.

Exercise 2.4. Prove that every isometric embedding is injective, that is, never maps two different points to the same one.

Let X be a metric space. We can consider the collection of all isometries from X to itself.

Exercise 2.5. The set of all isometries from X to Y is called $\mathrm{Isom}(X,Y)$. Prove that $\mathrm{Isom}(X,X)$ forms a group under function composition.

For sets, a group action was a homomorphism to the bijections group of the set – i.e. each group element corresponded to a bijection. For a metric space, we have the same concept, except with isometries instead of bijections.

Definition 2.13. A *group action*, or an *isometric action* of a group G on a metric space X is a homomorphism $\alpha:G\to\mathrm{Isom}(X,X)$. We write $g.x$ instead of $\alpha(g)(x)$ if the action is clear.

Example 2.14. The additive group $\mathbb{Z}$ acts on the real line by shifting it. To be precise, $g.x=x+g$ defines a group action.

2.5. Problems to think about

2.1. Prove that G acts on itself in the following way: Fix a $g \in G$, then $g.x = gx$.

2.2. Use the above problem to prove *Cayley's theorem.* Any finite group G can be realised as a subgroup of the symmetric group S_n for some $n \in \mathbb{N}$.

2.3. Find all isometries of $\mathbb{R}$. (Hard extra: which group is $\mathrm{Isom}(\mathbb{R})$?)

2.4. Which cyclic groups admit non-trivial isometric actions on $\mathbb{R}$?

2.5 (Very hard)**.** For those who know some topology: Prove that an isometric embedding from a compact metric space to itself is always an isometry.

3. A metric on groups?

First we should talk a bit about what a graph is. Graphs have probably shown up in a lot of olympiad questions, but it won't hurt to look at a proper definition.

Definition 3.1. An *undirected simple graph* G is a pair (V, E), where V is a set, and E is a subset of $\binom{V}{2}$, that is, two-element subsets of V. We call V the *vertices* of G, and E the *edges.*

We imagine G as a collection of points, of which some are connected. Note that in our definition above, a vertex is never connected to itself.

Definition 3.2. A *walk* on G is a sequence of vertices, of which two consecutive ones are connected by an edge. We say it is a walk *between* the first vertex and the last vertex of the sequence.

From this point of view, it's not hard to connect a metric space to a graph. Take the vertices as the elements of the metric space, and define the distance of two vertices to be the length of the shortest walk between them, which exists due to well-ordering. The only problem here is that a walk might not exist – in this case define the distance to be infinite.

Definition 3.3. A graph $G = (V, E)$ is called *connected* if there is a walk between any pair of vertices.

Exercise 3.1. The above defined distance function is appropriately called the *shortest-path distance.* Prove that for a connected graph, the shortest-path distance turns it into a metric space.

3.1. The Cayley graph

Now comes the fundamental observation connecting the world of groups to the world of metric spaces. For any group G, given a generating set, we are able to turn it into a graph, and thus into a metric space. Let us see how.
An important note to make here is that we pick a generating set S for G, and then only are we able to make this jump, and the result depends heavily on the choice of S. So it is not right to say that given a group, there is a unique metric space associated to it.
The intuition is as follows – because S is a generating set for G, we can write every possible $g \in G$ as a finite product of elements from S. Then it is possible to see how some elements are in some

sense „close“to others. Namely, if they differ by few multiples of elements in S. So we would say that $s_1s_2s_3$ and s_1s_2 are very close, while $s_1s_2s_3$ and s_4 aren't necessarily. Notice that we have made an arbitrary choice here – we multiply on by the *right* when considering distance, i.e. we care if the differing elements of S are on the right. We are now ready to present the definition.

DEFINITION 3.4. Let G be a group and $S \subseteq G$ such that $\langle S \rangle = G$. Then the *Cayley graph* $\text{Cay}(G,S)$ of G with respect to S is the graph that has as vertices the elements of G, and for the edge set we allow (g,h) as an edge if and only if $g^{-1}h$ or its inverse $h^{-1}g$ is in S, and further it is not the identity.

Why do we make the somewhat awkward distinction in the end about $g^{-1}h$ not being the identity? It is so the following holds.

Exercise 3.2. Given a group G with a generating set S, $\text{Cay}(G,S)$ is a simple connected graph.

It would be very useful if the reader would make clear to herself where in this definition the arbitrary choice of „measuring distance on the right“comes in. All of our theory could be developed in the same way if we measured from the left instead, but we must make the choice one way or the other, which will lead to some asymmetries – see the next section.

Exercise 3.3. Show that $\text{Cay}(G,S)$ is *regular*, that is, every vertex is in the same amount of edges.

3.2. THE WORD METRIC

We make the next logical step in our discussion. Now that we have associated a graph to our group, and a metric to our graph, we close the circle and find a metric for our group.

DEFINITION 3.5. Let G be a group and $S \subseteq G$ such that $\langle S \rangle = G$. Then the *word metric* d_S on G with respect to S is the shortest-path distance on $\text{Cay}(G,S)$.

Exercise 3.4. Generating sets make a difference!

(i) If $G = \mathbb{Z}$ under addition and $S = \{1\}$, what is $d_S(0,5)$?

(ii) If $G = \mathbb{Z}$ under addition and $S = \{2,3\}$, what is $d_S(0,5)$?

Remember the next theorem well!

THEOREM 3.6. The word metric is *left-invariant,* that is, for any $h \in G$ and $g_1, g_2 \in G$ we have that

$$d_S(hg_1, hg_2) = d_S(g_1, g_2).$$

PROOF. Left as an exercise but is very easy, just unravel the definitions. □

Note that this is an example of the arbitrary choice we made earlier. If we measured distance from the left, the metric would have turned out right-invariant.

3.3. Interlude: End goal of Svarc-Milnor

In everything that we have done so far, something should stick out as a bit scary. In the beginning of this course, we promised that we would find a relation between the metric on a group and the metric on the metric space it acts on. But now we just saw that we don't have a clearly fixed metric on a group, it depends on the generating set we pick! How do you think this problem will be solved in the end?

3.4. Problems to think about

3.1. Think about Cayley graphs with few edges.

- Does there exist a group that has a Cayley graph with exactly 2023 vertices and exactly 2024 edges?
- Does there exist a group that has a Cayley graph with exactly 2023 vertices and exactly 2023 edges?

3.2. Show that if G is abelian and not cyclic, $\mathrm{Cay}(G,S)$ contains a cycle of length 4.

4. Quasi-isometry and bilipschitz equivalence

4.1. Definitions and examples

Definition 4.1. Let (X,d_X) and (Y,d_Y) be metric spaces, and $f: X \to Y$. We say f is a *bilipschitz embedding* if we can find constants $k, K \in \mathbb{R}^+$ such that $\forall x, x' \in X$ the following estimate holds:

$$k \cdot d_X(x,x') \leq d_Y(f(x), f(x')) \leq K \cdot d_X(x,x').$$

In other words, the error generated by the map f is bounded linearly.

Definition 4.2. Let (X,d_X) and (Y,d_Y) be metric spaces, and $f: X \to Y$ a bilipschitz embedding. We call it a *bilipschitz equivalence* if there is an inverse $g: X \to Y$ that is also a bilipschitz embedding.

Definition 4.3. Two metric spaces are called *bilipschitz equivalent* if there is a bilipschitz equivalence between them.

The following lemma comes in handy sometimes.

Lemma 4.4. A bilipschitz embedding is a bilipschitz equivalence if and only if it is bijective.

Proof. The one direction is trivial as the existance of an inverse implies bijectivity. Now assume $f: X \to Y$ is a bilipschitz embedding that is also bijective, and let g be the inverse that is given by bijectivity. It remains to prove that g is a bilipschitz embedding. Given two arbitrary $y, y' \in Y$ we have

$$k \cdot d_X(g(y), g(y')) \leq d_Y(f(g(y)), f(g(y'))) = d_Y(y, y')$$

proving that $d_X(g(y), g(y')) \leq \frac{1}{k} d_Y(y, y')$, which is the right part of what it means to be a bilipschitz equivalence. The left side works in exactly the same way. $\square$

While this concept of a bilipschitz equivalence is good and useful, we will find it a bit too restrictive for our purposes in future. Here is a different, more liberating idea.

Definition 4.5. Let (X,d_X) and (Y,d_Y) be metric spaces, and $f:X\to Y$. We say f is a *quasi-isometric embedding* if we can find constants $k,K,m,M\in\mathbb{R}^+$ such that $\forall x,x'\in X$ the following estimate holds:

$$k\cdot d_X(x,x')+m\leq d_Y(f(x),f(x'))\leq K\cdot d_X(x,x')+M.$$

In other words, the error generated by the map f is bounded linearly, and may differ by a fixed constant.

We would like to proceed in the same way as before, but the definition of a *quasi-isometry* gives us a few more difficulties. It is overkill to demand the existance of an inverse. As we will see, quasi-isometries will connect groups and metric spaces of completely different cardinalities and forcing them to be bijective would make our theory unnecessarily rigid. But nonetheless we want a sort of inverse to exist, so that we can capture what it means for two metric spaces to be quasi-isometric, i.e. the relation going in both ways. Thus we come up with the concept of a quasi-inverse.

Definition 4.6. We say two functions $f,g:X\to Y$ *have finite distance* if there is a fixed constant $c\in\mathbb{R}$ such that for any $x\in X$ it is true that $d_Y(f(x),g(x))\leq c$. Then $h:X\to Y$ is *quasi-inverse* to $k:Y\to X$ if both compositions $h\circ k$ and $k\circ h$ have finite distance to their respective identities.

Now we may proceed as before.

Definition 4.7. Let (X,d_X) and (Y,d_Y) be metric spaces, and $f:X\to Y$ a quasi-isometric embedding. We call it a *quasi-isometry* if there is a quasi-inverse $g:X\to Y$ that is also a quasi-isometric embedding..

Definition 4.8. Two metric spaces are called *quasi-isometric* if there is a quasi-isometry between them.

4.2. Word metrics are BL-equivalent

It turns out this is exactly the property we need to get control of the many word metrics on a group. All of them are bilipschitz equivalent!

Theorem 4.9. Let G be a group, and S,T finite generating sets of G. Then

$$id_G:\mathrm{Cay}(G,S)\to\mathrm{Cay}(G,T)$$

is a bilipschitz equivalence.

Proof. The identity is bijective and thus by Lemma 4.4 it suffices to show that it is a bilipschitz embedding. Take any $g,h\in G$. If $d_S(g,h)=n$, then we can write by the definition that

$$h^{-1}g=s_1s_2\cdots s_n$$

where these are elements of S. We want to now show that the distance of g and h with respect to the generating set T is bounded by a constant times n. Applying the triangle inequality for metrics gives

$$\begin{aligned}d_T(g,h) &= d_T(h^{-1}g, 1_G) = d_T(s_1 s_2 \cdots s_n, 1_G)\\ &\le d_T(1_G, s_1) + d_T(s_1, s_1 s_2) + \ldots + d_T(s_1 s_2 \cdots s_{n-1}, s_1 s_2 \cdots s_n)\\ &= d_T(1_G, s_1) + d_T(1_G, s_2) + \ldots + d_T(1_G, s_n).\end{aligned}$$

There are n terms in this sum, but what constant can we pick to bound them all? Well, just take $C := \max_{s\in S} d_T(1_G, s)$. As S is a finite generating set, this maximum exists, and we get

$$d_T(g,h) \le C \cdot n = C \cdot d_S(g,h).$$

This is again just the right part of the definition of a bilipschitz equivalence, but swapping S and T in the above argument gives the left part. □

It really cannot be stressed how important it is to remember that the finiteness of S and T is what makes all of this work. The above theorem *does not* hold in the case of infinite generating sets. This is why in geometric group theory, we mostly only ever talk about finitely generated groups.

4.3. Important properties

Try and prove the next two theorems by yourself. They should be a bit long, but not particularly hard.

Theorem 4.10. Every isometry is a bilipschitz equivalence, and every bilipschitz equivalence is a quasi-isometry.

Theorem 4.11. Being quasi-isometric is an equivalence relation on the set of all metric spaces, and so is being bilipschitz equivalent.

4.4. Non-properties

You should be careful with assumptions – isometries have many nice properties that quasi-isometries do not.

Exercise 4.1. Show that a quasi-isometry doesn't have to be

(i) injective

(ii) surjective

(iii) continuous

(iv) dimension-preserving.

4.5. Quasi-dense image and quasi-inverses

We had a nice little lemma that made our life easier when proving that something was a bilipschitz equivalence, namely Lemma 4.4. We would like a corresponding version for quasi-isometries, and there is one. First we prepare with the following definition and lemma.

DEFINITION 4.12. Let X be a metric space and $D \subseteq X$. Then we say D is *quasi-dense* in X if there is a constant $K \in \mathbb{R}$ such that every point in X has at most distance K to D, or in other words

$$\exists K \in \mathbb{R} : \forall x \in X \exists y \in D : d_X(x,y) \leq K.$$

LEMMA 4.13. Assuming the axiom of choice, a function having quasi-dense image is equivalent to it having a quasi-inverse.

PROOF. This is a fun exercise. Try doing it yourself! □

THEOREM 4.14. A quasi-isometric embedding $f : X \to Y$ is a quasi-isometry if and only if it has quasi-dense image.

PROOF. By the lemma, one direction is trivial. Now let $f : X \to Y$ be a quasi-isometric embedding with quasi-dense image. Again invoke the lemma to get a $g : Y \to X$ that is quasi-inverse to f. It remains to prove that g is a quasi-isometric embedding. Pick arbitrary $y, y' \in Y$. As f and g are quasi-inverses, $f \circ g$ has finite distance to the identity. To be precise, there exists a $C \in \mathbb{R}$ with the property that any expression of the form $d_X(f(g(x)), x)$ is bounded by C. So we get

$$\begin{aligned} k \cdot d_Y(g(y), g(y')) + m &\leq d_Y(f(g(y)), f(g(y'))) \\ &\leq d_Y(f(g(y)), y) + d_Y(y, y') + d_Y(y', f(g(y'))) \\ &\leq d_Y(y, y') + 2C \end{aligned}$$

which rearranges to a linear upper bound on $d_Y(g(y), g(y'))$. The lower bound again works analogously. □

4.6. PROBLEMS TO THINK ABOUT

The problems here have an olympiad feel, so have fun trying to solve them!

4.1. Prove that $\mathbb{N}$ and $\mathbb{Z}$ under the inherited metric from $\mathbb{R}$ are *not* quasi-isomorphic.

4.2. Prove that $\mathbb{R}$ and $\mathbb{R}^2$ are *not* quasi-isomorphic.

4.3. Find counterexamples to show that quasi-isometry does not imply bilipschitz equivalence, and bilipschitz equivalence does not imply isometry.

4.4. Is it true that every bijective quasi-isometry is a bilipschitz equivalence? Is this true for word metrics?

4.5. When are finite groups quasi-isometric?

5. QUASI-GEODESIC SPACES

The next section is only tangentially related to what came before, but you will notice parallels between the ideas. We shift our focus from maps between metric spaces to the metric space itself.

5.1. Definition and examples

It turns out that in the Svarc-Milnor lemma we will require the metric space to be a bit more special than just any metric space, namely it should also be a *quasi-geodesic* space. First, let's see what a regular geodesic space is. This definition is motivated by the idea of straight lines in the plane – these are lines that are isometric embeddings of real intervals.

Definition 5.1. Given a metric space X, a *geodesic* is an isometric embedding $\phi\colon [0,L] \to X$. We say it is a geodesic between $\phi(0)$ and $\phi(L)$.

Take some time to understand this definition. A geodesic is a line connecting two points that is in some sense the shortest possible path between them, and simultaneously for every point in between as well. What do you think are geodesics on a sphere, for example?

Definition 5.2. A *geodesic* space is a metric space in which any two points are connected by a geodesic.

Example 5.3. Show that $\mathbb{R}^2 \backslash \{0\}$ is *not* a geodesic space.

Now that we have a good sense of what a geodesic space is, we extend our definition to quasi-geodesic spaces, which is exactly what you expected.

Definition 5.4. Given a metric space X, a *quasi-geodesic* is a quasi-isometric embedding $\phi\colon [0,L] \to X$. We will often call it a (k,m,K,M)-quasi-geodesic if we want to specify the constants in the definition of a quasi-isometric embedding. (Definition 4.5)

Definition 5.5. A (k,m,K,M)-*quasi-geodesic* space is a metric space in which any two points are connected by a (k,m,K,M)-quasi-geodesic. If we don't want to specify the constants, we can also just call it a quasi-geodesic space.

5.2. Problems to think about

5.1. Show that $\mathbb{R}^2 \backslash \{0\}$ is a quasi-geodesic space, but not a geodesic space.

5.2. Show that the shortest-path metric on a graph (and thus the word metric on a group) transforms it into a quasi-geodesic space.

5.3. Let $f: X \to Y$ be a quasi-isometry.

- If X is geodesic, is Y also geodesic?
- If X is quasi-geodesic, Y also quasi-geodesic?

6. The Svarc-Milnor lemma

Slowly but surely we are ready to state and prove the Svarc-Milnor lemma. In short, we will have an isometric action of a group on a quasi-geodesic space with some conditions, which will imply the quasi-isometry of the group and the space.

6.1. Statement and proof

For easiness of notation, for a subspace $S \subseteq X$ of a metric space we will write $B_c(S) := \{x \in X \mid \exists s \in S : d_X(s,x) \leq c\}$.

Theorem 6.1 (Svarc-Milnor). Let G be a group that isometrically acts on a (k,m,K,M)-quasi-geodesic space. Further let there exist a finite diameter subspace $B \subseteq X$ such that

(i) $\bigcup_{g \in G} g.B = X$.

(ii) $B' = B_{2M}(B)$ has the property that

$$S = \{g \in G \mid g.B' \cap B' \neq \emptyset\}$$

is finite.

Then S generates G (finitely!) and G equipped with the word metric with respect to S is quasi-isometric to X via the following quasi-isometry:

$$\begin{aligned} \psi : G &\to X \\ g &\mapsto g.x \end{aligned}$$

where $x \in X$ is arbitrarily chosen.

Proof. We prove the claims one after the other.
Step 1: $\langle S \rangle = G$.
Let $g \in G$, and pick any $x \in X$. We will now provide a way of writing g as a product of elements of S. We use the quasi-geodesic property by considering a quasi-geodesic between x and $g.x$. This is nothing else but a (k,m,K,M)-quasi-isometric embedding ϕ of $[0,L]$ such that $\phi(0) = x$ and $\phi(L) = g.x$. The trick is now to dissect this quasigeodesic into many small parts. In particular, we dissect it into n parts with $\frac{L}{n} \leq \frac{M}{K}$, i.e. each part being smaller than $\frac{M}{K}$. Now label the dissection x_0 to x_n. We now have a bound on the distance between x_{i-1} and x_i, where in the following i is an arbitrary integer with $1 \leq 1 \leq n$.
Claim: $d_X(x_{i-1}, x_i) \leq 2M$.
This is proven quickly using the properties of quasi-isometry.

$$d_X(x_{i-1}, x_i) \leq K \cdot d_X(\phi^{-1}(x_{i-1}), \phi^{-1}(x_i)) + M \leq K \cdot \frac{M}{K} + M = 2M.$$

Now since $g.B$ cover the whole space (see (i) in the theorem) we can find $g_i \in G$ such that $x_i \in g_i.B$. We now realise that $g_i.B$ and $g_{i-1}.B$ are very close. By left-invariance of the word metric, this means B and $g_{i-1}^{-1} g_i.B$ are close. To be more clear, x_i is in both $g_i.B$ as well as in $2M$ distance of $g_{i-1}.B$. Using the notation B' as above, this means x_i is both $g_i.B'$ as well as in $g_{i-1}.B'$, and thus applying g_{i-1}^{-1} gives that there is an element (namely $g_{i-1}^{-1}.x$) that is in both B' as well as in $g_{i-1}^{-1} g_i.B'$. But by the definition of S this implies that $g_{i-1}^{-1} g_i \in S$. To finish this step off, write (noting that $g_0 = 1_G$ and $g_n = g$)

$$g = g_n = (g_0^{-1} g_1)(g_1^{-1} g_2) \cdots (g_{n-1}^{-1} g_n),$$

all of which by above are elements of S.
Step 2: ψ is a quasi-isometric embedding.

By left-invariance of the word metric it suffices to prove the bounds for $d_X(\phi(1_G),\phi(g))$ (Why?). Note that above we picked n such that $n \geq L\cdot\frac{K}{M}$. Picking it sharply also provides $n-1 \leq L\cdot\frac{K}{M}$, thus implying

$$\begin{aligned} d_X(\phi(1_G),\phi(g)) = d_X(x,g.x) &\geq k\cdot L + m \\ &\geq k\cdot\frac{(n-1)M}{K}\cdot k + m \end{aligned}$$

which is linear in n, and by definition we have $n \geq d_S(1_G,g)$. This is the lower bound, and the upper bound is $d_X(\phi(1_G),\phi(g)) \leq C\cdot d_S(1_G,g)$ where C is the maximum of all expressions of the form $d_X(x,s.x)$ in a fashion similar to when we proved that all word metrics on a group are bilipschitz equivalent. This follows from a brief induction argument over $d_S(1,g)$. If it is zero, $g=1$ and we are done. If it is n, write $g=s_1s_2\cdots s_n$ and let $h=s_1s_2\cdot s_{n-1}$. Then the statement is true for h, and we get

$$\begin{aligned} d_X(\phi(1_G),\phi(g)) = d_X(x,g.x) &= d_X(x,h.x)+d_X(h.x,g.x) = \\ &= d_X(x,h.x)+d_X(x,s_n.x) \leq (n-1)\cdot C + C \\ &= n\cdot C. \end{aligned}$$

Step 3: ψ has quasi-dense image.
Let $y\in X$. By (i) we have $y\in g.B$. Again by (i) we can take $x\in h.B$. But then $\phi(gh^{-1}) = gh^{-1}.x$ is also in $g.B$, so the maximum distance y can have from it is the diameter of B, which is finite by assumption.
Step 2 and 3 taken together imply by Lemma 4.14 that ψ is a quasi-isometry, thus proving the Svarc-Milnor lemma. □

6.2. Application to $\mathbb{Z}$ and $\mathbb{R}$

Let's immediately see how this can be applied. Recalling Example 2.14, there is an isometric action of $\mathbb{Z}$ on $\mathbb{R}$. As $\mathbb{R}$ is a geodesic space and thus a quasi-geodesic space, we have a chance of applying Svarc-Milnor. We need a suitable choice for B, and we guess that $[0,1]$ might work. It obviously has finite diameter, and $B' = B_c$ is just another closed interval which has the property that for only finitely many $x\in\mathbb{Z}$ there is an intersection of B' with its shift by x. Thus we can apply Svarc-Milnor, and we see that $\mathbb{Z}$ and $\mathbb{R}$ are quasi-isometric.

7. *Topological statement of the Svarc-Milnor lemma*

This section requires a good understanding of point-set topology and is entirely optional. If you look up the Svarc-Milnor Lemma online, it is likely you will find something completely different to what we wrote above. It might look more like this:

Theorem 7.1 (Svarc-Milnor, topological version)**.** Let G be a group that isometrically acts on a proper quasi-geodesic space X, such that the action is properly discontinuous and cocompact. Then G is finitely generated and G under any word metric and X are quasi-isometric.

We will now go through each unknown word in this theorem and define them.

Definition 7.2. A *proper* metric space is one in which every closed ball is compact.

Exercise 7.1. Show that proper metric spaces are exactly those in which the Heine-Borel theorem holds (every closed and bounded set is compact).

Definition 7.3. A *properly discontinuous* action of G on X is one in which for every compact subset $K \subseteq X$ the set $\{g \in G \mid g.K \cap K \neq \emptyset\}$ is finite.

Recall that we can quotient a metric space by a group action by collapsing the orbits.

Definition 7.4. A *co-compact* action of G on X is one in which the quotient space $\frac{X}{G}$ is compact.

As we have already proved the non-topological version, what we want to do in this section is show that the assumptions of the non-topological version of Svarc-Milnor are fulfilled under the assumptions of the topological version. We start with two small lemmas.

Lemma 7.5. The quotient map $q: X \to \frac{X}{G}$ under an isometric action is open.

Proof. Given an open set $U \subseteq X$, we want to show that $q(U)$ is open. Equivalently, $q^{-1}(q(U))$ is open. This is not equal to U, but in fact is equal to the union of all elements which map to the same ones as those in U, so in other words the union of all orbits of the elements in U, so

$$q^{-1}(q(U)) = \bigcup_{g \in G} g.U$$

But as G acts via isometries which are open (their inverse is also an isometry, which is continuous), each set $g.U$ is open. The union of open sets is also open, and we are done. ☐

Lemma 7.6. Proper metric spaces are locally compact.

Proof. We need to show that every point has a compact neighbourhood. But we can just take any closed ball with non-zero radius centered at the point and are done (the open ball with half the radius proves that it is a neighbourhood). ☐

Glancing at our non-topological version, the only thing we need to actually do is provide a finite diameter subspace of X that fulfills the two conditions. As our action is properly discontinuous and the space is proper, picking B compact will already give us (ii) as well as finite diameter. It remains to show that there exists a compact B with the property that $\bigcup_{g \in G} g.B = X$. We claim that it suffices to consider a locally compact space, which X is by Lemma 7.6.

Theorem 7.7. Let X be a locally compact space, and G a group acting on X. This action being co-compact is equivalent to the existance of a compact subspace $B \subseteq X$ with the property that $\bigcup_{g \in G} g.B = X$.

Proof. Let us first prove the backwards direction, which is the one we don't actually need for our goal of reducing the topological statement to the non-topological one. Given such a compact subspace B, we wish to show that $G \curvearrowright X$ is co-compact, i.e. $\frac{X}{G}$ is compact. Let $q: X \to \frac{X}{G}$ be the quotient map, and U_i an arbitrary open cover of $\frac{X}{G}$. As quotient maps are always continuous, the pre-images $q^{-1}(U_i)$ give an open cover of X and thus also of B. As B is compact, we get a finite open cover $q^{-1}(U_{i_j})$ of B. But the property of the $g.B$ covering X means that every element of $\frac{X}{G}$ has a pre-image in B, proving that the U_{i_j} are in fact a finite open cover of $\frac{X}{G}$.
It now remains to prove the forwards direction. We are given that the action is co-compact, and we must now conjure up a compact B that fulfills the condition. We claim the following construction works:

- Pick $x_i \in X$ as representatives for each orbit of the group action of G on X.
- By local compactness, pick a K_i as a compact neighbourhood of each x_i.
- Let K_{i_j} be a suitably chosen finite subset of the K_i such that $\bigcup_j q(K_{i_j})$ covers $\frac{X}{G}$.
- Set $B = \bigcup_j K_{i_j}$.

That B fulfills the property of $g.B$ covering all of X is clear by construction. Also, B is compact because it is the union of finitely many compact sets. It remains to show that the third step is possible, i.e. we can actually pick a finite subset of the K_i such that they cover $\frac{X}{G}$. To construct one, consider the $q(K_i^0)$, where A^0 designates the interior (the union of all open sets within) of A. These are all open sets as q is open by Lemma 7.5, and they cover $\frac{X}{G}$. By compactness we pick a finite subcover, and it is precisely these K_{i_j} which fulfill the required properties. □

7.1. Problems to think about

7.1. Show that every proper metric space is complete.

7.2. Show that the word metric turns a group into a proper metric space.

References

[1] `https://web.evanchen.cc/napkin.html`

[2] Löh, C. (2017). Geometric group theory, an introduction.